



ALMA MATER STUDIORUM  
UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN  
LINGUE, LETTERATURE E CULTURE MODERNE: DIVERSITÀ ED  
INCLUSIONE

Ciclo 38

**Settore Concorsuale:** 10/G1 - GLOTTOLOGIA E LINGUISTICA

**Settore Scientifico Disciplinare:** GLOT-01/A (ex L-LIN/01) - GLOTTOLOGIA E LINGUISTICA

LANGUAGE-MEDIATED ABSTRACTION: CONCEPTUAL CONCRETENESS AND  
CATEGORICAL SPECIFICITY IN LANGUAGE PROCESSING

**Presentata da:** Tommaso Lamarra

**Coordinatore Dottorato**

Serena Baiesi

**Supervisore**

Marianna Marcella Bolognesi

**Co-supervisore**

Anna Maria Borghi

## **ABSTRACT**

The present work investigates the relationship between conceptual concreteness and categorical specificity in accounting for language-mediated abstraction, advancing the claim that these variables constitute distinct semantic dimensions and examining their role in shaping language processing, conversational dynamics, and communicative strategies across three experimental studies. A key contribution of this work lies in adopting a methodology characterized by progressively increasing ecological validity, moving from traditional psycholinguistic paradigms to more interactive and communicative settings.

The first study examines the effects of conceptual concreteness and categorical specificity on word processing, both independently and in interaction, across lexical and semantic decision tasks. In addition to replicating the well-established processing advantage of concrete words over abstract words, the study reports a specificity effect: a behavioral advantage of specific words, which instantiate high-resolution representations, over general words in an abstract/concrete semantic decision task.

The second study extends the investigation to the perceived plausibility of simulated conversational excerpts. The study examines the role of concreteness and specificity in conversational dynamics, focusing on two linguistic constructs: uncertainty and curiosity (i.e., the need for clarification). The results show that abstract sentences were more likely to elicit expressions of uncertainty and curiosity. However, expressions signaling the need for clarification were also judged as plausible following concrete-general sentences, suggesting their referential indeterminacy.

The third study analyzes semi-naturalistic interactions in which interlocutors describe objects under varying conditions of referential availability (present vs. absent) and shared knowledge (known vs. unknown). The study assesses whether interlocutors systematically modulate the concreteness and specificity of their language under these experimental conditions. The main findings indicate that specificity functions as a linguistic cue for maximizing pragmatic informativeness.

Together, these studies support the perspective that categorical specificity plays a central role in understanding how language-mediated abstraction shapes both language processing and communicative behavior.

## RIASSUNTO

Il presente lavoro indaga la relazione tra la concretezza concettuale e la specificità categoriale nel descrivere l'astrazione mediata linguisticamente, sostenendo che tali variabili costituiscono dimensioni semantiche distinte ed esaminandone il ruolo nel modellare l'elaborazione linguistica, le dinamiche conversazionali e le strategie comunicative attraverso tre studi sperimentali. Un contributo centrale di questo lavoro risiede nell'adozione di una metodologia caratterizzata da una progressiva crescita della validità ecologica, che si sviluppa a partire da paradigmi psicolinguistici tradizionali verso contesti comunicativi ed interattivi.

Il primo studio esamina gli effetti della concretezza concettuale e della specificità categoriale sull'elaborazione delle parole attraverso compiti di decisione lessicale e semantica. Oltre a replicare il ben noto vantaggio di elaborazione delle parole concrete rispetto a quelle astratte, lo studio riporta un effetto di specificità: un vantaggio comportamentale delle parole specifiche, che istanziano rappresentazioni ad alta risoluzione, rispetto alle parole generali in un compito di decisione semantica.

Il secondo studio estende l'indagine alla plausibilità percepita di estratti conversazionali simulati. In particolare, lo studio esamina il ruolo della concretezza e della specificità nelle dinamiche conversazionali, concentrandosi su due costrutti linguistici: l'incertezza e la curiosità (ossia la necessità di chiarificazione). I risultati mostrano che le frasi astratte tendono a suscitare maggiormente espressioni di incertezza e curiosità. Tuttavia, anche le espressioni che segnalano la necessità di chiarificazione sono state giudicate plausibili come risposte a frasi concrete ma generali, suggerendo una loro indeterminatezza referenziale.

Il terzo studio analizza le interazioni seminaturalistiche in cui gli interlocutori descrivono oggetti in condizioni variabili di disponibilità referenziale (presente vs. assente) e di conoscenza condivisa (nota vs. ignota). Lo studio valuta se gli interlocutori modulino sistematicamente la concretezza e la specificità del loro linguaggio in funzione di tali condizioni sperimentali. I principali risultati indicano che la specificità funziona come una strategia linguistica volta a massimizzare la quantità di informazione pragmaticamente rilevante veicolata.

Nel loro insieme, questi studi supportano la prospettiva secondo cui la specificità categoriale svolge un ruolo centrale nella comprensione di come l'astrazione mediata

linguisticamente influenzi sia l'elaborazione linguistica sia il comportamento comunicativo.



# **LANGUAGE-MEDIATED ABSTRACTION**

Conceptual Concreteness and Categorical Specificity  
in Language Processing

TOMMASO LAMARRA

University of Bologna



**ABSTRACTION**

[www.abstractionproject.eu](http://www.abstractionproject.eu)

**Grant agreement**

**ERC-2021-STG-101039777**



**European Research Council**  
Established by the European Commission

Funded by the European Union (GRANT AGREEMENT: ERC-2021-STG-101039777). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

A Mari, con amore  
(To Mari, with love)

## **ACKNOWLEDGMENTS**

I am sincerely happy to write these acknowledgments and to dedicate them to the extraordinary people without whom this PhD journey would not have been possible.

I am extremely grateful to my supervisor and inspiration, Marianna Marcella Bolognesi. It is thanks to her constant encouragement, personal empathy, patience, and guidance that I am becoming the researcher that I aspire to be. Her enthusiasm and determination in her work are truly contagious and inspiring. She mentored me over the past three years as supervisor and P.I., demonstrating how meaningful and rewarding scientific research can be. I am also sincerely grateful for all the meetings, lunch breaks, and opportunities to discuss topics of mutual interest in language and cognition, as well as for her feedback, which helped highlight the strengths of this dissertation and improved it through her suggested revisions.

I am also deeply grateful to my co-supervisor Anna Maria Borghi, who welcomed me into the Ballab, her Lab, after my graduation and introduced me to scientific research. Her work and enthusiasm have been a great source of inspiration. I am sincerely thankful for the many interesting conversations and stimulating discussions, as well as for her insightful feedback and suggestions that improved this dissertation. Together with Anna, I would like to thank people in her lab, especially Claudia Mazzuca.

I would like to express my sincere gratitude to my external advisor, Gabriella Vigliocco, who invited me to join her Language and Cognition Lab at University College London. I had the opportunity to spend three months under her supervision. I greatly appreciated her meticulous attention to detail in both conducting and presenting experimental research. Her dedication and her work are truly inspiring. I am grateful for the encouragement and the many insightful discussions we shared, which have contributed to parts of this dissertation. Many thanks also to the people in her lab for stimulating discussions during Lab meetings.

Together, Marianna, Anna, and Gabriella taught me the importance of interdisciplinary research.

My sincere gratitude also goes to my colleagues of the ABSTRACTION research group. All of them are inspiring researchers. First, I would like to express my profound gratitude to Caterina Villani and Giulia Rambelli. As postdoctoral researchers in the group, they both supported me and helped me grow as a researcher, while also becoming dear friends. I am especially grateful to Caterina for her patience and empathy, and for her continuous encouragement throughout my PhD research. I am equally grateful to Giulia for her kindness and for becoming a close friend; whenever I had doubts, I knew I could rely on her.

I would also like to thank Claudia Collacciani for her spontaneity and for welcoming me to Bologna; Andrea Amelio Ravelli for the time spent working on iconicity ratings; Adele Loia for her sincerity and for the opportunities to exchange advice and perspectives; and Paola Vernillo and Laura Occhipinti for stimulating discussions on topics of mutual interest.

A special thanks goes to all the researchers with whom I collaborated, as well as to my PhD colleagues.

I am grateful to Professors Valentina Cuccio and Filippo Domaneschi for having the time to read and review my dissertation, as well as to Professors Caterina Mauri, Diego Frassinelli, and Francesca Strik Lievers for agreeing to be part of the final examination committee.

I am extremely grateful to all my friends, who have been a source of support, laughter, joy, and understanding. In particular, I would like to thank friends from the Ecobrace group for all the time we spent together. I especially thank Giulio and Bea for being part of my life for as long as I can remember, and Bea and Xiaoran for hosting me in London.

I would also like to mention my friend and colleague Luca, for all the discussions on language we have shared since our university studies.

An immense thank you to my parents, Alessandra and Giovanni, for the sacrifices they made to open up opportunities for me, and for their endless support and love, as well as to my family, especially my uncle Antonio, and to the family I am about to join.

My deepest gratitude goes to Mari, my best friend and future wife, to whom this dissertation is dedicated, for her constant support, encouragement, and endless kindness

and love. She is the person who has always reminded me to believe in myself. This dissertation is also, in part, hers.

Finally, my gratitude also goes to the participants who voluntarily agreed to participate in the experiments presented in this dissertation, as well as to the Experimental Lab of the Department of Modern Languages, Literatures and Cultures of the University of Bologna, where the majority of the experiments were conducted.



# CONTENTS

|                                                                                                                                     |           |
|-------------------------------------------------------------------------------------------------------------------------------------|-----------|
| Abstract                                                                                                                            |           |
| Riassunto                                                                                                                           |           |
| Acknowledgments                                                                                                                     |           |
| List of Figures.....                                                                                                                | xv        |
| List of Tables .....                                                                                                                | xvii      |
| <b>Introduction .....</b>                                                                                                           | <b>19</b> |
| <b>Part 1: Theoretical Background.....</b>                                                                                          | <b>25</b> |
| <b>Chapter 1 .....</b>                                                                                                              | <b>27</b> |
| <b>The Pathways of Abstraction.....</b>                                                                                             | <b>27</b> |
| 1.1. A Problem of Definition: Abstraction, the Great Unknown .....                                                                  | 27        |
| 1.1.1. Abstraction in Philosophy: The Problem of Abstract Objects .....                                                             | 28        |
| 1.1.2. Abstraction in Psychology: Towards Conceptualization .....                                                                   | 29        |
| 1.1.3. Abstraction in Linguistics: Language-Mediated Simplification and Inferential Structures.....                                 | 30        |
| 1.1.4. The Inferential Dimension of Abstraction: Mapping Between Domains .....                                                      | 31        |
| 1.1.5. The Multiple Levels of Abstraction Across Disciplines .....                                                                  | 33        |
| 1.2. Narrowing the Concept of Abstraction: Three Definitions .....                                                                  | 34        |
| 1.2.1. Abstraction as Conceptualization .....                                                                                       | 35        |
| 1.2.2. Abstraction as Schematization.....                                                                                           | 38        |
| 1.2.3. Abstraction as Subjectification .....                                                                                        | 40        |
| 1.3. Experimental Approaches to Understand Abstraction.....                                                                         | 45        |
| 1.4. The Polysemous Nature of Abstraction: The Senses of Abstract .....                                                             | 49        |
| 1.4.1. Dynamic Abstraction in Perceptual Symbol Systems: Six Senses of Abstraction                                                  | 53        |
| 1.4.2. Disentangling the Polysemous Nature of Abstraction: Conceptual Concreteness, Abstractness, and Categorical Specificity ..... | 58        |
| <b>Chapter 2 .....</b>                                                                                                              | <b>63</b> |

|                                                                                                              |           |
|--------------------------------------------------------------------------------------------------------------|-----------|
| <b>Conceptual Concreteness and The Complexity of Abstractness .....</b>                                      | <b>63</b> |
| 2.1. What are Concepts? From a Definitional Structure of Concepts to Multidimensional Conceptual Spaces..... | 63        |
| 2.1.1 The Classical View .....                                                                               | 64        |
| 2.1.2. Integrated and Flexible Conceptual Representations: Frames and Theories of Domains.....               | 67        |
| 2.2. Abstract vs. Concrete Concepts .....                                                                    | 70        |
| 2.2.1. Varieties of Dimensions for Abstract and Concrete Concepts.....                                       | 73        |
| 2.3. The Concreteness Effect: Differences in Language Processing and Representation .....                    | 76        |
| 2.4. The Symbol Grounding Problem.....                                                                       | 80        |
| 2.4.1. Embodied vs. Distributional Meaning.....                                                              | 83        |
| 2.4.2. Hybrid Theories of Word Meaning .....                                                                 | 85        |
| 2.5. The Challenge of Grounding Abstract Concepts .....                                                      | 88        |
| 2.5.1. Conceptual Metaphor Theory (CMT).....                                                                 | 89        |
| 2.5.2. Affective Embodiment Account (AEA) .....                                                              | 90        |
| 2.5.3. Language and Situate Simulation (LASS) .....                                                          | 91        |
| 2.5.4. Words as Social Tools .....                                                                           | 92        |
| 2.6. The Social Dimension of Abstractness .....                                                              | 94        |
| <b>Chapter 3 .....</b>                                                                                       | <b>98</b> |
| <b>The Neglected Variable: Categorical Specificity.....</b>                                                  | <b>98</b> |
| 3.1. Taxonomic Relations in Linguistics and Psycholinguistics .....                                          | 98        |
| 3.2. Taxonomic Structure of Conceptual Categories: Superordinate, Basic, Subordinate Levels .....            | 101       |
| 3.2.1 Basic Level Categories .....                                                                           | 102       |
| 3.2.2. Superordinate Categories.....                                                                         | 106       |
| 3.2.3. Subordinate Categories and Expertise.....                                                             | 107       |
| 3.3. Theories of Conceptual Categorization: The Prototype Theory .....                                       | 108       |
| 3.3.1. Exemplar Theory .....                                                                                 | 109       |
| 3.3.2. Theory Theory .....                                                                                   | 110       |
| 3.4. Disentangling Conceptual Concreteness and Categorical Specificity.....                                  | 112       |

|                                                                                                                    |            |
|--------------------------------------------------------------------------------------------------------------------|------------|
| 3.4.1. Human Generated Specificity Ratings and Relations with Other Semantic Variables .....                       | 115        |
| 3.5. Exploring the Impact of Abstraction and Categorical Specificity in Language: Recent Empirical Evidence.....   | 118        |
| 3.5.1. First Steps in Deepening the Relationship Between Conceptual Concreteness and Categorical Specificity ..... | 124        |
| <b>Methodological and Theoretical Rationale for the Experimental Section .....</b>                                 | <b>131</b> |
| From Representations to Conversations .....                                                                        | 131        |
| Overview of Experimental Section .....                                                                             | 135        |
| <b>Part 2: Experimental Section: Case Studies .....</b>                                                            | <b>142</b> |
| <b>Chapter 4 .....</b>                                                                                             | <b>144</b> |
| <b>Concreteness and Specificity as Isolated Words .....</b>                                                        | <b>144</b> |
| 4.1. Methodology.....                                                                                              | 145        |
| 4.2. Data Analyses and Results.....                                                                                | 149        |
| 4.3. Discussion.....                                                                                               | 156        |
| <b>Chapter 5 .....</b>                                                                                             | <b>160</b> |
| <b>Concreteness and Specificity in the Plausibility of Simulated Dialogues .....</b>                               | <b>160</b> |
| 5.1. Norming Studies.....                                                                                          | 161        |
| 5.2. Methodology.....                                                                                              | 164        |
| 5.3. Data Analyses and Results.....                                                                                | 167        |
| 5.4. Discussion.....                                                                                               | 174        |
| <b>Chapter 6 .....</b>                                                                                             | <b>177</b> |
| <b>Concreteness and Specificity in Naturalistic Conversations .....</b>                                            | <b>177</b> |
| 6.1. Methodology.....                                                                                              | 178        |
| 6.2. Data Analyses and Results.....                                                                                | 181        |
| 6.3. Discussion.....                                                                                               | 191        |
| <b>Part 3: Discussion .....</b>                                                                                    | <b>196</b> |
| <b>Chapter 7 .....</b>                                                                                             | <b>198</b> |
| <b>Towards a Comprehension Theory of Language-Mediated Abstraction .....</b>                                       | <b>198</b> |
| 7.1. Specificity Effect on Semantic Processing.....                                                                | 198        |

|                                                                          |            |
|--------------------------------------------------------------------------|------------|
| 7.2. The Pragmatic Ambiguity of General Sentences.....                   | 203        |
| 7.3. Specificity as a Linguistic Cue in Naturalistic Conversations ..... | 210        |
| 7.4. Conclusions .....                                                   | 215        |
| <b>References.....</b>                                                   | <b>218</b> |
| Appendix A.....                                                          | 246        |
| Appendix B.....                                                          | 253        |

## LIST OF FIGURES

|                                                                                                                                                                                          |     |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| <b>Figure 3.1.</b> Branching tree diagram schematizing the class inclusion relation of hyponymy/hypernymy of the word “animal”.....                                                      | 99  |
| <b>Figure 3.2.</b> The intersection of concreteness and specificity, resulting in four distinct quadrants: abstract-general, abstract-specific, concrete-general, concrete-specific..... | 115 |
| <b>Figure 4.1.</b> Overview of the experimental design and methodological structure of the three experiments described in this chapter.....                                              | 145 |
| <b>Figure 4.2.</b> Interaction effect between concreteness and specificity variables in predicting RTs in the Lexical Decision Task.....                                                 | 152 |
| <b>Figure 4.3.</b> Interaction effect between concreteness and specificity variables in predicting RTs in the Semantic Decision Task.....                                                | 155 |
| <b>Figure 5.1.</b> Overview of the experimental design and methodology of the study described in this chapter.....                                                                       | 161 |
| <b>Figure 5.2.</b> Schematic representation of the experimental procedure.....                                                                                                           | 165 |
| <b>Figure 5.3.</b> Predicted probabilities of plausibility judgments across all sentence types, as a function of the follow-ups.....                                                     | 169 |
| <b>Figure 5.4.</b> Predicted probabilities of plausibility judgments across all sentence types, as a function of the follow-ups.....                                                     | 171 |
| <b>Figure 5.5.</b> Rating scores for general and specific sentences across Confidence, Experts Need, and Social Metacognition dimensions.....                                            | 173 |
| <b>Figure 6.1.</b> Overview of the experimental design and methodological approach of the three analyses described in this chapter.....                                                  | 178 |
| <b>Figure 6.2.</b> The modulation of noun concreteness across the experimental conditions..                                                                                              | 183 |
| <b>Figure 6.3.</b> The modulation of noun specificity across the experimental conditions....                                                                                             | 184 |
| <b>Figure 6.4.</b> The interaction of <i>Session</i> and <i>Known</i> condition in the sentence concreteness.....                                                                        | 185 |
| <b>Figure 6.5.</b> The interaction of <i>Session</i> and <i>Known</i> condition in the sentence specificity.....                                                                         | 186 |
| <b>Figure 6.6.</b> Bubble plot indicating the distribution of occurrences of conversational strategies associated with specific linguistic variables across conditions.....              | 189 |

**Figure 6.7.** Forest plots of standardized fixed-effect estimates showing the magnitude and statistical significance of each predictor’s effect on individual noun concreteness (left plot) and sentence concreteness (right plot), while controlling for all other variables in the model.....190

**Figure 6.8.** Forest plots of standardized fixed-effect estimates showing the magnitude and statistical significance of each predictor’s effect on individual noun specificity (left plot) and sentence specificity (right plot), while controlling for all other variables in the model.....191

## LIST OF TABLES

|                                                                                                                                                                            |     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| <b>Table 1.1.</b> Comparative characteristics of Peircean logical operations in terms of breadth, depth, and information change, adapted from Peirce (1931-1958: 2.422)... | 43  |
| <b>Table 2.1.</b> Summary of dimensions and semantic variables for which concrete and abstract concepts might differ.....                                                  | 76  |
| <b>Table 3.1.</b> Differences among concepts at varying levels of abstraction in terms of informativeness and distinctiveness.....                                         | 105 |
| <b>Table 3.2.</b> Overview of semantic, lexical, and behavioral variables for which general and specific words may differ.....                                             | 128 |
| <b>Table 4.1.</b> Means and standard deviations for concreteness, imageability, and specificity ratings across the four word types.....                                    | 150 |
| <b>Table 4.2.</b> Summary of linear mixed models, each incorporating additional fixed effects in predicting reaction times.....                                            | 151 |
| <b>Table 4.3.</b> Lexical Decision Task: model comparison metrics with concreteness added first.....                                                                       | 153 |
| <b>Table 4.4.</b> Lexical Decision Task: model comparison metrics with specificity added first.....                                                                        | 153 |
| <b>Table 4.5.</b> Semantic Decision Task: model comparison metrics with concreteness added first.....                                                                      | 156 |
| <b>Table 4.6.</b> Semantic Decision Task: model comparison metrics with specificity added first.....                                                                       | 156 |
| <b>Table 5.1.</b> Means and Standard Deviations for Confidence, Social Metacognition, Experts Needs ratings across the two types of sentences.....                         | 173 |
| <b>Table 5.2.</b> Means and Standard Deviations for Confidence, Social Metacognition, Experts Needs ratings across the four types of sentences.....                        | 174 |
| <b>Table 6.1.</b> Odds Ratios values and <i>p</i> -values indicating significance of the effect size.....                                                                  | 187 |
| <b>Table A.1.</b> Full list of Italian words varying in concreteness and specificity used in the behavioral tasks presented in Chapter 4.....                              | 249 |
| <b>Table A.2.</b> Full list of nonwords used in the behavioral tasks presented in Chapter 4....                                                                            | 252 |
| <b>Table B.1.</b> Full list of Italian target sentences varying in concreteness and specificity used in the behavioral tasks presented in Chapter 5.....                   | 255 |

|                   |                                                                                                                                                                                                                                  |     |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| <b>Table B.2.</b> | Full list of Italian follow-up responses expressing uncertainty, certainty, need for additional information (i.e., curiosity), and willingness to end the conversation, used in the behavioral tasks presented in Chapter 5..... | 255 |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|

## INTRODUCTION

This work is organized into three main parts, with a methodological section placed between Parts 1 and 2. Part 1 provides the theoretical framework that underpins the experimental research presented in Part 2. Part 1 reviews key portions of the scientific literature on abstraction, conceptual representations, and taxonomic relations, emphasizing three case studies in which concreteness and specificity are examined: first, as isolated words, second, as embedded in conversational contexts, and third, in semi-naturalistic conversations. This progression illustrates an approach in which the ecological validity of the methodologies employed progressively increases.

Chapter 1 offers a cross-disciplinary review of the abstraction process, examining the primary traits of this notion and highlighting its resistance to a single, univocal definition. Different disciplines propose distinct traits of abstraction as potential definitions. As the chapter unfolds, it becomes apparent that two main senses of abstraction are often conflated: a conceptual and linguistic sense. These two senses form the foundation for the central variables of this work: conceptual concreteness and categorical specificity. According to recent proposals, these two variables, particularly specificity, are crucial for investigating language-mediated abstraction. The chapter concludes by outlining the rationale for studying these variables.

Chapter 2 focuses on conceptual concreteness and abstractness, summarizing both theoretical and empirical research on conceptual knowledge. It contrasts the classical view of concepts as fixed lists of features with contemporary accounts that portray concepts as multidimensional, dynamic, and flexible constructions. This multidimensionality is especially relevant in the study of abstract and concrete concepts, which are primarily characterized by different dimensions and types of experience (e.g., sensorimotor and linguistic). The chapter discusses the well-established concreteness effects as a prime example of divergence and offers a review of empirically supported theories suggesting that abstract concepts, despite lacking precise reference, can be grounded in various types of experience, with a particular emphasis on the social dimension.

Chapter 3 explores how the lexicon is organized at different levels of specificity and granularity through taxonomic and class inclusion relations. Specificity is framed as a form of language-mediated abstraction, distinguishing between general, inclusive words (i.e., hypernyms) and more specific, less inclusive words (i.e., hyponyms). The chapter also examines how conceptual knowledge mirrors this taxonomic organization of the lexicon, identifying three levels of categorization: subordinate, superordinate, and basic. Additionally, the chapter reviews theories that explain how categories are represented, either as schematic summaries of the best traits of category members or as collections of exemplars. Recent research on operationalizing and manipulating levels of abstraction is also discussed. This chapter underscores the lack of systematic research on the relationship between conceptual concreteness and categorical specificity in communicative contexts, highlighting a gap in the current literature.

Between Parts 1 and 2, a methodological section advocates for a hybrid methodology, balancing valid settings that bridge laboratory studies and real-life conversations. The section also outlines the experimental portion of this work, illustrating the methodological and conceptual rationale as well as the theoretical claims and research questions that motivated the three case studies that I conducted over the course of this research project.

Part 2 consists of three experimental chapters, each presenting methodologies, results, and preliminary discussions.

Chapter 4 presents the first case, in which concreteness and specificity are investigated as isolated words using traditional psycholinguistic paradigms (e.g., ratings, lexical decision task, and semantic decision task). The study replicated the well-established concreteness effect and introduced a specificity effect in the semantic decision task, showing that specific words, which instantiate high-resolution semantic representations, have an advantage in language processing, particularly in the semantic decision task.

Chapter 5 examines the role of concreteness and specificity in affecting conversational dynamics, embedding these variables in simulated conversational contexts. The findings replicated and expanded upon previous research, demonstrating that generality, like abstractness, exhibits a form of referential indeterminacy, which can trigger requests for clarification.

Chapter 6 discusses methodology, procedures, and results of a study in which concreteness and specificity are examined in the language produced during semi-naturalistic conversations. The study consisted of an automatic POS-tagging annotation and included three distinct analyses. The results suggest that concreteness and specificity can be understood as linguistic cues supporting the acquisition of knowledge. In particular, the informativeness of specificity can be modulated as a function of communicative strategies.

Finally, within the third and concluding part of this work, Chapter 7 provides a comprehensive discussion of the findings from the three experimental studies, synthesizing the results and emphasizing the contribution of specificity as a key measure of language-mediated abstraction. It argues that specificity conveys levels of informativeness that enhance word processing and contribute to the success of conversations.



“The difference between the almost right word and the right word is really a large matter. ’tis the difference between the lightning bug and the lightning.”

(Mark Twain)

“I like good strong words that mean something...”

(Louisa May Alcott)

“Words are alive, a wondrous euphony that lets our souls resonate.”



**PART 1**  
**THEORETICAL BACKGROUND**



# Chapter 1

## The Pathways of Abstraction

### 1.1. A Problem of Definition: Abstraction, the Great Unknown

A useful starting point for understanding the term “abstraction” is its Latin etymology. The word “abstraction” derives from the Latin verb *abstrahere*, composed of the prefix *ab*, meaning “out of”, and the root verb *trahere*, meaning “to pull”. Thus, the literal meaning of this term suggests an operation of pulling something out from something else.

Nonetheless, contemporary dictionaries provide a range of definitions. For example, the *Cambridge Dictionary* lists four meanings of abstraction (“Abstraction”):

1. The quality of existing as or representing an idea, a feeling, etc., and not a material object, or something that has this quality.
2. The quality of being very general and not based on real examples, or something that has this quality.
3. The action of removing or separating something from a place or context (the situation, facts, words, etc. that exist around something), or something that is removed like this.
4. An idea that develops by looking at or thinking about a number of different things.

Definition (3) aligns with its etymological meaning, describing a process of extracting something from individual objects or events. Definitions (1) and (2), however, suggest that abstraction can now also be conceived as two distinct qualities: existing as something immaterial, such as an abstract idea like “humanism” or a feeling like “freedom” (1); and expressing something detached from specific cases or concrete examples, as in a generalized argument lacking precision (2). The fourth definition (4) suggests a general idea encompassing multiple instances, such as the conceptual category represented by the word “animal”.

This heterogeneity, even in everyday usage, highlights the difficulties inherent in defining and dealing with this notion.

Abstraction is a central cognitive ability (Barsalou, 2003, 2005a; Burgoon et al., 2013), which helps explain why it has attracted scholarly attention across periods and scientific disciplines. Despite its significance, even in contemporary debates, the notion resists a stable, universally accepted definition, thereby varying across fields and theoretical perspectives.

This chapter undertakes a diachronic and interdisciplinary review of abstraction, tracing its conceptual evolution, identifying its key dimensions, examining experimental paradigms in which it has been studied, and clarifying its definitional scope as well as its practical and theoretical applications.

### **1.1.1. Abstraction in Philosophy: The Problem of Abstract Objects**

A primary challenge in interpreting the notion of abstraction already arises in ancient thought, where the definition of abstraction as a process was often conflated with what might be designated as “abstract object”, that is, the result of the process itself. In ancient philosophy, for instance, abstraction was considered an operation belonging to the domain of ontology. In some cases, it was identified as transcendental entities, such as Platonic ideas and mathematical objects, or immanent universal principles. These abstractions had their own ontological status and were conceived as objects *abstracted* from the chaotic reality of the sensible world.

In Aristotelian philosophy, abstraction was compared to the logical method of *subtraction* (Cleary, 1985; Bäck, 2014), namely, an objective and not subjective process through which the individual substances are isolated by subtracting, separating, and removing the proper attributes and types from the totality of sensible reality. Bäck (2014) proposed that the idea of *selective attention* better captures the concept of abstraction for Aristotle, as it emphasizes the importance of selecting relevant information (for a more detailed discussion on this topic, see Bäck, 2014). Nonetheless, we can observe that the process of abstraction fundamentally involves an operation of extracting specific elements to identify the abstracted or abstract objects.

In the later empiricist philosophy, the notion of abstraction similarly denoted an operation of identifying and selecting particular information; however, it was conceived as a psychological and epistemological process that produces general ideas and terms

from the multitude of concrete and individual experiences. These abstractions were likewise considered abstract objects that, nevertheless, do not correspond to the real structures of the sensible world, but rather consist of mental constructions that stand as signs for individual qualities or traits (for more details on the “abstract objects” topic, see Falguera et al., 2025). Thus, the different nature of the products of the process determines the nature of the process itself, which now becomes a cognitive and mental faculty aimed at defining general terms from single individual features potentially applicable to a multitude of cases.

So far, abstraction has been understood as an essential mental capacity that enables the organization of knowledge of the world. Notably, within psychology and linguistics, abstraction consistently retained the status of a cognitive faculty. However, scholarly interest in this notion gradually shifted from the nature of abstract objects to the clarification of the characteristics and potentialities of abstract thought, conceived as an advanced stage of cognitive and linguistic development.

### **1.1.2. Abstraction in Psychology: Towards Conceptualization**

In philosophy, abstractive processes were understood as being fundamentally related to the logical description of the world. In contrast, in psychology, emphasis is placed on the role of high-level abilities in cognition and, in some cases, in language in supporting the capacity for abstraction.

In his account of children’s development from concrete to abstract thought, Piaget distinguished two different types of abstraction: *empirical abstraction* and *reflective abstraction*. Empirical abstraction involves extracting information, qualities, and regularities from objects, leading to the formation of conceptual knowledge, a process akin to that emphasized by empiricists. Reflective abstraction, by contrast, underlies logical and mathematical knowledge and entails a transition to a higher level of reflection and reorganization of knowledge, through which novel structures of reality are discovered and brought to awareness — for instance, the ability to organize a sequence of actions in a causal order (Piaget, 1970, 1997).

Similarly, Dewey claimed that the value of abstraction lies not only in its discriminative power but also in its capacity to transcend immediate contexts and generate

new instruments for epistemological inquiry, such as words, concepts, and symbols, by revealing and establishing relations among objects, situations, and processes in the world (Dewey, 1910, 1938). Importantly, although closely related, Dewey distinguished abstraction and generalization, since the latter reflects the application of more general regularities and principles to novel individual cases (for further details on differences between abstraction and generalization, see ¶ 1.4.2.).

The psychologist Vygotsky emphasized the mediational role of language in supporting the capacity for abstraction, which underlies the full development of *conceptual thinking*, in Vygotsky's terminology. In particular, children's development involves a gradual refinement of their ability to abstract (Van Oers, 2001), that is, the ability to detect and retain relevant traits of concrete objects or contexts necessary for forming a genuine concept. A genuine concept, Vygotsky explained, represents a new formulation based on an abstracted synthesis of salient features. The *mastery* of abstraction processes marks the transition from *complex thinking* (i.e., the pre-conceptual stage in which children establish relations between objects but cannot yet form actual categories) to *conceptual thinking*.

Language and words play a crucial role in supporting conceptualization through abstraction. Through their symbolic nature, words enable children to internalize and master the information synthesized in abstracted concepts and to apply it across multiple contexts and situations. Indeed, the coordinated action of abstraction and language supports analytical thinking and generalization. In Vygotskian terms, the essence of word meaning corresponds to a "*generalized reflection of reality*" (Vygotsky, 1934/1986). Therefore, the use of language as a system of abstract and general signs contributes to organizing knowledge about the world in terms of general conceptual categories rather than isolated experiences.

### **1.1.3. Abstraction in Linguistics: Language-Mediated Simplification and Inferential Structures**

This cognitive and psychological perspective aligns with linguistic theoretical accounts that consider abstraction as fundamentally mediated by language and supporting inferential processes. For example, the linguistic anthropologist Enfield stressed the role of abstraction and language in simplifying the complexity of the world. As he argues:

We shall see that the two processes of abstraction — from reality to perception and from perception to language — serve two different functions. While perception reduces reality to make it easier to navigate the physical world, word meanings further reduce complexity to make it easier to navigate the *social* world. (Enfield, 2024, p. 5)

The perceptual system inherently reduces and abstracts the complexity of the world, since a cognitive system cannot process all information and must select relevant traits. Moreover, abstraction operates along with language to unify the multiplicity of different objects, events, and experiences by constructing generalized categories and shared signs among members of a linguistic community. This process enables social communication, cooperation, and coordinated interaction.

Understanding how abstraction simplifies the complexity of reality has been a major topic in linguistic research (a more detailed discussion is provided in ¶ 1.2.2.). Another relevant topic related to abstraction in linguistics is inferential processes. In particular, Jackendoff highlighted that abstract thought is structured over concrete conceptual knowledge through inferential mechanisms. He claimed that abstract objects, such as concepts associated with abstract words like “confidence” or more general words like “object”, are *conceptual structures* not based on *perceptual features*, but rather on *inferential features* (Jackendoff, 2003). In particular, these concepts are not described by a whole set of features directly grounded on perceptual experience (i.e., <round>, <yellow>, <heavy>). Instead, they are built over an inferential system composed of a collection of inferences that, in turn, potentially activate perceptual features by connecting the abstract object to more concrete objects (Jackendoff, 2003; Fortescue, 2017). For instance, in Jackendoff’s account, a concept like OBJECT is available in the conceptual system because it is understood as <something useful for a purpose> or <something that can be created or found>. Notably, Jackendoff thought that language and concepts interface in the same way.

Although Jackendoff did not provide an extensive account of abstraction and at times implicitly treated abstraction as based on schematic representations (see ¶ 1.2.2.), he nevertheless characterized abstraction, in the account provided above, as an inferential mechanism underlying the connection between concrete and abstract thought.

#### **1.1.4. The Inferential Dimension of Abstraction: Mapping Between Domains**

Expanding on the inferential nature of abstraction anticipated in the previous section, research on metaphors and analogies emphasized the relational character of abstraction and its role in mapping structures between one domain, typically concrete, and another, often abstract. In psychological research, the *Structure Mapping* account (Gentner, 1983; Gentner & Markman, 1997; Asmuth & Gentner, 2017) posits that comparative reasoning is based on relational structures, specifically, on the mapping between a base domain, i.e., the source of knowledge, and a target domain, the one to be explained (Gentner, 1983). According to this account, literal similarities, analogies, and abstractions represent different types of domain comparisons organized along a continuum. The emergent hierarchy among these types of comparison may arise from the level of sophistication of the predicates mapped between domains. For instance, literal similarity involves a comparison in which both object predicates (e.g., <red>) and relational predicates (e.g., <turn around>) are mapped from the base to the target. In contrast, analogies and abstractions involve mappings primarily based on relational predicates, while object predicates progressively became less relevant or left unmapped.

Analogies and abstractions (or *general* laws) differ because one of the arguments of an abstraction comparison is an abstract lexical entity, such as “attractive force”, rather than a word denoting a concrete object like “vacuum cleaner”. Both types of comparison map relational properties from the base domain to the target domain; however, in the case of abstractions, the base itself is a more abstract term.

Asmuth and Gentner (Asmuth & Gentner, 2017; Gentner & Asmuth, 2017) proposed that abstraction is a key process in conceptualization because of its relational nature. From this perspective, abstract concepts (e.g., CONNECTION) can be understood by figuratively extending the meaning and the conceptual representation of an existing concrete object (e.g., BRIDGE), suppressing immediate perceptual features, and selecting relevant relational properties (see Chapter 2 for further discussion). Nevertheless, this account appears to conceive abstractions as sophisticated comparisons grounded in analogical thinking and metaphorical extensions (Bowdle & Gentner, 2005; Wolff & Gentner, 2011; Asmuth & Gentner, 2017).

This close relation between metaphorical mapping and abstraction is consistent with broader evidence from cognitive linguistics and cognitive science. According to the *Conceptual Metaphor Theory*, abstract knowledge is grounded in concrete and

sensorimotor experiences (Lakoff & Johnson, 1980, 1999; Gibbs Jr, 2008). Metaphors are complex cognitive devices structured upon the sensorimotor system due to neural interconnectivity (Bolognesi, 2020). Consequently, the brain's architecture enables metaphorical extensions to develop from body-based, pre-linguistic, and generalized schematic representations derived from sensory and perceptual experience, that is, *image-schemas* (see Johnson, 1987). Within this framework, conceptual metaphors constitute a central mechanism for abstract conceptualization (Johnson & Tucker, 2021). The pervasiveness of metaphors in thought and language supports systematic mapping from a concrete source to an abstract target domain (for a detailed discussion, see Chapter 2). Thus, metaphorical extension can be viewed as an embodied process of abstraction, transferring information from concrete to abstract knowledge.

### **1.1.5. The Multiple Levels of Abstraction Across Disciplines**

An additional crucial dimension of abstraction concerns the possibility of organizing information at different levels of analysis. Starting from concrete experience, it is possible to select and extract features or elements to construct an abstraction at a specific level of analysis. This abstraction can serve as the basis for further abstractions at higher levels. From this perspective, abstraction can be considered a dynamic and creative process of constructing new levels. This phenomenon is widely discussed across computer science, philosophy of information, psychology, and linguistics.

Linguistics accounts for this phenomenon through a specific type of relation between lexemes, namely *taxonomic relations*. Words can be organized into taxonomies that reflect different levels of abstraction. In particular, words may stand in *hyponymic* or *hypernymic* relations to one another. For instance, a word functions as a *hyponym* when it denotes specific objects or individuals in relation to another term (e.g., “snow leopard”; “weeping willow”), and as a *hypernym* when it denotes a broader class of referents (e.g., “animal”; “plant”; “organism”; Lyons, 1977). As in natural taxonomies, linguistic taxonomies are structured according to a logic of inclusion (see the extended discussion in Chapter 3): the word “animal” is more general than the lexical item “snow leopard” because it includes the referents denoted by “snow leopard”; “organism” is more general than both “animal” and “plant” because it encompasses all their referents.

Rosch and colleagues demonstrated that conceptual knowledge of natural categories is organized in similar taxonomic relations (Rosch, 1978; Rosch et al., 1976). They identified three main levels of categorization: the *superordinate*, *basic*, and *subordinate* levels. Superordinate categories are highly general, whereas subordinate categories refer to specific objects or entities. Basic categories represent an intermediate level in which the most common and frequent concepts are stored (DOG, BIRD, FLOWER, CAKE). This topic will be discussed in detail in Chapter 3.

In computer science and in philosophy of information, similar layered structures are often analyzed through the so-called method of *Levels of Abstraction* (LoAs; Floridi, 2008; Angius et al., 2025). Levels of Abstraction (LoAs) indicate the set of types of information considered appropriate and relevant at a specific level of analysis, that is, for a specific Level of Abstraction (LoA). Floridi (2008) explains that different features become salient depending on the context, for instance, when one is purchasing a wine or tasting it. Thus, different types of information become relevant at those specific levels of analysis. Furthermore, the quantity and nature of relevant information vary based on the LoA.

Levels of Abstractions (LoAs) are closely connected to two other complementary notions: *Levels of Organization* (LoOs) and *Levels of Explanation* (LoEs). Levels of Organization (LoOs) refer to the structural organization — often, though not necessarily, hierarchical — of the system under analysis, much like taxonomies of conceptual and linguistic categories. Levels of Explanation (LoEs) denote the different explanatory perspectives from which a phenomenon can be described. For instance, the same computational function can be described from the programmer’s point of view or from the perspective of the software’s execution (see Floridi, 2008, and Angius et al., 2025, for a detailed analysis). Intuitively, these three kinds of levels are interrelated, since the relevant information (Level of Abstraction) depends on the organization of the system (Level of Organization) and on the explanatory perspective and goals adopted (Level of Explanation).

## **1.2. Narrowing the Concept of Abstraction: Three Definitions**

This extensive review accounts for the heterogeneity of the notion of abstraction, reflecting its complex multidimensional nature and the shifting emphasis placed on

different aspects across research domains. Overall, abstraction enables us to transcend immediate, concrete, and specific contexts by establishing relations among entities and achieving a deeper understanding of our experience in the world. It can be characterized as a sophisticated process of selecting relevant feature information. Cognitive and linguistic processes operate in parallel with the mechanism of abstraction, supporting the formation of conceptual knowledge and linguistic structures. By generalizing from concrete experiences, we form broader and more general categories and concepts that encompass multiple specific instances, structuring this knowledge at varying levels of abstraction and degrees of detail. Furthermore, the inferential dimension of abstraction enables the mapping of concrete and abstract domains through analogical reasoning and figurative extension. Language plays a crucial role in these processes, providing symbols and structures that facilitate generalization, simplification, and organization of information.

To make sense of abstraction's complex, multidimensional nature, its key dimensions can be organized into three overarching themes that capture the most crucial conceptual and functional aspects of the notion:

- The categorization and formation of concepts by creating single conceptual structures that encompass a variety of individuals.
- The organization of conceptual knowledge and lexical elements in schematic representations at varying levels of abstraction, specificity, and precision.
- The idea of abstraction as a procedure of feature selection capable of generating abstract entities, conceived as mental objects that do not exist in the material world.

In what follows, I present three distinct definitions and accounts of abstraction, each highlighting one of these themes:

- Abstraction as Conceptualization.
- Abstraction as Schematization.
- Abstraction as Subjectification.

### **1.2.1. Abstraction as Conceptualization**

As noted above, abstraction is a process that proceeds in parallel with and supports the formation of conceptual knowledge, that is, *conceptualization*. However, in recent

discussions, abstraction has ceased to be conceived as a mental mechanism connected to conceptualization. Instead, these two notions have been conflated and treated as equivalent (Johnson & Tucker, 2021).

From a naturalistic philosophical account grounded in neural evidence, abstraction is a cognitive ability that enables one to manage the complexity of information in the world. Johnson and Tucker (2021) write:

We have been arguing that the process of conceptualization *is* the process of abstraction. A concept captures the meaning of many specific experiences in a way that generalizes to new situations. This generalization requires abstraction, so that the concept is not tied to specific, concrete experiences. The abstractive process involves recognizing individual features that nevertheless remain integrated into a general recurring pattern of experience. (p. 260)

According to this naturalistic account, abstraction enables us to move beyond immediate, concrete, and specific cases, allowing us to achieve a more integrated and complex level of information organization and representation. In this view, abstraction is a process that operates by balancing two complementary yet distinct mechanisms: *differentiation* and *integration*. Differentiation makes it possible to detect essential differences and distinctive features among exemplars, contexts, or situations. In contrast, integration enables one to identify commonalities, relations, and similarities among individuals and events. This dual process reflects both the transition from immediate and direct perceptual experiences to increasingly hierarchically structured knowledge and the complexity of neural networks that give rise to higher-level representations connecting different perceptual, affective, and cultural inputs recorded and processed in distinct neural areas.

Johnson and Tucker (2021) define a concept as a pattern of neural connectivity and a probabilistic distribution of expectancies formed in response to expected *affordances* (Gibson, 1950, 1979), dynamically established during interaction with the physical, cultural, linguistic, and social world. In this perspective, a genuinely abstract and general concept thus results from abstraction, through which a multitude of differentiated kinds of information stemming from distinct objects, contexts, and situations are retained in an integrated neural network in the human brain that preserves this conceptual complexity. In their words:

Abstraction is then a kind of unity amid variety, where we recognize difference and yet find patterns that hold what is differentiated together in a relatively unified manner. (Johnson & Tucker, 2021, p. 232)

Although the authors emphasize that meaning arises not only from linguistic meaning but also in all kinds of experience (perceptual, affective, cultural), language and words provide crucial support for this abstraction/conceptualization. At the early stages of their development, children, for instance, often use *holistic* and *syncretic* words, in which meanings are not yet clearly differentiated into stable conceptual categories and are not fully distinguished in terms of perceptual similarity, functional properties, and contextual use (Werner, 1957). A child may use the word “ball” indiscriminately to identify any round-shaped, bouncing object. As development progresses, however, children become aware that the world is more differentiated and that the word “ball” can refer to a variety of objects, each with its own individual features. They then come to recognize different exemplars of the more general and abstract concept BALL, such as:

- SOCCER BALL, which has both the <round> and <bouncing> features;
- RUGBY BALL, which partially has the <bouncing> feature, but not the <round> one;
- and even SNOWBALL, which has the <round> feature but typically does not bounce, and it is made of an atypical material (e.g., <made of snow>).

Therefore, the semantic scaffolding provided by words and their semiotic function enables us to use linguistic signs to convey the integrated yet differentiated nature of conceptual complexity.

From this perspective, it is important to highlight another essential aspect of abstraction emphasized by Johnson and Tucker (2021): its relationship to *metacognition* (i.e., the ability to think about one’s own or others’ thinking). Drawing on Piaget’s account of the transition from concrete to abstract thought (see Piaget, 1936, and also ¶ 1.1.2.), Johnson and Tucker (2021) stress that this developmental shift also marks a transition from the egocentrism characteristic of concrete thinking in childhood to the more reflective dimension of abstract and hypothetical thinking typical of adolescence. Although I will not examine this specific topic in detail, what is relevant for the present discussion is that abstraction develops in parallel with the emergence of self-awareness and through intersubjective engagement. Consequently, the complexity of the world abstracted from experience seems to correspond not only to relations between oneself and

the world, but also to relations involving oneself and others. Conceptualization, understood as abstraction, therefore entails a degree of complexity that encompasses perceptual, affective, social, interpersonal, and culturally grounded shared symbolic structures, most notably, those mediated by language.

### 1.2.2. Abstraction as Schematization

In linguistics, however, the relevance of abstraction lies in its function as a mechanism of simplification. In particular, the denotative extension of lexical elements enables speakers to generalize across different instances (see ¶ 1.2.3. for a precise differentiation between intension and extension). A single term can refer to multiple referents, as in the case of a general word like “reptiles”. Thus, words and other lexical elements function as simplifications produced through an abstraction process that groups a multitude of referents under a single term by setting aside irrelevant differences and retaining only the salient, common features. This mechanism also operates at another level, as when a single morpheme conventionally encodes the general property associated with a class of entities. A derivational suffix, for example, can shift a grammatical class into another, as in the addition *-ness* to an adjective to form a noun that denotes the abstract property expressed by the adjective (e.g., “ill-*ness*”; “happy-*ness*”; “rich-*ness*”).

In cognitive linguistics and more specifically within the framework of *Cognitive Grammar* (Langacker, 1986, 1987, 2008, 2013), this dimension of abstraction is equated to the process of *schematization*. Indeed, listing some basic abilities essential for language, Langacker (1999, p. 93) conceived schematization as a specific case of abstraction:

A second basic phenomenon, abstraction, is the emergence of a structure through reinforcement of the commonality inherent in multiple experiences. By its very nature, this abstractive process “filters out” those facets of the individual experiences which do not recur. We will mostly be concerned with a special case of abstraction, namely schematization, involving our capacity to operate at varying levels of “granularity” (or “resolution”). Structures that appear very different when examined in fine-grained detail may nonetheless be quite comparable in a coarse-grained view. A schema is the commonality that emerges from distinct structures when one abstracts away from their points of difference by portraying them with lesser precision and specificity.

Therefore, schematization can be understood as a form of abstraction through which a higher-level conceptual representation emerges by extracting commonalities

across individual *instances* (Evans & Green, 2006; Langacker, 2013). A *schema* is a representational construct situated at a general, superordinate level within conceptual knowledge, capturing only the features shared by more specific, subordinate instances (Fortescue, 2017) and encoding less detailed information than those instances (M. Bolognesi, 2025). Schematization abstracts away the variety of individual and idiosyncratic differences among specific instances and particular usage events, prioritizing relevant, inherent similarities. This mechanism is cognitive and embodied, contributing to the formation of grammar structures while being grounded in language usage.

The relationship between schemas and instances is of a type-token nature, in which a specific token represents a concrete, specific manifestation of a more general type. (Langacker, 1999) uses the terms *instantiation* or *specification* to describe the process by which a particular exemplar, such as “oak tree” or “poplar”, instantiates a corresponding extracted schema, such as “tree”. This relation can be formalized as  $A \rightarrow B$ , where A represents the schema (e.g., “furniture”) and B denotes the instance that specifies and elaborates the schema with greater detail (e.g., “lawn chair”).

It is evident that schematization involves, on the one hand, categorization, which is captured, as well, by the  $A \rightarrow B$  relation (Category A *applies to* Exemplar B, or Exemplar B *is compatible with* Category A), and, on the other hand, a hierarchical, layered structure organized according to levels of abstraction, as described in ¶ 1.1.5. A schema at a given level can itself function as an instance of a more general schema, forming a hierarchical regression that presumably culminates in the most general schema. The following relation illustrates the progression:

“living being”  $\rightarrow$  “pluricellular organisms”  $\rightarrow$  “plant”  $\rightarrow$  “tree”  $\rightarrow$  “oak tree”

Here, “tree” can instantiate the general schema “plant”, which in turn is comprehended in the even more general schema “pluricellular organism”, and so on.

This type of relation applies to any lexical units and mirrors the taxonomies existing in linguistics (see ¶ 1.1.5 and Chapter 3). However, Langacker refers to this phenomenon as a complementary notion, namely, *schematicity* (Langacker, 1999, 2013; Tuggy, 2010; Bolognesi, 2025). Consequently, lexical units and even conceptual structures are organized along a hierarchically layered continuum (Evans & Green, 2006), spanning from the highest degrees of schematicity to the highest degrees of specificity

(further details on *specificity* in ¶ 1.4.2. and extensively in Chapter 3). These two poles, schematicity and specificity, represent opposite ends of the same relationship.

Langacker (1987: 132) illustrates that “schematicity pertains to level of specificity, i.e. the fineness of detail with which something is characterized”. As such, linguistic and conceptual structures are organized at both lower and higher levels of detail (Tuggy, 2010). Schemas are constructs characterized by less information, lower precision, and slighter resolution compared to their instantiations, but they apply to a broader set of elements (Bolognesi, 2025).

Abstraction as schematization in linguistics is a cognitive process that requires refined abilities and may involve distinct cognitive mechanisms, such as categorization. However, unlike the naturalistic philosophical account, which views abstraction as a process through which we gradually become aware of the world’s complexity, abstraction as schematization in linguistics represents a simplification of complexity, driven by principles of cognitive economy and communicative efficiency. Similarly, as discussed above in relation to Enfield’s account (Enfield, 2024), such simplification constitutes a fundamental aspect of linguistic abstraction in response to communicative demands (see ¶ 1.1.3.). Both the naturalistic philosophical account and the linguistic perspective explain how abstraction supports the organization of information and complexity; however, cognitive linguistics and Cognitive Grammar place greater emphasis on the process of extraction and selection of relevant, shared features.

### **1.2.3. Abstraction as Subjectification**

Thus far, abstraction can be understood as a form of conceptualization, integrating multiple features while preserving their distinctions into a single conceptual structure that emerges from dynamic neural connections. Abstraction can also refer to the linguistic process of schematization, in which similarities between instances are extracted to generate a more general linguistic unit representing those instances at a higher level of abstraction with reduced precision and resolution. In this view, abstraction denotes the ability to incorporate information in linguistic elements at different hierarchical levels along the schematicity-specificity continuum.

In addition, abstraction may also be defined as a process of subjectification, i.e., the ability inherent to mathematical thinking and operations to focus on a specific quality or feature of an object, while disregarding any idiosyncratic details, thereby treating it as an abstract object (e.g., conceiving the quality *whiteness* as an abstract concept).

The semiologist and logician Peirce narrowed the definition of abstraction in this sense, beginning by distinguishing it from other mental, logical, and epistemological operations. He states:

**Generalization** is an increase of breadth and a decrease of depth, without change of information. **Induction** is a certain increase of breadth without a change of depth, by an increase of believed information. **Abstraction** is a decrease of depth without any change of breadth, by a decrease of conceived information. **Specification** is commonly used (I should say unfortunately) for an increase of depth without any change of breadth, by an increase of asserted information. [...] **Determination**, for any increase of depth. **Restriction**, for any decrease of breadth; but more particularly without change of depth, by a supposed decrease of information. **Descent**, for a decrease of breadth and increase of depth, without change of information. (Peirce, 1931-1958: 2.422; note: the term Supposition is omitted, as it parallels *Specification*, involving only a hypothetical increase of information)

To fully understand Peirce's definition, it is crucial to clarify what he meant by *breadth* and *depth*. According to the author, breadth refers to "the aggregate of real substances of which alone a term is predicable" (Peirce, 1931-1958: 2.414), namely, the entire set of individuals or entities to which a term applies. For instance, the breadth of the concept DOG encompasses all its subordinate types as well as every individual dog. This aligns with the traditional concept of *extension* in logic and *denotation* in linguistics. In contrast, depth is "the real concrete form which belongs to everything of which a term is predicable" (Peirce, 1931-1958: 2.414), corresponding to the complete set of features possessed by an individual or an entity. For example, depth refers to the full set of features associated with the concept of DOG (e.g., <has fur>; <four legs>; <barks>, etc.). This notion aligns with the concept of *intension* in logic and, more loosely, *connotation* in linguistics, as the latter also includes subjective and culturally shaped nuances.

Therefore, according to Peirce's definitions, those mental operations are characterized as follows:

- **Generalization** is an operation that involves a transition from a specific entity, such as "dog," to a broader, more general category at a higher level of abstraction, such as ANIMALS. This process expands the range of referents (i.e., an increase in

breadth) and reduces the total number of possible features (i.e., a decrease in depth), sacrificing specificity in favor of a broader denotation.

- **Induction** refers to the application of newly gathered information from experience to a broader range of exemplars, such as “some dogs wag their tails, then all dogs wag their tails” or “dogs breathe, then animals breathe”. This process expands the range of reference (i.e., an increase in breadth) without altering the number of features (i.e., no change in depth), as those features are already inherent in the broader categories.
- **Abstraction** involves focusing attention on a specific feature or detail while disregarding others (i.e., a decrease in depth), without any change in the range of referents (i.e., the breadth). For instance, one might focus solely on the color of an object or the bitterness of a lemon.
- **Specification** is a process in which further details are provided, increasing the number of features (i.e., the depth), by adding new information, while the range of references remains unchanged. For example, the assertions “dogs can swim” or “dogs can hear frequencies inaudible to humans” specify novel features of the conceptual category DOGS, thereby deepening the understanding of the concept.
- **Determination** is similar to specification, but instead of adding asserted information, it may involve the postulation of a specific subset within a broader concept. For example, determining that a dog can possess all the features of a dalmatian enriches the representation of the concept DOG by specifying a more detailed subset of features, without expanding the range of referents.
- **Restriction** is a process akin to descent but focuses only on narrowing the scope of referents, without increasing depth. Peirce notes that this process may also involve a perceived decrease in information. For instance, restricting the scope from ANIMAL to RODENT may involve a perceived loss of information regarding non-rodent exemplars and their specific features.
- **Descent** may be considered the opposite of generalization, as it involves narrowing the scope of exemplars (i.e., decrease in breadth) while simultaneously deepening the details of the representation (i.e., increase in depth). For example, starting with the broad category of MAMMAL and narrowing it to DOG involves an increase in the number of features, such as specific traits associated with exemplars of dogs.

Table 1.1. summarizes the distinctions among these operations with respect to breadth, depth, and information change.

| Operation      | Breadth   | Depth     | Information change   |
|----------------|-----------|-----------|----------------------|
| Generalization | Increase  | Decrease  | No change            |
| Induction      | Increase  | No change | Increase (believed)  |
| Abstraction    | No change | Decrease  | Decrease (conceived) |
| Specification  | No change | Increase  | Increase (asserted)  |
| Determination  | No change | Increase  | No change            |
| Restriction    | Decrease  | No change | Decrease (assumed)   |
| Descent        | Decrease  | Increase  | No change            |

**Table 1.1.** Comparative characteristics of Peircean logical operations in terms of breadth, depth, and information change, adapted from Peirce (1931-1958: 2.422).

Although the definitions of these operations evolved across the volumes of the *Collected Papers* (Peirce, 1931-1958), what is relevant to this discussion is Peirce's clarification that abstraction is essentially a process of attention and feature selection. Peirce further distinguished two forms of abstraction: *precisive abstraction* and abstraction as *subjectifaction*. Precisive abstraction (or *prescission*) refers to the psychological act of separating or distinguishing certain details from others. For instance, when one describes a garden as <lush>, other features are deliberately left out. This process can also involve a form of categorization, whereby, in representing multiple instances within a synthesized construct, specific and idiosyncratic aspects are omitted. Peirce explains:

The term abstraction bears two utterly different meanings in philosophy. In one sense it is applied to a psychological act by which, for example, on seeing a theatre, one is led to call up images of other theatres which blend into a sort of composite in which the special features of each are obliterated. Such obliteration is called precise abstraction. (Peirce, 1931-1958: 4.463)

This process enables one to recognize existing categories by tracing commonalities while disregarding individual specific details (Zeman, 1982).

However, criticizing modern logic and aligning with the medieval interpretation of abstraction, Peirce claimed that abstraction primarily involves focusing on one aspect while neglecting all others. In this view, the process of abstraction is more consistent with the definition of abstraction as subjectification, rather than precise abstraction. As Peirce states:

Thus, if the definition usually given of abstraction, that it is attention to a part of an idea with neglect of the rest, be accepted, the term must be applied, no longer to precise abstraction but exclusively to subjectification. (Peirce, 1931-1958: 2.428)

This form of abstraction involves the operation of nominalization, in which an adjective, such as “sweet”, is converted into a noun, such as “sweetness”, thereby becoming the subject of a proposition. This process corresponds to the formation of an abstract noun, or *ens rationis* in logical terminology, an entity that exists as a construct of consciousness and thought and can serve as the subject of an actual predication. For example, in a sentence like “the sweetness of the honey is delightful”, the nominalized quality or the feature functions as grammatical and logical subject. For this reason, Peirce also refers to this type of abstraction as *hypostatic abstraction*, insofar as the abstract object assumes an independent nature and can function as the subject of possible propositions (in line with the argument on abstract objects outlined in ¶ 1.1.1.). This notion of abstraction, rooted in mathematics and geometry, enables us to represent complex phenomena by uniting collectivities of instances into a single, unified conceptual entity. For example, the collective term “humankind” refers to groups of individual humans and cultures as a singular, abstract entity, thereby illustrating the process by which a multiplicity of instances is conceptualized as a unified whole, both in language and in the mind.

The present section has narrowed the multidimensional nature of abstraction into three major themes, namely, conceptualization, schematization, and subjectification (or formation) of abstract concepts. However, although these notions are analytically distinct, they partially tend to conflate in certain respects. For instance, the definition of abstraction as precise abstraction appears closely aligned with the notion of conceptualization (see ¶ 1.2.1.) and, to some extent, with the idea of schematization, insofar as precise abstraction does not explicitly refer to a superordinate level (¶ 1.2.2.). Moreover, the account of conceptualization discussed in ¶ 1.2.1. involves both the formation of abstract concepts, conceived as mental entities elaborated indirectly from physical experience, and the

acquisition of more general recurring patterns or schemas that retain common traits applicable to multiple exemplars.

Therefore, although the definitions of abstraction can, in principle, be distinguished across different theoretical accounts or terminologies, the polysemous nature of this notion (Bolognesi et al., 2020) prevents a clear disambiguation of its core meaning. What emerges from these three perspectives is that abstraction primarily deals with the cognitive process of conceptualization, while at the same time appearing to conflate two distinct senses:

- A conceptual sense, referring to the formation of abstract concepts as mental entities that encode multiple and complex experiences;
- A linguistic sense, relating to the formation of general lexical items that function as linguistically derived schemas preserving commonalities across instances (Fortescue, 2017).

The following section illustrates how this conflation between conceptual and linguistic senses of abstraction is also reflected in the experimental paradigms and methodological approaches used to investigate abstraction and to assess abstractive abilities empirically.

### **1.3. Experimental Approaches to Understanding Abstraction**

Empirical research on abstraction has explored its impact on various cognitive processes and capacities, including memory and learning (Medin et al., 1983), information organization (Rosch et al., 1976), and social cognition (Semin & Fiedler, 1988), as well as through computational simulations and modeling approaches (Goldstone & Sakamoto, 2003). However, the heterogeneity that emerges from theoretical frameworks on abstraction, considered one of the core processes in human cognition, reflects a fragmented perspective on the methodologies and procedures used to empirically investigate abstraction.

Burgoon et al. (2013) provided a comprehensive review of the methods employed to study abstraction, listing the types of tasks implemented and elucidating the specific constructs or dimensions of interest that researchers aimed to examine. In this review, Burgoon et al. (2013) defined abstraction as “a process of identifying a set of invariant

central characteristics of a thing” (p. 502), highlighting the functionality of processing information at different levels of abstraction. Therefore, on the one hand, abstraction involves the operation through which a set of common, relevant features is isolated, allowing multiple instances, cases, individuals, situations, or events to be recognized as belonging to the same category. On the other hand, abstraction entails a multiplicity of levels, as previously discussed, in which the relevance and the total number of salient features vary. For example, at lower levels of abstraction, that is, in their view, at more concrete levels, more details and information are retained, requiring a larger number of features to identify a category. In contrast, at higher levels of abstraction, namely, at more abstract levels, fewer details remain, and thus a smaller set of features, which are stable across various individuals, is needed for an accurate categorization.

Furthermore, Burgoon et al. (2013) underlined the role of abstraction and abstract thought in the ability to transcend immediacy, contingency, and concrete experience, as well as in the development of hypothetical reasoning. In psychological research, this capacity of overcoming the “here and now” is typically associated with the ability to traverse the *psychological distance* (Trope & Liberman, 2010; Wang et al., 2021). Psychological distance encompasses four interconnected dimensions: social distance (i.e., from oneself to others), temporal distance (i.e., from the present moment and future or past situations), spatial distance (i.e., from nearby or distant locations), and hypothetical distance (i.e., from real experience and hypothetical ones). Notably, in linguistics, the ability to think about or discuss spatially distant entities, past or future events, and hypothetical situations is known as *displacement* (Hockett, 1960; Vigliocco, Perniss, et al., 2014), a fundamental property of human language grounded in abstractive processes.

Psychological distance is similarly linked to abstraction within the framework of *Construal-Level Theory* (CLT; Trope & Liberman, 2010). According to this view, more abstract entities are perceived as more distant than more concrete entities. Across all four dimensions, low levels of abstraction involve concrete *construals* rich in peripheral features and specific details, whereas higher, more abstract levels omit idiosyncratic information and retain central, common, and stable properties. For instance, the conceptual representation of WHEELCHAIR is characterized by the specific <has wheels> feature, its particular structure, and the limited contexts in which it is encountered or with which it is associated (e.g., “hospitals”). Conversely, the representation of CHAIR, situated at a higher level of abstraction, requires fewer and

more general salient features (e.g., <to be sat on>) that are stable across a wide range of contexts (e.g., “office”, “church”, “house”). Moreover, Trope & Liberman (2010) explain that higher construals like the category FUNNY THINGS or even the concept of FUN are more inclusive than the category PLAYING FOOTBALL, because they encompass a broader set of situations and contexts, and may include a wider range of interpersonal engaging experiences. Thus, organizing construals hierarchically enables the expansion of mental horizons and crossing the psychological distance. Consequently, the formation of more abstract concepts allows one to transcend the contingent experiences and mentally situate them in different spatio-temporal collocations or hypothetical scenarios.

Building on these perspectives, Burgoon et al. (2013) grouped the task types reviewed in their work into three macro-categories, each defined by the elements manipulated to elicit abstraction processes (Burgoon et al., 2013, p. 506):

- Abstraction of perceptual stimuli, i.e., focusing on a perceptual whole (abstract level) versus its parts (concrete level);
- Abstraction of objects, i.e., focusing on similarities among objects (abstract level) versus their differences (concrete level);
- Abstractions of actions: i.e., focus on the underlying reasons for an action (abstract level) versus on context-specific behaviors (concrete level).

Across all types of tasks, successful engagement in abstract thinking appears to depend on identifying more stable, common features or elements, as well as core reasons and intentions across multiple intentions, at a higher level of abstraction, rather than focusing on specific, contextually bounded details. For example, in a task manipulating stimulus perception, such as the Navon letter task (Navon, 1977), in which configurations of multiple small letters (e.g., “L”) form a larger, global letter (e.g., “N”), participants were asked to recognize whether the configurations of letters contain an “L” or an “N”. The capacity to abstract would be assessed by the ability to quickly and correctly recognize the emergent, global letter. Additionally, in tasks manipulating the abstraction of objects like a category versus examples task (Fujita et al., 2006), participants are presented with distinct objects (e.g., table, guitar) and asked either to generate a superordinate category (e.g., FURNITURE, MUSICAL INSTRUMENT) or to produce a specific subordinate concept (e.g., POOL TABLE, ELECTRIC GUITAR). Thus, eliciting

a superordinate, more inclusive category may reflect a more abstract way of thought, whereas generating a specific example would correspond to a more concrete level.

Although it is crucial to emphasize the role of relevant feature selection for abstraction (as discussed in ¶ 1.2.3.), and to acknowledge that abstract thinking can encompass more complex experiences compared to concrete thinking (see ¶ 1.1.2. and 1.2.1.), the perspectives emerging from CLT and Burgoon et al.'s (2013) review appear to conflate the two definitions of abstraction discussed in ¶ 1.2.2. and ¶ 1.2.3., namely, abstraction as schematization and abstraction as subjectification, corresponding to the linguistic and conceptual senses of abstraction. On the one hand, abstraction refers to the ability to generalize from specific instances, such as “electric bike”, to more general and inclusive categories like VEHICLES within taxonomic hierarchies and relations of inclusiveness (see Chapter 3). On the other hand, abstraction also involves the formation of abstract mental entities that are detached from direct physical experience, such as philosophical concepts like DIALECTIC or IMMANENCE.

As previously noted by Bolognesi (2020), Burgoon et al. (2013) describe abstraction as organized along a continuum of levels, ranging from high to low degrees of concreteness. They write:

Lower levels of abstraction (i.e., higher levels of concreteness) capture thoughts that are more specific, detailed, vivid, and imageable [...], often encompassing readily observable characteristics (e.g., furry dog, ceramic cup; Medin & Ortony, 1989). Higher levels of abstraction (i.e., lower levels of concreteness), on the other hand, include fewer readily observable characteristics and therefore capture thoughts that are less imageable (e.g., friendly dog, beautiful cup). It is noteworthy that as a thing gets more abstract, the multitude of its incidental characteristics receive less weight, and therefore the number of central invariant characteristics that are necessary to identify the thing gets smaller. (Burgoon et al., 2013, p. 503)

Thus, in this view, abstraction is primarily defined in terms of a scale of concreteness, typically corresponding to the degree of perceptibility and physical interaction with a referent (Connell & Lynott, 2012; Bolognesi et al., 2020; Vergallito et al., 2020; Löhr, 2022; Villani et al., 2022; Langland-Hassan & Davis, 2023). For instance, the concept SOFA denotes physical entities with which one can interact, whereas BELIEF refers to a mental state that is indirectly accessible through sensorial experience. At the same time, however, the hierarchical organization of these levels implies that abstraction is also understood in terms of schematization. For example, a FRIENDLY DOG

represents a general category that could encompass a broader range of instances, including a “furry dog”.

Furthermore, in a table listing examples of abstractions, Burgoon et al. (2013) did not distinguish between abstract in the linguistic sense of more schematic and general representations and the conceptual sense of abstract as being detached from physical experience. Indeed, in this classification, nominalized abstract concepts such as NUTRITION are placed at a higher level of abstraction than general categories like FRUITS, which nevertheless denote concrete, physical entities. Similarly, emotional abstract concepts such as LOVE and broad categories like HUMAN are grouped at the same level of abstraction (see Burgoon et al., 2013, p. 503).

A comparable conflation can be observed in experimental paradigms like the Abstraction of actions tasks, where non-sensory-based abstract principles, such as intention and personality traits, are treated as more general representations at a higher level of abstraction, applicable to more specific, concrete, and context-dependent situations and behaviors.

Empirical research on abstraction suggests that the conceptual and linguistic senses of abstraction are frequently conflated due to the polysemous nature of the notion itself (Bolognesi et al., 2020; Langland-Hassan & Davis, 2023). While abstraction as non-physical entities and as more general, linguistically derived schematizations can be theoretically distinguished (as explained in ¶ 1.2.2. and 1.2.3.), empirical research often treats these two senses of abstraction interchangeably.

The following section further explores the conflation between the conceptual and the linguistic senses of abstraction, focusing on accounts from linguistic and cognitive science that attempt to distinguish them, while often conflating or overlapping them with other dimensions.

#### **1.4. The Polysemous Nature of Abstraction: The Senses of Abstract**

In the previous sections, abstraction has been discussed as a polysemous and multidimensional notion, whose core meaning resists straightforward definition. As discussed in ¶ 1.2.1, abstraction can primarily refer to the process of conceptualization and concept formation, in which key features are selected through the integration of

shared and distinctive properties to delimit a concept. Additionally, abstraction can encompass both a conceptual and a linguistic sense, which, although theoretically distinguished, are often conflated in empirical research. Specifically, abstraction can refer to: (1) the formation of abstract concepts that are relatively detached from sensory experience (see ¶ 1.2.3.); and (2) the formation of linguistic items that function as schematic representations of general traits across multiple instances (see ¶ 1.2.2.).

These distinctions are consistent with specialized knowledge of abstraction in cognitive science and psycholinguistics. For example, an authoritative resource such as the APA Dictionary of Psychology (Van den Bos, 2007) reports three meanings for abstraction, the first two of which correspond to the linguistic and conceptual sense, respectively:

1. The formation of general ideas or concepts by extracting similarities from particular instances [...].
2. Such a concept, especially a wholly intangible one, such as *goodness* or *truth*.

Note: the third meaning of abstraction refers to *conditioning*, defined as “discrimination based on a single property of multicomponent stimuli” and redirects to the “abstract” lexical entry.

Building on these meanings, scholars have recently examined various theoretical conceptualizations of abstraction to clarify the key aspects of this notion. Reed (2016), for instance, proposed three central senses of abstraction: (1) as non-material abstract concepts; (2) as schematic representations generalizing across particular instances; and (3) as a key feature selection process. Reed (2016) further proposed a taxonomic analysis of abstraction, in which these three senses are mapped to three distinct levels of granularity (Bolognesi et al., 2020):

- *Instance* level, where abstract (i.e., non-physical and non-sensory-based representation) is considered the opposite of concrete (i.e., physical and sensory-based representation).
- *Attribute* level, in which abstract is the opposite of equivalent, emphasizing the selection of distinctive key attributes rather than the equal weighting of all components.
- *Category* level (for a distinction between concepts and categories, see Chapter 2), where abstract is opposed to particular, in other words, schematic representations as opposed to single exemplars.

While a taxonomic account of abstraction is highly pertinent, as discussed in relation to the conceptual and linguistic organization of knowledge (see ¶ 1.1.5.) and the schematicity-specificity continuum (see ¶ 1.2.2.), the account proposed by Reed (2016) appears to schematize an inherently more complex scenario. In particular, the levels of analysis introduced by Reed (2016) are used to contrast concrete and abstract representations, which typically refer to concepts directly and indirectly derived from perceptual experience, respectively. At the instance level, the distinction between concrete and abstract is primarily interpreted in terms of representational format. Concrete instances are characterized as sensory-based and modal representations (Barsalou, 1999; Reed, 2016), whereas abstract instances are described as symbolic, linguistic, and amodal representations (Barsalou, 1999; Reed, 2016). However, the representation of abstract concepts is not purely amodal. While abstract concepts are primarily grounded in linguistic experience (Reilly et al., 2025), they also involves sensorimotor experiences (Borghi, 2023; Villani et al., 2021), and recruits other perceptual processes, such as emotional experiences (Kousta et al., 2011; Vigliocco, Kousta, et al., 2014), introspection (Barsalou, 1999; Barsalou & Wiemer-Hastings, 2005), and interoception (Borghi et al., 2019; Connell et al., 2018; Marson et al., 2023; Villani et al., 2021, 2022), namely, the ability to detect inner bodily states (Borghi, 2023).

Moreover, while abstraction can be understood as a process of key feature selection (as discussed in ¶ 1.3.), the sense of abstract as opposed to particular aligns closely with the linguistic sense of abstraction as schematic representation, whereby common features are generalized across multiple instances. Consequently, although Reed (2016) distinguished between abstract as amodal and non-material and abstract as general and schematic representations, placing these senses at different taxonomic levels of analysis (i.e., instance and category levels), both ultimately refer to abstract representations as contrasted to concrete ones. This overlap results in a partial conflation of the linguistic and conceptual senses of abstraction. Indeed, as shown by Bolognesi et al. (2020), both abstract and concrete concepts can be organized across multiple levels of abstraction. Accordingly, both abstract and concrete concepts may be instantiated at different levels of granularity and precision (see ¶ 1.2.2.), either as more specific exemplars (e.g., DALMATIAN; PRIME NUMBER) as well as broader and more general categories that schematize a limited set of salient features in their representations (e.g., ANIMAL; QUANTITY).

Other works in cognitive science and linguistics build on or implicitly incorporate the conceptual and linguistic sense of abstraction outlined above. In this regard, Fortescue (2017, p. 1) highlighted the existing tension between two distinct uses of the term “abstract”:

1. “Disembodied from specific sensory features”, typically employed in cognitive science research;
2. “Simplified or generalized across instances”, used predominantly in linguistic research.

These two uses of the term “abstract” closely correspond to the conceptual and linguistic senses of abstraction, respectively.

To address this tension, Fortescue (2017) proposed a more general and integrative sense of abstraction that encompasses both. Drawing on Whitehead’s notion of *transmutation* (Whitehead, 1978), abstraction is framed as a process that integrates multiple perceptual and inferential features (see 1.1.3.) into a simplified conceptual unity. Such unities function as nodes that simultaneously encode and instantiate specific entities, their relations, and contextual information. Crucially, Fortescue (2017) further argued that such clusters of perceptual and inferential information are grounded in patterns of neural connectivity, thereby ensuring that abstractions remain perceptually derived and embodied.

At the same time, this integrative process also allows abstractions to encompass a potentially large number of instances and relations. Consequently, the resulting integrated unity may be experienced as increasingly detached from direct physical experience, despite remaining ultimately grounded in perceptual experience (with implications similar to those proposed by the CLT, see ¶ 1.3.). In this way, abstraction enables the management of experiential complexity by supporting conceptualization and lexical semantic organization at multiple levels of granularity. These range from concrete, perceptually grounded concepts (e.g., TULIP), through representations of categories of individuals and social and cultural conceptual constructs (e.g., MARRIAGE), to relational structures and schemas (e.g., “x is the opposite of y”), and highly abstract entities or principles, such as the “wave function collapse” or the concept of GOD.

From a broader cognitive perspective, as argued by Johnson & Tucker (2021), abstraction can be understood as a process of integrating and differentiating experiential

complexity that supports the emergence of self-awareness and, at more advanced levels, metacognitive processes (see ¶ 1.2.1.). Similarly, Fortescue (2017) emphasized the centrality of this third, integrative sense of abstraction in relation to the enactment of self-awareness. Moreover, the development of abstraction may be driven by the acquisition of lexical and semantic structures, and more broadly by language itself, which plays a crucial role in this process of conceptual simplification (Langacker, 1999; Jackendoff, 2003; Langacker, 2008; Enfield, 2024).

Nevertheless, clarification of the persistent conflation between the conceptual and linguistic senses of abstraction, as discussed above, remains incomplete. Although Fortescue's account proposes abstraction as a broader, integrative process encompassing both senses, the emphasis placed on simplification and integration, processes largely mediated by language, arguably appears to collapse this third sense back into the linguistic one. As a result, experiential complexity is primarily accounted for in terms of schematic representations of individuals, categories of individuals, and their relational schemas as instantiated through the linguistic abstraction. However, the apparent detachment from perceptual experience emerges largely as an epiphenomenon rather than as a theoretically independent dimension.

The following section, therefore, introduces an alternative account that further deepens this problem by examining abstraction from the perspective of Embodied (Varela et al., 1992; Shapiro, 2019) and Grounded cognition (Barsalou, 2008), and by analytically distinguishing six distinct senses of abstraction.

#### **1.4.1. Dynamic Abstraction in Perceptual Symbol Systems: Six Senses of Abstraction**

The Embodied and Grounded Cognition (EGC) framework posits that conceptual representations, linguistic meanings, and cognitive processes are built upon perceptual experience and sensorimotor processes and are situated within specific contexts and situations. As such, cognition is shaped by both the physical and cultural environment (Varela et al., 1992; Glenberg, 2010; Borghi & Pecher, 2011; Johnson & Tucker, 2021; Bolognesi, 2025). Specifically, a central claim of grounded cognition accounts concerns the role of *simulations* in cognition (Barsalou, 2008; Pulvermüller et al., 2014; Cuccio, 2015).

Simulations are conceived as the reactivation and reenactment of perceptual, introspective, and bodily states during the execution of various cognitive processes, including language comprehension, conceptual processing, and memory. For example, processing the word “grasp”, reenactment of motor systems is involved. Similarly, understanding a sentence such as “I changed the light bulb” may activate distinct and specific patterns of actions, affective responses, perceptual details, and contextual information. Thus, during the processing of such a sentence, it may be possible to imagine a specific context (i.e., a kitchen), activate contextually related concepts (e.g., LADDER), and involve sensory expectations such as the heat of the light bulb, alongside motor routines like climbing. These simulations reflect the different modalities through which the information is stored in memory and integrated into multimodal representations.

Barsalou (2003, 2005) proposed an account in which abstraction is grounded on simulations within the Perceptual Symbol Systems (PSS) framework (Barsalou, 1999). According to PSS, cognitive processes are based on the formation of perceptual symbols, which are representations recorded as patterns of neural activity arising from perceptual experience (Barsalou, 1999). These representations are multimodal, being distributed across sensorial modalities; schematic, in that they selectively represent specific aspects of perception; dynamic, as they are variable rather than fixed and stable; componential, in that they can be combined to form more complex representations; and situated, as they instantiate contextual, situational, and thematic information (Yeh & Barsalou, 2006; Barsalou, 2008). Notably, these representations also function as symbols, insofar as they schematically stand for features, objects, relations, and events in off-line cognitive processes through simulation activity.

Barsalou’s (Barsalou, 2003, 2005a) account proposed that abstraction can be analytically framed in six distinct senses:

- **Sense 1. Categorical knowledge.** Categories are formed by abstracting regularities from experience.
- **Sense 2. Behavioral ability to generalize across instances.** The properties of a category’s members can be summarized behaviorally, for example, through linguistic expressions, such as generic statements (i.e., “cougars live in the Americas”) and quantifications (i.e., “some plants eat insects”).

- **Sense 3. Summary representation.** Behavioral abstractions in Sense 2 may rely on summary representations that encode recurring patterns across individuals, rather than representing individuals separately.
- **Sense 4. Schematic representation.** Categories are represented schematically by selecting relevant properties to enhance distinctiveness (Posner & Keele, 1968; Palmeri & Nosofsky, 2001).
- **Sense 5. Flexible representation.** Abstraction supports the flexible adaptation of representations to accomplish specific cognitive tasks, “including categorization, inference, language comprehension, reasoning, and so on” (Barsalou, 2005, p. 390).
- **Sense 6. Abstract concepts.** Concepts become increasingly abstract as they are detached from direct physical experience and grounded in introspective states and different situational contexts, thereby increasing the complexity of their representations (Barsalou & Wiemer-Hastings, 2005).

Barsalou (2003, 2005) described the abstraction process within the PSS framework, specifically focusing on Sense 3, i.e., summary representations, and assuming three crucial properties of abstractions: interpretation, structured representation, and dynamic realization.

As a cognitive process, abstraction can be understood as an interpretive process (Pylyshyn, 1973) through which objects, events, and relations are elaborated in terms of *type-token* propositions. A specific object (a token) can thus be interpreted under multiple *types*, such as “table”, “furniture”, “office desk”, “negotiation table”, or even “a place where my cat should not stand”. Therefore, abstractions may be represented as a set of potential interpretations rather than as a single fixed representation. Moreover, such interpretations can be combined in structured representations that encode relations among them. For instance, the preposition “in” may define a conceptual space in which a structured relation is instantiated between “glass” (object x) and “liquid” (object y). Finally, abstractions should not be viewed as a limited or stable set of interpretations; instead, they are dynamic, varying across time and context.

In Barsalou’s (2003, 2005) account, abstractions are grounded in the same representational *constructs* that underpin perceptual symbols, namely, *simulators* and *simulations*. A simulator is a distributed representational system that results from the brain’s association and modality-specific areas. Simulators arise from repeated

encounters with multiple exemplars or instances of a category, enabling the development of categorical and conceptual knowledge. As different exemplars of a category (e.g., COW) are encountered, similar neural patterns are activated. Consequently, multimodal and sensorimotor features are stored in modality-specific areas and integrated across different instances, contexts, and relations by a population of conjunctive neurons, giving rise to the simulators. Once stored in long-term memory, simulators can reenact specific subsets of their contents, producing goal-dependent simulations of different instances. (for the sensorimotor reenactment, see Damasio, 1989, and Barsalou, 2005). Simulators function as *types* of multimodal representations of categories across different instances, whereas respective multimodal simulations serve as *tokens* representing offline “instances in their absence during language, memory, and thought” (Barsalou, 2005, p. 400). Crucially, simulators can represent both properties (e.g., <hand>) across instances and relations (e.g., <above>) among instances, depending on the information processed. Together, these simulators constitute the cognitive interpretations of a category’s members.

According to this account, abstractions may consist of subsets of properties and relations simulators applied across instances. Since only a small collection of property and relations is required to interpret an instance, abstraction is formed dynamically and context-dependently. Furthermore, abstractions interpret instances based on the type-token relationship and can be understood as structured representations, as simulations can be embedded within complex relation simulators.

In conclusion, the six senses of abstraction were reframed within the Dynamic Interpretation of Perceptual Symbol Systems (DIPSS) framework (Barsalou, 2005):

- **Sense 1. Categorical knowledge.** Categories’ knowledge is abstracted from experience, generating property and relation simulators that interpret perceived instances.
- **Sense 2. Behavioral ability to generalize across instances.** Linguistic expressions, such as generic statements (i.e., “tigers have stripes”), describe abstractions through language. For instance, these abstractions reflect the use of *stripe* simulators to interpret a variety of simulated instances of the item “tiger”.
- **Sense 3. Summary representation.** Temporary abstractions can become more established in memory, thereby increasing the likelihood of interpreting a category.

The interpretive system remains flexible, though recurring abstractions can alter its dynamic qualities.

- **Sense 4. Schematic representation.** Abstractions are constituted by only a subset of relations and property simulators rather than the entire set of potential simulators. As a result, abstractions are schematic, representing a relatively limited set of details. Furthermore, simulators themselves are schematic, as they capture and reenact only partial information derived from perception.
- **Sense 5. Flexible representation.** The flexibility of representations that enables the accomplishment of various cognitive tasks depends on the dynamic interpretive system.
- **Sense 6. Abstract concepts.** The detachment of abstract concepts from direct physical experience results from the complexity of their representations, which encode more relation simulators and introspective property simulators. This complexity characterizes abstract concepts, as their conceptual space is interpreted through relational propositions and contextual information (Barsalou & Wiemer-Hastings, 2005).

This framework has the merit of conceiving abstraction as an interpretive process, grounded in perceptual structures, that plays a key role in shaping conceptual knowledge. Through repeated encounters with specific instances, schematic representations of contextual and perceptual information are encoded into memory. Furthermore, this account contributes to explaining the detached nature of abstract concepts, which are characterized by complex representations that encode relations to mental and interoceptive states, as well as contextual and situational information.

Nevertheless, the six senses of abstraction proposed within this framework do not fully clarify the conflation between the conceptual and linguistic senses of abstraction discussed in the previous sections. A critical examination of the six senses suggests that Sense 1 can be subsumed under the definition of conceptualization outlined in ¶ 1.2.1., while Sense 6 identifies the conceptual sense of abstraction, referring to abstract concepts detached from direct physical experience. Sense 2, by contrast, pertains to a distinct semantic dimension, although related to the abstraction, namely, *genericity* (Carlson & Pelletier, 1995). Genericity is defined as the extent to which a linguistic label refers to members of a category across different contexts, thereby reflecting its inclusiveness in use (Bolognesi, 2025; for a further discussion, see ¶ 3.5.1.). Finally, the remaining senses,

as summary and schematic representations, correspond to the linguistic sense of abstraction, insofar as they refer to schemas that preserve relevant commonalities across individuals and instances.

What emerges from the DIPSS framework is that abstraction is an interpretive process that selects and encodes relevant information through schematic representations applicable to a variety of instances, closely aligning with the linguistic sense of abstraction (see ¶ 1.2.2.). Moreover, abstraction appears to be primarily involved in the process of conceptualization and conceptual formation, whereas abstract concepts appear to reflect the outcomes of these processes rather than constituting a component of them.

For the purpose of the present work, it is necessary to draw a clear distinction between the process of abstraction, characterized by schematic representations, and the nature of the abstract concepts. For this reason, the following section provides a terminological and conceptual clarification to more precisely define abstraction by disentangling its conceptual and linguistic senses, namely, abstract concepts and schematic, general representations, respectively. These should be treated as two distinct dimensions, each operationalized through two separate semantic and psycholinguistic variables: *conceptual concreteness* and *categorical specificity*.

#### **1.4.2. Disentangling the Polysemous Nature of Abstraction: Conceptual Concreteness, Abstractness, and Categorical Specificity**

As discussed throughout this chapter, abstraction is a crucial cognitive phenomenon that has been examined from multiple perspectives and empirically investigated using a range of methodologies and paradigms. However, the polysemous nature of abstraction (Bolognesi et al., 2020; Langland-Hassan & Davis, 2023) resists a clear definition of its core meaning. On the one hand, abstraction has been identified as the process of conceptualization (see ¶ 1.2.1.). On the other hand, abstraction has been defined as encompassing both a linguistic and a conceptual sense: the process of forming linguistically derived schematic representations that generalize across specific instances, and mental entities indirectly derived from physical experience, that is, abstract concepts. Specifically, the conflation of these senses results in the difficulties of investigating abstraction.

The major issue lies in the failure to treat the conceptual and the linguistic sense as independent variables. As discussed in the previous sections (in particular, see ¶ 1.4 and 1.4.1.), abstraction constitutes an interpretive cognitive process that isolates and selects salient information across multiple objects, individuals, and contexts. This mechanism gives rise to simplified representations that comprise schematic reconstructions of relevant common features as well as relations among instances across various contexts. In this sense, abstraction appears to correspond to its linguistic sense, namely, the linguistic process of schematization (Langacker, 1999, see ¶ 1.2.2.), whereby higher-order schematic representations are generated by identifying and extracting commonalities among individuals at different levels of granularity and specificity. Interpreting abstraction as schematization further involves recognizing it as both a key feature selection process, as highlighted by Burgoon et al. (2013), and as potentially grounded in a type-token relationship, as emerged from the discussion of the DIPSS account (Barsalou, 2003, 2005; for the type-token relationship in schematization, see ¶ 1.2.2.).

Within abstraction, the linguistic process of schematization contributes to the formation of conceptual knowledge, resulting in the equation of abstraction with conceptualization, as expressed by Johnson & Tucker (2021, p. 260): “the process of conceptualization is the process of abstraction”.

Together, these implications align with a recent collective endeavor of 54 authors, who recognize the theoretical urgency to converge on definitions for key semantic constructs. In this consensus paper, abstraction was defined as follows, with 100% endorsement and a mean confidence score of 82 out of 100:

The process of forming general ideas or concepts by extracting similarities and general tendencies from direct experience, language, or other concepts. (Reilly et al., 2025, p. 250)

Reilly et al. (2025) made a clear distinction between abstraction and other related notions, such as *generalization*. As discussed in the previous sections, other accounts on abstraction also revealed the importance of theoretically clarifying the distinction between these two notions (see ¶ 1.1.2. and ¶ 1.2.3.). Specifically, in response to Colunga & Smith (2003), who argued that abstraction is identical to generalization by similarity and associative learning, Reilly et al. (2025) clarified that, while the two are similar, abstraction involves the formation of high-order representations that preserve core

features across referents. In contrast, generalization refers to the mechanism by which information acquired from specific instances is applied to novel contexts (in line with Dewey's interpretation of abstraction and generalization, see ¶ 1.1.2.). As Son et al. (2008) also emphasized, abstraction is a simplification process in which the schematic representation, abstracted from instances, contributes to the mechanism of generalization, thus further reinforcing the distinction between these two notions.

While abstraction corresponds to both its linguistic sense and the process of schematization, the linguistically derived representations, as well as conceptual knowledge, are the outcome of the process of abstraction itself. This distinction resolves the conflation between the linguistic and the conceptual senses of abstraction: the former refers to the process of abstraction, while the latter pertains to its result. Therefore, the nature of the ascribed conceptual sense of abstraction, specifically as abstract concepts, should be explored through a distinct theoretical notion, namely, *abstractness* (Borghi, 2022, 2023; Reilly et al., 2025).

Recently, Borghi & Binkofski (2014) and Borghi et al. (2019) highlighted the distinction between abstraction and abstractness. In their view, abstraction refers to the process of concept formation, where a set of features is extracted from multiple instances to represent a concept. Borghi et al. (2019) argued that this process is particularly prominent in the case of more general representations placed at higher levels of abstraction (e.g., FOOD; SUBSTANCE). In contrast, abstractness pertains to a specific type of concepts, i.e., abstract concepts, concerning:

knowledge domains that are less spatio-temporally bounded and cannot be fully tracked with exteroceptive senses alone. Prototypical examples of abstract concepts are “fantasy”, “freedom” and “justice”. (Borghi et al., 2019, p. 121)

More recently, Reilly et al. (2025) also endorsed this distinction between abstraction and abstractness, converged on a definition of abstractness with a confidence mean of 84 out of 100, defining it as:

The quality of a concept (or word) whose meaning is understood primarily on the basis of language, but also draws from interoceptive experiences, including emotion, introspection, and metacognition. Abstract concepts are often exemplified by perceptually dissimilar associated actions and events. (Reilly et al., 2025, p. 248)

Conversely, *concreteness*, which characterizes concrete concepts, those for which experience-based knowledge is more prominent, has been defined as follows, with a confidence mean of 92 out of 100:

The extent to which a word or concept evokes a (multi)sensory experience encompassing both the classical basic senses but also extending to the chemical senses, interoception, and sense of self (e.g., body awareness and related phenomena). (Reilly et al., 2025, p. 253)

The difference between concrete and abstract concepts relies on their different nature, which in turn clarifies the theoretical distinction between abstractness and concreteness, although the two are often treated as complementary aspects of a single semantic variable (for further discussion, see Chapter 2). These two dimensions have typically been investigated in psycholinguistics and cognitive science as opposite poles of a single operationalized scale in ratings studies, where participants subjectively evaluate the concreteness/abstractness of words on a Likert scale, typically ranging from 1 to 5 or 7. One of the most comprehensive datasets employed in these disciplines is provided by Brysbaert et al. (2014), which includes ratings of the concreteness dimension for both single and combined English lemmas distributed across different parts of speech (e.g., nouns, adjectives, verbs).

The distinction between abstraction and abstractness resolves the conflation between the linguistic and conceptual senses discussed in the previous sections, particularly in ¶ 1.3., where abstraction was described in relation to experimental methodologies for empirically investigating this process. However, it remains crucial to determine how the notion of abstraction, as linguistic schematization and as recently defined (Reilly et al., 2025), can be empirically investigated.

Bolognesi et al. (2020) and Bolognesi & Caselli (2023) introduced a novel semantic and psycholinguistic variable to represent a dimension that can account for the process of linguistically derived abstraction as an operationalized variable, namely, *categorical specificity*.

Categorical specificity is a psycholinguistic dimension that measures the degree of semantic precision in a word's meaning, determined by the extent of its conceptual scope and the inclusiveness of its category (Villani & Bolognesi, 2025). Specificity reflects the degree of precision in a linguistic or conceptual representation, distinguishing words or concepts that denote highly general categories, which encompass a broader set

of referents (e.g., OBJECT; ANIMAL), from those that refer to more specific categories, which denote fewer or more narrowly defined referents (e.g., PNEUMATIC DRILL; KOMODO DRAGON). This notion closely aligns with Langacker's notion of schematicity (Langacker, 1987; Evans & Green, 2006; Bolognesi, 2025), insofar as both specificity and schematicity organize conceptual and linguistic categories at varying levels of semantic detail of their representation. Therefore, specificity also reflects the taxonomic structure of linguistic categories and conceptual knowledge (see ¶ 1.1.5.; for further elaboration, see Chapter 3).

While the literature generally suggests that concrete concepts or natural categories are typically organized in taxonomies (Borghi, 2022; Rosch et al., 1976), Bolognesi et al. (2020) and Bolognesi & Caselli (2023) demonstrated that both abstract and concrete words can denote broad, inclusive conceptual categories (e.g., ORGANISM; CONDITION) or narrower, more specific conceptual categories (e.g., RABBIT; FREEDOM OF SPEECH; as discussed in ¶ 1.4.). Moreover, Bolognesi et al. (2020) and Bolognesi & Caselli (2023) showed that conceptual concreteness and categorical specificity present a mild, yet positive, correlation ( $r = 0.29$ ), suggesting that these two variables reflect two separate dimensions related to abstraction.

This work further explores the relationship between these two variables, which reflect the clarified and disentangled senses of abstraction, to provide empirical evidence and theoretical insights into their separate contributions to language processing (for further elaboration, see *Methodological and Theoretical Rationale for the Experimental Section*). This investigation will contribute to providing valuable insights into the role of linguistically derived abstraction in cognition, and more importantly, its relationship with language (see Chapter 7). By examining these semantic and psycholinguistic variables, this work investigates the contribution of abstraction in shaping conceptualization and influencing linguistic processing within a cognitive and psycholinguistic framework.

The following chapters introduce and discuss the variables of conceptual concreteness and categorical specificity, respectively.

## Chapter 2

### Conceptual Concreteness and The Complexity of Abstractness

#### 2.1. What are Concepts? From a Definitional Structure of Concepts to Multidimensional Conceptual Spaces

Understanding what concepts are requires considering them metaphorically as “the glue” that holds our mental representations and knowledge about the world together (Borghi, 2023; Murphy, 2002). As discussed in ¶ 1.4.2., conceptual knowledge is the outcome of the abstraction process through which schematic representations are formed by integrating shared features of multiple instances across various situations and contexts. For example, the concept CAT can be formed by extracting relevant features from various individual instances, resulting in a mental representation that can be associated with the linguistic label “cat”. The process closely parallels category formation in that both involve identifying regularities across instances and grouping them into a coherent schema (e.g., a class of entities).

Categories and concepts are closely interrelated notions, and their distinction is not always straightforward. Following the influential account proposed in the *Big Book of Concepts* (Murphy, 2002), a concept can be conceived as a mental representation of a class of entities that encodes all experience-based knowledge about them. In contrast, a category corresponds to the class of entities itself. Therefore, a category corresponds to a class of entities represented by a concept, such that making inferences about a concept entails making predictions about the corresponding category. In a similar vein, a category represents the entire set of possible referents of the category, while a concept reflects the always varying global body of physically and culturally derived information about these referents stored in the conceptual system. In this sense, the distinction between category and concept can be linked to the distinction between denotation and connotation (see ¶ 1.2.3.), without reducing the former two to the latter pair. However, in this work, particularly in Part 2, concepts will be treated as preserving a referential scope, since they mentally represent categories.

In forming categories and concepts, we recognize and categorize individual instances (e.g., various birds) in the world as members of a specific category, thereby activating the corresponding conceptual knowledge (e.g., BIRD; Barsalou, 2012). This categorization process occurs at varying levels of abstraction and specificity (e.g., ROBIN; BIRD, VERTEBRATE; ANIMAL, ORGANISM; Rosch, 1978; Bolognesi, 2025), as discussed extensively in Chapter 1 (for further details, see Chapter 3).

Nevertheless, the debate surrounding the nature of the conceptual system and the mechanisms of categorization remains a major topic in cognitive science and psycholinguistics. As discussed in ¶ 1.2.1. and 1.4.1., contemporary theories of concepts and categorization emphasize the multidimensional, integrated, and context-sensitive nature of conceptual representations, challenging the so-called *Classical view* (Smith & Medin, 1981) or *Classical theory* of categorizations, which treats concepts and categories as a rigid set of defining features and properties.

### **2.1.1. The Classical View**

The classical theory has been the dominant perspective on concepts and categorization since the time of the ancient Greek philosopher Aristotle (Murphy, 2002; Evans & Green, 2006). According to this view, concepts and related linguistic categories are characterized by a definitional structure and a set of necessary and sufficient conditions. Therefore, an individual instance would be considered a member of a specific category only if it met the conditions specified in the category's definition. Consequently, the meaning of a category corresponds to the content of its definition. In the classical view, concepts are mentally represented as definitions.

This definitional approach implies that a concept's definition includes a set of properties or features (typically sensorial or perceptual; Evans & Green, 2006) that are individually necessary and jointly sufficient for an instance to be considered a valid member of the category. Furthermore, all features included in the definition should have been jointly possessed to ensure membership in the category. For example, the concept CAT might be defined as “an agile, four-legged animal with fur, a tail, claws, and the ability to meow”. Thus, the conceptual representation would consist of these defining features, and each potential member of the category must possess all of them (<agility>, <four-legged>, <animal>, <fur>, <tail>, <claws>, <meowing>).

However, this approach faces at least three significant challenges:

1. Definitional problem
2. Category and conceptual fuzziness
3. Typicality

First, establishing a clear definition of a concept and identifying the set of all potential necessary and sufficient features for a valid membership in the category presents a considerable challenge. A famous and traditional example of this difficulty arises from the philosopher Wittgenstein's account regarding the impossibility of providing a precise definition of the concept GAME:

Consider for example the proceedings that we call "games". I mean board-games, card-games, ball-games, Olympic games, and so on. What is common to them all?—Don't say: "There *must* be something common, or they would not be called 'games' "—but *look and see* whether there is anything common to all.—For if you look at them you will not see something that is common to *all*, but similarities, relationships, and a whole series of them at that. To repeat: don't think, but look!—Look for example at board-games, with their multifarious relationships. Now pass to card-games; here you find many correspondences with the first group, but many common features drop out, and others appear. When we pass next to ballgames, much that is common is retained, but much is lost.—Are they all 'amusing'? Compare chess with noughts and crosses. Or is there always winning and losing, or competition between players? Think of patience. In ball games there is winning and losing; but when a child throws his ball at the wall and catches it again, this feature has disappeared. Look at the parts played by skill and luck; and at the difference between skill in chess and skill in tennis. Think now of games like ring-a-ring-a-roses; here is the element of amusement, but how many other characteristic features have disappeared! (Wittgenstein, 1958: 66)

In this example, the challenge in establishing a rigid set of shared features for every member of the category GAME becomes evident. Among the various subtypes and specific instances of the concept GAME, one can identify certain shared features (e.g., <amusement>), along with others that are distinctive or unique to particular subtypes (e.g., <luck> for CARD GAMES; for further details on distinctiveness, see Chapter 3). This argument underscores the impossibility of defining a fixed set of features that every member of a category must share.

To provide another example, the ability to fly constitutes a feature generally shared by most members of the category BIRD. However, exceptions exist, such as OSTRICH, CASSOWARY, and KIWI BIRD, which are all birds that cannot fly.

Secondly, the definitional approach assumes that categories and concepts have clear, well-defined, and neat boundaries, where membership can be determined discretely, with instances belonging exclusively to one category or another. However, speakers often face difficulties in categorizing certain objects, revealing uncertainty in whether a particular instance belongs to one category or another (e.g., whether TOMATO or AVOCADO should be classified as a FRUIT or VEGETABLE).

Empirical observations support this intuition. For example, Hampton (1979) found that, when asked to rate whether various items were category members, participants did not provide clear-cut evaluations. In seven out of eight categories, a number of items were categorized as members and non-members of a category, with only a minimal difference in deviation between the two groups. Moreover, McCloskey and Glucksberg (1978) found that participants, when providing repeated category judgements (e.g., “Is a lion an animal?”) at different times (e.g., after two weeks from the first judgment), often changed their response, revealing a genuine uncertainty in categorizing an item as either a member or non-member of a category. These findings suggest that boundaries between categories and conceptual representations exhibit degrees of fuzziness.

Finally, classical theory fails to account for the typicality effect, where some category members are considered more representative than others. The pioneering work of Rosch and colleagues (Rosch, 1975, 1978; Rosch et al., 1976) demonstrated that conceptual knowledge is structured across multiple levels of categorization and that, within a category, members do not hold the same status, but are instead organized around a central prototype, conceived as the best representative exemplar of the category, (e.g., CARDINAL as prototypical member of the category BIRD). Consequently, members that share more features with the prototype, that is, those more similar to the prototypical exemplar, are positioned closer to the center of the category. These members are considered more typical representatives of the category. Conversely, members located at the periphery, being less similar to the prototype, are regarded as less typical exemplars due to their greater dissimilarity to the prototype. Chapter 3 provides a more detailed discussion on this topic.

These critiques highlight the limitations of the classical view, thereby paving the way for alternative perspectives that account for a less rigid and more flexible interpretation of categorization and conceptual representation. While Chapter 3 explores

other theories, particularly those focused on the taxonomic organization of conceptual representations, such as *Exemplar* and *Prototype Theory*, the current discussion introduces frameworks within cognitive semantics and psychology that emphasize a multidimensional, integrated, and context-sensitive interpretation of conceptual representation. These include *Frame Semantics* and theories of *Domains*, which offer a more dynamic understanding of how concepts are interconnected with one another and with background knowledge, as well as how they are understood in varying contexts.

### **2.1.2. Integrated and Flexible Conceptual Representations: Frames and Theories of Domains**

An attempt to address the limitations of the classical theory and its understanding of concepts as a fixed set of stable features was made within the cognitive linguistic framework. Rather than being isolated and neatly defined, concepts and their associated words are interconnected with other concepts and the words related to them, representing recurring structural elements and reflecting a network of beliefs, knowledge, and experiences that shape understanding of the concepts themselves. These are crucial tenets of Fillmore's Frame semantics (Fillmore, 1975, 1977, 1985; Fillmore & Atkins, 1992). Frames are schematic knowledge structures that account for both syntactic and semantic relations, representing a broad set of assumptions and background information necessary for understanding concepts and words. For example, the word "tall", when describing vertical distance, can be represented differently depending on its context, whether referring to a human, a building, or other objects. "Tall" lexically contrasts with "short" to describe human height, pairs with "high" as the opposite of "low" when referring to vertical distance from a baseline (e.g., "power lines"), and contrasts with "low" when describing the vertical height of buildings (Fillmore, 1977; Croft & Cruse, 2004). Thus, the complete understanding of "tall" and its contextual use, as well as its semantic relations with other words, depends on the specified structured knowledge or frame (e.g., frames for HUMANS, BUILDINGS, or OBJECTS LOCATED A CERTAIN DISTANCE FROM THE GROUND).

Frames also incorporate syntactic roles, as well as sequential and scene-related information, in the representation of concepts and words. This suggests that understanding a set of concepts or words requires access to a specific frame, which

structures how words interact within sentences by providing a structured knowledge of sequential configurations. For instance, the frame of COMMERCIAL EVENT is central to understanding concepts and words like “buy”, “sell”, “pay”, specifying the argument structures assumed by *participants roles* like <buyer>, <seller>, <goods>, and <money> (Fillmore & Atkins, 1992; Croft & Cruse, 2004). Thus, within the COMMERCIAL EVENT frame, the concept BUY implies two participants (i.e., <buyer> and <goods>), while PAY may require two or three participants (i.e., <buyer>, <seller>, <goods>) and may evoke a relation with an additional participant, <amount paid>, as in the sentence, “Alexis paid 500€ for the smartphone”. Such frames have played a crucial role in Construction Grammar (Goldberg, 1995; Croft & Cruse, 2004). Furthermore, the idea that conceptual representation reflects structural elements of language is not new in cognitive linguistics. Talmy (2000) also highlighted how language reflects conceptual structures by encoding cognitive representations of the world through structured knowledge. The distinction between closed-class and open-class elements highlights structural roles in conceptual representation, with closed-class words (e.g., “and”, “or”, “the”, “some”) providing grammatical constructions and scene-structuring frameworks, and open-class words (e.g., “dog”, “house”, “garden”) ensuring substantive content of a particular scene.

A concept similar to frames is Langacker’s theory of domains (Langacker, 1987), according to which words and concepts are understood through relations with other concepts, words, and in contexts of use. For example, the concept and linguistic unit “radius” are inseparable from the relative concept and linguistic unit “circle”. Similar to frames, domains are schematic knowledge structures that provide background information necessary for understanding lexicalized concepts. As in frame semantics, multiple domains can be accessed to specify different components of a conceptual representation. However, unlike frames, domains can be distinct in terms of basic and abstract domains. This distinction builds on embodied and grounded cognition (see ¶ 1.4.1.) with basic domains derived from sensorimotor experience, whereas abstract domains are more complex conceptual structures (e.g., LOVE, MARRIAGE) that may involve both physical experience and knowledge about social and cultural activities. Moreover, domains can be organized hierarchically, presupposing different possible relations with lower-level domains functioning as subdomains for higher-level concepts. For example, FINGER is understood in relation to larger domains of HAND or BODY

through a part-whole (*meronymic*) relation. Similarly, the lexical item “nephew” is understood within multiple potential domains, all related to the FAMILIAL RELATIONS conceptual network.

In the field of cognitive psychology, Barsalou (1992) and Barsalou & Hale (1993) propose an account of conceptual representation that incorporates frames and background knowledge structures. This approach builds on Bartlett’s (1932) theory of schemata, which, similar to linguistic theories, emphasizes the role of recurring patterns of schematic representations of experience stored in memory. Barsalou (1992) and Barsalou & Hale (1993) further specified three key components of frames in relation to Fillmore’s account: attributes, values, and structural invariants.

Frames are defined by Barsalou & Hale (1993, p. 126) as “a co-occurring set of multivalued attributes that are integrated by structural invariants”, and can potentially represent various kinds of categories (e.g., OBJECTS, ENTITIES, PHYSICAL EVENTS, MENTAL EVENTS; Barsalou, 1992). An attribute is a related concept or category that participates in describing an aspect of the frame, while values consist of subordinate concepts that specify a subtype of the attribute. In the classic example of the CAR frame, FUEL, WHEEL, and TRANSITION are attributes that co-occur in the representation of the frame and can be further specified by more specific concepts as their values, such as PETROL or LPG for FUEL, and AUTOMATIC and MANUAL for TRANSITION (Barsalou, 1992; Evans & Green, 2006). Since attribute-value relations are taxonomic, values can also function as attributes, with more specific concepts becoming their values. This hierarchical structure allows more specific terms to inherit properties (e.g., the property of being an aspect of the frame) from more general attributes (see Chapter 3 for further discussion on taxonomic inheritance).

Within a frame, attributes co-occur together across exemplars and are systematically correlated. Structural invariants, therefore, represent the range of relations (e.g., spatial, temporal, intentional; Barsalou, 1992; see also Miller & Johnson-Laird, 1976 for an extensive review) that are abstracted from multiple exemplars and instances, ignoring specific idiosyncratic values, and reflected as attributes in the more schematic and general frames.

Frames can also represent typical patterns and prototypical aspects when exemplars’ values are consistent across the integrated instances. However, frames can

also represent atypical members, whose attributes and values differ from the most recurring aspects but still contribute to the overall category structure.

The frames described in Barsalou (1992) and Barsalou & Hale (1993) have been integrated within the framework of Perceptual Symbol Systems (PSS; Barsalou, 1999; see ¶ 1.4.1.) as information extracted from simulators that integrates multimodal information and relations across multiple instances that contribute to form relatively stable representations of the conceptual content (for further details see ¶ 1.4.1.. Therefore, part of our knowledge of a concept can correspond to a schematic frame that dynamically represents the types of knowledge, relations, and contexts associated with a concept, as well as potential simulations of instances (Evans & Green, 2006).

Although the exact nature of what constitutes a concept remains an ongoing debate in cognitive science, for instance, Johnson & Tucker (2021) emphasized the predictive nature of concepts (see ¶ 1.2.1.) and human cognition (Clark, 2016). Constructs such as fields or domains suggest that conceptual representations are shaped by various types of experience. These representations correspond to flexible, though relatively stable structures, influenced by the specific experiences recorded and the contexts in which they are employed. This aligns with an *encyclopedic view* of meaning (Evans & Green, 2006; Murphy, 2002) in lexical semantics, which posits that meanings are not confined to fixed definitions but are shaped by a network of linguistic and extralinguistic experiences. In this view, meanings emerge from the complex interactions of different types of experiences (linguistic, sensorial, contextual) that relate to the potential concepts associated with them.

A similar perspective is particularly relevant when considering concepts that differ in terms of experiences and dimensions reflected in their representations, such as concrete and abstract concepts.

## **2.2. Abstract vs. Concrete Concepts**

Traditionally, the difference between concrete and abstract concepts was based on the physical accessibility of their potential referents. Concrete concepts (e.g., DESK) typically refer to physically available referents, while abstract concepts (e.g., DIPLOMACY) refer to mental entities lacking physical, tangible, single referents in the

external world. However, this approach has often led to a negative definition of abstract concepts, defining them by what they are not, and to an ontological assumption of a dichotomous opposition between concrete and abstract concepts based on the absence or presence of physical perceptibility. Thus, abstract concepts were understood as unrelated to physical experience, while concrete concepts were tied to it.

However, recent empirical findings have demonstrated that while concrete concepts are grounded in perceptual features and sensorimotor experience, abstract concepts are also grounded in sensory-based experiences, including interoception and introspective states, as well as sensorimotor components, albeit to varying extents (Barsalou, 1999; Kousta et al., 2011; Connell et al., 2018; Villani et al., 2021; Banks & Connell, 2022; Falcinelli, Fini, Mazzuca, & Borghi, 2024).

Therefore, the traditional dichotomy has been challenged, leading to a more nuanced understanding of the relationship between concrete and abstract concepts as existing on a continuum rather than a clear-cut distinction (Wiemer-Hastings et al., 2001; Barsalou & Wiemer-Hastings, 2005; Banks & Connell, 2022). Indeed, it has been argued that certain words may be interpreted as both abstract and concrete concepts depending on their meanings or connotations. For example, the word “government” can refer to both an abstract institution and a physical building (Wiemer-Hastings et al., 2001). Altarriba et al. (1999) also suggested considering emotional words like “fear” as a distinct class of concepts, separated from both abstract and concrete concepts.

Barsalou & Wiemer-Hastings (2005) found that, in analyzed clusters, as concepts move from concrete to abstract, they tended to become more hierarchically complex, with intermediate concepts occupying the space between the two extremes. Highly concrete concepts tended to access more physical elements (e.g., objects, locations), which become more salient in episodic memory and within specific situations. In contrast, highly abstract concepts were less likely to be tied to specific situations and tended to describe more contextual variability and refer to more social aspects (e.g., relations between people, social institutions), introspections (e.g., beliefs, intentions), and more complex relations (e.g., causal relations).

Supporting this view, Asmuth & Gentner (2017) observed that abstract relational concepts are easier to identify than concrete relational concepts, with the opposite being true for entity concepts. Consequently, abstract concepts tend to refer more to relational

properties and concepts, describing connections to other concepts and contextual information. In contrast, concrete concepts refer more to the intrinsic properties of referents (Gentner & Asmuth, 2017). As a result, relational abstract concepts are more influenced by contextual information and relational structures, resulting in a greater variability during encoding and a lower recognition accuracy compared to concrete entity concepts.

Barsalou et al. (2018) also argued that the traditional distinction between concrete and abstract concepts, particularly when studied as decontextualized words (see Chapter 4 for further discussion), is overly simplistic. According to the *Situated Conceptualization* framework (Barsalou, 2005b, 2016), both concrete and abstract concepts are encoded in broader situational networks (as discussed for frames and domains in ¶ 2.1.2.), which are distributed across various brain regions. These networks are connected to varying degrees, to both internal (e.g., goals, emotions, motivation, mentalizing activities) and external elements (e.g., settings, objects, agents) to support comprehension and action in real-world situations. Specifically, concrete concepts are more closely tied to external situational elements, while abstract concepts arise from the processing of internal elements and more subjective experiences. Furthermore, abstract concepts integrate both internal and external situational elements into relational representations, which are stored in both working and long-term memory. Notably, both concrete and abstract concepts can involve both types of situational elements. For instance, a concrete activity like GARDENING can be viewed more abstractly when considered in terms of goals and emotions. Similarly, an abstract concept like POLITICS may be understood more concretely, particularly when associated with a specific situation, such as voting in a voting booth.

Although the notion of a continuum better captures the relationship between concrete and abstract concepts than the dichotomous view, researchers and scholars have recently advanced the idea that abstract and concrete concepts can be represented across a “multidimensional space, defined by various features and dimensions” (Borghi, 2023, p. 155; Crutch et al., 2013; Villani et al., 2019, 2021). This multidimensional model is further supported by the definitions of abstractness and concreteness provided in the consensus paper by Reilly et al. (2025; see also ¶ 1.4.2.). According to these definitions, abstractness refers to the quality of concepts (or words) acquired primarily through language, emotions, interoceptive experiences, and metacognition, while concreteness

describes the extent to which concepts (or words) elicit self-awareness, sensory, and interoceptive experiences.

### 2.2.1. Varieties of Dimensions for Abstract and Concrete Concepts

Several dimensions and semantic variables characterize abstract concepts and abstractness, as well as concrete concepts and concreteness. As discussed in the previous section and, in part, in the previous chapter, both concrete and abstract concepts and words are grounded in multiple sensory dimensions, though to varying extents. Banks & Connell (2022) conducted a series of exploratory analyses, showing that both concrete and abstract concepts are grounded in vision, consistent with arguments about the encoding of situational elements in mental representations. However, concrete concepts (e.g., BOTTLE) are more strongly correlated with haptic experience, while abstract concepts (e.g., CONSCIOUSNESS) are more grounded in interoception (i.e., the bodily awareness of inner states; Connell et al., 2018; Villani et al., 2021; Barca et al., 2020).

Compared to concrete concepts, abstract concepts are typically less imageable (Paivio, 1971), meaning they are less likely to evoke a clear mental image, and tend to receive lower ratings in *body-object-interaction* (i.e., a measure for the ease of interaction between human body and words' referents; Tillotson et al., 2008; Pexman et al., 2019). Moreover, more valenced words (i.e., more emotionally charged words) tend to be more abstract (Vigliocco, Kousta, et al., 2014). According to an influential framework, abstract concepts are grounded more in emotional and linguistic experience, while concrete concepts are grounded in sensorimotor experience (Kousta et al., 2011; for an extensive discussion on the *Affective Embodiment Account*, see ¶ 2.5.2.).

Abstract concepts are typically characterized by *low-dimensionality*, while concrete concepts exhibit *high-dimensionality* (Borghi, 2022, 2023). A category is considered low-dimensional if it refers to multiple exemplars that share few salient features, whereas a high-dimensional category refers to a smaller number of exemplars that share more relevant features (Lupyan & Mirman, 2013; Langland-Hassan et al., 2021). For example, concepts like JUSTICE and HOPE encompass a more heterogeneous set of experiences, entities, and contexts compared to concepts like VASE and COOKIE JAR. Dimensionality is also relevant for distinguishing general from specific words (see Chapter 3 for further details). Furthermore, the notion of dimensionality parallels the

distinction between statistically sparse categories, which are defined by a limited number of relevant features that vary independently, and statistically dense categories, which possess a greater number of interrelated features relevant for category membership.

As abstract concepts are distributed across broader networks of more variable linguistic contexts and word-based associations (Kiefer & Pulvermüller, 2012), they tend to exhibit higher *semantic diversity* (Hoffman, 2016) than concrete concepts. Abstract concepts also have lower *contextual availability* than concrete concepts (Schwanenflugel et al., 1988, 1992). Specifically, abstract words tend to have weaker, more diffuse contextual associations, making the retrieval of relevant contextual information from semantic memory more difficult. In contrast, concrete concepts are connected to a more restricted network of contexts, facilitating the retrieval of semantic information (for further details, see ¶ 2.3.). Moreover, abstract concepts exhibit low *situational systematicity* (Davis et al., 2020), as they rely on a sparser, more variable set of elements across multiple situations and over time, requiring greater semantic control to selectively activate the correct meaning. In contrast, concrete concepts are characterized by high situational schematicity, being more closely tied to recurring and recognizable perceptual elements within given situations.

Concrete concepts are typically acquired early in the linguistic development, whereas abstract concepts are generally acquired later (Della Rosa et al., 2010; Villani et al., 2019). This difference is likely because a full understanding of abstract concepts requires advanced linguistic, semantic, and metacognitive competencies. Furthermore, while concrete concepts are primarily acquired through perceptual experience with potential referents, the lack of physical referents for abstract concepts results in an acquisition that is primarily based on linguistic experience (Della Rosa et al., 2010). However, recent proposals advocated for the broader role of language in the acquisition of lexicalized conceptual representations. Cuccio & Caruana (2019) suggest a reconsideration of the modality of word acquisition, regardless of the nature of the referents, since the latter cannot fully predict the role of sensory, linguistic, and social knowledge in conceptual development. Moreover, Sloutsky & Deng (2019) highlight that, although concepts can either initially be acquired through sensory experiences and later lexicalized or acquired through language and successively grounded in perception, a large proportion of words are acquired linguistically through interaction and reading.

Similarly, it has been proposed that the role of language extends beyond acquisition, playing a crucial role in semantic scaffolding and neuro-enhancement of language use (Dove, 2020; 2023). Borghi et al. (2018) emphasized the importance of social interaction and metacognition (Shea, 2018) in understanding abstract words, for which it is more likely that people rely more on interlocutors and mechanisms of social cohesion to fully grasp their meanings (for a more detailed discussion, see ¶ 2.5.4. and 2.6.).

A further dimension interacting with concreteness and abstractness is *iconicity*, defined as the cross-modal correspondence between aspects of a word's form and aspects of its meaning (Winter et al., 2024). Onomatopoeic words such as “boom” and “click”, as well as words like “bomb” and “cuckoo”, are considered iconic words because their form imitates or resembles in their meaning the traits of the referents they denote. Iconicity is positively correlated with sensory strength, particularly in the auditory and haptic modalities (Winter et al., 2017). In light of this relationship with perceptual experiences, Lupyan & Winter (2018) argued that iconic words should be highly concrete, as they are closely tied to a specific, concrete referent, limiting the flexibility for abstract extension, as occurring, for example, in metaphors. However, more recent findings showed an opposite pattern, with iconicity ratings being negatively correlated with concreteness (Winter et al., 2024). In fact, the more iconic a word is, the more abstract it tends to be. This result has been explained through the relationship between abstractness and auditory modality. Thus, the ratings might have been influenced by the auditory modality (Winter et al., 2024). Nonetheless, the relationship between abstractness and iconicity requires further investigation.

Furthermore, a recent study (de Varda, Lamarra et al., 2025) demonstrated that iconicity ratings are positively correlated with categorical specificity, suggesting that highly iconic words are also highly specific. Therefore, the link between iconic words and specific referents and contexts (Lupyan & Winter, 2018; Perry et al., 2021) may be a consequence of specificity rather than concreteness (see ¶ 3.5.1. for further details).

The following Table 2.1. summarizes how abstract and concrete concepts differ across relevant dimensions and semantic variables.

|                                      | <b>Abstract Concepts</b> | <b>Concrete Concepts</b> |
|--------------------------------------|--------------------------|--------------------------|
| <b>Dimensionality</b>                | Low                      | High                     |
| <b>Relational Character</b>          | High                     | Low                      |
| <b>Age of Acquisition (AoA)</b>      | Late                     | Early                    |
| <b>Modality of Acquisition (MoA)</b> | Primarily Linguistic     | Primarily Perceptual     |
| <b>Grounding</b>                     | Interoceptive            | Exteroceptive            |
| <b>Perceptual Strength</b>           | Weaker                   | Stronger                 |
| <b>Body-Object-Interaction (BOI)</b> | Lower                    | Higher                   |
| <b>Imageability</b>                  | Lower                    | Higher                   |
| <b>Iconicity</b>                     | Stronger                 | Weaker                   |
| <b>Contextual Availability</b>       | Lower                    | Higher                   |
| <b>Semantic Diversity</b>            | Higher                   | Lower                    |
| <b>Social Metacognition</b>          | Need others more         | Need others less         |

**Table 2.1.** Summary of dimensions and semantic variables for which concrete and abstract concepts might differ. Adapted from (Borghi, 2022).

This section has discussed relevant dimensions and semantic variables that distinguish abstract concepts from concrete concepts. These differences often contribute to a behavioral advantage of concrete words over abstract ones, a well-established phenomenon, widely observed in both behavioral, psycholinguistic, and neurocognitive studies known as the concreteness effect.

### **2.3. The Concreteness Effect: Differences in Language Processing and Representation**

As discussed in the previous section, abstract and concrete concepts differ along multiple dimensions. These distinctions include the quality and quantity of information they convey (Vigliocco et al., 2009), the processes (Villani et al., 2021), and the modalities involved in their acquisition (Della Rosa et al., 2010), as well as the richness of linguistic and perceptual experiences they evoke (Lynott et al., 2020). Moreover, words

denoting concrete concepts are, on average, acquired earlier (Della Rosa et al., 2010; Schwanenflugel, 2013; Borghi, 2022) and are more imageable than words denoting abstract concepts (Paivio, 1990). Furthermore, abstract and concrete words have structurally different representational frameworks, with concrete words organized categorically based on their semantic similarity, whereas abstract words are represented across a broader set of associative networks (Crutch & Warrington, 2005). Finally, concrete concepts are more strongly associated with perceptual aspects of meaning, whereas abstract concepts rely more heavily on linguistic and social ones, both in their acquisition and use (Borghi, 2022).

These differences are reflected in the *concreteness effect*, a behavioral advantage in which concrete words are recognized faster and more accurately than abstract words, across a variety of tasks, such as lexicon decision (e.g., “is this a word or nonword), word naming, and recall (Kounios & Holcomb, 1994; Romani et al., 2008; Jessen et al., 2000). This effect is considered a cornerstone in cognitive and psycholinguistic research on word processing and has been explained by two theoretical frameworks: the *Dual Coding Theory* (DCT, Paivio, 1991) and the *Contextual Availability Theory* (CAT, Schwanenflugel et al., 1988; Schwanenflugel & Shoben, 1983).

According to the Dual Coding Theory (DCT), word meaning is represented in two interconnected systems: a verbal system for linguistic information (*the logogen system*) and a non-verbal system that supports visual and sensory information (*the imagen system*). In the logogen system, associative networks of words are formed through word-to-word associations arising from language experience, typically activated through written and spoken language. In contrast, the imagen system encodes sensory and imagistic experiences, preserving information about the modalities from which they originated. A key aspect of this framework is that the two representational systems can operate both in parallel and in isolation during language processing, as they are linked together. Concrete concepts are assumed to activate both representational systems, whereas abstract concepts primarily engage the verbal system and are only indirectly represented in the imagen system. This representational asymmetry is thought to account for the concreteness effect.

In contrast, the Contextual Availability Theory (CAT) accounts for the interaction between external stimuli and the speaker’s knowledge. The central tenet of this

framework is that processing is facilitated by the retrieval of contextual information (Schwanenflugel, 2013). According to CAT, the distinction between concrete and abstract concepts lies not in the representational code itself, but in the degree of accessibility of relevant semantic and contextual information stored in the semantic memory. Abstract concepts, which are weakly associated with a wider range of contexts, require more effort to retrieve from memory, whereas concrete concepts are more readily processed due to their stronger associations with a limited set of more accessible contexts. Thus, the greater variability in contextual associations for abstract concepts (Hoffman, 2016) makes it more difficult to retrieve semantically relevant information from contextual background knowledge, resulting in longer decision latencies.

Löhr (2024) specified that Dual Coding Theory (DCT) distinguishes abstract and concrete words based on the type of content and format (e.g., perceptual vs. linguistic-amodal content) in which information is encoded, as well as how it is organized. In contrast, Contextual Availability Theory (CAT) posits that information is organized differently.

Numerous studies have investigated processing differences between abstract and concrete words at both behavioral and neural levels, generally supporting the concreteness effect. However, several psycholinguistic studies have reported difficulties in replicating it (Barca et al., 2002; Papagno et al., 2009), or findings remain inconsistent with expected patterns, revealing an unclear nature of this phenomenon.

For example, Schwanenflugel & Shoben (1983) and Schwanenflugel et al. (1988) found that when abstract and concrete words are embedded in a supportive, meaningful context, rather than processed as isolated words, the difference in lexical decision latencies disappears. Moreover, when words are controlled for contextual availability, the concreteness effect is no longer observed, suggesting that contextual availability, rather than concreteness per se, may be considered the primary factor influencing response times (Schwanenflugel et al., 1988).

More recently, Connell & Lynott (2012) found that in both lexical decision and word naming tasks, strength in the dominant perceptual modality better explains the variance of decision latencies and accuracy than both concreteness and imageability, a key variable related to DCT. These findings have been interpreted as evidence that perceptual strength better accounts for the concreteness effect than concreteness itself.

Therefore, perceptual strength ratings reflect specific perceptual modality involved during the experience and reenactment of the referent concept, thereby influencing language processing.

Some studies have reported evidence supporting both theories. For example, Jessen et al. (2000) found that participants recognized concrete words more accurately than abstract words. Brain activity measurements revealed that concrete words elicited greater activation in brain regions involved in spatial processing (rather than visual imagery) and areas linked to verbal processing. In contrast, abstract words primarily activated regions associated with speech production and more complex visual processing. These findings suggest that both DCT and CAT may contribute to explaining the concreteness effect.

In some cases, abstract concepts have even shown a processing advantage over concrete ones. This pattern has been observed in two lexical decision tasks, where the stimuli were controlled for imageability and contextual availability, i.e., key psycholinguistic variables associated with DCT and CAT, respectively (Kousta et al., 2011; Barber et al., 2013). Kousta et al. (2011) observed that the effects of concreteness and imageability disappear when valence is taken into account. Specifically, Kousta et al. (2011) conducted two regression analyses on reaction time data from the English Lexicon Project (ELP; Balota et al., 2007), examining a set of 1,446 words. They found that when word valence was included in the regression, neither concreteness nor imageability significantly predicted response latencies. In contrast, valence emerged as a significant predictor, indicating a processing advantage for emotional words.

Furthermore, Kousta et al. (2011) and Vigliocco, Kousta, et al. (2014) reported a reversed effect when concrete and abstract concepts are matched on multiple psycholinguistic variables, with abstract concepts exhibiting a processing advantage. This reversed effect has been attributed to the stronger affective associations that typically characterize abstract concepts compared to concrete ones (Vigliocco, Kousta, et al., 2014).

Overall, the concreteness effect remains a phenomenon with several unresolved aspects, and alternative explanations have provided additional insights into its underlying mechanisms. As highlighted above, a promising approach to investigating this subject involves shifting the focus from a single dimension, such as concreteness, to other

variables that may collectively or additively contribute to generating a potential processing advantage.

As discussed, Dual Coding Theory posits an unbalanced role for linguistic and perceptual experiences in contributing to the semantic grounding of abstract and concrete words. In contrast, Contextual Availability Theory claims that the different organization of contextual information shapes the representation of abstract and concrete words. Together, these theories, along with the heterogeneous evidence arising from the study of the concreteness effect, highlight an ongoing debate regarding whether and to what extent word meanings are grounded primarily in perceptual experience, in linguistic associations and recurring co-occurrences of words, or in the interaction between both sources of information. A debate that emerged from the questions raised around the well-known *symbol grounding problem* (Harnad, 1990).

## **2.4. The Symbol Grounding Problem**

In previous sections, conceptual knowledge and word meanings were discussed as grounded in perceptual experience and the sensorimotor system, in interaction with both physical and cultural environments. Moreover, language use can reenact specific perceptual experiences or contextual information stored in memory. These claims are typically associated with the Embodied and Grounded Cognition (EGC) account (Lakoff & Johnson, 1980; Gallese & Lakoff, 2005; Borghi & Cimatti, 2010; Glenberg et al., 2013; Bolognesi, 2025). Therefore, concepts and words' meaning consist of modal representations, as different types of sensory-based information contribute to their representational format.

However, during the early decades of cognitivism, mental operations and, more broadly, cognition were equated with formal operations akin to calculus. This view was based on the Turing Machine, a mathematical model of computation that manipulates symbols according to predefined rules (Turing, 1937). Within this framework, symbols represent objects and relations in the world but were not derived from perceptual experiences; rather, they were arbitrary and amodal representations, which could be combined and processed according to a set of formal rules (see Physical Symbol System Hypothesis; Newell & Simon, 1972). Theoretical frameworks rooted in this view often

adopt the mind-computer analogy, where the brain is likened to hardware and cognitive functions to software.

One notable example of this symbolic approach is Chomsky's *Generative Grammar*, which posits the existence of a *Universal Grammar* (UG), a set of formal, syntactic, and universal principles underlying an innate linguistic *competence* (i.e., the unconscious knowledge that speakers have of their language; Chomsky, 1957, 1968, 1995; Evans & Green, 2006). According to Chomsky, language is an independent module of the mind, with a structure that can be understood in computational terms. This internal language system (i.e., *I-language*) operates based on recursive mechanisms provided by Universal Grammar to generate an infinite number of combinations of symbols (e.g., sentences). In this framework, meaning does not emerge from perceptual experience or specific sensory modalities (e.g., auditory, visual, haptic), but rather from the combinatory rules of innate syntactic structures.

Another account that supports the idea of an amodal symbolic system is Fodor's theory of the *modularity of the mind* (Fodor, 1983, 1985). Fodor's theory argues that the mind is composed of distinct, domain-specific modules, each specialized for particular cognitive functions. This framework emphasizes the separation between perception and cognition, which operate according to different principles (Evans & Green, 2006). Within Fodor's theory, the mind is structured in the following mental mechanisms:

- Transducers, which receive inputs from the perceptual level and convert them into a different manipulable format for other cognitive systems;
- Central systems, where primary cognitive processes operate, such as memory and inferential reasoning;
- Modules, which mediate between transducers and central systems.

The modules are domain-specific systems within which the mind is organized. Language consists of an innate and encapsulated module processing information independently of other modules (Robbins & Drayson, 2025). More specifically, language receives transduced information from the perceptual level (e.g., auditory or visual information) and processes it according to its domain-specific syntactic rules. Crucially, modality-specific sensory information is transduced into propositional symbols that can be manipulated and combined according to formal rules within the linguistic module, while sensory details, such as traces of specific sounds or visual cues, are largely omitted.

This process aligns with Fodor's idea of the *language of thought* hypothesis or *Mentalese*, a symbolic system in which information is translated into amodal propositions that can be manipulated based on its own syntax (Fodor, 1975).

In these frameworks, meaning does not emerge from semantic or pragmatic properties or sensory experiences, but rather from the manipulation of amodal symbols according to syntactic rules and formal recursive mechanisms. For instance, information about rabbits in the world would be translated into the symbol "rabbit", with its meaning derived from combinations with other symbols (e.g., "jump", "fur", "long ears") within syntactic structures, without any reference to specific exemplars or specific sensory-derived information (e.g., the softness of the fur, the tone of the cry, or an image of them jumping).

Such theories were challenged by the questions raised within the *symbol grounding problem* (Harnad, 1990, p. 335):

How can the semantic interpretation of a formal symbol system be made intrinsic to the system, rather than just parasitic on the meanings in our heads? How can the meanings of the meaningless symbol tokens, manipulated solely on the basis of their (arbitrary) shapes, be grounded in anything but other meaningless symbols? [...] How is symbol meaning to be grounded in something other than just more meaningless symbols? This is the symbol grounding problem.

Harnad (1990) argued against traditional symbolic systems, positing that meanings cannot be derived solely from combinations of "meaningless" symbols but should instead be grounded in projections and analogical mechanisms tied to sensory experience. To support this claim, Harnad (1990) advanced two key arguments: the well-known Chinese Room argument (Searle, 1980) and the Chinese/Chinese Dictionary-Go-Round argument.

The Chinese Room is a thought experiment formulated by Searle to challenge the concept of strong AI, to demonstrate that AI cannot possess intrinsic intentionality, which involves certain conscious mental states arising from brain activity. The notion of intentionality, borrowed from Brentano's (Brentano, 1874/1973) concept of *aboutness*, refers to the ability of thoughts, beliefs, and other mental states to be directed at objects or others' mental states, often expressed through propositional content (e.g., "X believes that P"; "X believes that Y thinks that P"). Searle's central argument unfolds as follows: Searle imagines himself in a room, receiving written messages in Chinese from another room. Searle does not understand Chinese, but he follows a program with English

instructions to manipulate Chinese symbols according to syntactic rules in order to produce appropriate responses. The critical point raised by Searle, and later echoed by Harnad, is that manipulating symbols in accordance with syntax alone does not result in true understanding. Although Searle can generate appropriate sentences by following formal syntactic rules, the Chinese characters and logographs remain meaningless to him. Searle's conclusion against strong AI is that machines are in the same situation as he in the Chinese room: while they may manipulate symbols and produce sentences that appear syntactically correct, they do not possess a true understanding of their semantic content.

Harnad (1990) enforced this point by providing a related thought experiment. In this example, one is asked to learn Chinese, but the only source of information available is a Chinese/Chinese dictionary, in which meaningless symbols (the *definientes*) are paired with other meaningless symbols (*definienda*). According to Harnad, this scenario, along with the Chinese Room argument, illustrates the impossibility of acquiring meaning without grounding symbolic representation, either directly or indirectly, in perceptual reference.

#### **2.4.1. Embodied vs. Distributional Meaning**

The symbol grounding problem has recently been invoked in ongoing debates (see, for instance, de Vega et al., 2008) concerning the types of information that contribute to shaping word meaning and conceptual representation, particularly whether such meaning is perceptually based or contextually derived. Within this framework, supporters of grounded and embodied cognition, who argued that the meaning of symbols should be grounded in bodily systems involved in perception, action, and affective states, are opposed to supporters of *Distributional Semantics* (Lenci & Sahlgren, 2023), who claim that word meaning may arise from relations among words (Firth, 1957; Bolognesi, 2020).

Distributional semantics presents a disembodied and, therefore, amodal alternative to the embodied and grounded perspective. It is based on the distributional hypothesis (Harris, 1954), which posits that certain aspects of word meaning emerge from distributional properties, namely, associations among words, linguistic contextual information, and patterns of statistical co-occurrences across large corpora of texts. Firth (1957) encapsulated this idea in the principle that “you shall know a word by the company it keeps”.

Generally, semantic representations in distributional models are generated in a bottom-up manner from the automatic acquisition of contextual information from and across corpora (Bolognesi, 2020). One of the most popular distributional models for representing word meaning is *Latent Semantic Analysis* (LSA; Landauer & Dumais, 1997). LSA is a mathematical model based on the principle that similar words tend to occur in similar contexts. Word meaning, in this framework, is understood as emerging from the distribution across contexts and occurrences within a corpus.

LSA models represent the distributional, semantic, and contextual similarity among words as quantified measures of geometrical proximity within multidimensional spaces (e.g., spaces of one, two, or three dimensions). For instance, words like “orange”, “bread”, “knife” would cluster closely together in an  $n$ -dimensional space, while the word “lamp” would be positioned farther away. Furthermore, based on semantic similarity, the distance between “bread” and “knife” will be smaller than the distance between “bread” and “lamp” or between “knife” and “lamp”.

In these models, word meaning is mathematically represented as a set of numerical values that describe the relations among words across their contexts. Therefore, the position of a word within one of the possible dimensions is determined by numerical coordinates. The vector, instead, is the sequence of coordinates that indicates the position of a point in a multidimensional space. The proximity between two vectors, namely the semantic similarity between two words, is quantified by calculating the cosine between the two vectors. Thus, words are numerically represented as vectors, describing their positions in relation to each other, and reflecting the linguistic contexts in which words occur.

LSA models primarily rely on frequencies of word occurrences within a corpus, whereas new generations of models, such as *word embeddings*, operate based on probability measures of expected occurrences (see Bolognesi, 2020, for an extensive discussion).

Supporters of grounded and embodied cognition oppose this perspective by raising similar questions that form the basis of the symbol grounding problem: How can word meaning and conceptual representation emerge from patterns of linguistic associations and contextual information, rather than being grounded in the sensorimotor system and multimodal perception distributed across the brain’s regions? However, as

Lenci (2008) suggested, part of word knowledge is linguistically derived through word associations and linguistic contexts. For instance, while one can form a conceptual representation such as DOG based on perceptual experience and information extracted from multiple exemplars, the knowledge that “dog is an animal” is largely linguistic and acquired through language itself.

The contribution of both perceptually and linguistically derived information has been further illustrated by Marconi (1997) in his account of *referential* and *inferential competence* (Lenci, 2008; Lenci et al., 2018). Referential competence involves various referential abilities (e.g., classification, recognition), including the ability to map words to appropriate referents in the world. Inferential competence, on the other hand, involves the ability to relate words to one another within the semantic lexicon through a set of inferences and logical connections.

In recent years, several hybrid accounts have emerged, positing, albeit to varying extents, that word meaning and the semantic representation of concepts are shaped by both amodal linguistic experience and perceptual experience.

#### **2.4.2. Hybrid Theories of Word Meaning**

An account worth mentioning is the *Symbol Interdependency Hypothesis* (Louwerse, 2007, 2011, 2018; Louwerse & Jeuniaux, 2008), which posits that cognition is both symbolic (i.e., amodal) and embodied (i.e., modality-specific). According to this hypothesis, language comprehension depends on both the interdependencies among linguistic, amodal symbols and the grounding of these symbols in the perceptual system. Within this framework, language plays a crucial role in cognition by encoding relations in the world (as discussed in ¶ 2.1.2), including embodied relations (Louwerse, 2008) and perceptual information.

This idea draws on previous theoretical accounts, notably the classical Peircean distinction between icons, indexes, and symbols as distinct types of signs (Peirce, 1931: 2.92). It also builds on the discussion of the hierarchical contribution of iconic, indexical, and symbolic processes in language processing (Deacon, 1997). Deacon (1997) proposed that symbols, which are arbitrary and conventional in Peircean semiotics, ground their meaning through indexical and iconic relations, those that exhibit physical contiguity and

natural resemblance with referents, respectively, and through relationships with other symbols (Louwerse, 2018). This view aligns with De Saussure's (1916/1966) discussion about a sign's dependence on its relationships to other signs, metaphorically represented by a chess piece's position on the chessboard.

Interestingly, Harnad (1990) also proposed a similar hybrid account emphasizing the role of iconic representations as the perceptual basis for the emergence of distinct categories and symbolic representations, thus offering a solution to the symbol grounding problem. More recently, Cuccio & Gallese (2018) proposed that embodied simulations can be understood as an iconic, abstract schematization, with denotative relations arising from the encoded similarity to referents and instances.

According to the Symbol Interdependency Hypothesis, language evolved as a cognitive mechanism essential for communication (Louwerse, 2011; Enfield, 2024). Thus, language is understood as a shortcut to meaning through interdependencies among lexical items. Although grounding is essential in the formation of a symbolic system, language processing does not always require activating perceptual information and more cognitively expensive simulations. Language can bootstrap encoded perceptual information through statistical regularities among words. For instance, Louwerse & Connell (2011) demonstrated that linguistic contexts and co-occurrences predict categories of perceptual modality (e.g., auditory, visual, haptic, olfactory, gustatory), albeit with less precision if compared with human-generated modality norms. This finding supports the idea that meaning can be induced not only through direct grounding to the referents of lexical items, but also indirectly, through semantic approximations emerging from networks of linguistic contexts and related words (in a similar vein to the notions of referential and inferential competence, as discussed in ¶ 2.4.1.).

Moreover, linguistic information more accurately explained faster response times in a feature recognition task involving items sharing or shifting perceptual modalities, while perceptual information more accurately explained slower reaction times (Louwerse & Connell, 2011). These results suggested that representations based on semantic approximations enable linguistic information to precede perceptual information in language processing. As a matter of fact, deeper conceptual processing in the simulation system incurs a processing cost. These findings confirmed previous research. Louwerse & Jeuniaux (2010) found that the contribution of linguistic and perceptual information in

language and conceptual processing is task-dependent, with linguistic information contributing more in shallow processing and perceptual information playing a larger role in deeper conceptual processing.

A pluralist perspective, proposed by Zwaan (2014), further suggests that the contribution of both embodied and linguistic experiences to language comprehension depends on the extent to which language use is embedded in the environment. Each level of embeddedness corresponds to a distinct type of context: *demonstration*, *instruction*, *projection*, *displacement*, and *abstraction*. While both linguistic and embodied experiences are simultaneously activated during comprehension, the relative contributions of these two types of information vary across different types of contexts along a scale. Specifically, demonstration tends to rely more on embodied representations, whereas the abstraction level is primarily based on amodal ones.

Empirical evidence further supports the view that both perceptual and linguistic experiences serve as distinct sources of information, jointly contributing to a richer semantic representation of words and concepts. Andrews et al. (2009) developed three main models of semantic representation that differ in terms of types of data utilized: one based solely on perceptual information (the experiential model), one based solely on linguistic information (the language-based model), and a hybrid model that integrates both perceptual and linguistic data (the combined model). Qualitative comparisons of correlations reveal qualitative differences between the experiential and language-based models in terms of neighborhood structures, with a notably low correlation between the two. These findings suggested that semantic representations derived from perceptual-based and linguistic-based information differ in their nature. Furthermore, model comparisons were also conducted to assess the efficacy of these models in predicting semantic effects in behavioral tasks. Results indicated that the combined model provided a more accurate approximation of semantic effects, suggesting that semantic representations are likely shaped by a statistical integration of perceptual, affective, and linguistic experiences.

Finally, Vigliocco et al. (2009) proposed an integrative theoretical framework, according to which the representation of both concrete and abstract concepts and words is shaped by experiential information, derived from sensorimotor and affective experiences, as well as linguistic information, emerging from word associations, co-

occurrence patterns, and syntactic structures. The core assumption of this framework is that the semantic distinction between abstract and concrete words is primarily a function of the different statistical distribution and proportion of perceptual and linguistic information. Specifically, sensorimotor experiences are more predominant in the representation of concrete words and concepts, whereas affective and linguistic experiences play a larger role in the representation of abstract words and concepts.

In summary, both theoretical proposals and experimental evidence support the view that linguistic and perceptual information jointly contribute to conceptual representation and word meaning. The following section explores various theoretical approaches that investigate the extent to which these types of information are involved in the grounding and representation of abstract concepts.

## **2.5. The Challenge of Grounding Abstract Concepts**

As discussed in ¶ 2.1., words that denote abstract and concrete concepts typically differ in terms of accessibility and perceptibility of their referents. Abstract words, such as “linguistics” or “invincibility”, lack a precise and clearly perceivable referent. In contrast, concrete words more readily map onto a single potential physical referent, as in the case of words like “wolf” or “yacht”. However, abstract words do not constitute an anomaly in the lexicon; rather, they are ubiquitous (Lupyan & Winter, 2018) and reflect one of the most sophisticated cognitive abilities of the human species.

For these reasons, and due to the detachment that characterizes abstract words and related concepts, they have become the main topic of an ongoing debate within embodied approaches to meaning and conceptual representation. This debate has been aptly summarized as addressing “the challenge or challenges of abstract concepts” (Borghi et al., 2017; Dove, 2021). Embodied cognition posits that meaning derives from sensorimotor, emotional, and perceptual systems involved in the interaction with the physical and cultural world. The abstractness that defines abstract words and related concepts, therefore, raises a critical question concerning how and to what extent such concepts can be grounded in experience.

Within this chapter, accounts and experimental evidence have already been presented in support of the grounding of abstract concepts in the perceptual system.

Abstract concepts have been argued to possess a relational character (Asmuth & Gentner, 2017; Gentner & Asmuth, 2017), linking to other concepts, intrinsic elements of situations represented (Barsalou et al., 2018), and to rely more heavily on introspective states and interoceptive experience (Connell et al., 2018; Villani et al., 2021). Moreover, the previous section emphasized how semantic representations are shaped by both perceptual and linguistic experience, with Vigliocco et al. (2009) proposing that linguistically derived and affective experiences jointly contribute to the formation of abstract conceptual representations.

The present section discusses four of the most valuable theoretical proposals that account for how abstract concepts are grounded within perceptual and embodied systems.

### **2.5.1. Conceptual Metaphor Theory (CMT)**

*Conceptual Metaphor Theory* (CMT; Lakoff & Johnson, 1980), one of the most influential approaches in the study of embodied cognition, emerged from linguistics and the analysis of metaphorical language, providing a framework for understanding how abstract concepts are structured through perceptual and bodily experience. According to this theory, metaphors are not merely rhetorical expressions but fundamental cognitive devices that contribute to the formation of conceptual representations. Abstract concepts are indirectly grounded in the sensorimotor system and perception through mappings onto more concrete concepts that are accessible through perceptual experience (Lakoff & Johnson, 1999). Through metaphorical mappings, knowledge from concrete, bodily-based, and affective source domains is mapped onto abstract target domains, as illustrated by the classical metaphor “life is a journey”.

Metaphorical mapping may operate on meaningful schematic representations of recurring sensory and motor experiences, known as *image-schemas* (Johnson, 1987). These schemas are preconceptual, relational structures that arise from recurrent sensorimotor interactions and patterns of engagement with the environment. Common examples include CENTER-PERIPHERY, UP-DOWN, and SOURCE-PATH-GOAL.

Through metaphorical language, abstract concepts become more comprehensible by being structured in terms of concrete, bodily grounded concepts. According to the speculative *Neural Theory of Metaphor* (Lakoff, 2008), metaphors are also embodied

insofar as systematic correlations in experience (e.g., feeling warmth while experiencing affection), which are linguistically expressed through metaphors (e.g., “Affection Is Warmth”), give rise to neural connections between brain areas involved in sensorimotor processing and those supporting abstract reasoning. These connections are assumed to operate by simulation mechanisms. As a consequence, activating an abstract concept may partially and indirectly activate the neural structures associated with the corresponding embodied experience. However, Cuccio (2018) argued that embodied simulations are not always involved during metaphor processing (see also Bolognesi, 2020).

Within psycholinguistic research, one of the most extensively investigated metaphorical mappings concerns the conceptualization of the abstract domain of time in terms of the more concrete domain of space (Boroditsky & Ramscar, 2002; Casasanto & Boroditsky, 2008).

Nonetheless, as noted by Borghi, although conceptual metaphors may facilitate the understanding, and possibly the acquisition, of abstract concepts, they are not necessarily constitutive of the representation of abstract concepts.

### **2.5.2. Affective Embodiment Account (AEA)**

According to the *Affective Embodiment Account* (AEA), abstract words are more strongly associated with emotional states than concrete words. Kousta et al. (2011) proposed that different types of experiences contribute to conceptual representation to varying extents. While sensorimotor experiences play a predominant role in the representation of concrete words, affective experience and emotional development are assumed to be particularly relevant for abstract concepts. In support of this view, it has been argued that emotionally charged abstract words are acquired earlier than emotionally neutral abstract words. Accordingly, a central assumption of the AEA is that emotional experience facilitates the acquisition of abstract words (Vigliocco et al., 2013).

As already discussed in ¶ 2.3., empirical evidence indicates that, when potential confounding variables are controlled (e.g., imageability, contextual availability), emotional valence accounts for a processing advantage of abstract concepts over concrete concepts in lexical decision tasks (Kousta et al., 2011). Moreover, during lexical-semantic processing, highly valenced abstract words have been shown to elicit stronger activation

in the rostral anterior cingulate cortex (rACC) compared to concrete words. Given the regulatory role of the rACC in emotional processing, these findings have been interpreted as supporting the claim that the processing of abstract concepts relies more heavily on affective mechanisms (Vigliocco, Koustas, et al., 2014).

This influential theoretical account emphasized the role of emotions in bootstrapping the development of abstract concepts. Nevertheless, empirical findings suggest the need for further investigation. For example, recent cross-linguistic evidence indicates that only a subset of abstract words was rated as more emotionally loaded (either positively or negatively) than concrete words. Furthermore, additional analyses revealed that when the concreteness/abstractness dimension is not operationalized as ratings, concrete words are evaluated as more emotionally charged than abstract words (Winter, 2023).

### **2.5.3. Language and Situated Simulation (LASS)**

This theoretical framework proposes that conceptual representations are supported by multiple interacting systems. In particular, the Language and Situated Simulation (LASS) theory focuses on the simulation and linguistic system, which interact dynamically during language processing. As discussed in ¶ 2.4.2., in tasks involving relatively shallow processing, such as lexical decision tasks, the linguistic system is activated earlier than the simulation system, providing access to conceptual representations based on linguistic associations and information. By contrast, when deeper semantic processing is required, the simulation system becomes more strongly engaged, enacting sensorimotor states associated with interactions with referents (Barsalou et al., 2008; Connell & Lynott, 2013; Louwerse & Connell, 2011).

Therefore, LASS theory suggests that language and linguistic information can function as a shortcut to conceptual representation, particularly in tasks with limited semantic demands. Although the LASS was not originally formulated to account for the grounding of abstract concepts (Borghi et al., 2017), it offers a valuable multiple-representation framework in which both concrete and abstract concepts are represented through the flexible, task-dependent activation of linguistic and simulation-based information.

#### 2.5.4. Words as Social Tools

According to the *Words as Social Tools* (WAT) theory (Borghi & Binkofski, 2014; Borghi et al., 2019), conceptual representation is shaped by both perceptual/motor and linguistic experiences. In the case of abstract concepts, linguistic experience plays a crucial role. Moreover, the WAT theory posits that linguistic experience encompasses not only word associations but also the full spectrum of information conveyed through language, including, for instance, syntactic structure (as seen for frames and domains in ¶ 2.1.2.). Crucially, WAT emphasizes the role of social interaction and the broader social dimension in conceptual representation, particularly relevant for abstract concepts (for a derived interpretation of the WAT theory, see the *Word as Cultivators* proposal; Schilhab, 2015; see also Chapter 7).

A central assumption of WAT theory (Borghi & Binkofski, 2014; Borghi, 2023) is to examine the difference underlying the representation of concrete and abstract concepts with a particular focus on the heterogeneity of their semantic subcategories (Falcinelli et al., 2024; Villani et al., 2019; Borghi, 2023). While concrete concepts have traditionally been classified into semantic domains such as natural and artifact categories (Keil, 1992), recent analyses indicate that abstract concepts can also be grouped into distinct semantic domains, rather than being treated as a monolithic whole. For example, a cluster analysis of 425 words (Villani et al., 2019) identified four subcategories of abstract concepts: concepts referring to principles, disciplines or religious domains, *Philosophical and Spiritual* (PS, e.g., ETHICS); concepts encompassing emotionally charged states or institutions, *Emotional and Mental State* (EMS, e.g., SERENITY); concepts denoting psychological and physical characteristics of the self, or notions related to social institutions, *Social and Self* (SS, e.g., FRIENDSHIP); and concepts referring to quantifiable physical notions, *Physical Space Time and Quantity* (PSTQ, e.g., PERIMETER).

These subcategories differ in multiple dimensions, including abstractness, perceptual strength, and inner grounding (Villani et al., 2019, 2021). Additionally, Falcinelli et al. (2024) examined Ecological and Technological concepts, highlighting their “hybrid” semantic nature with respect to concreteness and abstractness.

The WAT theory is structured around three main tenets, each divided into relative subtenets (Borghi, 2023). The first tenet posits that both abstract and concrete concepts are embodied and grounded in simulations and situated experiences, although specific dimensions may be more crucial for certain subtypes. As discussed in ¶ 2.2., abstract concepts are grounded in sensorimotor experiences, though these are less prominent than in concrete ones. In contrast, interoceptive states and inner grounding contribute more strongly to the representation of abstract concepts.

The second tenet and relative subtenets emphasize the central role of language and social interaction in the acquisition, use, and brain representation of concepts. Linguistic and social experience are highly relevant for abstract concepts, which lack a single, denotable referent, compared to concrete ones. Although both sensorimotor and linguistic experiences contribute to conceptual representation, research on modality of acquisition (MoA; Wauters et al., 2003; see also ¶ 2.2.1.) suggests that linguistic experience is crucial in the case of abstract concepts.

With concrete concepts, it is relatively straightforward to abstract shared perceptual features from multiple referents to form a category (e.g., COUCH). In contrast, the complexity and heterogeneity of abstract concepts make this operation more difficult, increasing. This high dimensionality (Lupyan & Mirman, 2013) of abstract concepts and the difficulties in retrieving a single physical referent likely increase reliance on linguistic information. Moreover, the WAT theory is that the linguistic experience also promotes sophisticated social abilities (Tomasello & Akhtar, 1995). For example, children likely learn abstract words like “future” relying more on linguistic inputs, verbal explanations, and interaction within social contexts. This reliance on others to establish and converge on word meaning, evoking more social interaction, has been defined as social metacognition (Borghi, 2022; Borghi et al., 2018b; for a further discussion, see ¶ 2.6.), a dimension particularly relevant for the conceptual representation of abstract words.

Finally, the third tenet asserts that the heterogeneity of abstract concepts results in a greater semantic diversity (Hoffman, 2016), contextual and situational variability (Davis et al., 2020; as discussed in ¶ 2.2.1.), and linguistic variability across languages (Borghi & Binkofski, 2014; Mazzuca et al., 2020).

While each of these theories provides valuable insights into the grounding of abstract concepts, the AEA and WAT frameworks particularly emphasize the crucial role of affective and linguistic experiences in the conceptual representation of abstract concepts. Proponents of AEA have recently highlighted the role of communicative contexts and social interactions in the acquisition of concepts, suggesting that emotions shared between caregivers and children may support word learning. Notably, as a theory of conceptual grounding, the WAT theory posits that linguistic contextual information promotes greater engagement in social interaction and metacognitive processes, which are fundamental to the acquisition and use of abstract concepts.

In the following section, this social dimension of abstract concepts is discussed in detail.

## **2.6. The Social Dimension of Abstractness**

The complexity and heterogeneity of abstract concepts are closely related to the abstractness that characterizes them. As discussed in the previous section, the WAT theory proposes that this heterogeneity and complexity, along with high dimensionality, namely, the reference to multiple entities that share few relevant features, lead speakers to rely more heavily on linguistic information (e.g., word associations and contextual co-occurrences) and to engage more in social interaction (Borghi & Binkofski, 2014).

Proponents of the WAT framework (Borghi & Binkofski, 2014; Borghi, 2023) argue that the referential indeterminacy of abstract concepts fosters the development of metacognitive abilities. During the processing of abstract concepts, which lack a single perceptual referent, one may reenact linguistically based information stored during the acquisition of abstract concepts, thereby activating networks of associated words and related contexts. Borghi & Binkofski (2014) and Borghi (2023) proposed that this process may occur involving inner speech (Vygotsky, 1934/1986), understood as an internalized language-based mechanism which involves self-directed communication. When this process fails to yield a satisfactory meaning or definition for an abstract word, one may turn to external sources of information, such as spoken or written language. In this sense, uncertainty regarding an abstract word's meaning prompts greater reliance on others. This process has been described as social metacognition (Borghi et al., 2018a, 2019).

Empirical findings are consistent with this proposal. For example, abstract concepts have been shown to elicit more prominent activation of the mouth motor system than concrete concepts, suggesting the reenactment of linguistic information stored during acquisition (Borghi & Zarcone, 2016; Fernyhough & Borghi, 2023).

In spoken language and social interaction, dialogue constitutes a joint activity (Clark, 1996; Pickering & Garrod, 2021) in which interlocutors may negotiate the meaning of words. Importantly, such interaction may be asymmetric, with interlocutors differing in their degree of knowledge. Consequently, one develops metacognitive processes (Shea, 2018) to infer and understand others' mental states and converge in their representations and meanings.

Recent empirical evidence further supports the idea that abstract concepts engage social mechanisms to a greater extent than concrete concepts. In rating tasks, abstract words are typically associated with lower scores in *word confidence*, that is, lower certainty regarding a word's meaning, than concrete words (Mazzuca et al., 2022; Falcinelli et al., 2024). Moreover, abstractness has been found to correlate positively with both *socialness*, namely, the social relevance of a concept (Diveica et al., 2022; P. M. Pexman et al., 2023), and social metacognition (Villani et al., 2019). In addition, interlocutors tend to be more inclined to negotiate the meaning of abstract concepts and start a conversation with them compared to concrete concepts (Fini et al., 2023; Mazzuca & Santarelli, 2023; Falcinelli et al., 2024).

Given this close relationship between social interaction and abstractness, recent research has begun to examine mechanisms underlying online conversations involving concrete and abstract words in experimentally designed, situated contexts. One notable example comes from a free production task in which participants were asked to pretend to be involved in real conversations and to reply to 60 sentences varying in abstractness/concreteness (30 abstract and 30 concrete; e.g., "I bought strawberries"; "I found a remedy"). The nature of participants' responses differed as a function of the target sentences' abstractness/concreteness. Specifically, sentences including abstract concepts elicited more expressions indicating uncertainty and a greater need for social validation (e.g., "And then?"), as well as more questions related to causality and agency (e.g., *why* and *who* questions). In contrast, sentences with concrete concepts evoked more spatio-

temporal questions (e.g., *when* and *where* questions) and expressions signaling comprehension (e.g., “Ok”, “Me too”; Villani et al., 2022).

A subsequent study by Mazzuca, Villani et al. (2025) sought to replicate these findings on social mechanisms underlying engagement with abstract and concrete sentences, using a more controlled paradigm and collecting behavioral measures (i.e., reaction times). In this study, participants were asked to judge the plausibility of realistic linguistic exchanges composed of the same target sentences used in Villani et al. (2022), paired with follow-up expressions consisting of actual responses provided by participants in that earlier study. These follow-ups included *backchannels*, as well as *open* and *restricted repairs* (Dideriksen et al., 2019), and were annotated according to their communicative function: signaling uncertainty (e.g., “What do you mean?”) or certainty (e.g., “Good”), expressing a need for additional information and a desire to continue the conversation (i.e., curiosity; e.g., “Tell me more”), signaling a willingness to end the interaction (e.g., “Fine”), or posing causal (e.g., “for whom?”; “why?”) or spatio-temporal questions (e.g., “when?”; “where?”).

Across three different experiments, Mazzuca, Villani et al. (2025) explored whether abstract sentences were considered more plausible if paired with follow-ups signaling uncertainty, curiosity (i.e., the need for clarification), and causal questions. Conversely, concrete sentences were predicted to be judged as more plausible when followed by expressions signaling certainty, the willingness to conclude the conversation, and questions related to the spatio-temporal dimension. The results largely replicated the findings of Villani et al. (2022), showing that abstractness is associated with the linguistic constructs of uncertainty and curiosity. This pattern was partially supported by slower responses when participants judged abstract sentences paired with uncertainty and curiosity follow-ups as implausible, compared to concrete sentences. However, in this specific task, spatio-temporal questions were judged as more plausible than causal questions when paired with both abstract and concrete sentences, even though abstract sentences were generally evaluated as less plausible than concrete sentences.

As discussed by Mazzuca, Villani et al. (2025), this apparent discrepancy between the two studies may be attributable to the task-oriented context (Yee & Thompson-Schill, 2016). Whereas a production task allows a wider range of associations and response patterns to emerge, a plausibility judgement task may constrain the activation of semantic

information to a binary choice (e.g., “yes” or “no”), thereby restricting the potential range of relevant semantic information. An additional explanation is that, although abstract and concrete words differ across contexts and situations, with abstract words typically embedded in more heterogeneous situations and contexts than concrete words, words referring to concepts generally tend to evoke situations that could plausibly occur in real experience, that is, within space and time (Davis et al., 2020). Consequently, while abstract sentences in a production task may elicit *why* and *who* questions, reflecting the relevance of agency and causality for their representation, in a plausibility judgement task, where the focus is on the potential validity of a realistic conversation in which an event is described, *where* and *when* questions may be considered as more plausible responses because they are functional to describe events or situations within a spatio-temporal dimension.

Taken together, these findings support the view that abstractness constitutes a semantic variable characterized by intrinsic complexity, which in turn prompts more social engagement and metacognitive abilities compared to concreteness.

However, within this framework, words and sentences varying in concreteness and abstractness have not yet been examined with respect to categorical specificity. As noted by Bolognesi et al. (2020), specificity characterizes both concrete (e.g., “furniture”, “table”) and abstract words (e.g., “quantity”, “number”), highlighting the need to investigate specificity as an independent semantic dimension alongside concreteness and abstractness.

The following section, therefore, addresses the so far neglected variable of categorical specificity within cognitive and psycholinguistic research.

## Chapter 3

### The Neglected Variable: Categorical Specificity

#### 3.1. Taxonomic Relations in Linguistics and Psycholinguistics


Categorical specificity defines word meaning as belonging to a particular position within a hierarchical structure at different levels of semantic abstraction and precision. More specific words are placed at lower levels of this hierarchy, typically denoting categories that refer to fewer or single members. In contrast, more general words reflect categories that encompass multiple exemplars. This hierarchical structure is “imposed” in this sense relation underlying the connection between specific and general words (Lyons, 1977), that is, *hyponymy/hypernymy* (see also ¶ 1.1.5.).

Within semantics and cognitive linguistics, the term *hypernymy* is less commonly used than *hyponymy* (Cruse, 2017), and the term *superordination* is sometimes adopted (Lyons, 1977; Leech, 1981). In the present section, however, *hypernym* is used for two reasons: (1) it is the correlative term of *hyponym*, and (2) it avoids potential confusion with the term *superordinate*, which will be discussed in the following section (see ¶ 3.2.).

In linguistics, *hyponymy/hypernymy* is defined as a paradigmatic semantic relation conveying a logic of class inclusion between pairs of words (Lyons, 1977; Leech, 1981; Croft & Cruse, 2004). A lexeme like “daffodil” is a hyponym, a more specific lexeme, of “flower”, which in turn is its hypernym or more general lexeme. Importantly, this relation is not absolute, but relative. For example, “flower” is a hyponym of the more general word “plant”, which, in turn, has “living being” as its hypernym. As this relation is based on a logic of class inclusion, hyponyms are members or subclasses (Lyons, 1977; Leech, 1981) included in a broader class (or category in conceptual terms). Thus, the relation between hypernym and hyponym is asymmetrical. The class (or category) FLOWER includes ROSE; therefore, the proposition “all roses are flowers” is true, while the opposite, “all flowers are roses”, is false. In fact, the class FLOWER is more inclusive than ROSE and can include other types of FLOWER, such as DAFFODIL, DAISY, and TULIP.

Lyons (1977) explains that inclusion can be understood in nuanced senses, whether referring to extension (*denotation*) or intension (*connotation*). As discussed in ¶ 1.2.3., extension refers to the entire set of entities to which a class applies, while intension consists of the set of attributes or features for which a term is predicable. Therefore, considering extensional inclusion, the lexical item “flower” is more “inclusive than “tulip”; however, “tulip” might be more inclusive than “flower” because “tulip” includes all the defining properties or features of “flower” and certain distinctive properties of its class (Lyons, 1977). Nonetheless, as stated in ¶ 2.1., in the present work, inclusion is interpreted in its extensional sense as the extent of the scope of its extension or denotation, that is, a referential inclusion (Leech, 1981).

The relation of inclusion is encoded in the hierarchical structure, which can typically be represented with a branching tree diagram (Lyons, 1977; Cruse, 2017), as illustrated in Figure 3.1.



**Figure 3.1.** Branching tree diagram schematizing the class inclusion relation of hyponymy/hypernymy of the word “animal”.

The hierarchical structure encoded in taxonomic relations of hypernymy/hyponymy typically implies transitivity. Thus, if “dog” is a hyponym of “mammal” and “mammal” is a hyponym of “animal”, then “dog” is a hyponym of “animal”, presumably inhering aspects of its semantic representation (for further details, see ¶ 3.2.).

Taxonomic hierarchies based on class inclusion relations can potentially include sets of different entities. For example, Lyons (1977) stated that both abstract (e.g., “quality”) and concrete nouns (e.g., “fish”) can be organized in hierarchies within the lexicon (as demonstrated by Bolognesi et al., 2020, see ¶ 3.4.), as well as mass nouns (e.g., “material”). Furthermore, although the relation of hyponymy/hypernymy is more naturally applied to nouns, even verbs and adjectives can exhibit this kind of sense relation. However, to be framed as “X is a (kind of) Y”, they require prior nominalization (Lyons, 1977).

Additionally, in some cases, hyponymy can be expressed through compositionality, and, in particular, through adjectival modification of a noun as well as through attributive nouns, where the modifier specifies a subtype of the hypernym, as in “musical instrument” or “emperor penguin”. In other cases, the integration of a syntagmatic modifier gives rise to co-hyponymy, as in the case of “despotic ruler” as a co-hyponym of “tyrant”. As Lyons (1977, p. 294) observes:

Hyponymy is a paradigmatic relation of sense which rests upon the encapsulation in the hyponym of some syntagmatic modification of the sense of the superordinate lexeme.

As mentioned above, taxonomic relations can be expressed as “X is Y” or “X is a kind (or type) of Y”, reflecting an IS-A (a kind of) relation, where *kind* functions as a semantic primitive (Wierzbicka, 1984). However, Cruse (1986) and Croft & Cruse (2004) argued that “X is a kind/type of Y” serves as the diagnostic frame of a sub-species of hyponymy, namely, *taxonymy* (Cruse, 1986), which is based on the specification of a particular aspect of semantic representation that licenses the relation. For example, “lumberjack is a kind of manual worker” sounds more natural than “lumberjack is a kind of man” because it specifies the <type of work> aspect encoded in the lexical item “manual worker” (Croft & Cruse, 2004). Nonetheless, taxonymy might be influenced by the relative distance between terms in the taxonomic hierarchical structure, making “long-distance” categorizations (e.g., “dalmatian is an animal”) less intuitive than more proximate categorizations (e.g., “dalmatian is a dog”; Murphy, 2002).

Recent research suggests that precise brain regions are more active in processing and elaborating taxonomic relations. In particular, Davis & Yee (2018) propose that the anterior temporal lobe supports taxonomic processing, contributing both to distinctions among highly similar referents (i.e., hyponyms and co-hyponyms) and less confusable

referents (i.e., hypernyms). Additionally, connectivity with the prefrontal cortex may facilitate distinctions supporting the selection of an appropriate label for the referent.

From a cognitive linguistic perspective, these linguistic structures reflect underlying processes of categorization and taxonomic organization. While the hierarchical structure of hypernyms and hyponyms in linguistics encodes levels of abstraction within the lexicon, analogous relations organize conceptual knowledge. Lexical items occupy specific positions along a continuum of granularity, and conceptual representations can likewise be arranged at varying levels of abstraction, precision, and inclusiveness.

### **3.2. Taxonomic Structure of Conceptual Categories: Superordinate, Basic, Subordinate Levels**

The taxonomic and hierarchical organization observable within the lexicon is mirrored in conceptual knowledge. Conceptual representations can be structured at varying levels of abstraction and inclusiveness along a vertical dimension. Within the conceptual system, the correlative terms of hypernym and hyponym correspond to *superordinate* and *subordinate* categories, respectively. Superordinate categories are more inclusive than subordinate categories, with the relation between categories typically captured by the IS-A relation (Collins & Quillian, 1969), analogous to the relation between hypernyms and hyponyms.

As illustrated in Figure 3.1., categories are connected through a series of class inclusion relations. For example, the expression “hammer is a tool” indicates that HAMMER is a subordinate and more specific category (or concept) included within the more general, superordinate category (or concept) TOOL. As discussed in the previous section (see ¶ 3.1.), the class inclusion relation (i.e., IS-A) is asymmetrical (e.g., “all pines are trees, but not all trees are pines”) and transitive (e.g., “if an instrument is an object and a brush is an instrument, then a brush is an instrument”).

In addition to transitivity, category relations involve property inheritance. For instance, asserting that “all wolves are animals” means that WOLF inherits properties or features encoded in the semantic and conceptual representation of ANIMAL, such as the ability to breathe or the possession of a skeletal structure. Subordinate categories differ

from superordinate categories by possessing additional properties that distinguish them both vertically from the superordinate categories and horizontally from categories at the same level. On this view, property inheritance thus supports the relevance of distinctive features for categorization, consistent with the principle of *cognitive economy*, which posits that not the entire set of features is redundantly represented at each level of categorization, but only those that are distinctive. Other features may be inferred via inheritance from subordinate categories (Murphy, 2002; see further discussion in ¶ 3.2.1.). For example, the ability to howl may be closely associated with the category WOLF, whereas the ability to breathe or to move may be inferred from the relation “a wolf is an animal”.

However, transitivity and property inferential processes are not always straightforward and may be influenced by the typicality of category members. Rips et al. (1973) found that, in a sentence verification task, participants verified IS-A statements with typical members of a category (e.g., “a goat is a mammal”) faster than IS-A sentences with atypical members (e.g., “an orca is a mammal”). Similarly, Sloman (1993, 1998) reported that, in one-step logical inferences, participants perceived conclusions involving typical members (e.g., “birds have pneumatic bones, so canaries have pneumatic bones”) as stronger than conclusions with atypical members (e.g., “birds have pneumatic bones, so emus have pneumatic bones”; Murphy, 2002).

Nevertheless, the distinctiveness of relevant features and the principle of cognitive economy have been proposed as central to certain models of categorization, such as prototype-based accounts (Rosch, 1978), contributing to the identification of the *basic level*, the defined optimal level of conceptualization, positioned between the *subordinate* and *superordinate* levels.

### **3.2.1. Basic Level Categories**

A first account for a consistent intermediate level of categorization was provided by Brown (1958), who observed that, when speaking with children, parents generally preferred short, frequent names that more precisely represented the referent (Hajibayova, 2013). For example, parents tended to refer to an instance of a dog using the word “dog” rather than “animal” or “Labrador”. However, the formalization and systematic development of a basic level of abstraction is due to the work of Rosch and colleagues

(Rosch et al., 1976; Rosch, 1978). In a series of influential studies on natural and basic objects (e.g., chair and tree), Rosch (1978) and Rosch et al. (1976) demonstrated that the conceptual system is multilayered and that an intermediate level, i.e., the basic level, is the most efficient for categorization, as reflected in words such as “hammer” or “bird”.

Rosch et al. (1976) also proposed a list of operational definitions of the basic level of abstraction:

- Commonality of attribute
- Similarity in motor movements
- Similarity in shape
- Identifiability of averaged shape

Their findings revealed that basic-level categories encompass the greatest number of members sharing a large set of common attributes. In feature listing tasks, the number and type of features listed by participants varied across the three levels of categorization (e.g., VEHICLE, BIKE, ELECTRIC BIKE). Participants listed the fewest number of features for superordinate categories, more features for basic categories, and the greatest number of features for subordinate categories. Feature types also differed: noun and adjective features predominated at the basic level (e.g., <feather>, <button>), subordinate level elicited additional adjective attributes (e.g., <brown>), while functional properties were more common at the superordinate level (e.g., <wearable>).

Similar patterns emerged when participants indicated the movements they would perform when interacting with objects (e.g., grasping, raising hands) from each category level. Compared to superordinate categories, basic-level and subordinate categories elicited a greater number of associated movements. Furthermore, members of subordinate and basic-level categories were more similar to one another in overall appearance and shape than those of superordinate categories, suggesting that superordinate categories are more difficult to mentally represent (Murphy, 2002).

Basic-level categories also exhibited a behavioral processing advantage, the well-known *basic-level advantage* (Murphy & Smith, 1982). In particular, participants verify category membership of object pictures more quickly at the basic level than at superordinate and subordinate levels (Rosch et al., 1976; Murphy & Smith, 1982; Jolicoeur et al., 1984; Murphy & Wisniewski, 1989). For instance, participants were faster to verify that the picture of a beagle was a DOG than an ANIMAL or a BEAGLE.

Similarly, in free naming tasks, participants named a picture of a beagle faster as a DOG than as an ANIMAL or a BEAGLE.

Additionally, basic-level categories are typically monolexemic, composed of a single lexical unit (Evans & Green, 2006), and they are more frequent and learned earlier in the development than superordinate and subordinate categories (Rosch et al., 1976; Mervis & Crisafi, 1982; Mervis, 1983; see also Murphy, 2002, chapters 9-10).

Rosch (1978) and Rosch et al. (1976) explained the semantic salience and the privileged status of the basic-level categories through two main principles, underlying the mechanisms of categorization: *cognitive economy* and *cue validity*. Cognitive economy has a different sense from the one explained in the previous section. In this context, it does not imply that distinctive features are encoded at different levels of abstraction; rather, it refers to the efficiency of a category in including as many members as possible while maximizing distinctiveness from other categories. Basic-level categories are cognitively efficient because they encompass numerous exemplars sharing relevant distinguishing features (e.g., CAT includes multiple subtypes with many shared features).

Cue validity is a probabilistic measure of the relevance of an attribute (i.e., a feature) for a given category, defined by how consistently the attribute is associated with members of that category (Rosch, 1978; Rosch et al., 1976). High cue validity indicates a strong association between an attribute and a given category (e.g., <wing> with BIRD), while low cue validity reflects an attribute that is less diagnostic for a particular category, as it occurs with relatively high probability in other categories. The basic level maximizes cue validity, whereas both superordinate and subordinate categories generally exhibit lower cue validity. Superordinate categories include many heterogeneous members that share few features, and subordinate categories differ from one another by only a small number of distinctive features. For example, the feature <wing> is associated with superordinate categories like BIRD, INSECT, and some members of ANIMAL (e.g., BAT), which differ in other features, while subordinate levels like RAVEN, PELICAN, and SEAGULL share most features and differ only in a limited set of additional distinctive features.

Murphy (2002) and Murphy & Brownell (1985) proposed an alternative account for the preference for basic-level categories and concepts, grounded in a structural explanation known as *differential explanation*, which relies on two properties:

*informativeness* and *distinctiveness*. Informativeness refers to the amount of information encoded in the semantic representation of a category or concept. Distinctiveness reflects the extent to which a category differs from other categories at the same level of abstraction.

According to this account, summarized in Table 3.1., the basic level represents the optimal level of abstraction because basic-level concepts (e.g., DOG; CHAIR) are both highly informative and distinctive. They are associated with many features and differ from a relatively large number of other concepts at the same level. In contrast, superordinate concepts (e.g., ANIMAL; OBJECT) are highly distinctive, as they refer to a broad and heterogeneous set of members, making categories clearly distinct from one another, but they are relatively uninformative because these members share few features. Subordinate concepts (e.g., ICELANDIC SHEEPDOG; BARBER CHAIR), conversely, are highly informative, being more specific and possessing a larger number of features, but they are less distinctive from one another.

| Levels of Abstraction | Informativeness | Distinctiveness |
|-----------------------|-----------------|-----------------|
| Superordinate         | Low             | High            |
| Basic                 | High            | High            |
| Subordinate           | High            | Low             |

**Table 3.1.** Differences among concepts at varying levels of abstraction in terms of informativeness and distinctiveness. Adapted from (Murphy, 2002).

From these perspectives, basic-level categories and concepts exemplify the principle of cognitive economy by achieving an optimal balance between specification and differentiation (Bauer & Just, 2017), which may account for the processing advantage of this level of abstraction. Recent neurocognitive research has confirmed and refined this idea, demonstrating that basic-level distinctions tend to emerge earliest during object recognition (Greene & Rohan, 2025). Neural representations appeared to prioritize basic categories even when tasks required access to subordinate categories (Bolognesi, 2025).

However, typicality can modulate this advantage. When subordinate categories include atypical members, a subordinate-level advantage in categorization may emerge alongside the basic-level advantage for typical members (Jolicoeur et al., 1984).

This section has examined the basic level both on its own and in relation to the other two levels of abstraction. The following sections provide a more detailed discussion of the superordinate and subordinate levels.

### **3.2.2. Superordinate Categories**

One of the most influential examinations of superordinate categories is presented in Markman (1985), which provides evidence that superordinate category terms are prevalently mass nouns rather than count nouns. For example, “furniture”, “jewelry”, and “food” are all superordinate concepts and mass nouns, whereas basic-level concepts are generally count nouns (Croft & Cruse, 2004). Furthermore, Markman (1985) demonstrated that this pattern is consistent across approximately 20 languages spanning multiple language families and subgroups.

Markman (1985) attributed this syntactic pattern to the fact that superordinate categories typically refer to collections of multiple objects, rather than individual referents. This assumption is supported by Wisniewski & Murphy (1989), who found that in a textual analysis, basic-level nouns were more likely to be used to refer to single referents, whereas superordinate nouns more often refer to classes of objects. Moreover, Murphy & Wisniewski (1989) showed that the processing advantage of basic-level concepts over superordinate concepts diminishes when participants categorize objects within a scene.

Finally, as discussed in the previous section, feature listing tasks revealed that basic-level and subordinate categories tend to elicit more features expressed as nouns and adjectives. In contrast, superordinate categories elicit more functional properties. This focus on functional attributes and the representation and collection of individuals contributes to speakers’ perception of superordinate categories as more abstract (Murphy, 2002).

### 3.2.3. Subordinate Categories and Expertise

Outside the field of language acquisition, subordinate categories have been investigated less consistently than basic and superordinate categories. Subordinate categories are generally polymorphemic and, as discussed for hyponyms, often composed of modifier-head patterns (e.g., “teaspoon”, “butter knife”, “mountain bike”; Croft & Cruse, 2004). Compared to basic-level concepts, subordinate concepts provide greater specificity and additional details (e.g., POOL TABLE is a table for playing pool). Although cognitively less efficient than basic-level concepts, subordinate concepts become more relevant when clarification or further information is required, depending on the pragmatic context (Cruse, 1977).

Subordinate concepts have been more extensively examined in relation to speakers’ domain-specific expertise. Rosch et al. (1976) speculated that experts’ knowledge relies on fine-grained representation within subordinate categories. Tanaka & Taylor (1991) provided a more systematic investigation of the relationship between levels of abstraction and expertise, using a within-subject design to compare novices and experts in specific domains.

Participants completed a features listing task, a free-naming task, and a categorization task. Novices replicated the results reported by Rosch et al. (1976) across the tasks: they elicited more features for basic-level concepts than subordinate concepts, named objects in pictures with basic-level words more frequently than subordinate words, and recognized objects as members of basic-level categories more quickly than as members of subordinate and superordinate categories. In contrast, experts in their domains listed more features at the subordinate levels, were more likely to produce subordinate-level words in the free-naming task, and were equally fast in recognizing objects as members of basic-level and subordinate categories, while categorization at the superordinate level remained slower.

Tanaka & Taylor (1991) further observed that, for experts in their domains, feature listings showed that subordinate categories were both informative and distinctive, reflecting knowledge of unique features of specific exemplars. Consequently, subordinate concepts in expert knowledge can be as cognitively efficient as basic-level concepts (Murphy, 2002).

### 3.3. Theories of Conceptual Categorization: The Prototype Theory

As discussed in Chapter 2, the internal structure of a category is not necessarily represented as a definite list of features with sharp boundaries; rather, categories may be conceived as gradual structures, where membership depends on the degree of similarity among members and boundaries between categories are blurred. Building on Wittgenstein's notion of *family resemblances* (Wittgenstein, 1958: 67), Rosch & Mervis (1975) demonstrated that category membership is influenced by the similarity of an exemplar to other members. Typical exemplars share many features with other members of the same category but differ from members of other categories. In contrast, atypical members share fewer similarities with other members of the same category. Consequently, not every member is equally representative of the category (Geeraerts, 2008). This representativeness can be captured through *Goodness-Of-Exemplar* ratings (Croft & Cruse, 2004), with the most representative members identified as *prototypes*.

Prototypes are summary representations of a category structured as a unified whole based on features that are most likely associated with it (Murphy, 2002, 2016). Crucially, these features vary in semantic relevance, determined by how consistently a feature recurs within a category while remaining uncommon in other categories. For example, <bark> is more relevant to the category DOG than <has fur>, which is shared across multiple categories. Therefore, the more relevant features a category member possesses, the better it represents the category and the more similar it is to the prototype (Hampton, 1997; Croft & Cruse, 2004).

In contrast to the classical view (see ¶ 2.1.1.), the *prototype theory* does not posit that specific features are necessary for category membership. Instead, a feature's relevance is determined by its weight in the semantic representation, computed based on its frequency within the category and rarity outside it. Prototype theory also accounts for borderline cases, where a member may share an equal amount of similarities with multiple categories (e.g., AVOCADO with FRUIT and VEGETABLE). Moreover, because typical and prototypical members exhibit highly relevant features, they are categorized more rapidly than atypical members (Rosch & Mervis, 1975).

As schematic representations, prototypes can be interpreted as a schemata (Rumelhart & Ortony, 1977) which encode the relations between features and their relative importance within a conceptual and semantic representation. A schema organizes

features into dimensions, or *slots*, each associated with specific values, known as *fillers* of the slot (Murphy, 2002). For example, a feature can be represented as a dimension (e.g., <head color>, <number of eyes>) to which different values apply (e.g., <blue>, <yellow> for <head color>). Slots may constrain the set of possible values (e.g., birds typically have two eyes, or one if injured), and values can compete with each other (e.g., a dog's color of the fur can be brown or black or white). Slots can also be interrelated, restricting each other's values (e.g., flightless birds cannot migrate). As schemata, prototypes provide a structured, systematic representation of a category, capturing potential values, features, and relations.

A prototype thus constitutes an abstracted, summary representation that maximizes the relevance of recurring features and typical values within a given category. Prototype theory involves a probabilistic view of the conceptual structure, in which not all features are required for category membership; rather, the weighted set of relevant features is constitutive of category membership.

By contrast, exemplar theory offers an alternative account, proposing that conceptual structure is based on collections of stored instances, which are retrieved as needed rather than abstracted into a prototype.

### **3.3.1. Exemplar Theory**

According to the *exemplar theory*, conceptual representations are not based on abstracted summary representations derived from specific category members. Instead, categories are formed through information encoded from each specific instance encountered and stored in memory.

Similarities play a central role also in this approach. For example, when a Maine Coon, a specific exemplar or token of the category CAT, is encountered, recognizing it as a CAT involves a comparison between that specific exemplar and stored exemplars of CAT in memory. Category membership is determined by evaluating the similarity of that specific Maine Coon to the set of recorded instances of CAT (Medin & Schaffer, 1978). If the Maine Coon is highly similar to other store exemplars, it will be considered a typical member of the category. By contrast, a sphynx, which shares fewer similarities with other exemplars of CAT, would be regarded as an atypical exemplar.

Consequently, as in the prototype theory, exemplar theory predicts a typicality-based processing advantage in categorization tasks (Nosofsky & Palmeri, 1997). Typical exemplars tend to be recognized more quickly as category members because they are more similar to the stored instances, whereas atypical members are categorized more slowly due to their lower overall similarity to the exemplar set stored in memory.

Medin & Schaffer (1978) proposed that similarity is computed as a function of two factors: the importance of the dimensions on which exemplars match or differ, and the degree of similarity on each dimension (for further details, see Murphy, 2002).

A further issue concerns whether exemplars should be treated as types or tokens. In particular, concept formation may differ depending on whether individuals frequently encounter many distinct instances of an exemplar (e.g., different labradors) or repeatedly encounter the same instance. To address this question, Nosofsky (1988) conducted an experiment using color patches that varied in saturation and brightness, manipulating the frequency of stimulus presentation. The results showed that items presented more frequently were judged to be more typical than those presented less frequently. Moreover, computational simulations of exemplar models provided a better account of the data when each stimulus presentation was treated as a distinct exemplar (i.e., a token) rather than when each type was considered as an exemplar. These findings suggest that the frequency of encounters with exemplars plays a crucial role in category formation, potentially outweighing the diversity of exemplar type.

### **3.3.2. Theory Theory**

Similar to Langacker's theory of domains (see ¶ 2.1.2.), the *theory theory* posits that concepts are knowledge structures embedded in and supported by background knowledge. Murphy (2002) characterized this approach as a knowledge-based account, arguing that concepts are not learned in isolation but as part of a broader body of general knowledge that contributes to the representation of the concepts themselves. This background knowledge encompasses a wide range of information, including contextual and domain-specific knowledge associated with a concept. As one acquires new information related to a concept, this background knowledge is continuously updated; consequently, categorization can be understood as a form of reasoning that relies on inferential processes grounded in background knowledge.

Gopnik (2003) and Gopnik & Wellman (1994) traced the foundations of this theory to children's cognitive and linguistic development. Using the same cognitive devices as adults, children actively construct knowledge about the world by abstracting and reorganizing instances and relations, thereby forming background information that supports the development of intuitive "theories" regarding the world. On this view, children begin to make predictions, gather evidence, and engage in inductive and inferential reasoning based on their existing background knowledge. When predictions are falsified, children recalibrate their background knowledge is revised in light of new acquired information, leading to the refinement of these theories and enabling more accurate predictions.

From the perspective of conceptual structures, the theory theory holds that category formation is not driven by repeated encounters with specific exemplars or by similarity to an abstract summary representation that maximizes relevance of the most recurring features. Instead, category structure is grounded in the retrieval of relevant information from background knowledge. This view aligns with Barsalou's (1985) notion of *ideals*. Ideals are not summary representations of category instances but rather features that category members ought to possess to optimally fulfill goals or functional roles associated with the category. Categories structured around such ideals are referred to as *goal-derived categories* (Barsalou, 1985). The relevant features of these categories are retrieved from background knowledge that situates the category within a network of relations.

Accordingly, a central assumption of the theory theory is that the relevance of category features is evaluated with respect to general knowledge about the category itself, rather than solely on the basis of similarity or frequency of occurrence.

Together, these approaches and the taxonomic organization of conceptual representation have been highly influential, constituting cornerstones of cognitive and psycholinguistic research. However, because they drew largely on the study of naturalistic taxonomies, these frameworks of conceptual representation across different levels of abstraction have focused primarily on taxonomies of natural categories and object (artifact) categories, that is, concepts typically denoted by concrete words. As noted by Lyons (1977, see ¶ 3.1.), abstract concepts can also be hierarchically organized, and abstract words are capable of entering semantic class-inclusion relations (e.g., "Buddhism

IS-A religion”; “freedom IS-A condition”), although such taxonomic structures may be less pronounced (Borghi, 2022). If taxonomic relations are indeed fundamental to conceptual and linguistic organization, their relevance cannot be restricted to the concrete concepts alone.

Moreover, the traditional distinction between basic, subordinate, and superordinate levels has typically been assumed a priori, despite being empirically supported, and treated as qualitatively discrete, rather than being explicitly operationalized through quantitative measures. Within these frameworks, the basic level has consistently been privileged, whereas subordinate and superordinate levels have received comparatively less systematic investigation. To develop a more comprehensive account of conceptual representation and word meaning, abstraction and inclusiveness should therefore be reconceptualized as continuous dimensions rather than as discrete categorical types. Such reconceptualization may also lead to more nuanced hypotheses concerning the internal structure of lexicalized conceptual categories, for instance, whether they are best characterized as prototype-based, exemplar-based, theory-based, or as hybrid systems.

Adopting this perspective allows a shift of focus toward categories occupying the higher and lower levels of abstraction, namely, highly general and highly specific words. Because the processes of abstraction and reference are closely related to both categorical specificity and conceptual concreteness (see Chapter 1), a detailed examination of their relationship is required. The following section addresses this issue by introducing two studies that explicitly investigate the relationship between categorical specificity and conceptual concreteness, demonstrating that, although correlated, they constitute two distinct dimensions of abstraction.

### **3.4. Disentangling Conceptual Concreteness and Categorical Specificity**

Chapter 1 outlined the theoretical foundations supporting the view that conceptual concreteness and categorical specificity should not be conflated. Conceptual concreteness reflects a variable related to abstraction, characterizing the nature of the conceptual representations. In contrast, categorical specificity operationalizes the linguistically mediated process of abstraction. While concreteness and abstractness denote the outcomes of this process, the generality/specificity continuum captures the process itself.

Empirical evidence supporting this distinction was first provided by Bolognesi et al. (2020), who systematically explored the relationship between these two variables. They argued that concreteness and specificity reflect distinct aspects of the referential process. Concreteness pertains to the quality of conceptual and semantic representation of a referent, as being characterized by perceptual and sensory experience or, in the case of abstract concepts, heterogeneous experience that lack single reference. Categorical specificity, by contrast, indexes the process of abstraction, reflecting the degree of precision and inclusiveness of a linguistic category or lexicalized concept across taxonomic levels of granularity. More specific categories and words (e.g., “dolphin”) are less inclusive, referring to fewer exemplars with richer defining features, whereas more general categories and words (e.g., “vertebrate”) encompass a larger number of referents that share fewer salient features.

As discussed in Chapter 2, this aspect of category representation is defined as dimensionality, which underlies the widespread conflation of general words with abstract words and concreteness with specificity. Both general and abstract concepts are typically *low-dimensional* (Lupyan & Mirman, 2013; Borghi, 2022), referring to multiple category members that share relatively few semantically relevant features. Conversely, both concrete and specific concepts are *high-dimensional*, denoting fewer exemplars that share a greater number of distinguishing features.


To empirically investigate this conflation, Bolognesi et al. (2020) conducted three separate analyses. In the first analysis, they examined correlations between human concreteness ratings from Brysbaert et al. (2014) and three measures of specificity extracted from the lexical database *WordNet* (Miller, 1995) for 13,518 words. WordNet organizes lexical items through two main types of relation: synonymy and hypernymy/hyponymy. As discussed in § 3.1., the hypernym/hyponym relationship constitutes an IS-A relation (Wierzbicka, 1984), in which a category is included in another category as a specification or specific instance (e.g., “rake IS-A tool”). WordNet is based on a tree-like structure, with nodes connected in layers based on levels of granularity. The three specificity measures calculated a score extracted by examining a word’s position relative to other nodes in the hierarchy, considering its distance from the root and leaves, the number of hypernyms, and the overall depth of the taxonomy. Results indicated that correlations between concreteness and the three measures of specificity were positive and

statistically significant, but modest (ranging from  $r = 0.263$  to  $r = 0.361$ ), suggesting that the two variables capture related but distinct dimensions.

In the second analysis, Bolognesi et al. (2020) explored whether the set of 3,896 nouns categorized as abstract or concrete in WordNet, under the nodes ABSTRACTION and PHYSICAL\_ENTITY, respectively, differed in terms of concreteness and specificity. The nouns classified under PHYSICAL\_ENTITY were significantly more concrete than those under ABSTRACTION. Furthermore, nouns in the PHYSICAL\_ENTITY group were also more specific than ABSTRACTION nouns, although the magnitude of this difference was modest. These results suggested that more concrete words tend to be more specific, whereas more abstract words are likely more general.

The third analysis further clarified this complex relationship by crossing concreteness and specificity variables to examine the extreme poles of the distribution. In particular, within the 13,518 words analyzed in the first analysis, Bolognesi et al. (2020) examined words with concreteness and specificity scores falling within the lowest and highest 0.25 point ranges of the 1-5 Likert scale. Results showed that words were distributed across four quadrants: abstract-general, abstract-specific, concrete-general, and concrete-specific (see figure 3.2.). Although granularity levels were more pronounced in concrete words compared to abstract words, specificity characterized both types of concepts. For instance, VEHICLE is a concrete conceptual category that includes the more specific HYBRID CAR, just as ART is an abstract conceptual category encompassing the more specific POINTILLISM.

Overall, these findings demonstrated that human-generated concreteness and specificity, as extracted from WordNet's encyclopedic-based knowledge, are related but independent semantic variables, reflecting distinct dimensions of the abstraction and referential process. Concreteness primarily reflects the quality of conceptual representation, whereas specificity captures the granularity of linguistic categorization.



**Figure 3.2.** The intersection of concreteness and specificity, resulting in four distinct quadrants. abstract-general, abstract-specific, concrete-general, and concrete-specific. Note: boundaries between quadrants are not sharp. Adapted from Bolognesi et al. (2020).

### 3.4.1. Human-Generated Specificity Ratings and Relations with Other Semantic Variables

Bolognesi et al. (2020) demonstrated that concreteness and specificity are two distinct variables that contribute to the semantic representation of words. The third analysis revealed the emergence of four categories: abstract-general, abstract-specific, concrete-general, and concrete-specific. However, a potential limitation of that study lies in the fact that specificity measures were extracted from WordNet (Fellbaum, 1998; Miller, 1995). Although WordNet is a valuable linguistic resource, the encyclopedic knowledge it encodes may not fully reflect speakers' actual semantic and conceptual representations.

To address this limitation, Bolognesi & Caselli (2023) collected human-generated specificity ratings for different parts of speech (i.e., nouns, adjectives, and verbs) for

1,121 Italian words from the Italian adaptation of Affective Norms for English Words (ANEW; Bradley & Lang, 1999; Montefinese et al., 2014).

To obtain specificity ratings, Bolognesi & Caselli (2023) adopted the Best-Worst Scaling (BWS) method (Louviere et al., 2015). This methodology requires participants to evaluate a set of words and to identify the most and the least specific item in each set. In this study, 352 participants were presented with  $m$  tuples of four words belonging to the same part of speech. For each tuple, participants were asked to rate target words from the most specific to the least specific.

The use of the BWS method was theoretically motivated by the relational nature of specificity. Along the generality/specificity continuum, a word or category is not inherently specific or general; rather, its level of specificity is defined in relation to more general (i.e., hypernym or superordinate category) or more specific (i.e., hyponym or subordinate categories) words or categories. Thus, the referential scope of a category depends on its degree of inclusiveness relative to other categories. For instance, the word “berry” constitutes a relatively precise category because it is more specific than words like “fruit” or “food”, yet it is more general than “blueberry”, which denotes a narrower and more specific subtype. As argued by Bolognesi & Caselli (2023), the BWS effectively captures this relational nature of specificity by requiring participants to compare words and rank within hierarchies.

Across three separate studies, Bolognesi & Caselli (2023) investigated whether human-generated specificity norms correlate with WordNet-derived specificity measures employed by Bolognesi et al. (2020), whether human-generated concreteness and specificity ratings are correlated, and how specificity relates to other psycholinguistic variables, including imageability, age of acquisition (AoA), valence, arousal, dominance, word frequency, and familiarity. Age of acquisition norms were sourced from the itAoA dataset (Montefinese et al., 2019), while frequency measures were obtained from the SUBTLEX-IT corpus (Crepaldi et al., 2016), containing word frequencies extracted from subtitles. All other variables were collected from Montefinese et al. (2014). In a fourth study, Bolognesi & Caselli (2023) explored the impact of categorical specificity on reaction times of a lexical decision task (i.e., “is this string of letters a word or nonword?”), using data reported by Vergallito et al. (2020).

In the first study, 743 nouns extracted from the *Open Multilingual WordNet* (Bond & Paik, 2012) were correlated with an equivalent set of nouns from the Italian ANEW dataset. The analyses revealed a significant but relatively low positive correlation between the two types of specificity measures ( $r = 0.443$ , Spearman correlation coefficient). According to Mukaka (2012), correlations between 0.30 and 0.50 are considered low. This result was interpreted as suggesting that the encyclopedic and lexicographic knowledge encoded in WordNet does not fully reflect speakers' lexicon. As Bolognesi & Caselli (2023) noted, WordNet included levels of taxonomic depth and granularity that exceed those typically accessed by average speakers, who may categorize a German shepherd within a taxonomy as “dog → mammal → animal → organism”, rather than as “shepherd dog → dog → canine → carnivore → mammal → vertebrate → phylum Chordata → animal → eukaryote → organism”.

The second study revealed a low albeit positive correlation between human-generated specificity and concreteness ratings, collected by Montefinese et al. (2014) ( $r = 3.16$ ; Spearman correlation coefficient). This finding is consistent with the results from Bolognesi et al. (2020), further supporting the idea that, even when both variables are derived from human evaluations, concreteness and specificity reflect two distinct semantic dimensions involved in reference and abstraction. For example, “object” may be a highly concrete word despite being relatively general, whereas “photosynthesis” denotes a highly specific yet abstract word.

In the third study, correlations between human-generated specificity ratings and other psycholinguistic variables (i.e., age of acquisition, valence, arousal, dominance, familiarity, and imageability) were significant but generally low, with coefficients below 0.238. In particular, specificity showed negative correlations with valence, arousal, dominance, and familiarity. However, when controlling for the other variables, the relationship between dominance and specificity was no longer significant. These results suggest that the more specific a word is, the more it tends to be unfamiliar and associated with negative feelings. In contrast, general words tend to be more familiar and linked to positive feelings.

By contrast, imageability and age of acquisition showed a significant positive correlation with specificity, indicating that the more imageable words tend to be more specific, whereas less imageable words are more general. Moreover, general words tend

to be acquired earlier in language development, whereas more specific words are typically learned later (Mervis & Crisafi, 1982). Notably, specificity showed a moderate negative correlation with word frequency ( $r = -0.526$ ), suggesting that specific words tend to be less frequent, whereas general words occur more frequently in language use.

In the fourth study, linear regressions revealed a significant positive effect of specificity on lexical decision latencies: more specific words require longer processing times, whereas general words tended to be recognized more quickly. However, a quadratic model revealed a U-shape relationship, such that words with intermediate levels of specificity elicited faster response times than words with either very high or very low specificity scores in the lexical decision data reported by Vergallito et al. (2020). This pattern is consistent with the processing advantage typically observed for basic-level words (see ¶ 3.2.), which may correspond to words with medium specificity scores.

Taken together, these results provide a deeper understanding of the semantic representation of both general and specific words, confirming that categorical specificity constitutes an independent variable capturing a precise aspect of meaning, namely, the level of inclusiveness and granularity encoded in words and lexicalized concepts. Moreover, specificity was shown to exert an additive effect on reaction times, highlighting its role in language processing. Finally, the low correlation between human-generated concreteness and specificity ratings provides further evidence for the conceptual disentanglement of these two variables.

Overall, categorical specificity should be regarded as a crucial component of semantic representation, as it accounts for the taxonomic organization of conceptual knowledge across different levels of linguistically mediated abstraction. At the same time, the relationship between concreteness and specificity remains complex and warrants further investigation to clarify the mechanisms underlying the abstraction process.

### **3.5. Exploring the Impact of Abstraction and Categorical Specificity in Language: Recent Empirical Evidence**

Within the empirical investigation of linguistically mediated abstraction, Bolognesi et al. (2020) demonstrated that conceptual concreteness and categorical specificity constitute two distinct semantic dimensions, reflecting different aspects of the abstraction process.

In particular, categorical specificity operationalizes the process of schematization (see ¶ 1.2.2.), allowing referents to be processed at varying levels of granularity.

Further evidence is provided by Bolognesi & Caselli (2023), who found that categorical specificity exhibits low correlations with other semantic and lexical variables, including age of acquisition, arousal, dominance, valence, imageability, and word familiarity. Their results indicate that general words tend to be more familiar and associated with more positive feelings than specific words, although they are less imageable. In contrast, specific words are less familiar but more imageable. Moreover, specificity showed an inverse relationship with word frequency, suggesting that general words are frequent, whereas highly specific words are relatively rare. This finding aligns with the existing literature proposing that subordinate categories are associated with more specialized knowledge (Tanaka & Taylor, 1991; see also ¶ 3.2.).

In addition, Bolognesi & Caselli (2023) reported that specificity affects language processing: words at different levels of abstraction and precision require varying processing times. Notably, words with medium specificity scores, potentially corresponding to the basic level, were associated with faster response times in a lexical decision task (Vergallito et al., 2020).

Traditionally, general words (i.e., hypernyms or superordinate categories) differ from specific words (i.e., hyponyms or subordinate categories) with respect to distinctiveness, informativeness (Murphy, 2002), and inclusiveness (Lyons, 1977). As discussed in ¶ 3.1., inclusiveness can be understood differently, depending on whether it refers to extension or intension. However, in the present work, inclusiveness refers to the extension and scope of denotation. Accordingly, general words are more inclusive, encompassing a broader network of entities, contexts, and situations, whereas specific words refer to a narrower set of exemplars and contexts. In this sense, general words are highly distinctive, whereas specific words are less distinctive. In contrast, specific words are more informative than general words, as their representations are semantically richer (see ¶ 3.2.). Finally, general words are low-dimensional, as they refer to multiple instances that share relatively fewer relevant features, while specific words are high-dimensional, denoting a small number of referents that share a greater number of relevant features (Lupyan & Mirman, 2013; Borghi, 2022).

The findings reported by Bolognesi & Caselli (2023) contribute to a deeper understanding of the semantic representation of words across different levels of the taxonomic hierarchy, highlighting the relevance of studying categorical specificity as a semantic variable and as a potential quantifiable index of the abstraction process. Specificity helps explain how words at different levels of abstraction vary in their semantic representations and in their effects on language processing in behavioral tasks. Investigating these differences is therefore crucial for understanding how abstraction may shape linguistic behavior and cognition.

Given the importance of abstraction as a fundamental mechanism of semantic representation, recent research in cognitive science and psycholinguistics has focused on examining the cognitive and linguistic consequences of representing meaning at different levels of abstraction.

For example, Van Hoef et al. (2023) employed four word→picture categorization tasks to investigate whether the distance of sensorimotor (i.e., ratings of sensorimotor strength; Lynott et al., 2020) and linguistic (i.e., co-occurrence frequencies) experiences predict categorization performance and contribute to explaining basic-level advantage. Participants were presented with a series of word-picture pairs and were asked to evaluate (i.e., yes vs. no) whether the picture item provided was a potential member of the word-labelled category. Results showed that the sensorimotor distance between a word category (e.g., “flower” or “plant”) and a member (e.g., “tulip”) predicted categorization performance as effectively as categorical distinctions based on discrete traditional, taxonomic levels (i.e., superordinate, basic, subordinate). However, when sensorimotor and linguistic distributional distances were weighted to reflect the range of words typically produced to name a given object, without assuming an *a priori* preference for basic-level words, the model incorporating both linguistic and sensorimotor experiences more effectively predicted response times and accuracy than the model based on discrete taxonomic distinctions.

These results confirmed the emergence of a basic-level advantage: picture items were faster categorized when preceded by a basic-level word (e.g., “dog”) than by a more general word (e.g., “animal”). However, within this study, the processing advantage was explained in terms of sensorimotor and linguistic distributional similarity. Van Hoef et al. (2023) proposed that representation of concepts closer in a taxonomic structure (e.g., CAT

and SIAMESE) may be more similar in terms of sensorimotor and linguistic distributional experiences than concepts that are farther apart (e.g., ANIMAL and SIAMESE). This assumption is consistent with the implications derived from the notion of dimensionality (Lupyan & Mirman, 2013): concepts occupying closer positions within a taxonomy may exhibit less variation in dimensionality and thus share more relevant information than concepts that are taxonomically distant.

Support to this assumption comes from a recent study showing that a computational model of category verification incorporating both sensorimotor and linguistic distributional information performed similar word→picture categorization tasks more efficiently when there was a greater degree of overlap in sensorimotor experience and/or linguistic distributional knowledge between exemplar (i.e., picture) and word labelled category (Wingfield et al., 2024).

An additional study by Kobrock et al. (2024) employed a reference game paradigm to examine referential communication at different levels of abstraction. In this task, a speaker produced utterances to help an interlocutor identify two target pictures among a set of distractors. The manipulation consisted of varying the level of abstraction of the conceptual contexts (i.e., the set of pictures among distractors). Three conceptual contexts were constructed, in which target pictures represented instances belonging either to the same superordinate category (e.g., ANIMAL), the same basic-level category (e.g., DOG), or the same subordinate category (e.g., DALMATIAN).

Experimenters found that referential communication was tailored to the level of abstraction of the context. In particular, speakers produced words matching the appropriate level of reference and abstraction defined by the conceptual context: for instance, speakers employed the hyponym “dalmatian” to identify two dalmatians among pictures of other dog breeds, and the superordinate word “animal” to identify a robin and a wolf among pictures of objects. Notably, Kobrock et al. (2024) replicated the basic-level effect, showing that speakers were faster to produce basic-level terms than superordinate and subordinate terms when referring to target pictures at the corresponding level of abstraction.

Together, these studies provide further evidence and new insights into well-established phenomena such as the basic-level advantage, while enabling a more fine-grained analysis of semantic representations that vary in abstraction. More broadly, they

shed light on specific mechanisms and dimensions underlying the abstraction process itself. However, these studies did not operationalize categorical specificity, nor did they explicitly address the relationship between concreteness and specificity, the two dimensions that are assumed to capture distinct but complementary aspects of the referential process.

Within the literature reviewed in this work, only two studies across behavioral sciences, psycholinguistics, and computational linguistics have examined both concreteness and specificity, using concreteness norms from Brysbaert et al. (2014) and specificity measures extracted from WordNet database (Fellbaum, 1998; Miller, 1995), as operationalized by Bolognesi et al. (2020).

The first study, Iliev & Axelrod (2017), investigated the contribution of specificity and concreteness in predicting lexical decision data from the Balota et al. (2007) dataset. In particular, Iliev & Axelrod (2017) conceptualized levels of specificity within a taxonomic structure as reflecting informational precision. Within this framework, more specific words were associated with higher precision and informativeness, whereas general words reflected lower granularity. Iliev & Axelrod (2017) introduced two measures of lexical precision derived from WordNet (Miller, 1995): *semantic inclusiveness* (P-inclusiveness) and *semantic depth* (P-depth). Semantic inclusiveness captures a word's level of generality by quantifying the number of subordinate nodes it possesses within a taxonomic tree: a greater number of subordinate nodes corresponds to higher generality (i.e., inclusiveness) and lower precision. Semantic depth, by contrast, measures the distance between a word and the top-level node in the taxonomy, with greater depth corresponding to higher specificity and informativeness.

Employing these measures, Iliev & Axelrod (2017) modeled concreteness and specificity (referred to as *precision*) as predictors of word recognition latencies, controlling for word length and frequency. The findings revealed opposing effects of concreteness and specificity: concrete words were processed more quickly, replicating the concreteness effect (Jessen et al., 2000), whereas more specific (i.e., precise) words were associated with slower processing times. These results have been interpreted as indicating that concrete words were more easily recognized due to their clear, defined representations. In contrast, although specific words have well-defined representations,

their lower level of generality and higher informational precision may impose greater cognitive demands during lexical access.

The second study (Schulte Im Walde & Frassinelli, 2022) explicitly drew on the distinction between concreteness and specificity pointed out by Bolognesi et al. (2020) and examined distributional measures potentially affecting each variable. The authors distinguished concreteness/abstractness from generality/specificity as two separate aspects of abstraction, employing human-generated concreteness ratings from Brysbaert et al. (2014) and deriving specificity scores for words in hypernym-hyponym relations and relative subsets in WordNet (Miller, 1995). Across a series of experiments, they identified reliable vector-space measures based on word co-occurrence patterns for both concreteness and specificity.

Their results showed that these measures differed between concreteness and specificity. In particular, frequency and *word entropy* were more reliable measures in distinguishing more general from more specific words, whereas neighborhood density variants better accounted for differences between abstract and concrete words. These findings suggested that general words are used more frequently and are more predictable than specific words, while abstract words tend to occur in a broader set of semantically diverse contexts compared to concrete words, which appear in fewer but more tightly connected contexts. These results align with previous findings on contextual variability (Hoffman, 2016) and availability (Schwanenflugel et al., 1988, 1992) characterizing abstract and concrete words.

Together, these two studies further deepen the understanding of the semantic representation of general and specific words and, in particular, their potential different contributions to predicting response latencies. However, both studies relied on specificity measures extracted from WordNet that, as noted by Bolognesi & Caselli (2023), may not accurately reflect the lexical organization or semantic distinctions represented by average language users, given their reliance on dictionary-based knowledge. Moreover, Schulte Im Walde & Frassinelli (2022) appear to treat concreteness and specificity as parallel dimensions reflecting the abstraction process, whereas Chapters 1 and 2 have discussed that concreteness and abstractness should be better understood as outcomes of the abstraction process itself. Notably, although Iliev & Axelrod (2017) did not refer to abstractness as a semantic construct, they clearly distinguished concreteness and

specificity as reflecting, respectively, the degree of perceptibility of referents and the scope of inclusiveness and levels of granularity of words' meaning.

Finally, the analyses conducted by Bolognesi et al. (2020) revealed that crossing concreteness and specificity gives rise to four types of categories: abstract-general, abstract-specific, concrete-general, and concrete-specific (see ¶ 3.4.). The emergence of these four categories further complicates the relationship between conceptual concreteness and categorical specificity, as well as the interpretation of their underlying semantic representation.

### **3.5.1. First Steps in Deepening the Relationship Between Conceptual Concreteness and Categorical Specificity**

Recently, the ERC-funded research group ABSTRACTION (Grant agreement: ERC-2021-STG-101039777) has investigated the relationship between conceptual concreteness and categorical specificity, examining their semantic representations and respective roles in language processing and word categorization, through paradigms drawn from psycholinguistics and computational linguistics.

Building on Bolognesi & Caselli (2023), Ravelli et al. (2025) released human-generated specificity ratings in English for the 1,034 words included in the ANEW dataset (Bradley & Lang, 1999). The authors adopted the same methodology employed by Bolognesi & Caselli (2023), namely, the Best-Worst Scaling (BWS) method. As discussed in ¶ 3.4.1., the BWS constitutes a valid methodology for capturing the relational nature of specificity. Across three studies, Ravelli et al. (2025) examined the relationship between abstractness/concreteness (Brysbaert et al., 2014) and specificity, as well as their associations with the same psycholinguistic variable considered by Bolognesi & Caselli (2023): valence, arousal, dominance, word familiarity, age of acquisition (AoA), and word frequency. Overall, Ravelli et al. (2025) confirmed for the English language the pattern of results previously reported by Bolognesi & Caselli (2023).

Furthermore, Ravelli et al. (2025) investigated the contribution of both concreteness and specificity in predicting response latencies of lexical decision tasks (i.e., is this string of letters a word or a nonword?) across multiple datasets (Balota et al., 2007; Keuleers et al., 2012; Hutchison et al., 2013; Tucker et al., 2019; Mandera et al., 2020),

as well as response latencies in an abstract/concrete semantic decision task (i.e., is this word abstract or concrete?) using a recent dataset (Pexman et al., 2017).

Results showed that specificity exerted a significant positive effect on reaction times in lexical decision tasks: the more specific a word was, the longer the response latency, whereas general words were associated with faster responses. This pattern is consistent with findings reported by Bolognesi & Caselli (2023) and Iliev & Axelrod (2017). Notably, Ravelli et al. (2025) observed a stronger effect of specificity in semantic decision task, with an opposite direction: more specific words required less processing time, whereas more general words were associated with longer response latencies.

In addition, Ravelli et al. (2025) identified a significant interaction between concreteness and specificity in predicting reaction times in both lexical decision and semantic decision tasks, including datasets in which specificity showed particularly robust effects (Balota et al., 2007; Keuleers et al., 2012; Pexman et al., 2017; Mander et al., 2020). The interaction exhibited a negative direction and appeared to be primarily driven by specificity, as concreteness did not show any significant independent effect on chronometric data.

Other research has contributed to clarifying the semantic representations of abstract-general, abstract-specific, concrete-general, and concrete-specific words. For instance, Villani et al. (2024) conducted a fine-grained semantic annotation of words included in the Italian adaptation of the ANEW dataset (Montefinese et al., 2014). Their annotation was organized according to a hierarchical structure comprising three levels of granularity and semantic precision: macro-category (e.g., ENTITY, QUALITY), category (e.g., LIVING, NON-LIVING, BODILY, NON-PERCEPTUAL), and subcategory (e.g., ANIMALS, FLOWER AND PLANTS, EMOTIONS, SOCIAL ARTIFACTS).

Villani et al. (2024) examined whether semantic content and categorical specificity contributed to predicting degrees of words' conceptual concreteness. While semantic content alone was a good predictor of concreteness, adding categorical specificity to the linear regression models significantly improved the predictive power of the model. Thus, incorporating information about specificity resulted in more efficient prediction of conceptual concreteness, particularly semantic content defined at a coarse level. For example, when the only semantic information for the word "weapon" was that

it belonged to the macro-category ENTITY, the inclusion of specificity information improved the prediction of its high concreteness.

Moreover, Villani et al. (2024) analyzed correlations between semantic content and the four different categories of words emerging from the intersection of concreteness and specificity (i.e., abstract-general, abstract-specific, concrete-general, and concrete-specific; Bolognesi et al., 2020). The results revealed patterns consistent with established differences in the semantic representation of abstract and concrete concepts. For instance, inner experiences and non-perceptual qualities were strongly associated with abstract concepts, whereas living and non-living entities were strongly associated with concrete concepts. Importantly, these differences also occurred across levels of abstraction and specificity, with the strongest contrast emerging between abstract-general and concrete-specific words.

Further evidence comes from Rambelli & Bolognesi (2023, 2024), who investigated the role of conceptual concreteness and categorical specificity in explaining the distributional properties of Italian and English words included in the ANEW datasets (Bradley & Lang, 1999; Montefinese et al., 2014). Their analyses focused on distributional measures of contextual variability, including word entropy, word-context associations, word-word associations, and semantic neighbors (Hoffman, 2016; Schulte Im Walde & Frassinelli, 2022). Overall, results revealed highly consistent patterns across the two languages, showing that differences in contextual variability are more strongly driven by specificity than by concreteness. Regardless of their concreteness, more specific words tended to occur in narrower and more similar contexts, whereas more general words were associated with a broader range of less tightly related contexts, suggesting a greater variability.

In addition, a significant interaction between concreteness and specificity emerged, indicating that the semantic representation of highly abstract-specific words is characterized by more predictable contexts. In particular, *entropy* (i.e., a measure of contextual uncertainty) was strongly and negatively associated with specificity, suggesting that the more specific a word is, the more predictable its contexts are. This effect was particularly pronounced for abstract-specific words.

Specificity and concreteness have also been investigated in relation to linguistic notions involved in reference. Collacciani et al. (2024), for example, investigated a

linguistic mechanism related to the perceived specificity of words as shaped by their contextual use, known as *genericity* (for a difference between generality and genericity, see also ¶ 3.5.1.). Genericity refers to the extent to which a linguistic expression applies to members of a category within a given context. For instance, the sentence “mosquitoes fly” typically refers to all exemplars of the category; “mosquitoes bite” applies only to female exemplars; and “mosquitoes carry malaria” refers to a narrower subset. Thus, the referential scope of a category appears to vary as a function of the predicate.

Collacciani et al. (2024) introduced a novel paradigm and dataset operationalizing genericity as a graded phenomenon along continuous dimensions of inclusiveness (i.e., generality/specificity) and abstractness/concreteness. Their results indicate that a binary distinction between generic and non-generic use in contexts fails to capture the graded nature of meaning (Bolognesi, 2025) and that both inclusiveness and abstractness/concreteness shape how broadly or narrowly a category’s referential scope is perceived.

Finally, as anticipated in ¶ 2.2.1., specificity has been shown to positively correlate with iconicity (i.e., the resemblance between a word’s form and meaning; de Varda, Lamarra et al., 2025), suggesting that more specific words tend to be more iconic. Iconic words are typically associated with specific referents and contexts, encoding salient perceptual or semantic properties of those referents in their meaning. Since specificity indexes semantic precision and category inclusiveness, its relationship with iconicity suggests that highly iconic words activate a relatively precise set of perceptual and semantic features, indicative of high-resolution semantic representations.

As discussed in previous sections, categorical specificity has emerged as a crucial variable that operationalizes linguistically mediated abstraction and independently contributes to the semantic representation of concepts across degrees of abstractness/concreteness. Multiple metrics have been proposed to operationalize this variable, including human ratings, WordNet-based measures, and the BWS method. Villani & Bolognesi (2025) reviewed these approaches and introduced a novel, promising metric for collecting specificity data: *Word Ladders*. Word Ladders is a mobile application developed by the Abstraction research group to collect specificity values through a generative, free production task embedded in a gamified environment. Participants generate taxonomic word ladders (i.e., taxonomic relations) by producing more general

(e.g., “fruit) and specific words (e.g., “Granny Smith) by adding steps above and below the initial prompt (e.g., “apple”). Data collected through Word Ladders is currently undergoing analysis to construct an ontology of hierarchical, taxonomic (though not exclusively) semantic relations, primarily encompassing class inclusion (i.e., IS-A) relations, using an ecologically valid methodology.

Together, these recent studies provide an empirically grounded theoretical integration of specificity as a key semantic dimension underlying linguistically mediated abstraction. By overcoming the limitations of traditional frameworks, which focus primarily on concrete categories and basic levels of abstraction, this body of work highlights the relationship between conceptual concreteness and categorical specificity, substantially enriching our understanding of the semantic representations of general and specific words. Table 3.2. summarizes potentially relevant semantic, lexical, and behavioral variables associated with these representations.

|                                 | <b>General Words</b> | <b>Specific Words</b> |
|---------------------------------|----------------------|-----------------------|
| <b>Dimensionality</b>           | Low                  | High                  |
| <b>Affective Content</b>        | Positive             | Negative              |
| <b>Age of Acquisition (AoA)</b> | Early                | Late                  |
| <b>Word Familiarity</b>         | Higher               | Lower                 |
| <b>Word Frequency</b>           | Higher               | Lower                 |
| <b>Contextual Variability</b>   | Higher               | Lower                 |
| <b>Contextual Entropy</b>       | Less predictable     | More predictable      |
| <b>Iconicity</b>                | Lower                | Higher                |
| <b>Inclusiveness</b>            | High                 | Low                   |
| <b>Informativeness</b>          | Low                  | High                  |
| <b>Distinctiveness</b>          | High                 | Low                   |
| <b>Response Duration in LDT</b> | Shorter              | Longer                |
| <b>Response Duration in SDT</b> | Longer               | Shorter               |

**Table 3.2.** Overview of semantic, lexical, and behavioral variables for which general and specific words may differ.

In the present work, I adopt the definition proposed by Villani & Bolognesi (2025), according to which categorical specificity refers to the degree of semantic precision in a word's meaning, determined by the extent of its conceptual scope or category inclusiveness. Words with higher specificity denote more narrowly defined words and concepts, characterized by a greater number of distinguishing features and, consequently, higher-resolution semantic representations. In contrast, words with lower specificity refer to broader, more inclusive categories, encompassing a wider range of referents that share fewer relevant features, and thus exhibiting lower-resolution semantic representations. Categorical specificity is closely associated with hierarchical taxonomic structures, in which more specific terms occupy lower levels and represent subsets of broader categories.

Despite these advances, the relationship between conceptual concreteness and categorical specificity has not yet been systematically investigated in conversational contexts or social interactions, nor in paradigms that extend beyond classical psycholinguistic research, which typically examines isolated, decontextualized words. The potential role of these two variables in shaping interaction remains largely unexplored, particularly in relation to conversations and pragmatic aspects of meaning. The following section highlights the relevance of studying semantic variables as they are embedded in naturalistic contexts and conversations, employing more ecologically valid stimuli and experimental paradigms.



## Methodological and Theoretical Rationale for the Experimental Section

### From Representations to Conversations

In the previous theoretical chapters, it has been discussed how conceptual representations and related word meanings are unlikely to be structured as dictionary-like definitions, in which meanings and concepts are associated with a fixed and *a priori* set of different kinds of properties and features. Rather, both word meanings and conceptual representations are better understood within an encyclopedic view of meaning (Evans & Green, 2006; Murphy, 2002), according to which background knowledge and multiple sources of information (linguistic and extralinguistic) contribute to shaping meaning.

The relationship between concepts and word meanings is complex, and a comprehensive discussion of this topic lies beyond the scope of the present work. Traditionally, this relationship has often been described as a linear mapping, whereby a conceptual core is directly associated with a word's meaning. For example, the label “cow” activates the relative corresponding conceptual representation COW. However, such a one-to-one relation fails to capture the more intricate interactions between words and concepts. The case of lexical ambiguity is emblematic: a single word may be associated with multiple related (*polysemy*) or unrelated (*homonymy*) meanings, thereby entering into one-to-many relations with conceptual representations.

Without taking a definitive stance on the precise nature of the distinction between concepts and word meanings, it is nevertheless reasonable to assume that the two are mutually influential. Conceptual representations and word meanings are grounded in a background of diverse experiences, including linguistic, perceptual, and contextual information. For instance, changes in linguistic and semantic experience primarily affect word meaning and, in turn, update the associated conceptual representations; conversely, changes in perceptual and sensorimotor experience reshape conceptual representations and thereby induce corresponding modifications in word meanings.

Over the past decades, cognitive and psycholinguistic research has increasingly focused on identifying and isolating the types and the number of different information, dimensions, and variables that contribute to shaping the word meaning and representation

of lexicalized concepts. To this end, a wide range of measures has been developed to collect data for large sets of words, often comprising thousands of items, resulting in the creation of several publicly available datasets. For example, human participants have been asked to provide subjective ratings on Likert scales for words along semantic dimensions such as concreteness, imageability, contextual availability, emotional valence, body-object-interaction (BOI), age of acquisition (AoA), iconicity, and perceptual strength across different modalities (visual, auditory, olfactory, tactile, gustatory). In parallel, additional dimensions and variables reflecting linguistically derived information have been extracted from dictionary-based, computational resources, and large corpora, including semantic relations (e.g., *synonymy*, *hypernymy*, *hyponymy*, *meronymy*), number of senses, co-occurrence statistics, word frequency, semantic neighborhood density, and other measures of contextual variability (e.g., word entropy).

Taken together, this wide range of dimensions and variables constitutes a semantic network that reflects a pluralistic approach to meaning and conceptual representations (Vigliocco et al., 2009; Borghi & Binkofski, 2014; Pexman, 2020). According to these pluralist accounts, meaning and semantic representations are structured along multiple dimensions, corresponding to different types of experience, whose relative contribution may vary across word types, as extensively discussed in relation to concrete and abstract concepts (see ¶ 2.2. and 2.2.1.). A central assumption of these pluralist approaches is that such semantic dimensions influence language processing, as well as lexical and semantic access to meaning. One framework that builds on this assumption is the *semantic richness* framework (Pexman, 2020), which posits that richer semantic representations elicit greater semantic activation. This increased activation, in turn, supports faster and more accurate word processing, in line with the principle that in lexical processing, “more is better” (Balota et al., 1991).

Consequently, the study of language processing, and by extension, language use, has relied on a wide range of laboratory-based experimental paradigms in which participants perform different tasks (e.g., lexical decision, semantic decision, categorization, text reading, picture naming, feature production tasks). These paradigms typically collect behavioral measures (e.g., reaction times, accuracy, eye movements). In addition to behavioral data, language processing has also been investigated using neuropsychological and neuroimaging techniques, including EEG and event-related potentials (e.g., the N400), as well as functional neuroimaging methods, which provide

complementary insights into the temporal and neural dynamics underlying meaning processing.

Across these paradigms, words are most often presented in isolation and in decontextualized form, displayed under controlled conditions for a fixed duration or until a response is produced. Accordingly, the explanatory and predictive power of semantic variables is typically assessed in terms of the proportion of variance they account for in behavioral outcomes (e.g., slower/faster or more/less accurate responses).

These traditional psycholinguistic methods offer several important advantages for studying word meaning and conceptual representations. For instance, rating-based measures and feature-listing tasks provide valuable insights into the nature of semantic representations by identifying which variables or dimensions are considered central. Moreover, highly controlled experimental paradigms provide standardized methods to investigate language processing during word recognition, categorization, production, and reading. These paradigms have been widely employed to assess the contribution of the different semantic dimensions.

Despite these strengths, traditional methods also present notable limitations. A primary concern is their focus on isolated words, which may obscure how meanings are constructed and modulated in more naturalistic contexts. Semantic dimensions that appear highly influential when words are presented in isolation may become less relevant, or interact differently, when words are embedded in sentences, discourse, or situational contexts (Banks et al., 2023). For instance, as discussed in ¶ 2.3., the concreteness effect, that is, the behavioral processing advantage of concrete words over abstract words, is attenuated or no longer observed when words are embedded in supportive contexts (Schwanenflugel et al., 1988; Schwanenflugel & Shoben, 1983; Schwanenflugel, 2013) or when other semantic variables, such as contextual availability, imageability, valence, are controlled (Schwanenflugel et al., 1988; Kousta et al., 2011).

The concreteness/abstractness dimension further illustrates the limitations of traditional psycholinguistic methods. For instance, Connell & Lynott (2012) argued that concreteness ratings do not reliably capture the perceptual components of conceptual representations, at least not to the same extent as modality-specific perceptual strength ratings. Reijnierse et al. (2019) additionally demonstrated that polysemy can influence concreteness ratings, with metaphorical meanings decreasing the perceived concreteness

of nouns. These findings contribute to indicating that meanings are not activated as fixed conceptual and semantic cores but are flexibly shaped by contextual information, task demands, stimulus modality, and individual-specific factors (Lebois et al., 2015; Yee & Thompson-Schill, 2016). In response, a recent study has collected ecologically valid concreteness ratings for Italian and English words embedded in a set of different contexts (i.e., CONcreTEXT norms; Montefinese et al., 2023).

Another limitation of studying language solely through isolated and decontextualized words is the neglect of the social dimensions that characterize natural language use. Chapter 2 extensively discussed the pivotal role of the social dimension in the representation, acquisition, and use of abstract words (Borghi & Binkofski, 2014; Borghi et al., 2019). Abstract words generally rely on social cohesion, social validation, metacognitive abilities, prompting conversations, and interactive turns (Borghi, 2023).

Building on this perspective, language should be investigated, reconciling semantics and pragmatics, consistent with the encyclopedic view of meaning (Murphy, 2002; Croft & Cruse, 2004). Interactive and naturalistic methods, as well as experimental pragmatic approaches (Noveck, 2018; Domaneschi & Bambini, 2020), which treat language as a form of joint action (Clark, 1996; Pickering & Garrod, 2021), provide critical insights by capturing how speakers use words within realistic situated communicative contexts and conversations.

Two additional claims emerge from adopting more naturalistic methodologies: (1) the use of ecologically valid paradigms, naturalistic stimuli, and life-like tasks allows researcher to study language (including multimodal aspects) and cognition in real-world contexts, capturing the dynamics of social interactions that contribute to shaping word meanings and conceptual representations (Vigliocco et al., 2024); and (2) the integration of traditional psycholinguistic methods with interactive, ecologically valid approaches is essential. Although laboratory paradigms have limitations, they provide precise experimental control and enable the systematic investigation of relevant semantic dimensions. In contrast, interactive and real-world-based methods capture dynamic, context-dependent, and socially embedded properties of language use but do not provide rigorous control of laboratory-based experimental settings and are susceptible to confounding variables (Banks et al., 2023; Vigliocco et al., 2024). Together, these complementary approaches offer a more comprehensive understanding of language.

## Overview of Experimental Section

The following sections present three case studies investigating the relationship between categorical concreteness and conceptual specificity in understanding language-mediated abstraction. These studies employ paradigms and methodologies characterized by progressively increasing ecological validity:

1. The first case study (Chapter 4) explored the influence of isolated words varying in conceptual concreteness and categorical specificity alone and in interaction on word processing. Three traditional psycholinguistic tasks were employed, such as rating, lexical decision (e.g., “is this string of letters a word or a nonword?”), and abstract/concrete semantic decision task (e.g., “is this word abstract or concrete?”). Nonwords were generated for the lexical decision task. Reaction times and accuracy were collected.
2. The second case study (Chapter 5) investigated how conceptual concreteness and categorical specificity impact conversational dynamics, focusing on two linguistic constructs associated with abstractness, that is, uncertainty and curiosity. The study consisted of two sessions. In the first session, participants completed a plausibility judgement task in which they evaluated the plausibility of short conversation excerpts composed of target sentences varying in both abstractness/concreteness and specificity, paired with follow-up expressions reflecting different communicative functions. In the second session, participants completed a rating task in which the same target sentences were evaluated across three semantic dimensions: word confidence, social metacognition, and the need for experts. Ecological validity was enhanced by embedding stimuli within simulated conversational contexts and by deriving the follow-up expressions from the follow-up expressions from participant’s responses collected in a previous free-production study (Villani et al., 2022). Reaction times were collected.
3. The third case study (Chapter 6) analyzed a dataset of semi-naturalistic interactions where a knowledgeable speaker describes objects to an interlocutor under four different conditions, defined by the physical presence or absence of the object, and by whether the object was known or unknown to the interlocutor (Gu et al., 2025).

The study examined whether the specificity and concreteness of the nouns and sentences produced, as well as the linguistic and abstraction strategies employed (e.g., analogical categorization via co-hyponyms; reference to objects' perceptual features), varied as a function of these conditions. Ecological validity of the study was supported by the analysis of previously recorded and transcribed interactions conducted in a real-life setting embedded within an experimental paradigm.

Overall, the experimental part of the present work is structured to progress from an examination of the effect of concreteness and specificity on language processing of the semantic representations at the level of isolated words, to an investigation of their perceived role in shaping conversational dynamics within simulated conversations, and finally to an exploration of whether these variables reflect communicative strategies in semi-naturalistic contexts. In particular, the final study addresses how speakers adapt their language in response to potentially real-life conditions, such as the physical absence or presence of a referent and asymmetries in shared knowledge between interlocutors.

The theoretical claims that motivated these studies can be summarized as follows:

- Conceptual concreteness and categorical specificity are distinct semantic variables that reflect different aspects of the abstraction process and reference in word meaning.
- Differences in resolution and referential scope of semantic representations between general and specific words, together with the cognitive effort required to abstract a referent from instantiated single or multiple exemplars, may result in a relatively greater or smaller processing cost.
- Looser reference characterizing more general words in the specific communicative contexts may be perceived as indeterminate, whereas an appropriate degree of informativeness and precision encoded in specific words may be interpreted as more pragmatically exhaustive.
- The higher degrees of informativeness and semantic resolution associated with specific words may be consciously or unconsciously employed as abstraction-derived linguistic cues aimed at maximizing communicative success and transfer of knowledge.

Building on these claims and the existing scientific literature, the present work advances the hypothesis that language-mediated abstraction plays a systematic role in

shaping communicative strategies by regulating how informativeness and precision are encoded and interpreted. Within this theoretical perspective, the following research questions are formulated to investigate how categorical specificity influences the strategic deployment of linguistically derived abstraction in communicative behaviors.

### **Case study 1: conceptual concreteness and categorical specificity in word processing**

The well-established concreteness effect posits that concrete words display a processing advantage over abstract words (Jessen et al., 2000; Kounios & Holcomb, 1994; Romani et al., 2008). This effect appears to be consistent across both word recognition tasks (lexical decision) and categorization tasks (semantic decision; Pexman et al., 2017). In contrast, categorical specificity has been shown to exhibit opposite effects across these tasks. Differences in the degree of precision and variability encoded in semantic representations of general and specific words tend to produce task-dependent effects. In particular, general words are processed faster than specific words in lexical decision tasks (Iliev & Axelrod, 2017; Bolognesi & Caselli, 2023; Ravelli et al., 2025), whereas specific words show a processing advantage in semantic decision tasks (Ravelli et al., 2025).

However, previous studies have not investigated the effect of these variables on language processing using a controlled set of carefully selected and normed stimuli varying along both specificity and concreteness, and balanced for relevant lexical categories, such as word length and frequency.

#### **RQ1: Do concreteness and specificity affect lexical and semantic access to meaning independently and interactively?**

It is expected that the concreteness effect will be replicated across the tasks. In contrast, a consistent specificity effect across tasks is hypothesized, given the careful control of stimuli. In particular, due to their high-resolution, precision, and informativeness (Murphy, 2002), more specific words are expected to be processed faster than more general words, which are lower-dimensional (Lupyan & Mirman, 2013) and characterized by higher contextual variability (Rambelli & Bolognesi, 2023, 2024).

### **Case study 2: conceptual concreteness and categorical specificity in the plausibility of linguistic exchanges**

Within communicative contexts, abstract sentences are generally perceived as eliciting more uncertainty and requests for additional information and conversational turns (i.e., curiosity) than concrete sentences, which tend to elicit expressions of assertiveness and conversational closure. Moreover, abstract sentences are associated with increased requirements for social interaction, coordination, and deference to interlocutors' or experts' knowledge. However, the stimuli employed in previous studies typically reflected only a basic level of abstraction, without accounting for differences in categorical specificity.

The relationship between categorical specificity and conceptual concreteness has not yet been systematically investigated in conversational contexts. Gandolfi et al. (2023) proposed that abstract words are more likely to prompt social interaction and coordination because the absence of an easily accessible referent makes it more difficult for interlocutors to align their representations across different levels (e.g., situation models, phonology, syntax, semantics; Pickering & Garrod, 2004, 2021). In contrast, with concrete words, the retrieval of physical referents is more straightforward, allowing interlocutors to align their representations more easily. Similarly to abstract words, general words are characterized by high contextual variability and low-dimensionality, as they refer to multiple exemplars and contexts that share relatively few relevant features. As a result, interlocutors may find it more difficult to align their representations with general words. Conversely, specific words are highly informative and high-dimensional, since members of the category share a greater number of relevant features; consequently, interlocutors may align their representations more easily in sentences containing specific words that refer to single or fewer exemplars.

## **RQ2: How does categorical specificity shape uncertainty and curiosity in conversational dynamics?**

It is expected that general sentences will be judged as more plausible when paired with uncertainty and curiosity follow-ups, whereas specific sentences will be evaluated as more plausible when followed by expressions indicating comprehension and conversational closure. Moreover, general sentences are expected to receive higher scores in Social Metacognition and Word Confidence ratings than specific sentences, but lower scores in perceived need of experts, given that specific terms are typically mastered by specialists (Tanaka & Taylor, 1991).

### **Case study 3: conceptual concreteness and categorical specificity as communicative strategies in semi-naturalistic conversations**

Social interactions play a central role in the development of semantic and conceptual knowledge (Borghi et al., 2019, 2025; De Felice et al., 2023). Through language and communicative strategies, interlocutors establish a common ground (Clark & Brennan, 1991), align their representations (Pickering & Garrod, 2004; Gandolfi et al., 2023), and reactivate multimodal conceptual and semantic information (Lupyan & Thompson-Schill, 2012; Schilhab, 2015). However, the roles of concreteness and categorical specificity have not yet been examined in ecological, semi-naturalistic conversations.

Vigliocco et al. (2019) showed that caregivers employ iconic gestures and indexical cues more frequently when referents are physically absent and novel to children. These communicative strategies exploit imagistic properties to evoke referents, support information transfer, and facilitate knowledge acquisition. In line with *Word as Cultivators* (WAC) of others' mind hypothesis (Schilhab, 2015) and *label feedback hypothesis* (Lupyan & Thompson-Schill, 2012), linguistic cues may similarly reactivate perceptual and semantic information, particularly in communicative contexts aimed at knowledge transfer (Schilhab, 2015). Accordingly, speakers may strategically employ more concrete and specific language to maximize imagery and informativeness depending on specific conditions, such as novel or absent referents.

#### **RQ3: Do the concreteness and the specificity level of the speakers' language change in function of the four conditions?**

It is expected that speakers will increase the concreteness and specificity of their language when referents are both unknown and physically absent

The subsequent section presents empirical analyses addressing these research questions.

All the behavioral experiments reported in this dissertation were conducted with the Ethical approval of the Ethics Committee of the University of Bologna (Protocol n. 0121435). All participants received instructions, provided informed consent, and agreed to participate in the experiments voluntarily. The behavioral tasks and data collection

were conducted at the Experimental Lab of the Department of Modern Languages, Literatures and Cultures of the University of Bologna (“Views and opinions expressed are however, those of the author(s) only and do not necessarily reflect those of the Experimental Lab”); or online using Qualtrics and Prolific, in accordance with EU data protection regulations.

The case studies presented and analyzed in the following chapters draw on research output at different stages of publication: case study 1 (Chapter 4) is derived from (Lamarra et al., 2026); case study 2 (Chapter 5) is based on a manuscript currently submitted (Lamarra, Villani, Mazzuca, Borghi & Bolognesi, *submitted*); Case study 3 (Chapter 6) forms the basis of a manuscript currently in preparation (Lamarra, Ravelli, Bolognesi & Vigliocco, *in preparation*).



**PART 2**  
**EXPERIMENTAL SECTION:**  
**CASE STUDIES**



## Chapter 4

### Concreteness and Specificity as Isolated Words


This study investigates the influence of two semantic psycholinguistic variables, namely conceptual concreteness and categorical specificity, on word processing. These variables were examined using isolated words within classical psycholinguistic experimental paradigms.

Experiment 1 consisted of a rating task in which participants judged the Italian words used as stimuli for imageability, concreteness, and specificity. These ratings were subsequently used to norm the stimuli and served as factors in the analyses of the following two experiments.

Experiment 2 employed a lexical decision task (LDT) to assess the effects of the variables of interest on lexical access to meaning. Following the procedure introduced by Meyer & Schvaneveldt (1971), participants were asked to determine whether each stimulus was an Italian word (e.g., *tavolo*) or a nonword (e.g., *tolavo*).

Experiment 3 was an abstract/concrete semantic decision task (SDT), based on the task described in Pexman et al. (2017). In this task, participants categorized each presented word according to whether it referred to a concrete (e.g., “watermelon”) or an abstract entity (e.g., “coherence”). The use of both an LDT and an SDT aligns with the *semantic richness* framework (Balota et al., 1991; McClelland & Rumelhart, 1981; P. M. Pexman et al., 2017; Yap et al., 2015). According to this framework, both lexical and semantic information contribute to performance in word processing tasks; however, the extent of their respective contributions may depend on the specific demands of the task.

A schematic overview of the experimental design and methodology is provided in Figure 4.1.



**Figure 4.1.** Overview of the experimental design and methodological structure of the three experiments described in this chapter.

## 4.1. Methodology

For each of the three experiments, a total of 120 Italian words were selected from a previous study (Villani et al., 2022; see Appendix A). This dataset consisted of 30 abstract and 30 concrete words, each of them divided into three semantic domains: Animals ( $N = 10$ ), Food ( $N = 10$ ), and Tools ( $N = 10$ ), related to concrete concepts; PSTQ (*physical-spatio-temporal-quantitative*;  $N = 10$ ), PS (*philosophical-spiritual*;  $N = 10$ ), and EMS (*emotional-moral-social*;  $N = 10$ ), related to abstract words. In line with the existing literature (Banks et al., 2023), the abstract-concrete distinction is not considered a dichotomy, but rather as a continuum. Indeed, it has been proposed that the distribution of multidimensional experiences can vary across different kinds of abstract and concrete concepts. Therefore, the division into subcategories or semantic domains was implemented to guarantee greater granularity in the selection of the stimuli, enabling a more accurate interpretation of the results.

For each of the words included in Villani et al. (2022), a hypernym and a hyponym were identified based on the conceptual-semantic and lexical relations (i.e., taxonomies) included in the lexical database WordNet (Miller, 1995), specifically in the *MultiWordNet*

database (<https://multiwordnet.fbk.eu>). When lexical entries were not available in the resource, more general (i.e., hypernym or superordinate) and more specific (i.e., hyponym or subordinate) words were selected by consulting the Treccani Italian dictionary (<https://www.treccani.it/enciclopedia/dizionario>). The selection included all those terms that explicitly refer to a potential hypernym or hyponym in the description. The use of dictionaries as an integrative resource to WordNet was motivated by the assumption that both reflected encyclopedic knowledge. In cases when the original word from the previous study was already relatively specific (e.g., “hen”) or general (e.g., “insect”), the original word was retained as a stimulus.

Consequently, the final set of 120 words employed across the three experiments comprised 30 concrete and general words (e.g., “object”), 30 concrete and specific words (e.g., “beach umbrella”), 30 abstract and general words (e.g., “relationship”), and 30 abstract and specific words (e.g., “fellowship”), balanced in terms of word length (i.e., number of letters) and word frequency. The frequency of the words was estimated from SketchEngine based on the occurrences observed in the itTenTen20 corpus (Jakubíček et al., 2013), and subsequently log-transformed. A one-way analysis of variance (ANOVA) implemented in R through the `aov()` function demonstrated that the four kinds of words did not vary in frequency,  $F(3, 116) = 0.506, p = .679$  (abstract-general,  $M = 8.30$ ,  $SD = 1.88$ ; abstract-specific,  $M = 8.13$ ,  $SD = 2.27$ ; concrete-general,  $M = 8.03$ ,  $SD = 2.54$ ; and concrete-specific,  $M = 7.73$ ,  $SD = 1.98$ ;) and length,  $F(3, 116) = 0.357, p = .784$  (abstract-general,  $M = 5.09$ ,  $SD = 0.50$ ; abstract-specific,  $M = 5.02$ ,  $SD = 0.73$ ; concrete-general,  $M = 5.01$ ,  $SD = 0.75$ ; and concrete-specific,  $M = 4.89$ ,  $SD = 0.42$ ).

**Ratings (Experiment 1).** In this first Experiment, a total of 90 participants took part in three distinct rating tasks. Thirty participants rated the final set of stimuli in terms of concreteness (i.e., how much the word referred to something perceivable through the five senses, versus something hardly perceivable; 1 = “very abstract”, 5 = “very concrete”). A different group of thirty participants evaluated the words for imageability (i.e., how easily a word could evoke a mental image; 1 = “difficult to imagine”, 7 = “easy to imagine”). Finally, thirty participants rated the stimuli along the specificity dimension (i.e., how much a word referred to something highly specific and precise or a wider group of entities/objects or situations; 1 = “not at all specific”, 5 = “very specific”).

Each of the three rating tasks was implemented as a questionnaire on Qualtrics (<https://www.qualtrics.com/>). For the concreteness and the imageability ratings, participants were recruited using the Prolific crowdsourcing platform (<https://www.prolific.com/>). Participants who completed the specificity ratings were recruited through personal and professional networks. Each survey included two blocks of 60 items. Before rating the actual stimuli, participants completed a set of practice trials.

This experiment also served to operationalize each dimension as a continuous variable for implementation in the following two experiments. Imageability, in particular, was included as previous studies have demonstrated its validity as a stronger predictor of lexical decision latencies compared to concreteness (Vergallito et al., 2020).

**Lexical Decision Task (Experiment 2).** For this experiment, in which participants were asked to recognize whether the presented string of letters was a word or a nonword (Meyer & Schvaneveldt, 1971), a total of 120 nonwords were created using an online tool (<https://www.trainingcognitivo.it/GC/nonparole/>) and were employed as stimuli in addition to the 120 Italian words used in the ratings. The nonwords were constructed to have an equivalent number of syllables and length as the real words (see Appendix A).

The participants in this experiment were 43 students (32 female, 11 males;  $M$  age = 22.9,  $SD$  age = 3.14) of the University of Bologna who voluntarily agreed to participate in the lexical decision task. All participants were native Italian speakers and naïve regarding the purpose of the experiment.

The experiment was implemented on a PC via *E-Prime 3.0*. (Psychology Software Tools, Pittsburgh, PA). Each participant completed the task in a quiet, low-light room at the Department of Modern Languages, Literatures and Cultures, University of Bologna. Participants were seated in front of a monitor, and stimuli were presented in the middle of the screen. The experimental session began with a fixation cross appearing on the screen for 1000 *ms* (i.e., milliseconds), followed by a blank screen presented for 500 *ms*. The target string of letters was then displayed on the screen for 500 *ms*, after which participants had an additional 1000 *ms* to respond. To provide their answers and recognize trials as real words or nonwords, participants pressed two distinct keys on the keyboard (i.e., “x” and “m”), with response keys counterbalanced across participants. In the experiment, the final set of 240 items was distributed across two lists, each comprising 60 words and 60 nonwords. The 60 words incorporated 15 concrete-general words (5

foods, 5 tools, 5 animals), 15 concrete-specific words (5 foods, 5 tools, 5 animals), 15 abstract-general words (5 PSTQ, 5 PS, 5 EMS), 15 abstract-specific words (5 PSTQ, 5 PS, 5 EMS). Each list was divided into two blocks, each composed of 30 words (15 abstract, 15 concrete) and 30 nonwords. The stimulus presentation was randomized within the four blocks. To prevent potential semantic bias, words belonging to the same taxonomic relation were not presented within the same block. For instance, the more general term “mammal” and the more specific “dromedary” were assigned to separate blocks. Moreover, short breaks were included between the blocks, while a longer break was implemented after the completion of the first list, during which participants completed simple arithmetic calculations.

Before beginning the lexical task, participants were familiarized with the procedure by categorizing eight practice trials (4 words and 4 nonwords).

**Semantic Decision Task (Experiment 3).** This experiment implemented an abstract/concrete semantic decision task, as introduced by Pexman et al. (2017), in which each word appearing on the screen was categorized as either concrete or abstract. The stimuli were the same set of 120 words used in the two previous experiments.

Consistent with the previous experiment, the sample size included an equivalent number of participants. Specifically, 44 students from the University of Bologna voluntarily took part in the semantic decision task. All participants were native Italian speakers and were unaware of the experiment’s purpose.

The location and conditions of the experimental session were identical to those employed in the lexical decision task. However, the task began with a fixation cross displayed on the screen for 1000 *ms*, followed by a blank screen lasting for 500 *ms*. The target word was then presented on the screen for up to 3000 *ms* or until the participant responded. Participants categorized each word as abstract or concrete by pressing one of the two keys on the keyboard (i.e., “x” and “m”). The assignment of the key responses was counterbalanced across participants. The 120 words were distributed across two separate blocks, each comprising 30 concrete and 30 abstract words. These words varied in specificity and were balanced across the semantic domains, as in Experiment 2. Similarly, hypernyms and hyponyms related to the same taxonomy were prevented from appearing in the same block. Before the main task, participants completed a practice session consisting of 8 trials: 4 abstract words (2 general and 2 specific) and 4 concrete

words (2 general and 2 specific). A single break was included after 60 trials to prevent potential semantic priming effects, as in the lexical decision task. During this pause, participants performed simple arithmetic calculations.

## 4.2. Data Analyses and Results

All data analyses and visualization for the experiments were conducted using R (R Core Team, 2023) and RStudio (v. 4.2.3.). Rating scores and reaction times (RTs) were analyzed with Linear mixed-effects models using the *lmer()* function from the “lme4” package on R (Bates et al., 2015). Post-hoc contrasts on the rating scores were performed with the “emmeans” package (Lenth, 2021), employing Tukey’s adjustment for multiple comparisons.

**Ratings (Experiment 1).** Firstly, Pearson’s correlation revealed an expected strong positive correlation between concreteness and imageability ( $r = 0.96$ ), a moderate positive correlation between specificity and imageability ( $r = 0.35$ ), and a low positive correlation between concreteness and specificity ( $r = 0.29$ ).

Subsequently, three separate linear mixed models were performed, each featuring a specific rating score as the dependent variable (i.e., concreteness, imageability, and specificity). The models included word type as an independent variable (abstract-general vs. abstract-specific vs. concrete-general vs. concrete-specific), and participants and words as random intercepts. The Results demonstrated a significant overall effect of word type on concreteness,  $F(3, 116) = 345.47, p < .001$ ; imageability,  $F(3, 116) = 167.67, p < .001$ ; and specificity ratings,  $F(3, 116) = 85.821, p < .001$ .

Post-hoc contrasts revealed significant differences between abstract-general and concrete-general, as well as between abstract-specific and concrete-specific, all  $ps < .001$ , in both the concreteness and imageability models. Moreover, no significant differences were found between abstract-general and abstract-specific in both imageability and concreteness, all  $ps > .44$ . In contrast, unexpected significant differences were observed between concrete-general and concrete-specific in terms of both imageability and concreteness, all  $ps < .0006$ . Additionally, significant and equally surprising differences emerged between abstract-specific and concrete-specific, as well as abstract-general and concrete-general, within the specificity models, all  $ps < .031$  (see Table 4.1. for the means

and standard deviations for all types of words across the dimensions of concreteness, imageability, and specificity).

|                   | Concreteness |           | Imageability |           | Specificity |           |
|-------------------|--------------|-----------|--------------|-----------|-------------|-----------|
|                   | <i>Mean</i>  | <i>SD</i> | <i>Mean</i>  | <i>SD</i> | <i>Mean</i> | <i>SD</i> |
| Abstract-General  | 1.78         | 0.988     | 2.23         | 1.16      | 2.9         | 1.30      |
| Abstract-Specific | 1.97         | 1.18      | 2.41         | 1.25      | 3.43        | 1.25      |
| Concrete-General  | 4.46         | 0.819     | 4.30         | 0.957     | 2.55        | 1.23      |
| Concrete-Specific | 4.99         | 0.185     | 4.86         | 0.488     | 4.48        | 0.868     |

**Table 4.1.** Means and standard deviations for concreteness, imageability, and specificity ratings across the four word types. Reproduced from Lamarra et al. (2026).

**Lexical Decision Task (Experiment 2).** During pre-processing, practice trials, nonwords, and errors (4.8%) were excluded from subsequent analyses. All participants demonstrated an accuracy rate exceeding 80% and were therefore eligible for inclusion in the analyses. Additionally, 3 stimuli (i.e., *prodotto agricolo* “agricultural product”; *elisio* “elysium”; *dromedario* “dromedary”) were removed from the analyses due to an accuracy rate falling below 70%.

The high correlation observed between concreteness and imageability suggested a potential collinearity among predictors (Zuur et al., 2010). To assess this, Variance Inflation Factors (VIF) were computed (Fox, 2005) via the `vif()` function from R’s package “car” (Fox & Weisberg, 2019). RTs (i.e., reaction times) were modeled as a function of concreteness, imageability, and specificity, which were included as fixed effects. Random intercepts were specified for participants and words to account for variability at the individual and item levels. As expected, concreteness and imageability exceed the threshold score of 10 (11.96 and 12.57, respectively), indicating substantial collinearity between these two predictors. In contrast, specificity showed the lowest VIF value (1.20). Given that the primary interest of this work is the relationship between concreteness and specificity, and that imageability yielded the highest VIF value, imageability was excluded from subsequent models. However, models including both


specificity and imageability as fixed effects were also analyzed. All scripts and materials are available on the following OSF repository link: <https://osf.io/j2smv/>.

A stepwise regression was conducted to analyze the models. Following this procedure, a series of linear mixed-effects models was performed using the `lmer()` function from the “lme4” package in R (Bates et al., 2015). First, a baseline model was specified with only dependent variables (i.e., RTs) and random intercepts (Subject and words). Concreteness was then introduced in the model as a predictor. This was followed by a model including only specificity as fixed effects. Subsequently, a model incorporating both concreteness and specificity was fitted. Finally, a model tested the interaction between these two variables (see Table 4.2.).

| Model   | Dependent variable | Fixed Effects              | Random Effects                   |
|---------|--------------------|----------------------------|----------------------------------|
| Model 0 | RT                 | /                          | (1   Participants) + (1   Words) |
| Model 1 | RT                 | Concreteness               | (1   Participants) + (1   Words) |
| Model 2 | RT                 | Specificity                | (1   Participants) + (1   Words) |
| Model 3 | RT                 | Concreteness + Specificity | (1   Participants) + (1   Words) |
| Model 4 | RT                 | Concreteness * Specificity | (1   Participants) + (1   Words) |

**Table 4.2.** Summary of linear mixed models, each incorporating additional fixed effects in predicting reaction times.

No significant differences in concreteness or specificity, nor significant interaction between them, were observed in the models (see Figure 4.2.) However, concreteness almost reached significance in both Model 1 and Model 3 ( $p = 0.06$  and  $p = 0.056$ , respectively).



**Figure 4.2.** Interaction effect between Concreteness and Specificity variables in predicting RTs in the Lexical Decision Task. Blue line refers to low specificity scores (i.e., general words) and red line refers to high specificity scores (i.e., specific words). Reproduced from Lamarra et al. (2026).

A comparative analysis of the model fits was further conducted using the `anova()` function. Model comparisons suggested that the inclusion of concreteness in Model 1 improved model fit relative to Model 0 in explaining variance in response latencies, although the improvement was not statistically significant. Neither the addition of specificity in Model 3 nor the interaction between specificity and concreteness in Model 4 significantly improved the model fit (see Table 4.3.).

|         | AIC   | BIC   | logLik | deviance | Chisq  | df | Pr(>Chisq) |
|---------|-------|-------|--------|----------|--------|----|------------|
| Model 0 | 60699 | 60725 | -30346 | 60691    |        |    |            |
| Model 1 | 60698 | 60730 | -30344 | 60688    | 35.580 | 1  | 0.059      |

|         |       |       |        |       |        |   |      |
|---------|-------|-------|--------|-------|--------|---|------|
| Model 3 | 60699 | 60738 | -30344 | 60687 | 0.2055 | 1 | 0.65 |
| Model 4 | 60701 | 60746 | -30344 | 60687 | 0.3220 | 1 | 0.57 |

**Table 4.3.** Lexical Decision Task: model comparison metrics with Concreteness added first. Reproduced from Lamarra et al. (2026).

An additional model comparison revealed that including specificity as the first predictor in Model 2 did not improve the model fit relative to Model 0. Model 3, which added concreteness alongside specificity, showed a marginal improvement. Finally, the inclusion of the interaction between the two variables did not result in a significant improvement in the model fit (see Table 4.4.).

|         | AIC   | BIC   | logLik | Deviance | Chisq  | df | Pr(>Chisq) |
|---------|-------|-------|--------|----------|--------|----|------------|
| Model 0 | 60699 | 60725 | -30346 | 60691    |        |    |            |
| Model 1 | 60701 | 60734 | -30346 | 60691    | 0.0209 | 1  | 0.885      |
| Model 3 | 60699 | 60738 | -30344 | 60687    | 3.7426 | 1  | 0.053      |
| Model 4 | 60701 | 60746 | -30344 | 60687    | 0.3220 | 1  | 0.57       |

**Table 4.4.** Lexical Decision Task: model comparison metrics with Specificity added first. Reproduced from Lamarra et al. (2026).


**Semantic Decision Task (Experiment 3).** Practice trials and errors were excluded from subsequent analyses. All participants reached an accuracy rate exceeding 80% and were therefore retained for further analyses. However, participants struggled to accurately categorize 8 stimuli (i.e., *dollaro* “dollar”; *valuta* “currency”; *paga* “payment”; *diagramma* “diagram”; *copertura* “coverage”; *elisio* “elysium”; *sezione* “section”; *rivolta* “revolt”) as concrete or abstract, with accuracy rates falling below 70%. Consequently, these items were excluded from the analyses.

As in Experiment 2, linear mixed-effects models were employed to analyze decision latencies from the semantic decision task via the `lmer()` function from the “lme4” package in R (Bates et al., 2015). RTs were modeled as a function of the three predictors (i.e., concreteness, imageability, and specificity). Random intercepts were specified for participants and words. Given the previously noted collinearity between concreteness and

imageability in the lexical decision task, `vif()` function from R's package "car" (Fox & Weisberg, 2019) was used to assess potential collinearity among predictors in the regression. Variance Inflation Factors (VIF; Fox, 2005) values indicated substantial collinearity with concreteness and imageability exceeding the conventional threshold of 10 (12.65 and 13.27, respectively). In contrast, the VIF value associated with specificity was much lower (1.19). Due to the highest VIF, imageability was excluded from the final models. For transparency, models including both specificity and imageability are available in the OSF repository.

In line with the previous experiment, a stepwise regression approach was used to analyze the final models. Reaction times (RTs) were included as the dependent variable, concreteness and specificity as fixed effects, and participants and words as random intercepts. The model inclusion sequence began with the baseline Model 0, which incorporated only random effects. Concreteness was added in Model 1, while specificity was the sole predictor considered in Model 2. Model 3 included both concreteness and specificity as independent variables, and Model 4 examined their interaction (see Table 4.2. above).

Results from Model 1 revealed a significant effect of concreteness on the dependent variable,  $F(1, 108.35) = 69.864, p < .001$ . Model 2 reported a significant effect on reaction times,  $F(1, 109.15) = 20.958, p < .001$ . A significant effect of both concreteness and specificity was observed in Model 3,  $F(1, 107.47) = 54.05, p < .001$ , and  $F(1, 107.55) = 9.726, p = .002$ , respectively. However, Model 4 revealed a significant effect of concreteness,  $F(1, 107.14) = 5.965, p = .0162$ , but no significant effect of specificity,  $F(1, 107.45) = 3.15, p = .0787$ , or interaction,  $F(1, 107.11) = 0.84, p < .36$  (see Figure 4.3.).



**Figure 4.3.** Interaction effect between Concreteness and Specificity variables in predicting RTs in the Semantic Decision Task. Blue line refers to low specificity scores (i.e., general words) and red line refers to high specificity scores (i.e., specific words). Reproduced from Lamarra et al. (2026).

Model fit comparisons indicated that the inclusion of concreteness in Model 1 led to a significant improvement in model fit relative to the baseline Model 0. Crucially, Model 3 demonstrated that adding specificity to the model significantly improved model fit compared to Model 2, which included only concreteness. However, Model 4, incorporating the interaction between the variables, did not provide a significant enhancement to the model fit (see Table 4.5.).

|         | AIC   | BIC   | logLik | deviance | Chisq   | df | Pr(>Chisq) |
|---------|-------|-------|--------|----------|---------|----|------------|
| Model 0 | 67471 | 67497 | -33732 | 67463    |         |    |            |
| Model 1 | 67419 | 67451 | -33704 | 67409    | 54.6250 | 1  | <0.001     |
| Model 3 | 67411 | 67450 | -33700 | 67339    | 9.5065  | 1  | 0.002      |
| Model 4 | 67412 | 67457 | -33699 | 67398    | 0.8645  | 1  | 0.3525     |

**Table 4.5.** Semantic Decision Task: model comparison metrics with Concreteness added first. Reproduced from Lamarra et al. (2026).

The second comparative analysis, in which specificity was added first, demonstrated that including specificity in Model 2 significantly improved model fit relative to the baseline Model 0. Model 3 yielded a similar significant improvement. Nonetheless, the inclusion of the interaction between predictors in Model 4 did not result in a significant improvement in model fit (see Table 4.6.).

|         | AIC   | BIC   | logLik | Deviance | Chisq   | df | Pr(>Chisq) |
|---------|-------|-------|--------|----------|---------|----|------------|
| Model 0 | 67471 | 67497 | -33732 | 67463    |         |    |            |
| Model 2 | 67454 | 67486 | -33722 | 67444    | 194.029 | 1  | <0.001     |
| Model 3 | 67411 | 67450 | -33700 | 67339    | 44.7286 | 1  | <0.001     |
| Model 4 | 67412 | 67457 | -33699 | 67398    | 0.8645  | 1  | 0.3525     |

**Table 4.6.** Semantic Decision Task: model comparison metrics with Specificity added first. Reproduced from Lamarra et al. (2026).

### 4.3. Discussion

Results from Experiment 1 (Ratings) align with previous research. This study confirmed the positive but mild correlation between specificity and concreteness observed by Bolognesi et al. (2020). Conversely, as reported in previous studies (Richardson, 1976; Connell & Lynott, 2012; Khanna & Cortese, 2021), concreteness and imageability showed a higher correlation. Finally, as a novel finding, imageability and specificity exhibited a moderate positive correlation.

The ratings provide valuable evidence regarding the taxonomic organization of speakers' conceptual knowledge. These findings demonstrate that both concrete and

abstract concepts can be characterized by different levels of specificity, although taxonomic hierarchies are less pronounced for abstract concepts, in line with what has been advanced by Borghi (2022). Moreover, a significant difference in concreteness was observed between concrete-general (e.g., FOOD) and concrete-specific concepts (e.g., TART), whereas no difference in concreteness was found between abstract-general (e.g., ATTRIBUTE) and abstract-specific concepts (e.g., EUPHORIA). A plausible explanation for this distinction lies in their different denotative ranges and referential structures. General words and concepts typically refer to broader categories that may denote multiple potential members or exemplars. In contrast, specific words and concepts refer to less extensive but more narrowly defined categories. The broader denotative potential of general concepts may influence their representation in terms of perceived concreteness. For instance, a concrete general-concept like SUBSTANCE may be rated as slightly less concrete due to its ability to instantiate a wider set of exemplars, while a more specific and defined concept like ASPIRIN may be rated as more concrete. These differences appear not to likely emerge between abstract concepts, because both abstract-general and abstract-specific words typically lack univocal and tangible referents and refer to multiple situations.

Results from Experiment 2 (Lexical Decision Task) showed that concreteness and specificity did not affect response latencies at the lexical level. This may be due to the semantic nature of these variables. In lexical decision tasks, shallow processing is involved, with participants relying more on word forms rather than engaging with a full semantic understanding (Balota et al., 1991; Yap et al., 2011; Pexman et al., 2017). Semantic variables, such as concreteness and specificity, typically require deeper semantic processing of the stimuli. These findings are consistent with models of word processing that propose an interactive and cascading sequence of processes (McClelland & Rumelhart, 1981; Balota et al., 1991; Yap et al., 2015). According to this framework, orthographic inputs are initially processed at the letter level, subsequently at the lexical level, where structural and formal properties of the words are accessed (e.g., frequency). Semantic level is the deepest level and involves the activation of conceptual meaning. Information flows bidirectionally among these levels (interactive activation), which can be activated simultaneously without requiring full completion of the preceding level (cascading activation). Thus, information can flow between semantic units (*meaning-level units*) and lexical units (*word-level units*). Crucially, the contribution of lexical and

semantic information can vary depending on the task demands. Although semantic information can support visual word recognition by providing feedback activation to the lexical level, it generally plays a secondary role compared to lexical-level properties, such as word length and frequency (Yap et al., 2011; Pexman et al., 2017).

In the present task, the absence of significant effects for both concreteness and specificity may also be attributed to the short stimulus exposure times. In contrast, previous works reported longer presentation durations, such as 2000 *ms* employed in the study by Vergallito et al. (2020). Nevertheless, the exploratory models available in the OSF repository revealed a significant effect of imageability, suggesting two key implications: first, that the presentation durations were sufficient to engage the semantic level; and second, that distinct semantic variables may require varying exposure times to activate their specific semantic information and affect word recognition in lexical decision tasks. This interpretation is also supported by the effect of concreteness that approached significance.

From a visual inspection of the interaction between concreteness and specificity, some notable trends emerge (see Figure 4.2.). Concepts with lower concreteness but high specific scores (i.e., abstract-specific concepts) appear to be the slower concepts to be processed, revealing a difference in processing between the two kinds of abstract concepts. A similar difference does not seem to emerge between concrete-general and concrete-specific concepts. Finally, concepts with lower specificity scores (i.e., general concepts, both abstract and concrete) did not differ in processing times.

The semantic decision task (Experiment 3) replicated the concreteness effect, showing a processing advantage for concepts with higher concreteness scores. Concrete concepts required shorter processing responses compared to abstract concepts. Notably, higher specificity scores were associated with shorter reaction times, revealing what can be defined as a “specificity effect”. Specifically, general concepts were processed more slowly than specific concepts. This may be due to the broader range of category members denoted by general concepts and contexts associated with them (Rambelli & Bolognesi, 2023, 2024). One plausible interpretation is that more specific words facilitate access to semantic representations by activating narrower and denser networks, whereas more general concepts may evoke broader and less defined networks, potentially slowing processing.

The absence of an interaction between concreteness and specificity suggests that these variables contribute independently to the prediction of the dependent variables (i.e., reaction times), despite both exhibiting similar negative effects. Furthermore, their additive effects on response times in the semantic decision task, along with their low correlation ( $r = 0.29$ ), provide additional evidence that concreteness and specificity should not be conflated, as argued by Bolognesi et al. (2020). Rather, they represent two different dimensions of semantic representation and referential process.

## **Chapter 5**

### **Concreteness and Specificity in the Plausibility of Simulated Dialogues**


In two preregistered experiments (<https://osf.io/erxd6>), this study investigates the impact of specificity on conversational dynamics and its relation to concreteness, adopting the methodology used by Mazzuca, Villani et al. (2025), while using sentences that are more specific and more general than those abstract and concrete sentences employed in that study. In each experiment, the relationship between specificity and concreteness is further explored through the implementation of simulated conversations.

During the first session, participants were presented with target sentences varying in both concreteness and specificity, followed by two different types of follow-up expressions. Participants were asked to evaluate whether or not the linguistic exchanges could plausibly occur in a real conversation.

In Experiment 1, the follow-ups were expressions signaling uncertainty (e.g., “What do you mean?”) or certainty (e.g., “Well done”), while in Experiment 2, follow-ups reflected expressions conveying the need for clarification and the willingness to continue the conversation (i.e., curiosity; e.g., “Tell me more”) or the willingness to end the conversation (e.g., “Thanks”; see Appendix B).

Unlike Mazzuca, Villani et al. (2025), the present study did not include a third experiment on question types. This decision was based on the rationale that causal (“why”, “who”) and spatio-temporal questions (“where”, “when”) do not directly address requests of exemplification, making them less relevant for the specificity variable. Plausibility judgments and reaction times (RTs) are collected.

In the second session, following each experiment, participants rated the target sentences on a 7-point Likert scale for three semantic variables: Confidence, Social Metacognition, and Experts Need (see Figure 5.1. for a schematic overview of the experimental design and methodology).



**Figure 5.1.** Overview of the experimental design and methodology of the study described in this chapter.

While the target sentences were constructed ad hoc for the experiments, the follow-up expressions were retrieved from participants’ responses in a previous free production study (Villani et al., 2022), thereby enhancing the ecological validity of the present study (Banks et al., 2023; Vigliocco et al., 2024).

### 5.1. Norming studies

For each of the 42 sentences (21 abstract, 21 concrete) used in Mazzuca, Villani et al. (2025), a hyponym and a hypernym of the original target concept were identified, and two corresponding sentences, varying in specificity, were generated. For instance, given the concrete sentence “I bought some strawberries”, where “strawberries” represents the basic concept, a more specific sentence, “I bought some wild strawberries”, and a more general sentence, “I bought some food”, were created. The same procedure was applied to abstract sentences. Most of the stimuli were retrieved from WordNet (Miller, 1995); however, for concepts not available there, Italian dictionaries (i.e., Treccani Dictionary, <https://www.treccani.it/enciclopedia/dizionario>) were consulted. This process resulted in

a total of 84 sentences, varying in both concreteness and specificity (e.g., “I made a tart/dessert”; “I calculated the percentage/quantity”).

To generate and validate the target sentences to be used in the two experiments, two distinct norming studies were conducted as online questionnaires on Qualtrics (<https://www.qualtrics.com/it/>).

**Validation of general-specific pairs.** The first norming study aimed to assess whether the level of specificity of specific/general word pairs retrieved from WordNet and dictionaries aligned with speakers’ judgments. A sample of 25 participants (10 male, 13 female, 2 unspecified; M age = 24.4; SD age = 4.50) was presented with each pair of specific and general words and asked to evaluate which word was the more specific or the more general, selecting a binary response (i.e., yes/no). For instance, given the pair “dalmatian/animal”, participants were asked whether “dalmatian” was more specific than “animal” and whether “dalmatian” was more general than “animal”, as well as the same questions with terms reversed. The order of specific and general word pairs was counterbalanced (i.e., 21 pairs specific-general and 21 pairs general-specific). The percentage of accuracy for each item was calculated, and five pairs of words with an accuracy < 70% were excluded. A total of 74 target sentences was obtained ( $n = 21$  pairs of general/specific concrete sentences and  $n = 16$  pairs of general/specific abstract sentences).

**Validation of target sentences.** Following the methodology employed by Mazzuca, Villani et al. (2025), a second norming study was conducted to balance the selected sentences across key dimensions. A sample of 39 participants (12 male, 27 female; M age = 29; SD age = 12.3) rated the resulting 74 sentences on a 7-point Likert scale in terms of:

- Concreteness (1= concrete; 7= abstract);
- Context Support (i.e., the extent to which they thought each sentence required a context to be understood: 1= not at all; 7= extremely);
- Naturalness (i.e., the extent to which they thought the sentence was “natural”: 1 = not at all; 7 = extremely natural);

- Familiarity (i.e., the extent to which they thought the sentence was familiar; 1 = not at all familiar; 7 = extremely familiar).

Additionally, participants were asked to indicate whether they knew the meaning of the target word (yes/no). Words known by fewer than 70% of participants were considered not eligible stimuli.

Each dimension was presented in a separate block. The blocks were randomized across participants, and the sentences were randomized within each block. The analysis revealed that 18 participants (46.15%) did not know the word *cacatua* (i.e., “cockatoo”). Consequently, the corresponding pair “cockatoo/feathered” was excluded from the analysis. To determine which sentences were suitable for inclusion in the study, a score of 4 on the Naturalness ratings was applied as a cut-off criterion. Data points where participants indicated they did not know the word were excluded from the analysis ( $n = 107$ ). Descriptive statistics showed that two sentences were rated, on average,  $< 4$  in Naturalness (i.e., “I tempered the tool”,  $M = 3.8$ , and “I searched for a dialectic”,  $M = 3.6$ ). Consequently, the corresponding pairs “I tempered the pastel pencil” and “I searched for a method” were also excluded. This resulted in a total of 68 target sentences (15 abstract-specific, 15 abstract-general, 20 concrete-specific, 20 concrete-general). To balance the number of sentences for sub-types, four additional pairs of concrete sentences with the lowest scores in naturalness ( $< 5.5$ ) were excluded: *bovino/mucca da latte*, *piccione viaggiatore/uccello*, *parapioggia/ombrellone pieghevole*, *alimento/carota viola* (i.e., “bovine/dairy cow”, “carrier pigeon/bird”, “rain cover/folding umbrella”, “food/purple carrot”). The final set of stimuli consisted of 60 sentences (15 abstract-specific, 15 abstract-general, 15 concrete-specific, 15 concrete-general).

Subsequently, specific and general sentences were balanced in terms of abstractness/concreteness, familiarity, naturalness, and context support. Linear mixed effect models were employed to analyze rating scores for each dimension as a function of the Level of the sentence (specific vs. general) with random intercepts for participants and sentences. Results showed that specific and general sentences do not differ for naturalness,  $X^2(1) = 0.005$ ,  $p = .94$ ; ( $EMM = 5.82$ ;  $SE = .218$ ,  $EMM = 5.80$ ;  $SE = .218$ , respectively), abstractness,  $X^2(1) = 0.1992$ ,  $p = .65$ , ( $EMM = 3.26$ ;  $SE = .257$ ,  $EMM = 3.39$ ;  $SE = .257$ , respectively) and context support,  $X^2(1) = 2.3318$ ,  $p = .13$ , ( $EMM = 2.92$ ;  $SE = .281$ ,  $EMM = 3.46$ ;  $SE = .281$ , respectively). However, specific sentences were rated

as less familiar ( $EMM = 4.51$ ;  $SE = .249$ ) than general sentences ( $EMM = 5.23$ ;  $SE = .249$ ),  $X^2(1) = 5.1577$ ,  $p = .023$ ). As established in scientific literature (Johnson & Mervis, 1997; Tanaka & Taylor, 1991; Bolognesi & Caselli, 2023), this result is likely due to the specific words' more specialized or unusual lexicon. Lastly, a final dataset of 60 sentences (15 abstract-specific, 15 abstract-general, 15 concrete-specific, 15 concrete-general) was retained for the subsequent analyses (see Appendix B).

## 5.2. Methodology


For each of the two Plausibility judgment tasks and Ratings, the 60 target sentences selected from the norming studies were used as materials (15 abstract-specific, 15 abstract-general, 15 concrete-specific, 15 concrete-general).

### Session 1

**Specificity and Uncertainty (Experiment 1).** Participants were presented with sentences varying in both concreteness and specificity, followed by two different types of follow-up expressions. Subsequently, they were asked to judge whether the linguistic exchanges shown could plausibly occur in a real conversation. The follow-ups reflected expressions signaling uncertainty and certainty. The follow-up expressions employed in this experiment were the same as those used in Mazzuca, Villani et al. (2025), as validated in their norming study (see Mazzuca et al., 2025, *Norming studies* section). These expressions corresponded to the two highest (i.e., “I did not understand”; “What do you mean?”) and two lowest values (i.e., “Good job”; “Well done”) on the Certainty-Uncertainty rating scale provided during their second norming study.

Stimuli were distributed across four blocks to ensure that each sentence was presented twice and paired with each type of follow-up within the same block. Furthermore, within each block, sentences were balanced to include an equal number of abstract and concrete sentences, as well as general and specific sentences. As in the experiments in Chapter 5, to avoid priming effects, words belonging to the same taxonomy were not presented within the same block. For instance, if “butterfly” was assigned to a block, then “flying animal” was placed in a distinct block. During the task, each target sentence was paired with each follow-up across the four blocks.

A total of 26 Italian native speakers (15 females, 11 males;  $M$  age = 21.5;  $SD$  = 1.57; age range = 19-25) were recruited at the University of Bologna and voluntarily agreed to participate in the experimental session. All participants were naïve to the purpose of the experiment. They were tested in a quiet, low-light environment at the Department of Modern Languages, Literatures and Cultures, University of Bologna. The experiment was implemented on a PC using E-Prime 3.0 (Psychology Software Tools, Pittsburgh, PA). The procedure followed that of Mazzuca, Villani et al. (2025; see Figure 5.2.).



**Figure 5.2.** Schematic representation of the experimental procedure.

In the first session, participants were presented with sentence/follow-up pairings on a computer screen and asked to judge whether the occurring linguistic exchanges could plausibly occur in an everyday conversation (yes vs. no). The order of sentences was randomized across participants, with a short break introduced between each block. The stimulus presentation consisted of a central black fixation cross displayed for 500 ms, followed by a target sentence, which remained on the screen for 1500 ms. The follow-up was then presented until a response was provided or for a maximum of 3000 ms. Reaction times (RTs) were recorded once the follow-up item was displayed. Participants responded by pressing two keys on the keyboard (i.e., “x” and “m”) to indicate whether the linguistic exchange was “plausible” or “not plausible”. The assignment of keys was counterbalanced across participants.

In the second session, participants, seated in front of the same PC, completed a rating task in which they were asked to evaluate the target sentences on a 7-point Likert scale across three distinct semantic dimensions (see *Ratings* section):

- Confidence (i.e., the extent to which participants feel confident about the meaning of the sentence; Mazzuca et al., 2022);
- Experts Need (i.e., the extent to which participants rely on experts' knowledge to understand the sentence; Falcinelli et al., 2024);
- Social Metacognition (i.e., the extent to which they think they need others' contribution to understand the sentence; Villani et al., 2019).

**Specificity and Curiosity (Experiment 2).** The procedure, design, and materials were identical to those used in Experiment 1. Participants were asked to evaluate the plausibility of target sentences that varied in both concreteness and specificity when paired with different follow-up expressions. In this experiment, however, the follow-up expressions conveyed either the willingness to continue the conversation and seek further information (i.e., curiosity) or the willingness to end and conclude the conversation. These expressions matched the two highest (i.e., “Describe”; “Tell me more”) and the two lowest scores (i.e., “Ok”; “Thanks”) on the End-Curiosity rating scale from the second norming study conducted by Mazzuca, Villani et al. (2025).

Twenty-seven Italian native speakers (19 females, 5 males, 1 identified as queer, 2 unspecified;  $M$  age = 22;  $SD$  = 1.85; age range = 19-25) were recruited at the University of Bologna and voluntarily agreed to participate in the experiment.

## **Session 2**

### **Ratings**

During the second session of each experiment, all participants ( $n = 53$ ) rated the 60 target sentences across three semantic dimensions related to metacognition: Confidence, Social Metacognition, and Experts Need. Specifically, participants evaluated each sentence on a 7-point Likert scale (1 = not at all; 7 = very much) in terms of:

- “How confident they were in having understood the sentence” (i.e., Confidence; Mazzuca et al., 2022)
- “How much they felt the need for other people’s input to understand the sentence” (i.e., Social Metacognition; Villani et al., 2019; Borghi, 2022);
- “How much they would have relied on experts’ knowledge to better understand the sentence” (i.e., Experts Need; Falcinelli et al., 2024; Mazzuca et al., 2022).

In each session, the order of both sentences and semantic dimensions was randomized across participants. Sentences were presented individually on the same computer screen in a separate session conducted via Qualtrics.

### **5.3. Data Analyses and Results**

Data analysis and visualization for all the experiments were performed with R (R Core Team, 2023) and RStudio (v. 4.4.1). To identify outliers, data points falling outside  $\pm 2SD$  from the average RTs of both “plausible” and “not plausible” responses were excluded from subsequent analyses.

#### **Session 1**

In both Experiment 1 and Experiment 2, mixed-effects binomial logistic regressions were performed using the `glmer()` function from the “lme4” package in R (Bates et al., 2015) to examine plausibility judgments (1 = yes vs. 0 = no) as a function of Level (general vs. specific), Type of follow-up (Experiment 1: uncertainty vs. certainty; Experiment 2: curiosity vs. end), and their interaction as fixed effects. Participants and sentences were modeled as random intercepts.

Wald Chi-Squared tests examine main effects and the significance of interactions via the `Anova()` function from the “car” package in R (Fox & Weisberg, 2019). In all experiments, RTs for “plausible” and “not plausible” responses were log-transformed due to detected skewness and then analyzed separately using linear mixed models implemented with the `lmer()` function from the “lme4” package (Bates et al., 2015). The models followed the structure as those used to assess plausibility judgements. Pairwise

differences were assessed using Tukey's adjustment in post-hoc contrasts performed using the "emmeans" package in R (Lenth, 2021).

To further explore the relationship between concreteness/abstractness and specificity/generalizability in conversational dynamics, additional exploratory models were conducted that were not included in the preregistration plan. In these models, Level (general vs. specific) was replaced with Quadrant, which contained information about all four types of sentences (abstract-specific vs. abstract-general vs. concrete-specific vs. concrete-general). This modification applied to models used for both plausibility judgements and RTs.


### **Specificity and Uncertainty (Experiment 1)**

Following the data exclusion criterion, 2.64% of data points ( $n = 165$ ) were removed from the analyses.

#### **Plausibility judgements**

**Level.** A significant main effect of Type of Follow-up was observed,  $X^2(1) = 1034.31$ ,  $p < .001$ . No other main effect or interaction reached significance, all  $ps > .20$ .

**Quadrant.** A significant main effect of Quadrant,  $X^2(3) = 18.615$ ,  $p < .001$ , and Type of Follow-up,  $X^2(1) = 936.948$ ,  $p < .001$ , were found. Additionally, a significant interaction between Quadrant and Type of Follow-up emerged from the model,  $X^2(3) = 155.149$ ,  $p < .001$ . Post-hoc contrasts indicated that, for all types of sentences, participants judged certainty follow-ups as more plausible compared to uncertainty follow-ups, all  $ps < .0001$ . However, abstract sentences, both general and specific, were considered more plausible when paired with uncertainty follow-ups compared to concrete sentences, both general and specific, all  $ps < .001$ . No significant difference was found within either abstract or concrete sentences, all  $ps > .27$ , nor within certainty follow-ups, all  $ps > .17$  (see Figure 5.3.).



**Figure 5.3.** Predicted probabilities of plausibility judgments across all sentence types, as a function of the type of follow-ups.

### Response times (RTs)

**Plausible. Level.** A significant main effect of Type of Follow-up was found,  $F(1, 2722.77) = 185.63, p < .001$ . Participants were overall slower with follow-ups signaling uncertainty ( $M = 1320, SD = 466$ ) compared to certainty follow-ups ( $M = 1033, SD = 441$ ). No other main effect or interaction reached significance, all  $ps > .30$ .

**Plausible. Quadrant.** Similarly, the model revealed a significant main effect of Type of Follow-up,  $F(1, 2979.7) = 157.82, p < .001$ . No other main effect or interaction was detected.

**Implausible. Level.** A significant main effect of Type of Follow-up was observed,  $F(1, 1916.6) = 23.096, p < .001$ , with participants being overall slower to judge certainty follow-ups as implausible ( $M = 1219, SD = 446$ ) than uncertainty follow-ups ( $M = 1113, SD = 465$ ). No other main effect or interaction was observed.

**Implausible. Quadrant.** The model revealed a significant main effect of Type of Follow-up,  $F(1, 1963.6) = 21.2309, p = < .001$ . Additionally, a significant interaction between Quadrant and Type of Follow-up was observed,  $F(3, 2043.72) = 4.2158, p = .0055$ . No other main effect reached significance (Quadrant,  $p = .86$ ). Post-hoc contrasts revealed that participants were faster to judge uncertainty follow-ups as implausible for both concrete-general,  $t(2275) = 3.162, p = .0016$ , and concrete specific sentences  $t(2503) = 4.96, p = < .0001$ , while no difference was found between type of follow-ups across abstract-general,  $t(1641) = 0.556, p = .578$ , and abstract-specific sentences,  $t(1532) = 0.972, p = .331$ . No differences were found within certainty follow-ups, all  $ps > 0.50$ . In contrast, within uncertainty follow-ups, a slight difference emerged between both types of abstract sentences (general and specific) and concrete-specific sentences, with abstract sentences-uncertainty pairings being responded to more slowly than concrete-specific sentences-uncertainty pairings, all  $ps < .033$ . Notably, no differences were observed between both types of abstract sentences and generic-concrete sentences, all  $ps > 0.54$ .

## Specificity and Curiosity (Experiment 2)


In this experiment, 2.68% of data points ( $n = 174$ ) were excluded from subsequent analyses.

### Plausibility judgments

**Level.** A significant main effect of Type of Follow-up was observed,  $X^2(1) = 290.09, p < .001$ , whereas no significant main effect of Level was found,  $X^2(1) = 2.23, p = .135$ . However, a significant two-way interaction between Level and Type of Follow-up emerged,  $X^2(1) = 118.25, p < .001$ . Post-hoc contrasts revealed that participants considered curiosity follow-ups more plausible than end follow-ups when paired with both general,  $z = 19.572, p < .0001$ , and specific sentences,  $z = 5.080, p < .0001$ . General sentences were judged as more plausible when paired with curiosity follow-ups compared to specific sentences,  $z = 6.472, p < .0001$ , while general-end pairings were judged as less plausible than specific-end pairings,  $z = -3.587, p < .001$ .

**Quadrant.** The model revealed a significant main effect of Quadrant,  $X^2(3) = 21.425, p < .001$ , a significant main effect of Type of Follow-up,  $X^2(1) = 214.276, p < .001$ , and a significant interaction between Quadrant and Type of Follow-up,  $X^2(3) = 745.321, p < .001$ .

.001. Post-hoc contrasts further revealed that, within the type of follow-ups, participants considered curiosity follow-ups more plausible than end follow-ups when paired with abstract sentences (both general and specific) and concrete-general sentences, all  $ps < .0001$ . However, concrete-specific sentences were judged as more plausible with end follow-ups,  $z = -12.456, p < .0001$ . Moreover, abstract sentences (general and specific) were considered more plausible when paired with curiosity follow-ups, all  $ps < .0001$ , and less plausible when paired with end follow-ups, all  $ps < .0223$ , compared to concrete sentences (both general and specific). No significant difference was found within abstract sentences. In contrast, within concrete sentences, concrete-general sentences were judged as more plausible when associated with curiosity follow-ups and less plausible when followed by end follow-ups, compared to concrete-specific sentences, all  $ps < .0001$  (see Figure 5.4.).



**Figure 5.4.** Predicted probabilities of plausibility judgments across all sentence types, as a function of the type of follow-ups.

### Response times (RTs)

**Plausible. Level.** No main effect or interaction was found.

**Plausible Quadrant.** No main effect or interaction was found.

**Implausible. Level.** A significant main effect of Type of Follow-up was observed,  $F(1,1332.8) = 47.7647, p < .001$ , while Level approached significance,  $F(1, 55.56) = 3.78, p = .05689$ . No interaction was detected. Overall, participants were slower to judge curiosity follow-ups as implausible ( $M = 1031; SD = 385$ ) compared to end follow-ups ( $M = 962; SD = 374$ ).

**Implausible. Quadrant.** The model revealed a main effect of Quadrant,  $F(3, 76.07) = 5.5274, p = .0017$ , and Type of Follow-up,  $F(1, 2527.75) = 58.5887, p < .001$ , but no significant interaction. Participants were generally faster to judge concrete-specific sentences as implausible ( $M = 978, SD = 358$ ) compared to both abstract sentences ( $M$  abstract-general = 995,  $SD$  abstract-general = 390;  $M$  abstract-specific = 995,  $SD$  abstract-specific = 401), and marginally faster than concrete-general sentences ( $M = 990, SD = 378$ ).

## Session 2

Ordinal mixed effects models were performed to analyze ratings on each semantic dimension using the `clmm()` function implemented via the “ordinal” package in R (Christensen, 2023). In each model, the rating score was treated as the dependent variable, the type of sentence (general vs. specific) as the independent variable, and participants and sentences were included as random intercepts. Similarly to the analysis of plausibility judgments and reaction Times (RTs), additional models were run. In these models, the type of sentence (general vs. specific) was replaced with Quadrant (abstract-specific vs. abstract-general vs. concrete-specific vs. concrete-general), while the rest of the model structure remained identical.


## Ratings

**Level.** The models revealed a significant difference between general and specific sentences on Social Metacognition variable,  $b = -0.6632, SE = 0.2636, z = -2.1516, p = .0119$ . General sentences ( $EMM = -1.19, SE = 0.267$ ) were associated with significantly higher Social Metacognition compared to specific sentences ( $EMM = -1.85, SE = 0.270$ ).

However, no statistically significant differences were found for Confidence and Experts Need, all  $ps > .416$  (see Table 5.1. for all means and the standard deviations and Figure 5.5. for distribution of ratings scores across dimensions and conditions).

|          | Confidence  |           | Social Metacognition |           | Experts Need |           |
|----------|-------------|-----------|----------------------|-----------|--------------|-----------|
|          | <i>Mean</i> | <i>SD</i> | <i>Mean</i>          | <i>SD</i> | <i>Mean</i>  | <i>SD</i> |
| General  | 4.93        | 1.55      | 3.05                 | 2.38      | 2.44         | 2.18      |
| Specific | 5.01        | 1.51      | 2.67                 | 2.30      | 2.51         | 2.21      |

**Table 5.1.** Means and Standard Deviations for Confidence, Social Metacognition, Experts Need ratings across the two types of sentences.



**Figure 5.5.** Rating scores for general and specific sentences across Confidence, Experts Need, and Social Metacognition dimensions.

**Quadrant.** A significant effect of Quadrant (abstract-specific vs. abstract-general vs. concrete-specific vs. concrete-general) was found on all three semantic variables. As expected, abstract sentences (general and specific) received higher scores in Social Metacognition and Experts Need and lower scores in Confidence, compared to concrete sentences (general and specific). Post-hoc contrasts showed that these differences were statistically significant across all models, all  $ps < .0243$ , with the smallest  $ps < .0001$ . However, there was no difference between the two types of abstract sentences. For concrete sentences, instead, participants provided higher Social Metacognition ratings for concrete-general sentences compared to concrete-specific sentences,  $z = 7.383, p < .0001$ . Therefore, the significant effect found in the model with Level (general vs. specific) was driven by concrete sentences. In addition, concrete-specific sentences were rated higher on Confidence than concrete-general sentences,  $z = -3.080, p = .011$ . Notably, no statistically significant difference between the two types of concrete sentences was observed in Experts Need (see Table 5.2. for all means and the standard deviations).

|                   | Confidence  |           | Social Metacognition |           | Experts Need |           |
|-------------------|-------------|-----------|----------------------|-----------|--------------|-----------|
|                   | <i>Mean</i> | <i>SD</i> | <i>Mean</i>          | <i>SD</i> | <i>Mean</i>  | <i>SD</i> |
| Abstract-General  | 4.71        | 1.64      | 3.31                 | 2.29      | 2.75         | 2.16      |
| Abstract-Specific | 4.55        | 1.64      | 3.49                 | 2.31      | 3.17         | 2.23      |
| Concrete-General  | 5.16        | 1.43      | 2.80                 | 2.45      | 2.14         | 2.16      |
| Concrete-Specific | 5.47        | 1.20      | 1.86                 | 1.99      | 1.85         | 1.97      |

**Table 5.2.** Means and Standard Deviations for Confidence, Social Metacognition, Experts Need ratings across the four types of sentences.

## 5.4. Discussion

Experiment 1 replicated findings from Mazzuca, Villani et al. (2025) with sentences varying in concreteness/abstractness and generality/specificity. Consistently, participants

considered certainty follow-ups more plausible than uncertainty follow-ups, as indicated by faster response times for judging certainty follow-ups as plausible.

Regardless of specificity levels, abstract sentences were considered more plausible when paired with uncertainty follow-ups compared to concrete sentences (see Figure 5.3.). No differences in plausibility were observed between levels of specificity (general vs. specific) within the same type of sentence (abstract or concrete). This supports the idea that uncertainty is a linguistic construct associated with the abstractness of a sentence (Borghi, 2022; Villani et al., 2022; Borghi et al., 2025; Mazzuca, Villani et al., 2025).

Additionally, participants exhibited slower response times when judging abstract sentences-uncertainty pairings as implausible. However, this trend was particularly evident when comparing abstract sentences with concrete-specific sentences, while no significant difference was observed when comparing abstract sentences with concrete-general sentences.

Experiment 2 demonstrated that, overall, participants considered curiosity follow-ups more plausible than end follow-ups. This finding was partially supported by slower response times in judging curiosity follow-ups as implausible across all types of sentences. However, general sentences paired with curiosity follow-ups were judged as more plausible than specific sentences-curiosity pairings, while specific sentences were considered more plausible than general sentences when paired with end follow-ups. The model featuring information about all types of sentences revealed that this effect was primarily driven by concrete sentences. Abstract sentences (general and specific) paired with curiosity follow-ups were considered more plausible than when associated with end follow-ups. Likely, abstract sentences-curiosity follow-ups were judged more plausible compared to concrete sentences-curiosity follow-ups.

Similar to uncertainty, curiosity (i.e., the need for clarification) is a dynamic associated with abstract sentences (Villani et al., 2022; Mazzuca, Villani et al., 2025), regardless of the specificity level. However, within concrete sentences, we observed that concrete-general sentences were judged more plausible when followed by curiosity follow-ups, while concrete-specific sentences were considered more plausible when associated with end follow-ups. Thus, like abstract sentences, concrete-general sentences were considered more plausible when suggesting a need for further interaction. In

contrast, concrete-specific sentences were considered more pragmatically exhaustive (Grice, 1989).

Results from ratings replicated the findings from Mazzuca, Villani et al. (2025), showing that abstractness elicits more social engagement and less confidence regarding the understanding of sentence meaning. In the rating sessions, this pattern was consistent across all levels of abstraction for abstract sentences. In contrast, results demonstrate that concrete sentences varied in terms of social interaction and confidence depending on the level of specificity. Concrete-general sentences elicited more Social Metacognition and less Confidence than concrete-specific sentences. The greater plurality of exemplars instantiated by concrete-general sentences may have introduced a degree of ambiguity (Magni, 2024), which may have increased the need for clarification. Conversely, the higher informativeness and specificity of concrete-specific sentences may have provided greater certainty regarding the meaning of the sentences.

## Chapter 6

### Concreteness and Specificity in Naturalistic Conversations

This study investigates how participants' language in naturalistic conversations varies in terms of concreteness and specificity as a conversational strategy to convey more vivid and informative representations. It examines the ECOLANG corpus (Gu et al., 2025), a dataset of multimodal, semi-naturalistic interactions between 31 dyads of English-speaking adults, where a knowledgeable *speaker* describes objects to an *addressee* under four experimental conditions:

- The object is physically situated in the scene (i.e., *Present* condition);
- The object is physically displaced from the scene (i.e., *Absent* condition);
- The object is familiar (e.g., “giraffe”) to the addressee (i.e., *Known* condition);
- The object is unfamiliar (e.g., “okapi”) to the addressee (i.e., *Unknown* condition).


Since the corpus recorded multimodal interactions with the stimuli, all objects used as stimuli represented highly concrete and specific referents of four semantic categories: fruits, musical instruments, tools, and animals.

In three separate analyses, the study illustrated in the present chapter investigates how the semantic and psycholinguistic variables of interest, namely concreteness and specificity, vary in the language produced by speakers under the four conditions. Moreover, this study explores the impact of experimental conditions on specific conversational strategies, as encoded in different linguistic variables: abstraction processes (i.e., analogical comparisons and categorization at the superordinate level), reference to objects' features (i.e., physical and encyclopedic; McRae et al., 2005), and reference to one's own emotional and personal experience with objects.

The first analysis examines whether the presence/absence or familiarity/unfamiliarity of the objects influences the degree of concreteness or the level of specificity employed by speakers in describing the objects.

The second analysis investigates whether the likelihood of employing the different conversational and linguistic strategies depends on the presence/absence or familiarity/unfamiliarity of the objects.

The third analysis explores which linguistic variable has the largest effect on the concreteness and specificity of language employed by speakers (see Figure 6.1. for a schematic overview of the experimental design and methodological structure).



**Figure 6.1.** Overview of the experimental design and methodological approach of the three analyses described in this chapter.

Overall, this study, for the first time, investigates concreteness and specificity, along with distinct abstraction strategies and conversational dynamics, in the context of ecologically valid data collected throughout experimental settings that more closely reflect natural language use (Banks et al., 2023; Vigliocco et al., 2024).

## 6.1. Methodology

This methodology section begins with an overview of the annotation process employed to extract all single nouns and multi-word combinations from the complete set of sentences produced by all adult speakers from the ECOLANG corpus (Gu et al., 2025).

## **Annotation**

An automatic Part-Of-Speech (POS) tagging annotation was performed on the total set of 22,581 adult speakers' sentences in the ECOLANG corpus using the `udpipe_annotate()` function from the UDPipe package (Wijffels, 2023) in R (R Core Team, 2023). A pre-trained model for the English language was utilized for tokenization and lemmatization, resulting in Universal Dependencies-style annotations (De Marneffe et al., 2021).

Given the absence of the complete conversational context (i.e., the addressees' sentences) and the spontaneous nature of the speech, the output was manually revised to ensure accuracy. As a result, all single nouns and multi-word combinations (Adjective + Noun; Noun + Noun) uttered by the speakers were extracted. These combinations were automatically identified based on syntactic dependency relations and linear token order. Specifically, pairs were extracted when the second token (i.e., a noun) served as the syntactic head of the preceding token (i.e., an adjective or noun). For instance, “dinner table” was selected because “table” syntactically followed “dinner” and functioned as its syntactic head. This procedure ensured that the extracted pairs reflected genuine syntactic dependencies rather than mere linear adjacency. The final dataset consisted of a total of 4,209 items, i.e., single nouns and multi-word combinations, that will be treated as nouns in subsequent sections.

## **Concreteness and Specificity across Conditions (Analysis 1)**

From the annotated dataset, only those nouns for which concreteness scores were available in both Brysbaert et al. (2014) and Muraki et al. (2022) were retained. The resulting final dataset for the first analysis consisted of 1,612 nouns produced in 6,612 sentences. In all such instances, a concreteness score for the noun was available in both Brysbaert et al. (2014) and Muraki et al. (2022); the score from Brysbaert et al. (2014) was retained, as it was the initial source consulted.

The Best-Worst Scaling method (Louviere et al., 2015) was employed to collect specificity ratings. This approach is currently considered the most effective methodology for capturing the relational nature of this variable (Bolognesi & Caselli, 2023; see Chapter 3). For example, “wind instrument” is generally considered a quite general word;

however, it is more specific than the even more general “instrument”, yet less specific than “flute” or “clarinet”.

The BWS was carried out using GPT-4, which has been shown to be a reliable proxy, given the strong correlation between human assessments and the responses generated by the LLM ( $r = 0.772$ ; Ravelli & Bolognesi, 2024). The model was configured according to the instructions provided by Ravelli et al. (2025), which were used as a system prompt to introduce each request, in line with the procedure described in Ravelli & Bolognesi (2024).

Since BWS scores fall within a -1 to 1 range, the values were rescaled to a 1-5 range to ensure consistency with the concreteness score reported in both Brysbaert et al. (2014) and Muraki et al. (2022).

### **Linguistic Variables across Conditions (Analysis 2)**

The entire set of 22,581 sentences produced by speakers was analyzed. In this second analysis, the distribution of several previously annotated linguistic variables was examined across the entire dataset. These variables included:

- A specific abstraction process, such as categorization through analogical correspondence between co-hyponyms (**Analogy**, e.g., “It’s like an orange”);
- A conversational dynamic revealing reference to speakers’ personal experiences with the object (**Social Self**, e.g., “I used to play it when I was at primary school”);
- A conversational dynamic signaling reference to the speakers’ own emotional dimension about the object (**Emotion Opinion**, e.g., “I like it”);
- Focus on objects’ perceptual or physical features (**Object**, e.g., “it’s spiky”);
- Focus on encyclopedic knowledge (McRae et al., 2005) that encompasses all types of knowledge about an object (**Situation**, e.g., “It was used in the sixteenth century era”).

Additionally, an ad hoc category was automatically derived to assess a second type of abstraction process, namely, categorization to a more general term at the superordinate level (**Superordinate categorization**). This was achieved through the identification of occurrences of the target objects’ hypernyms, previously extracted from WordNet (Miller,

1995), in strings that indicated reference to the superordinate level (e.g., “It’s an ancient instrument”; “As an animal”; “It’s a kind of fruit”).

### **Concreteness and Specificity across Linguistic Variables**

This analysis was carried out on a subset of 6,483 speakers’ sentences, obtained from the intersection of the two datasets employed in Analysis 1 and Analysis 2. This approach guaranteed that the analysis was restricted to sentences in which both linguistic variables had been annotated and nouns with available concreteness and specificity scores were included.

## **6.2. Data analyses and Results**

All the subsequent data analyses and visualizations were conducted using R (R Core Team, 2023) and RStudio (v. 4.4.1).

### **Concreteness and Specificity across Conditions (Analysis 1)**

Two separate linear mixed models were used to model concreteness and specificity scores of the nouns as continuous dependent variables via `lmer()` function from the “lme4” package in R (Bates et al., 2015). The *Session* condition (present vs. absent) and *Known* condition (known vs. unknown) served as categorical independent variables (0 = absent/unknown; 1 = present/known), analyzed independently and in interaction. *Word length* was included as a control variable.


The speakers were instructed to describe the objects. It is plausible to assume that, during the initial moments of stimulus presentation, relevant information may have been emphasized, leading to an increase in the concreteness or specificity of the language. Therefore, although the order of presentation of the *Session* condition was originally counterbalanced across dyads, the order of presentation (i.e., *Order*) was also included in the model as a control variable. Additionally, the individual information for each dyad was treated as a random intercept.

Two additional linear mixed models were performed, in which the concreteness and specificity scores of nouns were replaced as the dependent variable by the mean concreteness and specificity scores of the sentences, calculated from the average score of all nouns within each sentence. In these models, word length was excluded from the fixed effects, while sentence length was incorporated as a random intercept.

The concreteness and specificity scores were calculated at the sentence level to monitor how these variables fluctuated throughout the progression of the conversations. This approach offered a more dynamic perspective on variations in concreteness and specificity, both across sentences and within each sentence, in contrast to relying solely on the overall distribution of the nouns' concreteness and specificity scores across the entire interaction.


To assess potential conflicts among the variables, correlations were computed between the two experimental conditions (i.e., *Session* and *Known*) and between the dependent variables, concreteness and specificity. A negligible correlation was observed between *Session* and *Known* conditions ( $r = 0.09$ ), while a moderate, expected correlation was found between concreteness and specificity ( $r = 0.48$ ). Collinearity among all predictors and control variables in all linear mixed models was further assessed via the `vif()` function from the “car” package in R (Fox & Weisberg, 2019). Variance Inflation Factor (VIF; Fox, 2005) values for all predictors were below 2, indicating no evidence of problematic multicollinearity and suggesting low intercorrelation among predictors.

**Noun concreteness model.** Results revealed a significant negative main effect of the *Known* condition,  $F(1, 9308.9) = 15.9542, p < .001$  (see Figure 6.2.), word length,  $F(1, 9454.1) = 23.3456, p < .001$ , and Order,  $F(1, 9454.7) = 15.0493, p < .001$ , while the *Session* condition,  $F(1, 8978) = 1.1678, p = 0.279$ , and the interaction were not significant,  $F(1, 9366.6) = 0.0191, p = 0.889$ .




**Figure 6.2.** The modulation of noun concreteness across the experimental conditions.

**Noun specificity model.** The *Known* condition exhibited a significant negative main effect on specificity,  $F(1, 9235.6) = 18.5864, p < .001$  (see Figure 6.3.). No significant interaction between experimental conditions emerged,  $F(1, 9318.4) = 2.2525, p = .133$ . However, the *Session* condition approached significance,  $F(1, 8797) = 3.0449, p = .081$ . Among the other fixed effects, Order was not significant,  $F(1, 9452.5) = 0.0746, p = .785$ , whereas word length showed a statistically significant effect,  $F(1, 9455) = 875.7498, p < .001$ .




**Figure 6.3.** The modulation of noun specificity across the experimental conditions.

**Sentence concreteness model.** As in the model where concreteness scores of individual nouns served as the dependent variable, we observed a significant negative main effect of the *Known* condition,  $F(1, 9329.6) = 27.2786, p < .001$ , and Order,  $F(1, 9438.2) = 15.8244, p < .001$ , whereas neither the *Session* condition,  $F(1, 9036.7) = 0.9094, p = .34$ , nor the interaction was significant,  $F(1, 9366.9) = 0.0076, p = .931$  (see Figure 6.4.).



**Figure 6.4.** The interaction of *Session* and *Known* condition in the sentence concreteness.

**Sentence specificity model.** Notably, we observed a significant negative main effect of both Session,  $F(1, 8921.8) = 5.3072, p = .0213$ , and the *Known* condition,  $F(1, 9292) = 15.6094, p < .001$ , as well as a significant interaction between the two conditions,  $F(1, 9346.3) = 4.7869, p = .029$ , (see Figure 6.5.). Order variables did not display any significant effect.



**Figure 6.5.** The interaction of *Session* and *Known* condition in the sentence specificity.

### Linguistic Variables across Conditions (Analysis 2)

Separate logistic regressions were conducted using binomial generalized linear models for each selected linguistic variable, including the ad hoc category (i.e., Superordinate categorization), with the `glm()` function in R (R Core Team, 2023). The model structure was consistent across all analyses, treating each selected linguistic variable as the dependent variable, while *Session* and *Known* condition, and their interaction were included as predictors.

Wald Chi-Squared tests were employed to assess the significance of the main effects and interactions, using the `Anova()` function from the “car” package in R (Fox & Weisberg, 2019). Effect sizes were estimated by extracting odds ratios, 95% confidence intervals, and *p*-values via the `model_parameters()` function (see Table 6.1.) from the “parameters” package in R (Lüdtke et al., 2020).

**Absent/Present**

**Unknown/Known**

|                   | OR   | P         | OR   | p         |
|-------------------|------|-----------|------|-----------|
| <b>Analogy</b>    | 0.90 | .198      | 0.32 | < .001*** |
| <b>SupCategor</b> | 0.66 | .008**    | 0.39 | < .001*** |
| <b>Object</b>     | 0.86 | < .001*** | 0.61 | < .001*** |
| <b>Situation</b>  | 0.85 | < .001*** | 0.97 | .45       |
| <b>Emotion</b>    | 1.10 | .227      | 1.57 | < .001*** |
| <b>SocialSelf</b> | 0.91 | .117      | 2.52 | < .001*** |

**Table 6.1.** Odds Ratios values and *p*-values indicating significance of the effect size. The reported odds ratios (OR) reflect the likelihood of the variable occurring under the conditions of interest. An OR greater than 1 suggests a higher probability of the variable occurring when the condition is 1 (i.e., present/known), while an OR less than 1 indicates a greater likelihood of the variable occurring when the condition is 0 (i.e., absent/unknown). Asterisks indicate significance level (\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ ).

**Analogy.** The model revealed a significant negative main effect of the *Known* condition  $\chi^2(1) = 110.845, p < .001$ , but no significant effect for the *Session* condition or interaction, all  $ps > .198$ . The odds ratio for the *Known* condition was 0.32 (SE = 0.04, 95% CI [0.25, 0.40],  $z = -9.81$ ), confirming a statistically significant effect,  $p < .001$ , and suggesting a 68% reduction in the odds of the outcome relative to the reference condition.

**Superordinate categorization.** A significant negative main effect of the *Session* and *Known* condition emerged,  $\chi^2(1) = 7.0940, p = .0077$ , and  $\chi^2(1) = 23.4830, p < .001$ , respectively. However, the interaction did not reach significance,  $\chi^2(1) = 0.006, p = .937$ . The odds ratio for the *Session* condition was statistically significant,  $p = .008$  (OR = 0.66, SE = 0.10, 95% CI [0.48, 0.89],  $z = -2.67$ ), suggesting a 34% reduction in the likelihood compared to the reference condition. The odds ratio for the *Known* condition was 0.39 (SE = 0.08, 95% CI [0.26, 0.58],  $z = -4.59$ ), suggesting a significant effect,  $p < .001$ , and a 61% reduction in likelihood compared to the reference condition.

**Emotion Opinion.** A significant positive main effect of the *Known* condition was observed,  $\chi^2(1) = 34.564, p < .001$ . No other main effect or interaction reached significance, all  $ps > .38$ . Odd ratios confirmed that the *Known* condition was the only significant predictor,  $p < .001$  (OR = 1.57, SE = 0.12, 95% CI [1.35, 1.82],  $z = 5.82$ ),

corresponding to a 57% increase in the odds of the outcome relative to the reference condition.

**Object.** The Wald Chi-Squared test revealed a significant negative main effect of both the *Session* and *Known* condition,  $\chi^2(1) = 14.352, p < .001$ , and  $\chi^2(1) = 121.343, p < .001$ , respectively, whereas the interaction was not significant. Odds ratios indicated a 14% decrease in the odds of the outcome for the *Session* condition compared to the reference condition (OR = 0.86, SE = 0.03, 95% CI [0.79, 0.93],  $z = -3.79$ ),  $p < .001$ , and a 39% reduction for the *Known* condition (OR = 0.61, SE = 0.03, 95% CI [0.56, 0.66],  $z = -10.95$ ),  $p < .001$ .

**Situation.** *Session* condition was the only significant predictor,  $\chi^2(1) = 13.0444, p < .001$ . No other main effect or interaction was significant, all  $ps > .19$ . Odds ratio for the *Session* condition was 0.85 (SE = 0.04, 95% CI [0.78, 0.93],  $z = -3.62$ ),  $p < .001$ , indicating a significant effect and a 15% reduction in the odds of the outcome relative to the reference condition.

**Social self.** The model revealed a significant positive main effect of the *Known* condition on the dependent variable,  $\chi^2(1) = 261.332, p < .001$ . In contrast, neither the *Session* condition nor the interaction was significant. The *Known* condition was associated with a 152% increase in the likelihood of the outcome compared to the reference condition (OR = 2.52, SE = 0.15, 95% CI [2.24, 2.83],  $z = 15.67$ ),  $p < .001$ .

Figure 6.6. illustrates the distribution of occurrences of conversational strategies associated with specific linguistic variables across conditions.




**Figure 6.6.** Bubble plot indicating the distribution of occurrences of conversational strategies associated with specific linguistic variables across conditions. The size of bubbles in this plot reflects the total number of occurrences distributed across the four conditions. Larger bubbles represent a higher number of occurrences, while smaller bubbles correspond to fewer occurrences, as denoted by the legend. Note. In the bubble plot, the labels *Session* and *Known* indicate the experimental *Sessions* and *Known* condition, respectively. Values of 0 and 1 correspond to absent/unknown (0) and present/known (1) conditions.

### Concreteness and Specificity across Linguistic Variables (Analysis 3)

This analysis replicates the same linear mixed models used in the first analysis, replacing the *Session* and *Known* condition with the linguistic variables as independent predictors. Control variables and random intercepts were included in all models. However, due to the absence of specific hypotheses regarding potential interactions and the number of predictors considered, no interaction terms were included.

**Concreteness model.** A significant main effect was observed for Superordinate Categorization, Emotion Opinion, Object, Situation, word length, and Order, all  $ps < .006$ . In contrast, analogy and Social self were not significant, all  $ps > .183$ .


**Sentence concreteness model.** Superordinate Categorization, Emotion Opinion, Object, Situation, and Order exhibited a significant main effect, all  $ps < .0062$ . in comparison to the model predicting the concreteness of individual nouns, Analogy emerged as a significant predictor in the current model,  $F(1, 9278.7) = 6.2511$ ,  $p = .0124$ . Notably, Social Self was not a significant predictor in either model,  $p = .718$  (see Figure 6.7.).



**Figure 6.7.** Forest plots of standardized fixed-effect estimates showing the magnitude and statistical significance of each predictor's effect on individual noun concreteness (left plot) and sentence concreteness (right plot), while controlling for all other variables in the model. Asterisks indicate significance level (\*  $p < .05$ , \*\*  $p < .01$ , \*\*\*  $p < .001$ ).

**Specificity model.** Analogy, Superordinate categorization, Emotion Opinion, Situation, and word length were significant predictors of specificity level, all  $ps < .0035$ . No other main effect was significant, all  $ps > .178$ , although Social Self approached significance, all  $ps = .085$ .

**Sentence specificity model.** Significant main effects were observed for Analogy, Superordinate categorization, Emotion Opinion, and Situation, all  $ps < .0286$ . Object, Social Self, and Order were not significant predictors, all  $ps > .0583$ . However, the patterns of results differed from the model in which the specificity of individual nouns was treated as the dependent variable. Notably, the effect of Superordinate categorization decreased in significance ( $b = -0.1324$ ,  $SE = 0.06045$ ,  $t = -2.19$ ,  $p = .0286$ ), whereas the effect of Situation slightly increased ( $b = 0.11642$ ,  $SE = 0.02058$ ,  $t = 5.656$ ,  $p < .001$ ). Furthermore, unlike the previous model, Object emerged as the predictor approaching significance,  $p = .0583$  (see Figure 6.8.).



**Figure 6.8.** Forest plots of standardized fixed-effect estimates showing the magnitude and statistical significance of each predictor's effect on individual noun specificity (left plot) and sentence specificity (right plot), while controlling for all other variables in the model. Asterisks indicate significance level (\*  $p < .05$ , \*\*  $p < .01$ , \*\*\*  $p < .001$ ).

### 6.3. Discussion

Analysis 1 revealed that the *Session* condition was not a significant predictor of either concreteness or specificity, suggesting that the physical absence or presence of the referent did not, in itself, influence the concreteness and specificity of the nouns produced by the speakers in their descriptions. In contrast, the *Know* condition exhibited a

significant negative effect on both concreteness and specificity, indicating that when describing an unknown or unfamiliar object, speakers tended to produce more concrete and specific nouns. One plausible interpretation is that the novelty of unfamiliar objects may have prompted speakers to produce a more concrete and specific language to convey denser semantic representation, which may facilitate alignment with the interlocutor (Garrod & Pickering, 2009; Pickering & Garrod, 2021).

Among the control variables, Word length displayed an expected negative effect on concreteness, indicating that longer nouns tend to be less concrete (Piantadosi et al., 2011; Reilly et al., 2017). In contrast, a positive relationship between Word length and specificity suggests that the more specific nouns tended to be generally longer (Murphy & Smith, 1982). The significance of the Order variable in the concreteness models indicates that descriptions tended to be more concrete at the beginning of the experimental session. This effect may be attributed to the fact that early descriptions focused on the most salient and relevant features of the objects, resulting in higher degrees of concreteness. This interpretation is consistent with the experimental design, in which speakers were trained on information about the objects before the task. Notably, the Order variable did not affect the specificity of the nouns or sentences used in the descriptions. This discrepancy further suggests that, although the objects were both highly concrete and specific, concreteness primarily influenced immediately related information (e.g., perceptual features such as “it has claws”), whereas specificity was not similarly constrained. A plausible explanation is that essential information can be conveyed even at relatively low levels of specificity, for instance, through categorization using superordinate categories or hypernyms (e.g., “it’s a bird”).

In the model analyzing sentence concreteness, the *Known* condition was a significant predictor, suggesting that speakers produced more concrete sentences when describing unfamiliar objects compared to familiar ones, thus mirroring the pattern observed in the noun concreteness model.

Notably, in the model predicting sentence specificity, both conditions were significant, and a significant interaction emerged, suggesting that speakers produced more specific sentences when objects were both unknown and physically absent from the scene. This pattern may indicate that speakers increased sentence specificity under conditions of object novelty and referential absence to maximize the informativeness of their language

(Murphy, 2002), thereby potentially supporting the acquisition of novel knowledge (Schilhab, 2015).

Results from Analysis 2 revealed that, among abstraction strategies, analogical comparisons were more likely to occur in the *Unknown* condition, whereas superordinate categorizations were more frequent when the objects were physically absent and unknown. This finding suggests that both abstraction processes may function as complementary strategies for enhancing the retrieval of relevant information by comparing the object with other stored exemplars or instantiating it in a category (Murgiano et al., 2021).

Speakers focused more on perceptual and physical features when the objects were unknown and absent, whereas references to encyclopedic features were more likely solely in the absent condition. By contrast, speakers also refer more to their own personal experiences and emotional dimension when the objects were known (e.g., “avocado”). Interestingly, no difference was observed between the absent and present conditions. It is plausible to assume that novelty, rather than physical presence/absence, influenced the common ground between interlocutors (Clark & Brennan, 1991; Clark & Wilkes-Gibbs, 1986), shifting the focus of conversations from perceptually relevant information when objects were unfamiliar, to speakers’ personal and emotional experiences once objects became familiar and key traits had already been acquired.

Finally, Analysis 3 provided further support for the role of emotional experience as a key predictor of concreteness, indicating that when speakers introduced their own emotional experience related to the objects, the concreteness of their language decreased. Thus, emotional experience was negatively associated with concreteness, in line with the existing literature (Kousta et al., 2011; Vigliocco, Kousta, et al., 2014). The effect was consistent across noun and sentence concreteness. By contrast, references to personal experience did not affect concreteness.

References to objects’ features (both perceptual and encyclopedic) were significant predictors of noun and sentence concreteness. When speakers focused on objects’ features, they lightly increased the concreteness of their language.

Notably, the superordinate categorization variable emerged as the second-strongest predictor of both noun and sentence concreteness. When speakers categorized objects at the superordinate level, that is, at a more general and less specific level of

abstraction, they produced nouns that retained highly concreteness scores. This finding is consistent with the idea that specificity and concreteness are two separate variables (Bolognesi et al., 2020; Rambelli & Bolognesi, 2024) that reflect distinct dimensions of referential process.

The two abstraction processes, i.e., superordinate categorization and analogical comparison, exerted the largest effect in predicting noun specificity. Categorization at the superordinate level led to a decrease in noun specificity, whereas analogical comparison between co-hyponyms increased specificity. Although the directions of these effects were consistent across noun and sentence specificity models, their statistical significance differed. Specifically, the negative effect of superordinate categorization over specificity was substantially reduced in the sentence specificity model, whereas the effect of analogical comparison showed a slight increase in significance. This divergence may reflect differences in the quantity and type of related information elicited at the sentence level between general and specific words. Referring to a general word characterized by low-resolution representation (Villani et al., 2024) and low-dimensionality (Lupyan & Mirman, 2013) may instantiate a broader and more variable set of potential related exemplars, contexts, and features, which vary in terms of their specificity level. In contrast, analogical comparisons between specific terms may require the instantiation of highly specific shared properties.

In conclusion, emotional experience emerged as a reliable predictor of both noun and sentence specificity, suggesting that speakers tended to employ less specific language when referring to objects through their emotional dimension. This finding aligns with previous literature indicating that, overall, reporting emotional experience is associated with variability in terms of granularity and precision (Barrett, 2006). Conversely, references to objects' perceptual features did not significantly predict specificity in either of the two models. However, references to encyclopedic features (e.g., "it's from South Asia") showed a positive effect on both noun and sentence specificity, indicating that encyclopedic knowledge tends to include more fine-grained and precise information.



**PART 3**  
**DISCUSSION**



## **Chapter 7**

### **Towards a Comprehension Theory of Language-Mediated Abstraction**

#### **7.1. Specificity Effect on Semantic Processing**

Chapter 4 examines the role of categorical specificity in language processing within classical paradigms in psycholinguistics (i.e., rating task, lexical decision task, and semantic decision task). Moreover, the study examined in Chapter 4 explored the contribution of specificity in explaining the well-known, yet still debated, concreteness effect, namely, the behavioral processing advantage of concrete words over abstract ones (see ¶ 2.3., for a detailed discussion). The existing literature on the concreteness effect reports inconsistent findings, particularly when stimuli are controlled for other relevant semantic variables, such as contextual availability, imageability, and valence. Crucially, conceptual concreteness and categorical specificity are positively but only weakly correlated: both abstract and concrete words can be defined at different levels of abstraction, ranging from highly general and inclusive conceptual categories (e.g., COGNITION; FOOD) to more specific and narrowly defined conceptual categories (e.g., DECISION-MAKING; JACKFRUIT; Bolognesi et al., 2020). However, studies investigating processing differences between abstract and concrete words have not controlled their stimuli for levels of categorical specificity. As a result, the potential contribution of specificity to the concreteness effect has remained an open question.

In addition, the study presented in Chapter 4 aims to clarify the independent impact of categorical specificity on language processing. Recent research has reported the effects of categorical specificity on chronometric data. However, these studies reported opposite patterns with specific words processed either faster or slower than general words, depending on the type of task performed (e.g., lexical decision and semantic decision tasks) and on the criteria used for stimulus selection (e.g., human-generated data or WordNet-derived measures).

Accordingly, the primary objective of the study reported in Chapter 4 was to investigate the role of categorical specificity in accounting for differences in cognitive effort required during the processing of words varying in levels of abstraction. To achieve

this objective, three experiments (i.e., a rating task, a lexical decision task, and a semantic decision task) were conducted using a carefully controlled set of stimuli that varied orthogonally in concreteness and specificity and balanced for word length and frequency.

The results support and expand previous findings. In particular, experiment 1 (rating task) confirmed that speakers perceive both abstract and concrete words at different levels of specificity. However, compared to concrete words, differences in categorical specificity and taxonomic organization are less pronounced for abstract words (Borghi, 2022). Moreover, experiment 1 and the rating data provided further evidence in support of the asymmetry between speakers' taxonomic conceptual representation and dictionary-based knowledge of WordNet. Despite being a valuable resource, WordNet does not reflect taxonomic knowledge and the organization of speakers' lexicon.

In experiment 2 (lexical decision task), neither concreteness nor specificity exhibited any effect in predicting reaction times. Therefore, during lexical access to meaning, it did not emerge a concreteness effect nor any specificity effect. This result may be explained by the relatively shallow process involved in word recognition tasks, in which lexical variables are more influential than semantic variables (Yap et al., 2011). For instance, Yap et al. (2011) pointed out that, when controlling and balancing stimuli for relevant lexical dimensions, such as word length and frequency, semantic variables explained 2% of variance in a lexical decision task. However, as briefly noted in ¶ 4.3., a qualitative and visual examination of the effect of concreteness and specificity on predicted reaction times (see Figure 4.2., Chapter 4), although non-significant, suggests that, regardless their concreteness, words with lower scores in specificity (i.e., more general words) may require relatively less effort to be recognized, in line with previous findings (Iliev & Axelrod, 2017; Bolognesi & Caselli, 2023; Ravelli et al., 2025). In these previous studies, specificity exhibited a positive effect on lexical decision latencies (see ¶ 3.4.1. and ¶ 3.5.1.), with general words associated with faster response latencies than specific words. Moreover, this tendency appears to be more evident for words perceived as highly specific and low concrete (i.e., abstract-specific words), exhibiting the slowest times and presumably greater effort in being processed. This marked difference between abstract-general and abstract-specific words in lexical processing is intriguing and deserves further investigation.

The semantic decision task (experiment 3) replicated the concreteness effect (i.e., concrete words were categorized faster than abstract words) and reported a significant effect of specificity on decision latencies: specific words were recognized faster than more general ones during the abstract/concrete categorization task, regardless of their concreteness. Notably, in this experiment, with a controlled set of stimuli, a significant interaction between concreteness and specificity did not emerge, unlike the one reported by Ravelli et al. (2025). However, a plausible explanation for this inconsistency may be that the sample of stimuli used in experiment 3 and Ravelli et al. (2025) differ. In experiment 3, a controlled set of concrete and abstract words for which hypernyms and hyponyms have been generated was employed, while Ravelli et al. (2025) examined words in the ANEW dataset, which are characterized by highly affective content and were not previously controlled for concreteness and specificity. However, the absence of interaction in the semantic decision task demonstrates that specificity exhibits an independent effect in predicting response times when deeper semantic processing is required.

This specificity effect found in experiment 3 closely aligns with previous research. In particular, Ravelli et al. (2025) found a similar significant negative effect when comparing semantic decision response (Pexman et al., 2017) and specificity scores (Bolognesi & Caselli, 2023; see also ¶ 3.5.1.). Thus, words at lower levels of abstraction (more specific) required less effort to be categorized as abstract or concrete than words at higher levels of abstraction (more general).

A plausible interpretation of this finding is that categorization of both abstract and concrete words may be influenced by the cognitive effort required to abstract and instantiate a single or multiple referents. This hypothesis aligns with exemplar-based accounts of category representations (Nosofsky, 1986, 1988; Daelemans & Van den Bosch, 2005; Chandler, 2017), which propose that participants engaging in semantic decision tasks tend to instantiate and represent referents or exemplars. Consistently, Heit & Barsalou (1996) showed that in other semantics-based categorization tasks, such as typicality judgments and exemplar production, participants instantiate specific subordinate exemplars or specific aspects and features even when processing basic-level categories. Therefore, when engaging with general words like “quality” or “organism”, speakers may likely instantiate multiple exemplars belonging to the relative categories, resulting in a greater effort and, consequently, longer categorization times. Conversely,

more specific words (e.g., “empathy”; “frog”) tend to denote a more constrained reference, allowing for relatively effortless instantiation of referents and faster processing. This interpretation is also consistent with seminal findings by Rosch (1978) and Rosch et al. (1976), who demonstrated that, within concrete categories, superordinate terms are processed more slowly than words at the other two levels of abstraction (i.e., basic and subordinate levels). Importantly, the semantic decision task discussed in Chapter 4 (experiment 3) shows that this processing advantage of more specific words also emerges within abstract concepts.

Typically, specific and general words differ in their degree of informativeness and distinctiveness (see also ¶ 3.2.1. and ¶ 3.5.1.). In particular, specific words tend to be highly informative and less distinctive, whereas general words are more distinctive but less informative (Murphy, 2002; Murphy & Brownell, 1985). Accordingly, it is plausible to hypothesize that, relative to general words, specific words activate narrower and denser semantic representations. These representations are associated with the instantiation of a smaller set of referents or exemplars, for which salient and relevant semantic information is more readily accessible. As a result, the high-resolution semantic content characterizing specific words (Villani & Bolognesi, 2025) may enhance performance in semantic decision tasks by facilitating the retrieval of task-relevant information, such as determining whether a word is concrete or abstract.

The results reported in Chapter 4 provide crucial insights into the distinction between conceptual concreteness and categorical specificity. In particular, they show: (i) a confirmation of weak correlation between the two variables ( $r = 0.29$ ), consistent with (Bolognesi et al., 2020); (ii) the absence of collinearity; (iii) additive effect over response latencies in behavioral tasks (i.e., semantic decision task); and (iv) task-dependent patterns observed across lexical and semantic decision tasks. Although the effect of the variables was not significant in the lexical decision task, the analyses revealed a consistent effect of concreteness across the tasks, alongside an opposite pattern for specificity, as noted above. The linear regression coefficients indicate that concrete words are associated with faster responses in both lexical and semantic decision tasks. By contrast, when included in the model, specificity exerted a positive effect on lexical decision latencies and a negative effect on semantic decision latencies. Thus, a reverse pattern emerged: words with higher specificity (i.e., specific) were processed more slowly than words with lower specificity (i.e., general) in the lexical decision task, in line with previous findings

(Iliev & Axelrod, 2017; Bolognesi & Caselli, 2023; Ravelli et al., 2025), whereas highly specific words were processed faster than words with lower specificity in the semantic decision task. Overall, these results support the idea that conceptual concreteness and categorical specificity are distinct dimensions that independently influence word processing. They also suggest that the impact of semantic variables may vary as a function of task demands (Goh et al., 2016; Pexman, 2012; Pexman et al., 2017).

These findings align with the *semantic richness* framework (Pexman, 2020; Yap et al., 2015), according to which greater feedback activation from the semantic level provides support to recognition at the lexical level. During lexical decision tasks, semantic variables are generally less influential than lexical ones. Nevertheless, it is plausible that general words, being associated with a broader range of contexts and referents, may benefit from diffuse semantic activation. Conversely, when the task requires a deeper understanding of word meaning, as in semantic decision tasks, the variability of semantic information associated with general words, along with the instantiation of competing referents, may hinder performance. Similar patterns have been reported in relation to the *number of senses*: words with multiple senses (i.e., ambiguous words) elicit faster responses in lexical decision tasks, but poorer performance in semantic categorizations (Eddington & Tokowicz, 2015).

In the semantic decision tasks (experiment 3) discussed in Chapter 4, where statistically significant effects emerged, general words showed a clear processing disadvantage. It is plausible that the larger number of contexts and competing referents associated with general words likely interfered with performance when deeper semantic access to meaning was required. In contrast, specific words, by narrowing the range of possible meanings and providing more readily accessible relevant information through higher-resolution semantic representations, supported faster and more accurate semantic categorizations. These findings suggest that informativeness plays a task-dependent role in language processing. While higher levels of detail and precision may hinder lexical recognition, as suggested by Iliev & Axelrod (2017), they appear to enhance performance in semantic decision tasks, where a more qualitative processing is involved, as demonstrated by results from the study examined in Chapter 4 and from Ravelli et al. (2025). Future studies should further investigate how informativeness interacts with language processing across different tasks.

Finally, the processing disadvantage observed for general words in the semantic decision task may also be explained by their typically low dimensionality (Lupyan & Mirman, 2013; Borghi, 2022). By denoting more inclusive categories whose members share relatively few salient features, general words may have been more difficult to evaluate and categorize as abstract or concrete. Given that low dimensionality is commonly associated with abstract concepts (see ¶ 2.2.1.), the latency differences observed between highly abstract-general words and highly abstract-specific words may reflect the more precise and constrained representation of the latter.

This interpretation is supported by recent findings from Rambelli & Bolognesi (2023, 2024; see also ¶ 3.5.1.), who showed that abstract-specific words tend to exhibit low contextual variability (Hoffman, 2016). Thus, such words are associated with a relatively restricted and predictable set of semantically related and accessible contexts. This is consistent with Contextual Availability Theory (Schwanenflugel et al., 1988, 1992; see also ¶ 2.3.), according to which word processing benefits from the ease of retrieving relevant contextual information from memory.

When investigated as isolated, decontextualized words within traditional psycholinguistics paradigms, conceptual concreteness and categorical specificity exhibit separate and additive contributions to chronometric data. The findings discussed in Chapter 4 suggest that the precision and granularity of information associated with words at different levels of abstraction, as well as the resolution of semantic representations, may yield either processing advantages or disadvantages depending on task demands. In particular, the high-resolution semantic content of more specific words gives rise to a specificity effect, namely, a processing advantage over general words when deeper semantic understanding is required.

## **7.2. The Pragmatic Ambiguity of General Sentences**

When treated as isolated variables rather than as embedded within contextual information, conceptual concreteness and categorical specificity exert measurable effects on language processing during semantic categorization. In particular, the semantic decision task discussed in Chapter 4 replicated the well-established concreteness effect (i.e., faster response associated with concrete words compared to abstract words; Jessen et al., 2000) and revealed a specificity effect (i.e., a behavioral advantage in the processing of specific

words over general words). However, the flexibility of semantic and conceptual representations may lead to the activation of different meanings as a function of context, stimulus presentation, and task demands (Yee & Thompson-Schill, 2016), as also reflected in the differential effects of semantic variables across lexical decision and semantic decision tasks.

As argued in *Methodological and Theoretical Rationale for the Experimental Section*, a comprehensive investigation of semantic variables requires examining their effects in interactive and contextualized settings that more closely reflect natural language use (Banks et al., 2023; Vigliocco et al., 2024). Although traditional psycholinguistic paradigms offer several advantages, such as controlled environments for identifying key dimensions of word meaning, linguistic experience is fundamentally a form of joint action (Clark & Brennan, 1991; Clark, 1996; Pickering & Garrod, 2021) and is therefore best captured through analysis of conversational dynamics.

As discussed in Chapter 2 (see ¶ 2.5.4. and 2.6.), the social dimension is particularly relevant in the study of abstractness and the meanings of abstract concepts (see also *Methodological and Theoretical Rationale for the Experimental Section*). However, the relationship between this dimension and categorical specificity remains largely unexplored. Referential scope, informativeness (Davies & Arnold, 2019), and precision of information (Grice, 1989) at different levels of abstraction may modulate communicative strategies and pragmatic dynamics during interactions and conversations.

The study reported in Chapter 5 explored the role of categorical specificity alongside abstractness/concreteness in conversational dynamics, focusing on two linguistic constructs associated with abstractness: uncertainty and curiosity (i.e., the need for clarification, additional information, and more interactive turns). The study adopted the same methodology employed by Mazzuca, Villani et al. (2025), asking participants to evaluate the plausibility of conversation excerpts composed of target sentences varying in both abstractness/concreteness and specificity, along with follow-up expressions. The experimental setting resembled traditional psycholinguistic paradigms: participants read and responded to sentences displayed on a computer, with behavioral measures such as reaction times collected. However, the ecological validity of the study was ensured by embedding words varying in abstractness/concreteness and categorical specificity within simulated conversations. In addition, follow-up expressions were derived from

participants' responses from a previous free production study (Villani et al., 2022) and coded for their communicative function, including signaling uncertainty or comprehension, and indicating willingness to end the conversations or request additional information.

Mazzuca, Villani et al. (2025) demonstrated that abstract sentences were associated with linguistic constructs of uncertainty and curiosity, whereas concrete sentences elicited responses signaling comprehension and a lack of need for additional conversational turns. Building on this work, the study discussed in Chapter 5 investigated whether the same linguistic constructs are associated with categorical specificity, given that general and specific words differ in the resolution of their semantic representations and in the scope of their reference (Lyons, 1977; Bolognesi et al., 2020; Bolognesi & Caselli, 2023). The study also explored whether the patterns observed for abstract and concrete sentences vary as a function of changes in their level of abstraction. Results replicate and expand previous findings.

Experiment 1 (session 1) revealed no significant interaction between the level of categorical specificity of target sentences and type of follow-ups (i.e., uncertainty vs. certainty). However, when considering the intersection between the specificity and concreteness/abstractness of target sentences (i.e., abstract-general, abstract-specific, concrete-general, concrete-specific), participants judged abstract sentences as more plausible than concrete sentences, regardless of level of specificity. These results replicate findings from Mazzuca, Villani et al. (2025), providing further evidence that uncertainty is a linguistic construct primarily associated with abstractness (Borghi et al., 2021; Villani et al., 2022). This pattern was further supported by participants' readiness in judging both types of concrete sentences (general and specific) as implausible when paired with uncertainty follow-ups, compared to when paired with certainty follow-ups. In contrast, no differences emerged between uncertainty and certainty follow-ups within abstract sentences.

Experiment 2 demonstrated that, when paired with curiosity follow-ups, general sentences were judged as more plausible than specific sentences, whereas the opposite pattern emerged with end (of conversation) follow-ups, with specific sentences being judged as more plausible. General sentences appear to require additional information and social interaction, whereas specific sentences are more pragmatically exhaustive.

However, when analyzing the four quadrants, this effect was primarily driven by general-concrete sentences; no significant differences were found between abstract-general and abstract-specific sentences. These findings suggest that curiosity, like uncertainty, is a linguistic dimension that characterizes abstractness, independent of specificity.

In line with previous research (Mazzuca et al., 2022), the ratings supported the idea that abstractness elicits lower confidence but higher social engagement, as well as the need for experts' contribution to achieve a full understanding of the meaning of abstract concepts (Borghi, 2023). Notably, concrete-general sentences scored higher in Social Metacognition and lower in Confidence ratings than concrete-specific sentences, suggesting that interlocutors might require more social validation, clarification, and additional information to establish reference with concrete-general sentences. Surprisingly, no differences were observed between concrete-specific and concrete-general sentences in Experts Need ratings.

Taken together, these findings support the theoretical framework suggesting that the variability of abstract concepts generates more uncertainty and increases the need for interaction to negotiate their meanings (Borghi, 2023; Mazzuca & Santarelli, 2023). These results further indicate that interlocutors may require additional conversational turns and greater mutual monitoring to align their representation with abstract concepts (Gandolfi et al., 2023; Mazzuca, Villani et al., 2025). Notably, the study reported in Chapter 5 provides new evidence that uncertainty is a marker of abstractness across different taxonomic levels, occurring with both abstract-general (e.g., “opinion”) and abstract-specific words (e.g., “deliberation”).

Experiment 2 also revealed that the alignment process can be effortful with concrete-general sentences (e.g., “I saw a mammal”), likely because their denotation is more indeterminate, encompassing multiple exemplars and members of the category (Murphy, 2002; Bolognesi et al., 2020; Bolognesi & Caselli, 2023). When interlocutors instantiate this variety of potential referents (Heit & Barsalou, 1996), additional conversational turns and requests for clarification may be required to disambiguate the appropriate referent. In contrast, concrete-specific words (e.g., “I saw a dromedary”), which trigger fewer well-defined referents, allow interlocutors to align their representations more easily.

Overall, findings from the study reported in Chapter 5 provide crucial insights into the relationship between conceptual concreteness, categorical specificity, and conversational dynamics. In particular, uncertainty appears to be a linguistic construct inherently linked to abstractness but not to specificity. This pattern is likely due to the absence of a single, physically identifiable reference and the intrinsic variability of abstract words' semantic representations, which generates an indefiniteness in their meanings. This interpretation aligns with the WAT proposal (Borghi & Binkofski, 2014; Borghi, 2023; see ¶ 2.5.4.), which posits that the indefiniteness of abstract words' semantic and conceptual representation triggers two processes: first, an initial *inner monitoring*, relying on one's own linguistic knowledge and contextual information to capture the meanings of an abstract word; and second, if inner monitoring fails, a *mutual monitoring*, which involves relying on others through social interaction or deference on experts (Borghi, 2022, 2023; Borghi & Binkofski, 2014). The results of the ratings support this view: abstract sentences elicited high scores in Social Metacognition and Experts Need, and low scores in Confidence in the sentence meaning.

A recent theoretical perspective further links vagueness to abstractness (Borghi et al., 2025). Vagueness refers to expressions whose conceptual boundaries are blurred or indeterminate, with an indefinite denotation that is not easily identifiable (Magni, 2024). Vagueness is especially relevant for abstract concepts, whose meanings are often debated and indeterminate, requiring more Social Metacognition, social validation, and negotiation to be fully understood.

This indefiniteness or vagueness does not appear to be explainable solely by instantiation at different levels of abstraction. Although concrete-general sentences encompass a broader set of potential instances than concrete-specific sentences, participants judged concrete-general and concrete-specific sentences as equally plausible when paired with the two types of follow-ups. One plausible interpretation is that the uncertainty associated with abstract sentences in interactions and conversations stems from their greater contextual and situational variability (Hoffman, 2016; Davis et al., 2020; Schulte Im Walde & Frassinelli, 2022). This variability increases the difficulties of retrieving a proper context or situation to frame the event described in the sentence. For example, a sentence like “I thought of destiny” could refer to multiple meanings (e.g., future plans, a required action, or the inescapability of human existence), each applicable across a wide network of potential contexts and real-world situations. This ambiguity

appears not to happen with concrete sentences, even when general, for which physical referents are more readily accessible. When engaging with a sentence like “I bought a basket of fruits”, interlocutors can more easily identify and represent particular situations that frame the event, compared to abstract sentences.

This interpretation is consistent with research on contextual availability (Schwanenflugel et al., 1988) and situational systematicity (Davis et al., 2020), which posits that abstract words tend to be less tightly associated with a range of multiple potentially applicable contexts and real-world situations, which may also vary across space and time (Davis et al., 2020). The resulting uncertainty and indefiniteness may arise from this greater variability and complexity (Barsalou & Wiemer-Hastings, 2005), presumably leading interlocutors to misalign in their representations at the earliest stages of conversation, namely, the *situation model*, which encodes the key elements of the situation under discussion (Zwaan & Radvansky, 1998; Pickering & Garrod, 2004, 2021).

Curiosity and need for clarification appear to be associated not only with abstractness but also with categorical specificity, particularly with concrete-general sentences. Unlike uncertainty, curiosity (i.e., the need for clarification) relates more to the quantity rather than the quality of information provided. For example, in the case of concrete-general sentences, such as “I washed an object”, interlocutors can frame the event described in a potential situation or context and activate the lexical entry “object”, but they may not identify the specific referent, resulting in semantic misalignment (Pickering & Garrod, 2004, 2021). By instantiating multiple potential and competing referents (e.g., a brush, a fork, a vase), interlocutors are more likely to request clarification and produce restricted repairs (e.g., “which one?”; Dideriksen et al., 2019).

Gandolfi et al. (2023) observed that when engaging with abstract words, interlocutors may struggle to align their representations and require more mutual monitoring because the retrieval of a potential referent is less straightforward than with concrete sentences. Experiment 2 discussed in Chapter 5 further demonstrates that alignment and mutual monitoring can also be necessary in the case of concrete-general words, whose references are looser compared to those of concrete-specific sentences. Conversely, participants judged concrete-specific sentences as more plausible when paired with follow-ups signaling the willingness to end the conversation or the absence of a need for additional information. This suggests that concrete-specific sentences, which

provide more direct referential access and instantiate fewer potential referents, support alignment between interlocutors' representations more readily.

Ratings results corroborate these findings. Concrete-general sentences received higher scores in Social Metacognition but lower scores in Confidence compared to concrete-specific sentences, reflecting the increased need for social monitoring when referents are more indeterminate.

Taken together, these results suggest that concrete-specific sentences are pragmatically more exhaustive (Grice, 1989) due to their greater informativeness (Murphy, 2002; Davies & Arnold, 2019). Consistent with Grice's *Cooperative Principle* and the view of conversations as joint actions aimed at mutual understanding (Clark & Wilkes-Gibbs, 1986; Clark, 1996; Garrod & Pickering, 2009), interlocutors appear to modulate referential expressions based on informativeness, that is, the amount of information provided. Davies & Arnold (2019) further noted that informativeness depends on both linguistic and non-linguistic contextual information and the nature of the referring expression. Informative expressions facilitate referential identification by denoting a smaller set of eligible candidate referents within a given context. In contrast, concrete-general sentences may introduce potential indeterminacy or a form of ambiguity.

Generality can be understood as a form of semantic indeterminacy corresponding to referential indeterminacy (Magni, 2024), lacking specification and precision (Zhang, 1998; Magni, 2024). General words typically have broad denotations (i.e., wide referential scope) and low-resolution semantic representations. Consistently, in experiment 2, participants evaluated concrete-general sentences as more plausible when paired with curiosity follow-up expressions (e.g., signaling the need for clarification), whereas concrete-specific sentences were considered more plausible when followed by responses indicating no need for additional clarification. This pattern likely reflects differences in precision and informativeness at lower levels of abstraction as opposed to the greater indeterminacy at higher levels of abstraction.

Finally, the association of abstractness with linguistic constructs of uncertainty suggests that the representation of abstract words and sentences may be characterized by intrinsic vagueness. By contrast, the association between generality and curiosity may reflect the referential indeterminacy, which can be more broadly understood as a potential form of ambiguity, as arising from the need to select among potential competitive

referents. Notably, the fact that abstract words and sentences are associated with both uncertainty and curiosity indicates that their representations may involve different forms of indeterminacy (e.g., vagueness, ambiguity, generality) or may simply exhibit a more pronounced indeterminacy compared to concrete-general sentences.

Overall, the study discussed in Chapter 5 confirmed the association between abstractness and linguistic constructions of uncertainty and curiosity (i.e., the request for additional clarification) in conversational contexts. Notably, no differences were revealed at different levels of abstraction within abstract sentences. The indeterminacy in abstract words is such that interlocutors need to require more social interaction to converge on a mutual understanding. However, the main differences between general and specific sentences were driven by concrete sentences. In particular, the generality of concrete-general sentences does not appear to plausibly elicit uncertainty in an interlocutor. However, results suggest that participants found plausible linguistic excerpts in which this referential indeterminacy fosters requests for clarification. Ratings in Social Metacognition and Word Confidence supported this finding. In contrast, the high informativeness and the constrained reference of concrete-specific sentences have been evaluated as more pragmatically exhaustive.

### **7.3. Specificity as a Linguistic Cue in Naturalistic Conversations**

Chapter 6 discusses a study examining conceptual concreteness and categorical specificity at the level of language produced in semi-naturalistic conversations. In particular, the study analyzed the ECOLANG dataset, which comprises recorded and transcribed multimodal semi-naturalistic interactions between two adult interlocutors. The interlocutors occupied asymmetrical roles, whereby a more knowledgeable speaker described a set of objects to an addressee under different experimental conditions: (i) whether the objects were physically present or absent from the scene; and (ii) whether the objects were known or unknown to the addressee. The objects represented highly concrete and specific referents of four semantic categories

Language produced by speakers has been automatically annotated and analyzed. In addition, several conversational strategies encoding linguistic variables were examined and coded. These included two abstraction strategies reflecting superordinate-level categorization (e.g., “it’s a tool”) and categorization via analogical correspondence with

co-hyponyms (e.g., “it’s an ostrich”), as well as references to object’s perceptual features (e.g., “it’s yellow”), encyclopedic knowledge regarding the object (e.g., “it is from South Asia”), speakers’ personal experience (e.g., “I saw it at the zoo”), and speakers’ emotional dimension (e.g., “I don’t like it”).

Through three analyses, the study in Chapter 6 explored whether:

- speakers modulate the concreteness and specificity of their language across these four experimental conditions (Analysis 1).
- The likelihood of employing specific conversational strategies varies as a function of these conditions (Analysis 2).
- particular conversational strategies emerged as best predictors of concreteness and specificity of speakers’ language (Analysis 3).

This study presents a novel approach to examining the concreteness and specificity of nouns and sentences produced during semi-naturalistic interactions. Sentence-level concreteness and specificity were computed as the average scores of the nouns contained within each sentence. This approach enabled a more dynamic examination of variation in both variables, both within each sentence and across sentences.

A previous study by Vigliocco et al. (2019) investigated the distribution and frequency of nonverbal multimodal cues, in particular iconic and indexical cues, in interactions between children and caregivers under the same experimental conditions. The study reported a more frequent use of iconic and indexical cues in conditions in which objects were absent and unknown, suggesting that such multimodal strategies may enhance learning. In particular, Vigliocco et al. (2019) proposed that indexical cues contribute to language learning by directing attention to present entities, whereas iconic cues support learning by evoking perceptual properties of the referents through imagery, especially when the objects are absent.

The primary aim of the study discussed in Chapter 6 was to elucidate the role of concreteness and specificity as linguistic cues that enhance imagery and informativeness of language used for communicative purposes. The results support this claim.

Analysis 1 demonstrated that noun-level concreteness and specificity were not influenced by *Session* condition (i.e., object presence vs. absence). When analyzed both

as single nouns and composed nouns (e.g., “musical instrument), the physical presence or absence of the referent did not modulate degrees of concreteness or levels of specificity. By contrast, familiarity or unfamiliarity with the object significantly influenced speakers’ language, with unknown objects eliciting increased concreteness and specificity. In particular, speakers used more concrete and specific nouns when describing unknown objects. This finding partially supports the view that, similarly to multimodal cues, linguistic cues can function as communicative strategies aimed at conveying semantically denser representations. One plausible interpretation is that, when describing unfamiliar objects, speakers employ more concrete and specific nouns to facilitate representational alignment with the interlocutor (Gandolfi et al., 2023; Garrod & Pickering, 2009; Pickering & Garrod, 2004), plausibly directing attention to salient traits and features of the referents, as observed for indexical and iconic cues.

An additional interpretation derives from a theoretical proposal, the Word as Cultivators (WAC) of others’ minds hypothesis, which posits that words function as *linguistic handles* capable of reactivating neural components associated with a verbal cue during the *linguification* process (Schilhab, 2015). Particularly in learning contexts, such as when an object is unfamiliar, interlocutors may strategically use words to activate perceptual and semantic neural subsystems in order to facilitate information transfer. Accordingly, speakers in the present study may have consciously or unconsciously employed more specific and concrete nouns, characterized by high-dimensional and richer semantic representations, to support the transmission of the information.

The present findings are also consistent with the evocative role of language described by Lupyan & Thompson-Schill (2012), who demonstrated that words activate aspects of conceptual representation (e.g., visual shape) more effectively than nonverbal cues alone. In line with the *label feedback hypothesis*, verbal labels may modulate perceptuo-conceptual processing by activating diagnostic features of semantic and conceptual representations.

At the sentence level, speakers produced more concrete sentences when objects were unfamiliar, mirroring the pattern observed at the noun level, regardless of the presence or absence of the object. However, sentence-level specificity exhibited a different pattern: *Session* condition (absent vs. present) and interaction with *Known* condition (known vs. unknown) reached statistical significance. In particular, speakers

produced more specific sentences when objects were both unfamiliar and physically absent. As observed for iconic cues, increased specificity within and across sentences appears to function as a strategy for evoking relevant properties and information, plausibly reflecting an effort at maximizing communicative informativeness (Murphy, 2002; Davies & Arnold, 2019). Although the present study did not analyze the language of both interlocutors, it is reasonable to hypothesize that, as in the *Tangram game*, the use of condensed, precise, and informative words increases the semantic resolution of the sentence (e.g., “the ice skater”; Clark & Wilkes-Gibbs, 1986), thereby facilitating convergence on shared representations.

The divergent patterns observed for concreteness and specificity at the sentence level may reflect differences in the type of information conveyed. Although speculative, concreteness may primarily serve an ostensive function, directly referring to the novel and physically available referent, as for the indexical cues. In contrast, specificity, like iconic cues, may possess a more intrinsically representational character. The informativeness associated with categorical specificity may parallel the imagistic and resemblance-based properties of iconicity, conveying fine-grained information about referents and enhancing vividness. In support of this interpretation, de Varda, Lamarra et al. (2025) reported a positive correlation between iconicity and specificity, as discussed in Chapters 2 and 3.

Analysis 2 revealed that analogical comparisons were more frequent when referents were unfamiliar, in line with foundational work on analogical reasoning and learning (Gentner, 1989; Gentner et al., 2017). As proposed by Gentner et al. (2017), structural alignment emphasizes shared relational structure between exemplars, enabling the formation of abstraction schemas that can be stored and reused.

By contrast, superordinate categorization occurred with higher likelihood when referents were both unknown and physically absent. General words associated with broader categories may activate a wider range of potentially relevant information and associations. Together, these findings suggest that abstraction through either co-hyponyms (i.e., specific words) or superordinate (i.e., general words) categorization reflects complementary strategies aimed at enhancing access to diagnostic information, either by directly comparing an exemplar to stored representations of similar entities (Murgiano et al., 2021) or by situating it within a broader conceptual category.

In addition, speakers focused more on objects' perceptual features when the objects were both unfamiliar and physically absent. Notably, speakers also referred more to encyclopedic knowledge of the objects when they were absent from the scene. In contrast, when objects were familiar, speakers focused more on their own personal and emotional experience. Together, these findings suggest that object familiarity/unfamiliarity, rather than the physical presence or absence, plays a central role in shaping the ongoing focus of the conversation. In line with accounts emphasizing the crucial role of *common ground*, that is, the background knowledge, beliefs, and assumptions shared by interlocutors (Clark, 1996; Pickering & Garrod, 2004), in shaping conversational dynamics, the inferred level of shared expertise between the interlocutors appears to modulate conversational strategies. When objects were unfamiliar (e.g., “okapi”), speakers were more likely to describe perceptual features and convey key information regarding the object. Conversely, when objects were familiar (e.g., “giraffe”), and the level of shared expertise is assumed, such that central information is taken to be known, the established common ground prompts speakers to more likely internalize the conversation, shifting the focus toward personal experience and emotional dimension.

Finally, the third analysis revealed that references to emotional experience were the strongest predictor of both noun and sentence concreteness. In particular, when speakers referred to objects through their emotional experience, the concreteness of their language decreased, in line with the Affective Embodiment Account (AEA). However, this finding does not directly inform the ongoing debate concerning whether abstract or concrete concepts are more emotionally charged (Winter, 2022). Nonetheless, the present result demonstrated that, in semi-naturalistic conversations, references to emotional dimension are systematically associated with a reduction in concreteness of the language used.

References to objects' features, both perceptual and encyclopedic, also emerged as significant predictors, revealing increased noun and sentence concreteness when speakers focused on these aspects. Notably, superordinate categorization was a relatively strong predictor of speakers' language concreteness, with consistent effects at both the noun and sentence levels. This finding provides further evidence for the view that categorical specificity and conceptual concreteness constitute distinct semantic dimensions and that the linguistic and conceptual senses of abstraction discussed in Chapter 1 cannot be conflated. Importantly, an increase in abstraction through the use of

hypernyms of highly concrete and specific target objects did not lead to a decrease in concreteness; rather, it was associated with an increase. This pattern further supports the distinction between abstraction and abstractness/concreteness.

Abstraction strategies emerged as key predictors of noun and sentence specificity. When speakers employed analogical comparisons, their language became more specific, plausibly to instantiate diagnostic features or properties and to establish a relational structure (Gentner et al., 2017). By contrast, expressions containing superordinate categorization were associated with decreased specificity at both noun and sentence levels. However, as extensively discussed in ¶ 6.3., comparisons between noun and sentence specificity models indicated that, at the sentence level, the negative effect of superordinate categorization on specificity was attenuated.

Interestingly, references to emotional experience also emerged as a relatively strong predictor of categorical specificity, indicating a decrease in specificity when speakers referred to the emotional dimension. As discussed in ¶ 6.3., this pattern aligns with existing literature suggesting that emotional experience involves variation in levels of granularity (Barrett, 2006). Moreover, Bolognesi & Caselli (2023) showed that reduced specificity is associated with positive affective states.

Taken together, the findings of this third analysis indicate that even within the domain of pragmatic and communicative strategies, rather than on semantic representations alone, core dimensions of concreteness and specificity exhibit distinct patterns. Concreteness appears to be primarily associated with references to perceptual and inner content, although superordinate categorization also reliably predicted increases in concreteness. In contrast, categorical specificity reflects abstraction-related strategies operating at different levels of abstraction, although emotional experience also emerged as a reliable predictor of specificity.

## **7.4. Conclusions**

Consider the metaphor of a camera: one can zoom in to focus on precise, detailed, and specific components of the scene, or zoom out to obtain a broader view in which the environment appears more blurred but encompasses multiple settings. The paradox of categorical specificity as a language-mediated form of abstraction lies precisely in this

duality. On the one hand, it is possible to instantiate precise information that enhances cognitive processing and communicative strategies by focusing on specific details. On the other hand, adopting a wider perspective involves engaging with indeterminate reference, which nonetheless provides access to greater variability and flexibility in meaning.

This work offers several strengths, including the investigation of the relationship between categorical specificity and conceptual concreteness in communicative contexts and social interaction, and the employment of a progressively ecological methodology. However, this work also has limitations, such as focusing on referential aspects of language, such as categorical inclusiveness, and not having deepened other crucial aspects, such as multimodality. Moreover, although it has advanced several crucial notions within the field of pragmatics, it did not explore the role of categorical specificity in traditional pragmatic paradigms. Future research may encompass the analysis of categorical specificity within interactive contexts, but also focus on these other aspects of meaning and methodological approach.



## REFERENCES

- Altarriba, J., Bauer, L. M., & Benvenuto, C. (1999). Concreteness, context availability, and imageability ratings and word associations for abstract, concrete, and emotion words. *Behavior Research Methods, Instruments, & Computers*, 31(4), 578–602. <https://doi.org/10.3758/BF03200738>
- Andrews, M., Vigliocco, G., & Vinson, D. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological Review*, 116(3), 463–498. <https://doi.org/10.1037/a0016261>
- Angius, N., Primo, G., & Turner, R. (2025). The Philosophy of Computer Science. In E. N. Zalta & Nodelman, Uri (Eds.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/spr2025/entries/computer-science/>
- Asmuth, J., & Gentner, D. (2017). Relational Categories are More Mutable than Entity Categories. *Quarterly Journal of Experimental Psychology*, 70(10), 2007–2025. <https://doi.org/10.1080/17470218.2016.1219752>
- Bäck, A. (2014). *Aristotle's Theory of Abstraction* (Vol. 73). Springer International Publishing. <https://doi.org/10.1007/978-3-319-04759-1>
- Balota, D. A., Ferraro, F. R., & Connor, L. T. (1991). On the Early Influence of Meaning in Word Recognition: A Review of the Literature. In P. J. Schwanenflugel, *The Psychology of Word Meanings* (pp. 187–222).
- Balota, D. A., Yap, M. J., Hutchison, K. A., Cortese, M. J., Kessler, B., Loftis, B., Neely, J. H., Nelson, D. L., Simpson, G. B., & Treiman, R. (2007). The English Lexicon Project. *Behavior Research Methods*, 39(3), 445–459. <https://doi.org/10.3758/BF03193014>
- Banks, B., Borghi, A. M., Fargier, R., Fini, C., Jonauskaitė, D., Mazzuca, C., Montalti, M., Villani, C., & Woodin, G. (2023). Consensus Paper: Current Perspectives on Abstract Concepts and Future Research Directions. *Journal of Cognition*, 6(1), 62. <https://doi.org/10.5334/joc.238>
- Banks, B., & Connell, L. (2022). Multi-dimensional sensorimotor grounding of concrete and abstract categories. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1870), 20210366. <https://doi.org/10.1098/rstb.2021.0366>

- Barber, H. A., Otten, L. J., Kousta, S.-T., & Vigliocco, G. (2013). Concreteness in word processing: ERP and behavioral effects in a lexical decision task. *Brain and Language*, 125(1), 47–53. <https://doi.org/10.1016/j.bandl.2013.01.005>
- Barca, L., Burani, C., & Arduino, L. S. (2002). Word naming times and psycholinguistic norms for Italian nouns. *Behavior Research Methods, Instruments, & Computers*, 34(3), 424–434. <https://doi.org/10.3758/BF03195471>
- Barca, L., Mazzuca, C., & Borghi, A. M. (2020). Overusing the pacifier during infancy sets a footprint on abstract words processing. *Journal of Child Language*, 47(5), 1084–1099. <https://doi.org/10.1017/S0305000920000070>
- Barrett, L. F. (2006). Solving the Emotion Paradox: Categorization and the Experience of Emotion. *Personality and Social Psychology Review*, 10(1), 20–46. [https://doi.org/10.1207/s15327957pspr1001\\_2](https://doi.org/10.1207/s15327957pspr1001_2)
- Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11(4), 629–654. <https://doi.org/10.1037/0278-7393.11.1-4.629>
- Barsalou, L. W. (1992). Frames, Concepts, and Conceptual Fields. In A. Lehrer & E. F. Kittay (Eds.), *Frames, Fields, and Contrasts: New Essays in Semantic and Lexical Organization* (pp. 21–74). Routledge/Taylor & Francis Group.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22(4), 577–660. <https://doi.org/10.1017/S0140525X99002149>
- Barsalou, L. W. (2003). Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 358(1435), 1177–1187. <https://doi.org/10.1098/rstb.2003.1319>
- Barsalou, L. W. (2005a). Abstraction as Dynamic Interpretation in Perceptual Symbol Systems. In L. Gershkoff-Stowe & D. H. Rakison (Eds.), *Building object categories in developmental time* (pp. 389–431). Erlbaum.
- Barsalou, L. W. (2005b). Situated Conceptualization. In *Handbook of Categorization in Cognitive Science* (pp. 619–650). Elsevier Science Ltd. <https://doi.org/10.1016/B978-008044612-7/50083-4>
- Barsalou, L. W. (2008). Grounded Cognition. *Annual Review of Psychology*, 59(Volume 59, 2008), 617–645. <https://doi.org/10.1146/annurev.psych.59.103006.093639>

- Barsalou, L. W. (2012). The Human Conceptual System. In K. McRae, M. Joanisse, & M. Spivey (Eds.), *The Cambridge Handbook of Psycholinguistics* (pp. 239–258). Cambridge University Press. <https://doi.org/10.1017/CBO9781139029377.013>
- Barsalou, L. W. (2016). Situated conceptualization: Theory and applications. In Y. Coello & M. H. Fischer (Eds.), *Foundations of embodied cognition: Perceptual and emotional embodiment* (pp. 11–37). Routledge/Taylor & Francis Group.
- Barsalou, L. W., Dutriaux, L., & Scheepers, C. (2018). Moving beyond the distinction between concrete and abstract concepts. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1752), 20170144. <https://doi.org/10.1098/rstb.2017.0144>
- Barsalou, L. W., & Hale, C. (1993). Components of Conceptual Representation. From Feature Lists to Recursive Frames. In I. van Mechelen, J. Hampton, R. S. Michalski, & P. Theuns (Eds.), *Categories and Concepts: Theoretical Views and Inductive Data Analysis* (pp. 97–144). Academic Press.
- Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). Language and simulation in conceptual processing. In M. de Vega, A. Glenberg, & A. Graesser (Eds.), *Symbols and Embodiment: Debates on meaning and cognition* (pp. 245–284). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199217274.003.0013>
- Barsalou, L. W., & Wiemer-Hastings, K. (2005). Situating Abstract Concepts. In D. Pecher & R. A. Zwaan (Eds.), *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thinking* (pp. 129–163). Cambridge University Press. <https://doi.org/10.1017/CBO9780511499968.007>
- Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. Cambridge University Press.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using **lme4**. *Journal of Statistical Software*, 67(1). <https://doi.org/10.18637/jss.v067.i01>
- Bauer, A. J., & Just, M. A. (2017). A brain-based account of “basic-level” concepts. *NeuroImage*, 161, 196–205. <https://doi.org/10.1016/j.neuroimage.2017.08.049>
- Bolognesi, M. (2020). *Where words get their meaning: Cognitive processing and distributional modelling of word meaning in first and second language*. John Benjamins Publishing Company. <https://doi.org/https://doi.org/10.1075/celcr.23>

- Bolognesi, M. (2025). Beyond concreteness: Why word specificity is the missing piece in theories on embodied language comprehension. *Frontiers in Psychology*, 16, 1718655. <https://doi.org/10.3389/fpsyg.2025.1718655>
- Bolognesi, M., Burgers, C., & Caselli, T. (2020). On abstraction: Decoupling conceptual concreteness and categorical specificity. *Cognitive Processing*, 21(3), 365–381. <https://doi.org/10.1007/s10339-020-00965-9>
- Bolognesi, M. M., & Caselli, T. (2023). Specificity ratings for Italian data. *Behavior Research Methods*, 55(7), 3531–3548. <https://doi.org/10.3758/s13428-022-01974-6>
- Bond, F., & Paik, K. (2012). A survey of wordnets and their licenses. In *Proceedings of the 6th global WordNet conference (GWC 2012)*, 64–71.
- Borghi, A. M. (2022). Concepts for Which We Need Others More: The Case of Abstract Concepts. *Current Directions in Psychological Science*, 31(3), 238–246. <https://doi.org/10.1177/09637214221079625>
- Borghi, A. M. (2023). *The Freedom of Words: Abstractness and the Power of Language*. Cambridge University Press. <https://doi.org/10.1017/9781108913294>
- Borghi, A. M., Barca, L., Binkofski, F., Castelfranchi, C., Pezzulo, G., & Tummolini, L. (2019). Words as social tools: Language, sociality and inner grounding in abstract concepts. *Physics of Life Reviews*, 29, 120–153. <https://doi.org/10.1016/j.plrev.2018.12.001>
- Borghi, A. M., Barca, L., Binkofski, F., & Tummolini, L. (2018a). Abstract concepts, language and sociality: From acquisition to inner speech. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1752), 20170134. <https://doi.org/10.1098/rstb.2017.0134>
- Borghi, A. M., Barca, L., Binkofski, F., & Tummolini, L. (2018b). Varieties of abstract concepts: Development, use and representation in the brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1752), 20170121. <https://doi.org/10.1098/rstb.2017.0121>
- Borghi, A. M., & Binkofski, F. (2014). *Words as social tools: An embodied view on abstract concepts* (Vol. 2). Springer New York.
- Borghi, A. M., Binkofski, F., Castelfranchi, C., Cimatti, F., Scorolli, C., & Tummolini, L. (2017). The challenge of abstract concepts. *Psychological Bulletin*, 143(3), 263–292. <https://doi.org/https://doi.org/10.1037/bul0000089>
- Borghi, A. M., & Cimatti, F. (2010). Embodied cognition and beyond: Acting and sensing the body. *Neuropsychologia, The Sense of Body*, 48(3), 763–773. <https://doi.org/10.1016/j.neuropsychologia.2009.10.029>

- Borghgi, A. M., Fini, C., & Tummolini, L. (2021). Abstract Concepts and Metacognition: Searching for Meaning in Self and Others. In M. D. Robinson & L. E. Thomas (Eds.), *Handbook of Embodied Psychology: Thinking, Feeling, and Acting* (pp. 197–220). Springer International Publishing. [https://doi.org/10.1007/978-3-030-78471-3\\_9](https://doi.org/10.1007/978-3-030-78471-3_9)
- Borghgi, A. M., Mazzuca, C., & Tummolini, L. (2025). The role of social interaction in the formation and use of abstract concepts. *Nature Reviews Psychology*, 4(7), 470–483. <https://doi.org/10.1038/s44159-025-00451-z>
- Borghgi, A. M., & Pecher, D. (2011). Introduction to the Special Topic Embodied and Grounded Cognition. *Frontiers in Psychology*, 2. <https://doi.org/10.3389/fpsyg.2011.00187>
- Borghgi, A. M., & Zarcone, E. (2016). Grounding Abstractness: Abstract Concepts and the Activation of the Mouth. *Frontiers in Psychology*, 7. <https://doi.org/10.3389/fpsyg.2016.01498>
- Boroditsky, L., & Ramscar, M. (2002). The roles of body and mind in abstract thought. *Psychological Science*, 13(2), 185–189. <https://doi.org/10.1111/1467-9280.00434>
- Bowdle, B. F., & Gentner, D. (2005). The career of metaphor. *Psychological Review*, 112(1), 193–216.
- Bradley, M. M., & Lang, P. J. (1999). *Affective norms for English words (ANEW): Instruction manual and affective ratings*. Technical report C-1, the center for research in psychophysiology.
- Brentano, F. (1874). *Psychology from an Empirical Standpoint* (O. Kraus & McAllister Linda L., Eds.; A. C. Rancurello, D. B. Terrell, & McAllister Linda L., Trad.). Routledge & Kegan Paul Ltd.
- Brown, R. (1958). How shall a thing be called? *Psychological Review*, 65(1), 14–21. <https://doi.org/10.1037/h0041727>
- Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46(3), 904–911. <https://doi.org/10.3758/s13428-013-0403-5>
- Burgoon, E. M., Henderson, M. D., & Markman, A. B. (2013). There Are Many Ways to See the Forest for the Trees: A Tour Guide for Abstraction. *Perspectives on Psychological Science*, 8(5), 501–520. [https://doi.org/https://doi.org/10.1177/1745691613497964](https://doi.org/10.1177/1745691613497964)
- Carlson, G. N., & Pelletier, F. J. (Eds.). (1995). *The Generic Book: The Semantics of Generics*. University of Chicago Press.
- Casasanto, D., & Boroditsky, L. (2008). Time in the mind: Using space to think about time. *Cognition*, 106(2), 579–593. <https://doi.org/10.1016/j.cognition.2007.03.004>

- Chandler, S. (2017). The analogical modeling of linguistic categories. *Language and Cognition*, 9(1), 52–87. <https://doi.org/10.1017/langcog.2015.24>
- Chomsky, N. (1957). *Syntactic Structures*. Mouton.
- Chomsky, N. (1968). *Language and Mind*. Harper and Row.
- Chomsky, N. (1995). *The Minimalist Program*. The MIT Press.
- Christensen, R. (2023). *Ordinal-Regression Models for Ordinal Data*. R package version 2023.12-4.1 [Software]. <https://CRAN.R-project.org/package=ordinal>
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780190217013.001.0001>
- Clark, H. H. (1996). *Using Language*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511620539>
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition*. (pp. 127–149). American Psychological Association. <https://doi.org/10.1037/10096-006>
- Clark, H. H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22(1), 1–39. [https://doi.org/10.1016/0010-0277\(86\)90010-7](https://doi.org/10.1016/0010-0277(86)90010-7)
- Cleary, J. J. (1985). On the Terminology of «Abstraction» in Aristotle. *Phronesis*, 30(1), 13–45.
- Collacciani, C., Ravelli, A. A., & Bolognesi, M. (2024, maggio). Specifying Genericity through Inclusiveness and Abstractness Continuous Scales. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING).
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 8(2), 240–247. [https://doi.org/10.1016/S0022-5371\(69\)80069-1](https://doi.org/10.1016/S0022-5371(69)80069-1)
- Colunga, E., & Smith, L. B. (2003). The emergence of abstract ideas: Evidence from networks and babies. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 358(1435), 1205–1214. <https://doi.org/10.1098/rstb.2003.1306>
- Connell, L., & Lynott, D. (2012). Strength of perceptual experience predicts word processing performance better than concreteness or imageability. *Cognition*, 125(3), 452–465. <https://doi.org/10.1016/j.cognition.2012.07.010>

- Connell, L., & Lynott, D. (2013). Flexible and fast: Linguistic shortcut affects both shallow and deep conceptual processing. *Psychonomic Bulletin & Review*, 20(3), 542–550. <https://doi.org/10.3758/s13423-012-0368-x>
- Connell, L., Lynott, D., & Banks, B. (2018). Interoception: The forgotten modality in perceptual grounding of abstract and concrete concepts. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1752), 20170143. <https://doi.org/10.1098/rstb.2017.0143>
- Crepaldi, D., Amenta, S., Mandera, P., Keuleers, E., & Brysbaert, M. (2016). Frequency estimates from different registers explain different aspects of visual word recognition. *International Meeting of the Psychonomic Society*, 5–8.
- Croft, W., & Cruse, D. A. (2004). *Cognitive Linguistics*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511803864>
- Cruse, D. A. (1977). The pragmatics of lexical specificity. *Journal of Linguistics*, 13(2), 153–164. <https://doi.org/10.1017/S0022226700005363>
- Cruse, D. A. (1986). *Lexical Semantics*. Cambridge University Press.
- Cruse, D. A. (2017). The Lexicon. In M. Aronoff & J. Rees-Miller (Eds.), *The Handbook of Linguistics* (pp. 235–254). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781119072256.ch12>
- Crutch, S. J., Troche, J., Reilly, J., & Ridgway, G. R. (2013). Abstract conceptual feature ratings: The role of emotion, magnitude, and other cognitive domains in the organization of abstract conceptual knowledge. *Frontiers in Human Neuroscience*, 7. <https://doi.org/10.3389/fnhum.2013.00186>
- Crutch, S. J., & Warrington, E. K. (2005). Abstract and concrete concepts have structurally different representational frameworks. *Brain*, 128(3), 615–627. <https://doi.org/10.1093/brain/awh349>
- Cuccio, V. (2015). Embodied Simulation and Metaphors: On the Role of the Body in the Interpretation of Bodily-Based Metaphors. *Epistemologia: rivista italiana di filosofia della scienza*, XXXVIII(1), 99–113. <https://doi.org/10.3280/EPIS2015-001007>
- Cuccio, V. (2018). *Attention to Metaphor: From neurons to representations* Attention. <https://doi.org/10.1075/milcc.7>
- Cuccio, V., & Caruana, F. (2019). Rethinking the abstract/concrete concepts dichotomy. *Physics of Life Reviews*, 29, 157–160. <https://doi.org/10.1016/j.plrev.2019.04.007>

- Cuccio, V., & Gallese, V. (2018). A Peircean account of concepts: Grounding abstraction in phylogeny through a comparative neuroscientific perspective. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1752), 20170128. <https://doi.org/10.1098/rstb.2017.0128>
- Daelemans, W., & Van den Bosch, A. (2005). *Memory-based language processing*. Cambridge University Press.
- Damasio, A. R. (1989). Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition, Special Issue Neurobiology of Cognition*, 33(1), 25–62. [https://doi.org/10.1016/0010-0277\(89\)90005-X](https://doi.org/10.1016/0010-0277(89)90005-X)
- Davies, C., & Arnold, J. E. (2019). Reference and Informativeness. In C. Cummins & N. Katsos (Eds.), *The Oxford Handbook of Experimental Semantics and Pragmatics* (p. 0). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780198791768.013.6>
- Davis, C. P., Altmann, G. T. M., & Yee, E. (2020). Situational systematicity: A role for schema in understanding the differences between abstract and concrete concepts. *Cognitive Neuropsychology*, 37(1–2), 142–153. <https://doi.org/10.1080/02643294.2019.1710124>
- Davis, C. P., & Yee, E. (2019). Features, labels, space, and time: Factors supporting taxonomic relationships in the anterior temporal lobe and thematic relationships in the angular gyrus. *Language, Cognition and Neuroscience*, 34(10), 1347–1357. <https://doi.org/10.1080/23273798.2018.1479530>
- De Felice, S., Hamilton, A. F. D. C., Ponari, M., & Vigliocco, G. (2023). Learning *from* others is good, *with* others is better: The role of social interaction in human acquisition of new knowledge. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1870), 20210357. <https://doi.org/10.1098/rstb.2021.0357>
- De Marneffe, M.-C., Manning, C. D., Nivre, J., & Zeman, D. (2021). Universal Dependencies. *Computational Linguistics*, 1–54. [https://doi.org/10.1162/coli\\_a\\_00402](https://doi.org/10.1162/coli_a_00402)
- De Saussure, F. (1916). *Course in General Linguistics* (C. Bally & A. Sechehaye, Eds.; W. Baskin, Trad.). McGraw-Hill Book Company.
- de Varda, A. G., Lamarra, T., Ravelli, A. A., Saponaro, C., Giustolisi, B., & Bolognesi, M. (2025). IconicITA: Iconicity ratings of the Italian affective lexicon. *PLOS One*, 20(12), e0337947. <https://doi.org/10.1371/journal.pone.0337947>
- de Vega, M., Glenberg, A., & Graesser, A. (Eds.). (2008). *Symbols and Embodiment: Debates on meaning and cognition*. Oxford University Press.

- Deacon, T. W. (1997). *The Symbolic Species: The Coevolution of Language and the Brain*. WW Norton & Company.
- Della Rosa, P. A., Catricalà, E., Vigliocco, G., & Cappa, S. F. (2010). Beyond the abstract-concrete dichotomy: Mode of acquisition, concreteness, imageability, familiarity, age of acquisition, context availability, and abstractness norms for a set of 417 Italian words. *Behavior Research Methods*, 42(4), 1042–1048. <https://doi.org/10.3758/BRM.42.4.1042>
- Dewey, J. (1910). *How We Think*. D.C. Heath & Co Publishers.
- Dewey, J. (1938). *Logic: The Theory of Inquiry*. Henry Holt and Company.
- Dideriksen, C., Fusaroli, R., Tylén, K., Dingemanse, M., & Christiansen, M. H. (2019). *Contextualizing Conversational Strategies: Backchannel, Repair and Linguistic Alignment in Spontaneous and Task-Oriented Conversations*. 261–267. <https://doi.org/10.31234/osf.io/fd8y9>
- Diveica, V., Pexman, P. M., & Binney, R. J. (2022). Quantifying social semantics: An inclusive definition of socialness and ratings for 8388 English words. *Behavior Research Methods*, 55(2), 461–473. <https://doi.org/10.3758/s13428-022-01810-x>
- Domaneschi, F., & Bambini, V. (2020). Pragmatic Competence. In *The Routledge Handbook of Philosophy of Skill and Expertise* (pp. 419–430). Routledge.
- Dove, G. (2020). More than a scaffold: Language is a neuroenhancement. *Cognitive Neuropsychology*, 37(5–6), 288–311. <https://doi.org/10.1080/02643294.2019.1637338>
- Dove, G. (2021). The Challenges of Abstract Concepts. In M. D. Robinson & L. E. Thomas (Eds.), *Handbook of Embodied Psychology: Thinking, Feeling, and Acting* (pp. 171–195). Springer International Publishing. [https://doi.org/10.1007/978-3-030-78471-3\\_8](https://doi.org/10.1007/978-3-030-78471-3_8)
- Dove, G. O. (2023). Rethinking the role of language in embodied cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1870), 20210375. <https://doi.org/10.1098/rstb.2021.0375>
- Eddington, C. M., & Tokowicz, N. (2015). How meaning similarity influences ambiguous word processing: The current state of the literature. *Psychonomic Bulletin & Review*, 22(1), 13–37. <https://doi.org/10.3758/s13423-014-0665-7>
- Enfield, N. J. (2024). *Language vs. Reality: Why Language Is Good for Lawyers and Bad for Scientists*. MIT Press.
- Evans, V., & Green, M. (2006). *Cognitive Linguistics: An Introduction*. Edinburgh University Press.

- Falcinelli, I., Fini, C., Mazzuca, C., Alessandri, G., Alivernini, F., Baiocco, R., Chirico, A., Filosa, L., Palombi, T., Pistella, J., Tavolucci, S., Lucidi, F., & Borghi, A. M. (2024). What does “Internet” mean to us as we age? A multi-task investigation on the conceptualization of the technological domain across generations. *Computers in Human Behavior Reports*, 16, 100531. <https://doi.org/10.1016/j.chbr.2024.100531>
- Falcinelli, I., Fini, C., Mazzuca, C., & Borghi, A. M. (2024). *The TECo Database: Insights on The Semantic Organization of The Ecological Domain. Vol. 46*.
- Falguera, J. L., Martínez-Vidal, C., & Rosen, G. (2025). Abstract Objects. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2025 Edition). <https://plato.stanford.edu/archives/fall2025/entries/abstract-objects/>
- Fellbaum, C. (1998). *WordNet: An Electronic Lexical Database*. The MIT Press. <https://doi.org/10.7551/mitpress/7287.001.0001>
- Fernyhough, C., & Borghi, A. M. (2023). Inner speech as language process and cognitive tool. *Trends in Cognitive Sciences*, 27(12), 1180–1193. <https://doi.org/10.1016/j.tics.2023.08.014>
- Fillmore, C. J. (1975). An alternative to checklist theories of meaning. In *Proceedings of the first annual meeting of the Berkeley Linguistics Society*, 1, 123–131.
- Fillmore, C. J. (1977). Scenes-and-frames semantics. In A. Zampolli (Eds.), *Linguistic structures processing* (Vol. 5, pp. 55–88). North Holland Publishing.
- Fillmore, C. J. (1985). Frames and the semantics of understanding. *Quaderni di Semantica*, 6, 222–254.
- Fillmore, C. J., & Atkins, B. T. (1992). Toward a Frame-Based Lexicon: The Semantics of RISK and Its Neighbors. In A. Lehrer & E. F. Kittay (Eds.), *Frames, Fields and Contrasts: New Essays in Semantic and Lexical Organization* (pp. 75–102). Routledge/Taylor & Francis Group.
- Fini, C., Falcinelli, I., Cuomo, G., Era, V., Candidi, M., Tummolini, L., Mazzuca, C., & Borghi, A. M. (2023). Breaking the ice in a conversation: Abstract words prompt dialogs more easily than concrete ones. *Language and Cognition*, 15(4), 629–650. <https://doi.org/10.1017/langcog.2023.3>
- Firth, J. R. (1957). A Synopsis of Linguistic Theory, 1930-55. In *Studies in Linguistic Analysis* (pp. 1–31). Blackwell.
- Floridi, L. (2008). The Method of Levels of Abstraction. *Minds and Machines*, 18(3), 303–329. <https://doi.org/10.1007/s11023-008-9113-7>
- Fodor, J. A. (1975). *The Language of Thought*. Harvard University Press.

- Fodor, J. A. (1983). *The Modularity of Mind*. The MIT Press.  
<https://doi.org/10.7551/mitpress/4737.001.0001>
- Fodor, J. A. (1985). Précis of The Modularity of Mind. *Behavioral and Brain Sciences*, 8(1), 1–5.  
<https://doi.org/10.1017/S0140525X0001921X>
- Fortescue, M. (2017). *The Abstraction Engine: Extracting Patterns in Language, Mind and Brain* (Vol. 00094). John Benjamins Publishing Company.
- Fox, J. (2005). Linear models, problems. In *Encyclopedia of social measurement* (Vol. 2, pp. 515–522).
- Fox, J., & Weisberg, S. (2019). *An R Companion to Applied Regression* (3rd ed.). Sage.
- Fujita, K., Trope, Y., Liberman, N., & Levin-Sagi, M. (2006). Construal Levels and Self-Control. *Journal of personality and social psychology*, 90(3), 351–367. <https://doi.org/10.1037/0022-3514.90.3.351>
- Gallese, V., & Lakoff, G. (2005). The Brain's concepts: The role of the Sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology*, 22(3), 455–479.  
<https://doi.org/10.1080/02643290442000310>
- Gandolfi, G., Pickering, M. J., & Garrod, S. (2023). Mechanisms of alignment: Shared control, social cognition and metacognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1870), 20210362. <https://doi.org/10.1098/rstb.2021.0362>
- Garrod, S., & Pickering, M. J. (2009). Joint Action, Interactive Alignment, and Dialog. *Topics in Cognitive Science*, 1(2), 292–304. <https://doi.org/10.1111/j.1756-8765.2009.01020.x>
- Geeraerts, D. (2008). Chapter 4 Prototype theory. In D. Geeraerts (Eds.), *Cognitive Linguistics: Basic Readings* (pp. 141–166). De Gruyter Mouton.  
<https://www.degruyterbrill.com/document/doi/10.1515/9783110199901.141/html>
- Gentner, D. (1983). Structure-Mapping: A Theoretical Framework for Analogy\*. *Cognitive Science*, 7(2), 155–170. [https://doi.org/10.1207/s15516709cog0702\\_3](https://doi.org/10.1207/s15516709cog0702_3)
- Gentner, D. (1989). The mechanisms of analogical learning. In A. Ortony & S. Vosniadou (Eds.), *Similarity and Analogical Reasoning* (pp. 199–241). Cambridge University Press.  
<https://doi.org/10.1017/CBO9780511529863.011>
- Gentner, D., & Asmuth, J. (2017). Metaphoric extension, relational categories, and abstraction. *Language, Cognition and Neuroscience*, 34(10), 1298–1307.  
<https://doi.org/https://doi.org/10.1080/23273798.2017.1410560>

- Gentner, D., Maravilla, F., Ball, L. J., & Thompson, V. A. (2017). Analogical Reasoning. In *International Handbook of Thinking and Reasoning* (pp. 186–203). Routledge. <https://doi.org/https://doi.org/10.4324/9781315725697>
- Gentner, D., & Markman, A. B. (1997). Structure Mapping in Analogy and Similarity. *American Psychologist*, 52(1), 45–56. <https://doi.org/https://doi.org/10.1037/0003-066X.52.1.45>
- Gibbs Jr, R. W. (Eds.). (2008). *The Cambridge Handbook of Metaphor and Thought*. Cambridge University Press.
- Gibson, J. J. (1950). *The perception of the visual world*. Houghton Mifflin.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.
- Glenberg, A. M. (2010). Embodiment as a unifying perspective for psychology. *WIREs Cognitive Science*, 1(4), 586–596. <https://doi.org/10.1002/wcs.55>
- Glenberg, A. M., Witt, J. K., & Metcalfe, J. (2013). From the Revolution to Embodiment: 25 Years of Cognitive Psychology. *Perspectives on Psychological Science*, 8(5), 573–585. <https://doi.org/10.1177/1745691613498098>
- Goh, W. D., Yap, M. J., Lau, M. C., Ng, M. M. R., & Tan, L.-C. (2016). Semantic Richness Effects in Spoken Word Recognition: A Lexical Decision and Semantic Categorization Megastudy. *Frontiers in Psychology*, 7. <https://doi.org/10.3389/fpsyg.2016.00976>
- Goldberg, A. (1995). *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago University Press.
- Goldstone, R. L., & Sakamoto, Y. (2003). The transfer of abstract principles governing complex adaptive systems. *Cognitive Psychology*, 46(4), 414–466. [https://doi.org/10.1016/S0010-0285\(02\)00519-4](https://doi.org/10.1016/S0010-0285(02)00519-4)
- Gopnik, A. (2003). The Theory Theory as an Alternative to the Innateness Hypothesis. In L. M. Antony & N. Hornstein (Eds.), *Chomsky and His Critics* (pp. 238–254). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9780470690024.ch10>
- Gopnik, A., & Wellman, H. M. (1994). The theory theory. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture* (pp. 257–293). Cambridge University Press. <https://doi.org/10.1017/CBO9780511752902.011>
- Greene, M. R., & Rohan, A. M. (2025). The brain prioritizes the basic level of object category abstraction. *Scientific Reports*, 15(1), 31. <https://doi.org/10.1038/s41598-024-80546-4>
- Grice, P. H. (1989). *Studies in the Way of Words*. Harvard University Press.

- Gu, Y., Donnellan, E., Grzyb, B., Brekelmans, G., Murgiano, M., Brieke, R., Perniss, P., & Vigliocco, G. (2025). The ECOLANG Multimodal Corpus of adult-child and adult-adult Language. *Scientific Data*, 12(1), 89. <https://doi.org/10.1038/s41597-025-04405-1>
- Hajibayova, L. (2013). Basic-level categories: A review. *Journal of Information Science*, 39(5), 676–687. <https://doi.org/10.1177/0165551513481443>
- Hampton, J. A. (1979). Polymorphous concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 18(4), 441–461. [https://doi.org/10.1016/S0022-5371\(79\)90246-9](https://doi.org/10.1016/S0022-5371(79)90246-9)
- Hampton, J. A. (1997). Psychological representation of concepts. In M. A. Conway (Eds.), *Cognitive models of memory* (pp. 81–110). The MIT Press.
- Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1), 335–346. [https://doi.org/10.1016/0167-2789\(90\)90087-6](https://doi.org/10.1016/0167-2789(90)90087-6)
- Harris, Z. S. (1954). Distributional Structure. *WORD*, 10(2–3), 146–162. <https://doi.org/10.1080/00437956.1954.11659520>
- Heit, E., & Barsalou, L. W. (1996). The Instantiation Principle in Natural Categories. *Memory*, 4(4), 413–452. <https://doi.org/10.1080/096582196388915>
- Hockett, C. F. (1960). The Origin of Speech. *Scientific American*, 203, 88–111.
- Hoffman, P. (2016). The meaning of ‘life’ and other abstract words: Insights from neuropsychology. *Journal of Neuropsychology*, 10(2), 317–343. <https://doi.org/10.1111/jnp.12065>
- Hutchison, K. A., Balota, D. A., Neely, J. H., Cortese, M. J., Cohen-Shikora, E. R., Tse, C.-S., Yap, M. J., Bengson, J. J., Niemeyer, D., & Buchanan, E. (2013). The semantic priming project. *Behavior Research Methods*, 45(4), 1099–1114. <https://doi.org/10.3758/s13428-012-0304-z>
- Iliev, R., & Axelrod, R. (2017). The Paradox of Abstraction: Precision Versus Concreteness. *Journal of Psycholinguistic Research*, 46(3), 715–729. <https://doi.org/10.1007/s10936-016-9459-6>
- Jackendoff, R. (2003). *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford University Press.
- Jakubíček, M., Kilgariff, A., Kovář, V., Rychlý, P., & Suchomel, V. (2013). The TenTen corpus family. In *7th international corpus linguistics conference CL, 2013*.
- Jessen, F., Heun, R., Erb, M., Granath, D. O., Klose, U., Papassotiropoulos, A., & Grodd, W. (2000). The Concreteness Effect: Evidence for Dual Coding and Context Availability. *Brain and Language*, 74(1), 103–112. <https://doi.org/10.1006/brln.2000.2340>

- Johnson, K. E., & Mervis, C. B. (1997). Effects of varying levels of expertise on the basic level of categorization. *Journal of Experimental Psychology: General*, 126(3), 248–277. <https://doi.org/10.1037/0096-3445.126.3.248>
- Johnson, M. (1987). *The body in the mind: The bodily basis of meaning, imagination, and reason*. University of Chicago Press.
- Johnson, M. (2007). *The Meaning of the Body: Aesthetics of Human Understanding*. University of Chicago Press. <https://doi.org/10.7208/chicago/9780226026992.001.0001>
- Johnson, M. L., & Tucker, D. M. (2021). *Out of the Cave: A Natural Philosophy of Mind and Knowing*. The MIT Press.
- Jolicoeur, P., Gluck, M. A., & Kosslyn, S. M. (1984). Pictures and names: Making the connection. *Cognitive Psychology*, 16(2), 243–275. [https://doi.org/10.1016/0010-0285\(84\)90009-4](https://doi.org/10.1016/0010-0285(84)90009-4)
- Keil, F. C. (1992). *Concepts, Kinds, and Cognitive Development*. The MIT Press.
- Keuleers, E., Lacey, P., Rastle, K., & Brysbaert, M. (2012). The British Lexicon Project: Lexical decision data for 28,730 monosyllabic and disyllabic English words. *Behavior Research Methods*, 44(1), 287–304. <https://doi.org/10.3758/s13428-011-0118-4>
- Khanna, M. M., & Cortese, M. J. (2021). How well imageability, concreteness, perceptual strength, and action strength predict recognition memory, lexical decision, and reading aloud performance. *Memory*, 29(5), 622–636. <https://doi.org/10.1080/09658211.2021.1924789>
- Kiefer, M., & Pulvermüller, F. (2012). Conceptual representations in mind and brain: Theoretical developments, current evidence and future directions. *Cortex, Language and the Motor System*, 48(7), 805–825. <https://doi.org/10.1016/j.cortex.2011.04.006>
- Kobrock, K., Uhlemann, C., & Gotzner, N. (2024). Superordinate referring expressions in abstraction: Introducing the concept-level reference game. *In Proceedings of the Annual Meeting of the Cognitive Science Society*, 46. <https://doi.org/10.31219/osf.io/cv74u>
- Kounios, J., & Holcomb, P. J. (1994). Concreteness effects in semantic processing: ERP evidence supporting dual-coding theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(4), 804–823. <https://doi.org/10.1037/0278-7393.20.4.804>
- Kousta, S.-T., Vigliocco, G., Vinson, D. P., Andrews, M., & Del Campo, E. (2011). The representation of abstract words: Why emotion matters. *Journal of Experimental Psychology: General*, 140(1), 14–34. <https://doi.org/10.1037/a0021446>

- Lakoff, G. (2008). The neural theory of metaphor. In *The Cambridge handbook of metaphor and thought* (pp. 17–38). Cambridge University Press.  
<https://doi.org/10.1017/CBO9780511816802.003>
- Lakoff, G., & Johnson, M. (1980). *Metaphors We Live by*. University of Chicago Press.
- Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh: The embodied mind and its challenge to Western thought*. Basic Books.
- Lamarra, T., Villani, C., & Bolognesi, M. M. (2026). Specificity effect in concrete/abstract semantic categorization task. *Cognitive Processing*, 27, 25–43.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2), 211–240. <https://doi.org/10.1037/0033-295X.104.2.211>
- Langacker, R. (2008). *Cognitive Grammar: A Basic Introduction* (1<sup>a</sup> ed.). Oxford University Press New York. <https://doi.org/10.1093/acprof:oso/9780195331967.001.0001>
- Langacker, R. W. (1986). An Introduction to Cognitive Grammar. *Cognitive Science*, 10(1), 1–40.  
[https://doi.org/10.1207/s15516709cog1001\\_1](https://doi.org/10.1207/s15516709cog1001_1)
- Langacker, R. W. (1987). *Foundations of cognitive grammar, Volume 1: Theoretical prerequisites*. Stanford University Press.
- Langacker, R. W. (1999). *Grammar and Conceptualization*. De Gruyter Mouton.  
<https://doi.org/10.1515/9783110800524>
- Langacker, R. W. (2013). *Essentials of Cognitive Grammar*. Oxford University Press.
- Langland-Hassan, P., & Davis, C. P. (2023). A context-sensitive and non-linguistic approach to abstract concepts. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1870), 20210355. <https://doi.org/10.1098/rstb.2021.0355>
- Langland-Hassan, P., Faries, F. R., Gatyas, M., Dietz, A., & Richardson, M. J. (2021). Assessing abstract thought and its relation to language with a new nonverbal paradigm: Evidence from aphasia. *Cognition*, 211, 104622. <https://doi.org/10.1016/j.cognition.2021.104622>
- Lebois, L. A. M., Wilson-Mendenhall, C. D., & Barsalou, L. W. (2015). Are Automatic Conceptual Cores the Gold Standard of Semantic Processing? The Context-Dependence of Spatial Meaning in Grounded Congruency Effects. *Cognitive Science*, 39(8), 1764–1801.  
<https://doi.org/10.1111/cogs.12174>
- Leech, G. (1981). *Semantics: The Study of Meaning* (Second edition). Penguin Books.

- Lenci, A. (2008). Distributional semantics in linguistic and cognitive research. *Italian Journal of Linguistics*, 20(1), 1–31.
- Lenci, A., Lebani, G. E., & Passaro, L. C. (2018). The Emotions of Abstract Words: A Distributional Semantic Analysis. *Topics in Cognitive Science*, 10(3), 550–572. <https://doi.org/10.1111/tops.12335>
- Lenci, A., & Sahlgren, M. (2023). *Distributional Semantics*. Cambridge University Press. <https://doi.org/10.1017/9780511783692>
- Lenth, R. (2021). *emmeans: Estimated marginal means, aka least-squares means* (Version R package version 1.6.0) [Software]. <https://CRAN.R-project.org/package=emmeans>
- Löhr, G. (2022). What Are Abstract Concepts? On Lexical Ambiguity and Concreteness Ratings. *Review of Philosophy and Psychology*, 13(3), 549–566. <https://doi.org/10.1007/s13164-021-00542-9>
- Löhr, G. (2024). Does the mind care about whether a word is abstract or concrete? Why concreteness is probably not a natural kind. *Mind & Language*, 39(5), 627–646. <https://doi.org/10.1111/mila.12473>
- Louviere, J. J., Flynn, T. N., & Marley, A. A. J. (2015). *Best-Worst Scaling: Theory, Methods and Applications* (1<sup>a</sup> ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9781107337855>
- Louwerse, M., & Connell, L. (2011). A Taste of Words: Linguistic Context and Perceptual Simulation Predict the Modality of Words. *Cognitive Science*, 35(2), 381–398. <https://doi.org/10.1111/j.1551-6709.2010.01157.x>
- Louwerse, M., & Jeuniaux, P. (2008). Language comprehension is both embodied and symbolic. In M. de Vega, A. Glenberg, & A. Graesser (Eds.), *Symbols and Embodiment: Debates on meaning and cognition*. Oxford University Press.
- Louwerse, M. M. (2007). Symbolic or embodied representations: A case for symbol interdependency. In T. K. Landauer, D. S. McNamara, S. Dennis, & Kintsch (Eds.), *Handbook of latent semantic analysis* (pp. 107–120). Lawrence Erlbaum Associates Publishers.
- Louwerse, M. M. (2008). Embodied relations are encoded in language. *Psychonomic Bulletin & Review*, 15(4), 838–844. <https://doi.org/10.3758/PBR.15.4.838>
- Louwerse, M. M. (2011). Symbol Interdependency in Symbolic and Embodied Cognition. *Topics in Cognitive Science*, 3(2), 273–302. <https://doi.org/10.1111/j.1756-8765.2010.01106.x>

- Louwerse, M. M. (2018). Knowing the Meaning of a Word by the Linguistic and Perceptual Company It Keeps. *Topics in Cognitive Science*, 10(3), 573–589. <https://doi.org/10.1111/tops.12349>
- Louwerse, M. M., & Jeuniaux, P. (2010). The linguistic and embodied nature of conceptual processing. *Cognition*, 114(1), 96–104. <https://doi.org/10.1016/j.cognition.2009.09.002>
- Lüdecke, D., Ben-Shachar, M., Patil, I., & Makowski, D. (2020). Extracting, Computing and Exploring the Parameters of Statistical Models using R. *Journal of Open Source Software*, 5(53), 2445. <https://doi.org/10.21105/joss.02445>
- Lupyan, G., & Mirman, D. (2013). Linking language and categorization: Evidence from aphasia. *Cortex*, 49(5), 1187–1194. <https://doi.org/10.1016/j.cortex.2012.06.006>
- Lupyan, G., & Thompson-Schill, S. L. (2012). The evocative power of words: Activation of concepts by verbal and nonverbal means. *Journal of Experimental Psychology: General*, 141(1), 170–186. <https://doi.org/10.1037/a0024904>
- Lupyan, G., & Winter, B. (2018). Language is more abstract than you think, or, why aren't languages more iconic? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1752), 20170137. <https://doi.org/10.1098/rstb.2017.0137>
- Lynott, D., Connell, L., Brysbaert, M., Brand, J., & Carney, J. (2020). The Lancaster Sensorimotor Norms: Multidimensional measures of perceptual and action strength for 40,000 English words. *Behavior Research Methods*, 52(3), 1271–1291. <https://doi.org/10.3758/s13428-019-01316-z>
- Lyons, J. (1977). *Semantics (Vol. 1)*. Cambridge University Press.
- Magni, E. (2024). The role of ambiguity and vagueness in language change. In I. Fiorentini & C. Zanchi (Eds.), *Vagueness, Ambiguity, and All the Rest* (Vol. 347, pp. 10–28). John Benjamins Publishing Company. <https://doi.org/10.1075/pbns.347.02mag>
- Mandera, P., Keuleers, E., & Brysbaert, M. (2020). Recognition times for 62 thousand English words: Data from the English Crowdsourcing Project. *Behavior Research Methods*, 52(2), 741–760. <https://doi.org/10.3758/s13428-019-01272-8>
- Marconi, D. (1997). *Lexical competence*. MIT Press.
- Markman, E. M. (1985). Why superordinate category terms can be mass nouns. *Cognition*, 19(1), 31–53. [https://doi.org/10.1016/0010-0277\(85\)90030-7](https://doi.org/10.1016/0010-0277(85)90030-7)
- Marson, F., Paoletti, P., Naor-Ziv, R., Carducci, F., & Ben-Soussan, T. D. (2023). Embodied empathy and abstract concepts' concreteness: Evidence from contemplative practices. In *Progress in Brain Research* (Vol. 277, pp. 181–209). Elsevier. <https://doi.org/10.1016/bs.pbr.2022.12.005>

- Mazzuca, C., Falcinelli, I., Michalland, A.-H., Tummolini, L., & Borghi, A. M. (2022). Bodily, emotional, and public sphere at the time of COVID-19. An investigation on concrete and abstract concepts. *Psychological Research*, 86(7), 2266–2277. <https://doi.org/10.1007/s00426-021-01633-z>
- Mazzuca, C., Majid, A., Lugli, L., Nicoletti, R., & Borghi, A. M. (2020). Gender is a multifaceted concept: Evidence that specific life experiences differentially shape the concept of gender. *Language and Cognition*, 12(4), 649–678. <https://doi.org/10.1017/langcog.2020.15>
- Mazzuca, C., & Santarelli, M. (2023). Making it abstract, making it contestable: Politicization at the intersection of political and cognitive science. *Review of Philosophy and Psychology*, 14(4), 1257–1278. <https://doi.org/10.1007/s13164-022-00640-2>
- Mazzuca, C., Villani, C., Lamarra, T., Bolognesi, M. M., & Borghi, A. M. (2025). Abstractness impacts conversational dynamics. *Cognition*, 258, 106084. <https://doi.org/10.1016/j.cognition.2025.106084>
- McClelland, J. L., & Rumelhart, D. E. (1981). McClelland, J. L., & Rumelhart, D. E. (1981). An Interactive Activation Model of Context Effects in Letter Perception: Part I. An Account of Basic Findings. *Psychological Review*, 88(5), 375–407.
- McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, 37(4), 547–559. <https://doi.org/10.3758/BF03192726>
- Medin, D. L., Dewey, G. I., & Murphy, T. D. (1983). Relationships between item and category learning: Evidence that abstraction is not automatic. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9, 607–625.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85(3), 207–238. <https://doi.org/10.1037/0033-295X.85.3.207>
- Mervis, C. B. (1983). Acquisition of a lexicon. *Contemporary Educational Psychology*, 8(3), 210–236. [https://doi.org/10.1016/0361-476X\(83\)90015-2](https://doi.org/10.1016/0361-476X(83)90015-2)
- Mervis, C. B., & Crisafi, M. A. (1982). Order of acquisition of subordinate-, basic-, and superordinate-level categories. *Child Development*, 53(1), 258–266. <https://doi.org/10.2307/1129660>
- Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90(2), 227–234. <https://doi.org/10.1037/h0031564>

- Miller, G. A. (1995). WordNet: A lexical database for English. *Communications of the ACM*, 38(11), 39–41. <https://doi.org/10.1145/219717.219748>
- Miller, G. A., & Johnson-Laird, P. N. (1976). *Language and perception*. Harvard University Press.
- Montefinese, M., Ambrosini, E., Fairfield, B., & Mammarella, N. (2014). The adaptation of the Affective Norms for English Words (ANEW) for Italian. *Behavior Research Methods*, 46(3), 887–903. <https://doi.org/10.3758/s13428-013-0405-3>
- Montefinese, M., Gregori, L., Ravelli, A. A., Varvara, R., & Radicioni, D. P. (2023). CONcreTEXT norms: Concreteness ratings for Italian and English words in context. *PLOS ONE*, 18(10), e0293031. <https://doi.org/10.1371/journal.pone.0293031>
- Montefinese, M., Vinson, D., Vigliocco, G., & Ambrosini, E. (2019). Italian Age of Acquisition Norms for a Large Set of Words (ItAoA). *Frontiers in Psychology*, 10. <https://doi.org/10.3389/fpsyg.2019.00278>
- Mukaka, M. M. (2012). Statistics corner: A guide to appropriate use of correlation coefficient in medical research. *Malawi Medical Journal: The Journal of Medical Association of Malawi*, 24(3), 69–71.
- Muraki, E. J., Abdalla, S., Brysbaert, M., & Pexman, P. M. (2022). Concreteness ratings for 62,000 English multiword expressions. *Behavior Research Methods*, 55(5), 2522–2531. <https://doi.org/10.3758/s13428-022-01912-6>
- Murgiano, M., Motamedi, Y., & Vigliocco, G. (2021). Situating Language in the Real-World: The Role of Multimodal Iconicity and Indexicality. *Journal of Cognition*, 4(1), 38. <https://doi.org/10.5334/joc.113>
- Murphy, G. L. (2002). *The Big Book of Concepts*. The MIT Press. <https://doi.org/10.7551/mitpress/1602.001.0001>
- Murphy, G. L. (2016). Is there an exemplar theory of concepts? *Psychonomic Bulletin & Review*, 23(4), 1035–1042. <https://doi.org/10.3758/s13423-015-0834-3>
- Murphy, G. L., & Brownell, H. H. (1985). Category differentiation in object recognition: Typicality constraints on the basic category advantage. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11(1), 70–84. <https://doi.org/10.1037/0278-7393.11.1.70>
- Murphy, G. L., & Smith, E. E. (1982). Basic-level superiority in picture categorization. *Journal of Verbal Learning and Verbal Behavior*, 21(1), 1–20. [https://doi.org/10.1016/S0022-5371\(82\)90412-1](https://doi.org/10.1016/S0022-5371(82)90412-1)

- Murphy, G. L., & Wisniewski, E. J. (1989). Categorizing objects in isolation and in scenes: What a superordinate is good for. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(4), 572–586. <https://doi.org/10.1037/0278-7393.15.4.572>
- Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology*, 9, 353–383.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Prentice-Hall.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, 115(1), 39–57. <https://doi.org/10.1037/0096-3445.115.1.39>
- Nosofsky, R. M. (1988). Similarity, frequency, and category representations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(1), 54–65. <https://doi.org/10.1037/0278-7393.14.1.54>
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104(2), 266–300. <https://doi.org/10.1037/0033-295X.104.2.266>
- Noveck, I. (2018). *Experimental Pragmatics: The Making of a Cognitive Science*. Cambridge University Press. <https://doi.org/10.1017/9781316027073>
- Paivio, A. (1971). *Imagery and Verbal Processes*. Holt, Rinehart and Winston.
- Paivio, A. (1990). *Mental representations: A dual coding approach*. Oxford University Press.
- Paivio, A. (1991). Dual coding theory: Retrospect and current status. *Canadian Journal of Psychology / Revue canadienne de psychologie*, 45(3), 255–287. <https://doi.org/10.1037/h0084295>
- Palmeri, T. J., & Nosofsky, R. M. (2001). Central Tendencies, Extreme Points, and Prototype Enhancement Effects in Ill-Defined Perceptual Categorization. *The Quarterly Journal of Experimental Psychology Section A*, 54(1), 197–235. <https://doi.org/10.1080/02724980042000084>
- Papagno, C., Fogliata, A., Catricalà, E., & Miniussi, C. (2009). The lexical processing of abstract and concrete nouns. *Brain Research*, 1263, 78–86. <https://doi.org/10.1016/j.brainres.2009.01.037>
- Peirce, C. S. (1931). *The Collected Papers of Charles Sanders Peirce* (Voll. 1–8). Harvard University Press.
- Perry, L. K., Custode, S. A., Fasano, R. M., Gonzalez, B. M., & Savy, J. D. (2021). What Is the Buzz About Iconicity? How Iconicity in Caregiver Speech Supports Children’s Word Learning. *Cognitive Science*, 45(4), e12976. <https://doi.org/10.1111/cogs.12976>

- Pexman, P. (2012). Meaning-based influences on visual word recognition. *Visual Word Recognition Volume 2: Meaning and Context, Individuals and Development, Vol. 2*, 24–43. <https://doi.org/10.4324/9780203106976>
- Pexman, P. M. (2020). How does meaning come to mind? Four broad principles of semantic processing. *Canadian Journal of Experimental Psychology / Revue canadienne de psychologie expérimentale*, 74(4), 275–283. <https://doi.org/10.1037/cep0000235>
- Pexman, P. M., Diveica, V., & Binney, R. J. (2023). Social semantics: The organization and grounding of abstract concepts. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1870). <https://doi.org/10.1098/rstb.2021.0363>
- Pexman, P. M., Heard, A., Lloyd, E., & Yap, M. J. (2017). The Calgary semantic decision project: Concrete/abstract decision data for 10,000 English words. *Behavior Research Methods*, 49(2), 407–417. <https://doi.org/10.3758/s13428-016-0720-6>
- Pexman, P. M., Muraki, E., Sidhu, D. M., Siakaluk, P. D., & Yap, M. J. (2019). Quantifying sensorimotor experience: Body–object interaction ratings for more than 9,000 English words. *Behavior Research Methods*, 51(2), 453–466. <https://doi.org/10.3758/s13428-018-1171-z>
- Piaget, J. (1936). *Origins of intelligence in the child*. Routledge & Kegan Paul.
- Piaget, J. (1970). *Genetic Epistemology* (E. Duckworth, Trad.). Columbia University Press.
- Piaget, J. (1997). *The Principles of Genetic Epistemology*. Routledge. <https://books.google.it/books?id=rr-avb4T8ksC>
- Piantadosi, S. T., Tily, H., & Gibson, E. (2011). Reply to Reilly and Kean: Clarifications on word length and information content. *In Proceedings of the National Academy of Sciences*, 108(20). <https://doi.org/10.1073/pnas.1103550108>
- Pickering, M. J., & Garrod, S. (2004). The interactive-alignment model: Developments and refinements. *The Behavioral and Brain Sciences*, 27(2), 212–225. <https://doi.org/10.1017/S0140525X04450055>
- Pickering, M. J., & Garrod, S. (2021). *Understanding dialogue: Language use and social interaction* (pp. xvii, 283). Cambridge University Press. <https://doi.org/10.1017/9781108610728>
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77(3, Pt.1), 353–363. <https://doi.org/10.1037/h0025953>
- Pulvermüller, F., Moseley, R. L., Egorova, N., Shebani, Z., & Boulenger, V. (2014). Motor cognition–motor semantics: Action perception theory of cognition and communication.

- Neuropsychologia, Special Issue in Honor of Marc Jeannerod*, 55, 71–84. <https://doi.org/10.1016/j.neuropsychologia.2013.12.002>
- Pylyshyn, Z. W. (1973). What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin*, 80(1), 1–24. <https://doi.org/10.1037/h0034650>
- R Core Team. (2023). *R: A Language and Environment for Statistical Computing* [R Foundation for Statistical Computing]. <https://www.R-project.org/>
- Rambelli, G., & Bolognesi, M. (2024, maggio). The Contextual Variability of English Nouns: The Impact of Categorical Specificity beyond Conceptual Concreteness. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING).
- Rambelli, G., & Bolognesi, M. M. (2023). Contextual Variability depends on Categorical Specificity rather than Conceptual Concreteness: A Distributional Investigation on Italian data. In *Proceedings of the 15th International Conference on Computational Semantics*, 16–27.
- Ravelli, A. A., & Bolognesi, M. M. (2024). Yet another approximation of human semantic judgments using LLMs... But with quantized local models on novel data. *IJCoL. Italian Journal of Computational Linguistics*, 10(10–2).
- Ravelli, A. A., Bolognesi, M. M., & Caselli, T. (2025). Specificity ratings for English data. *Cognitive Processing*, 26(2), 283–302. <https://doi.org/10.1007/s10339-024-01239-4>
- Reed, S. K. (2016). A Taxonomic Analysis of Abstraction. *Perspectives on Psychological Science*, 11(6), 817–837. <https://doi.org/10.1177/1745691616646304>
- Reijnierse, W. G., Burgers, C., Bolognesi, M., & Krennmayr, T. (2019). How Polysemy Affects Concreteness Ratings: The Case of Metaphor. *Cognitive Science*, 43(8), e12779. <https://doi.org/10.1111/cogs.12779>
- Reilly, J., Hung, J., & Westbury, C. (2017). Non-Arbitrariness in Mapping Word Form to Meaning: Cross-Linguistic Formal Markers of Word Concreteness. *Cognitive Science*, 41(4), 1071–1089. <https://doi.org/10.1111/cogs.12361>
- Reilly, J., Shain, C., Borghesani, V., Kuhnke, P., Vigliocco, G., Peelle, J. E., Mahon, B. Z., Buxbaum, L. J., Majid, A., Brysbaert, M., Borghi, A. M., De Deyne, S., Dove, G., Papeo, L., Pexman, P. M., Poeppel, D., Lupyan, G., Boggio, P., Hickok, G., ... Vinson, D. (2025). What we mean when we say semantic: Toward a multidisciplinary semantic glossary. *Psychonomic Bulletin & Review*, 32(1), 243–280. <https://doi.org/10.3758/s13423-024-02556-7>

- Richardson, J. T. E. (1976). Imageability and concreteness. *Bulletin of the Psychonomic Society*, 7(5), 429–431. <https://doi.org/10.3758/BF03337237>
- Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 12(1), 1–20. [https://doi.org/10.1016/S0022-5371\(73\)80056-8](https://doi.org/10.1016/S0022-5371(73)80056-8)
- Robbins, P., & Drayson, Z. (2025). Modularity of Mind. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2025 Edition). <https://plato.stanford.edu/entries/modularity-mind/>
- Romani, C., McAlpine, S., & Martin, R. C. (2008). Concreteness Effects in Different Tasks: Implications for Models of Short-Term Memory. *Quarterly Journal of Experimental Psychology*, 61(2), 292–323. <https://doi.org/10.1080/17470210601147747>
- Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104(3), 192–233. <https://doi.org/10.1037/0096-3445.104.3.192>
- Rosch, E. (1978). Principles of Categorisation. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and Categorisation* (pp. 27–48). Erlbaum.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7(4), 573–605. [https://doi.org/10.1016/0010-0285\(75\)90024-9](https://doi.org/10.1016/0010-0285(75)90024-9)
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8(3), 382–439. [https://doi.org/10.1016/0010-0285\(76\)90013-x](https://doi.org/10.1016/0010-0285(76)90013-x)
- Rumelhart, D. E., & Ortony, A. (1977). The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), *Schooling and the Acquisition of Knowledge*. Erlbaum.
- Schilhab, T. S. S. (2015). Words as cultivators of others minds. *Frontiers in Psychology*, 6. <https://doi.org/10.3389/fpsyg.2015.01690>
- Schulte Im Walde, S., & Frassinelli, D. (2022). Distributional Measures of Semantic Abstraction. *Frontiers in Artificial Intelligence*, 4, 796756. <https://doi.org/10.3389/frai.2021.796756>
- Schwanenflugel, P. J. (2013). Why are Abstract Concepts Hard to Understand? In Schwanenflugel, Paula J. (Eds.), *The Psychology of Word Meanings* (pp. 223–250). Psychology Press. <https://doi.org/10.4324/9780203772119-10>

- Schwanenflugel, P. J., Akin, C., & Luh, W.-M. (1992). Context availability and the recall of abstract and concrete words. *Memory & Cognition*, 20(1), 96–104. <https://doi.org/10.3758/bf03208259>
- Schwanenflugel, P. J., Harnishfeger, K. K., & Stowe, R. W. (1988). Context availability and lexical decisions for abstract and concrete words. *Journal of Memory and Language*, 27(5), 499–520. [https://doi.org/10.1016/0749-596x\(88\)90022-8](https://doi.org/10.1016/0749-596x(88)90022-8)
- Schwanenflugel, P. J., & Shoben, E. J. (1983). Differential context effects in the comprehension of abstract and concrete verbal materials. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(1), 82–102. <https://doi.org/10.1037/0278-7393.9.1.82>
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(3), 417–424. <https://doi.org/10.1017/S0140525X00005756>
- Semin, G. R., & Fiedler, K. (1988). The cognitive functions of linguistic categories in describing persons: Social cognition and language. *Journal of Personality and Social Psychology*, 54, 558–568.
- Shapiro, L. (2019). *Embodied Cognition* (2<sup>a</sup> ed.). Routledge. <https://doi.org/10.4324/9781315180380>
- Shea, N. (2018). Metacognition and abstract concepts. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1752), 20170133. <https://doi.org/10.1098/rstb.2017.0133>
- Sloman, S. A. (1993). Feature-Based Induction. *Cognitive Psychology*, 25(2), 231–280. <https://doi.org/10.1006/cogp.1993.1006>
- Sloman, S. A. (1998). Categorical Inference Is Not a Tree: The Myth of Inheritance Hierarchies. *Cognitive Psychology*, 35(1), 1–33. <https://doi.org/10.1006/cogp.1997.0672>
- Sloutsky, V. M., & Deng, W. (Sophia). (2019). Categories, concepts, and conceptual development. *Language, Cognition and Neuroscience*, 34(10), 1284–1297. <https://doi.org/10.1080/23273798.2017.1391398>
- Smith, E. E., & Medin, D. L. (1981). *Categories and Concepts*. Harvard University Press.
- Son, J. Y., Smith, L. B., & Goldstone, R. L. (2008). Simplicity and generalization: Short-cutting abstraction in children's object categorizations. *Cognition*, 108(3), 626–638. <https://doi.org/10.1016/j.cognition.2008.05.002>
- Talmy, L. (2000). *Toward a Cognitive Semantics* (Vols. 1–2). MIT Press.
- Tanaka, J. W., & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology*, 23(3), 457–482. [https://doi.org/10.1016/0010-0285\(91\)90016-H](https://doi.org/10.1016/0010-0285(91)90016-H)

- Tillotson, S. M., Siakaluk, P. D., & Pexman, P. M. (2008). Body—Object interaction ratings for 1,618 monosyllabic nouns. *Behavior Research Methods*, 40(4), 1075–1078. <https://doi.org/10.3758/brm.40.4.1075>
- Tomasello, M., & Akhtar, N. (1995). Two-year-olds use pragmatic cues to differentiate reference to objects and actions. *Cognitive Development*, 10(2), 201–224. [https://doi.org/10.1016/0885-2014\(95\)90009-8](https://doi.org/10.1016/0885-2014(95)90009-8)
- Trope, Y., & Liberman, N. (2010). Construal-Level Theory of Psychological Distance. *Psychological review*, 117(2), 440–463. <https://doi.org/10.1037/a0018963>
- Tucker, B. V., Brenner, D., Danielson, D. K., Kelley, M. C., Nenadić, F., & Sims, M. (2019). The Massive Auditory Lexical Decision (MALD) database. *Behavior Research Methods*, 51(3), 1187–1204. <https://doi.org/10.3758/s13428-018-1056-1>
- Tuggy, D. (2010). Schematicity. In D. Geeraerts & H. Cuyckens (Eds.), *The Oxford Handbook of Cognitive Linguistics*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199738632.013.0004>
- Turing, A. M. (1937). On Computable Numbers, with an Application to the Entscheidungsproblem. In *Proceedings of the London Mathematical Society*, s2-42(1), 230–265. <https://doi.org/10.1112/plms/s2-42.1.230>
- Van den Bos, G. R. (2007). *APA dictionary of psychology* (American Psychological Association, Eds.; 1<sup>a</sup> ed.). American Psychological Association.
- Van Hoef, R., Connell, L., & Lynott, D. (2023). The effects of sensorimotor and linguistic information on the basic-level advantage. *Cognition*, 241, 105606. <https://doi.org/10.1016/j.cognition.2023.105606>
- Van Oers, B. (2001). Contextualisation for Abstraction. *Cognitive Science Quarterly*, 1(3), 279–305.
- Varela, F. J., Rosch, E., & Thompson, E. (1992). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.
- Vergallito, A., Petilli, M. A., & Marelli, M. (2020). Perceptual modality norms for 1,121 Italian words: A comparison with concreteness and imageability scores and an analysis of their impact in word processing tasks. *Behavior Research Methods*, 52(4), 1599–1616. <https://doi.org/10.3758/s13428-019-01337-8>
- Vigliocco, G., Convertino, L., De Felice, S., Gregorians, L., Kewenig, V., Mueller, M. A. E., Veselic, S., Musolesi, M., Hudson-Smith, A., Tyler, N., Flouri, E., & Spiers, H. J. (2024). Ecological

- brain: Reframing the study of human behaviour and cognition. *Royal Society Open Science*, 11(11), 240762. <https://doi.org/10.1098/rsos.240762>
- Vigliocco, G., Kousta, S., Vinson, D., Andrews, M., & Del Campo, E. (2013). The representation of abstract words: What matters? Reply to Paivio's (2013) comment on Kousta et al. (2011). *Journal of Experimental Psychology: General*, 142(1), 288–291. <https://doi.org/10.1037/a0028749>
- Vigliocco, G., Kousta, S.-T., Della Rosa, P. A., Vinson, D. P., Tettamanti, M., Devlin, J. T., & Cappa, S. F. (2014). The Neural Representation of Abstract Words: The Role of Emotion. *Cerebral Cortex*, 24(7), 1767–1777. <https://doi.org/10.1093/cercor/bht025>
- Vigliocco, G., Meteyard, L., Andrews, M., & Kousta, S. (2009). Toward a theory of semantic representation. *Language and Cognition*, 1(2), 219–247. <https://doi.org/10.1515/LANGCOG.2009.011>
- Vigliocco, G., Motamedi, Y., Murgiano, M., Wonnacott, E., Marshall, C., Maillo, I. M., & Perniss, P. (2019). Onomatopoeia, gestures, actions and words: How do caregivers use multimodal cues in their communication to children? In *Proceedings of the Annual Meeting of the Cognitive Science Society (Vol. 41)*.
- Vigliocco, G., Perniss, P., & Vinson, D. (2014). Language as a multimodal phenomenon: Implications for language learning, processing and evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1651), 20130292. <https://doi.org/10.1098/rstb.2013.0292>
- Villani, C., & Bolognesi, M. M. (2025). Specificity: Metrics and Norms. In *Reference Module in Social Sciences*. Elsevier. <https://doi.org/https://doi.org/10.1016/B978-0-323-95504-1.00894-2>
- Villani, C., Loia, A., & Bolognesi, M. M. (2024). The semantic content of concrete, abstract, specific, and generic concepts. *Language and Cognition*, 16(4), 867–894. <https://doi.org/10.1017/langcog.2023.64>
- Villani, C., Lugli, L., Liuzza, M. T., & Borghi, A. M. (2019). Varieties of abstract concepts and their multiple dimensions. *Language and Cognition*, 11(3), 403–430. <https://doi.org/10.1017/langcog.2019.23>
- Villani, C., Lugli, L., Liuzza, M. T., Nicoletti, R., & Borghi, A. M. (2021). Sensorimotor and interoceptive dimensions in concrete and abstract concepts. *Journal of Memory and Language*, 116, 104173. <https://doi.org/10.1016/j.jml.2020.104173>

- Villani, C., Orsoni, M., Lugli, L., Benassi, M., & Borghi, A. M. (2022). Abstract and concrete concepts in conversation. *Scientific Reports*, 12(1), 17572. <https://doi.org/10.1038/s41598-022-20785-5>
- Vygotsky, L. S. (1934). *Thought and Language* (A. Kozulin, Trad.; ed. rev.). The MIT Press.
- Wang, S., Hurlstone, M. J., Leviston, Z., Walker, I., & Lawrence, C. (2021). Construal-level theory and psychological distancing: Implications for grand environmental challenges. *One Earth*, 4(4), 482–486. <https://doi.org/10.1016/j.oneear.2021.03.009>
- Wauters, L. N., Tellings, A. E. J. M., Bon, W. H. J. V., & Haaften, A. W. V. (2003). Mode of acquisition of word meanings: The viability of a theoretical construct. *Applied Psycholinguistics*, 24(3), 385–406. <https://doi.org/10.1017/S0142716403000201>
- Werner, H. (1957). *The comparative psychology of mental development* (Revised). International University Press.
- Whitehead, A. N. (1978). *Process and Reality* (D. R. Griffin & D. W. Sherburne, Eds.; Corrected Edition). The Free Press.
- Wiemer-Hastings, K., Krug, J., & Xu, X. (2001). Imagery, Context Availability, Contextual Constraint and Abstractness. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 23.
- Wierzbicka, A. (1984). apples are not a “kind of fruit”: The semantics of human categorization. *American Ethnologist*, 11(2), 313–328. <https://doi.org/10.1525/ae.1984.11.2.02a00060>
- Wijffels, J. (2023). *udpipe: Tokenization, Parts of Speech Tagging, Lemmatization and Dependency Parsing with the «UDPipe» «NLP» Toolkit* (Version R package version 0.8.11) [Software]. <https://CRAN.R-project.org/package=udpipe>
- Wingfield, C., Van Hoef, R., & Connell, L. (2024). A Linguistic–Sensorimotor Model of the Basic-Level Advantage in Category Verification. *Cognitive Science*, 48(12), e70025. <https://doi.org/10.1111/cogs.70025>
- Winter, B. (2022). Abstract concepts and emotion: Cross-linguistic evidence and arguments against affective embodiment. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1870), 20210368. <https://doi.org/10.1098/rstb.2021.0368>
- Winter, B., Lupyan, G., Perry, L. K., Dingemanse, M., & Perlman, M. (2024). Iconicity ratings for 14,000+ English words. *Behavior Research Methods*, 56(3), 1640–1655. <https://doi.org/10.3758/s13428-023-02112-6>

- Winter, B., Perlman, M., Perry, L. K., & Lupyan, G. (2017). Which words are most iconic?: Iconicity in English sensory words. *Interaction Studies*, 18(3), 443–464. <https://doi.org/10.1075/is.18.3.07win>
- Wisniewski, E. J., & Murphy, G. L. (1989). Superordinate and basic category names in discourse: A textual analysis. *Discourse Processes*, 12(2), 245–261. <https://doi.org/10.1080/01638538909544728>
- Wittgenstein, L. (1958). *Philosophical Investigations* (G. E. M. Anscombe, Trad.). Basil Blackwell.
- Wolff, P., & Gentner, D. (2011). Structure-Mapping in Metaphor Comprehension. *Cognitive Science*, 35, 1456–1488. <https://doi.org/https://doi.org/10.1111/j.1551-6709.2011.01194.x>
- Yap, M. J., Lim, G. Y., & Pexman, P. M. (2015). Semantic richness effects in lexical decision: The role of feedback. *Memory & Cognition*, 43(8), 1148–1167. <https://doi.org/10.3758/s13421-015-0536-0>
- Yap, M. J., Tan, S. E., Pexman, P. M., & Hargreaves, I. S. (2011). Is more always better? Effects of semantic richness on lexical decision, speeded pronunciation, and semantic classification. *Psychonomic Bulletin & Review*, 18(4), 742–750. <https://doi.org/10.3758/s13423-011-0092-y>
- Yee, E., & Thompson-Schill, S. L. (2016). Putting concepts into context. *Psychonomic Bulletin & Review*, 23(4), 1015–1027. <https://doi.org/10.3758/s13423-015-0948-7>
- Yeh, W., & Barsalou, L. W. (2006). The Situated Nature of Concepts. *The American Journal of Psychology*, 119(3), 349–384. <https://doi.org/10.2307/20445349>
- Zeman, J. J. (1982). Peirce on abstraction. *The Monist*, 65(2), 211–229.
- Zhang, Q. (1998). Fuzziness—Vagueness—Generality—Ambiguity. *Journal of Pragmatics*, 29(1), 13–31. [https://doi.org/10.1016/S0378-2166\(97\)00014-3](https://doi.org/10.1016/S0378-2166(97)00014-3)
- Zuur, A. F., Ieno, E. N., & Elphick, C. S. (2010). A protocol for data exploration to avoid common statistical problems. *Methods in Ecology and Evolution*, 1(1), 3–14. <https://doi.org/10.1111/j.2041-210X.2009.00001.x>
- Zwaan, R. A. (2014). Embodiment and language comprehension: Reframing the discussion. *Trends in Cognitive Sciences*, 18(5), 229–234. <https://doi.org/10.1016/j.tics.2014.02.008>
- Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2), 162–185. <https://doi.org/10.1037/0033-2909.123.2.162>

## APPENDIX A

The tables below list the full set of 120 Italian words and 120 nonwords used as experimental stimuli in the behavioral studies discussed in Chapter 4 (Ratings, Lexical Decision, and Semantic Decision Task).

| N_words | Words             | English_version      | Subcategory | Type             |
|---------|-------------------|----------------------|-------------|------------------|
| 1       | Frutto            | Fruit                | FOOD        | Concrete-General |
| 2       | Ortaggio          | Vegatable            | FOOD        | Concrete-General |
| 3       | Bacca             | Berry                | FOOD        | Concrete-General |
| 4       | Cibo              | Food                 | FOOD        | Concrete-General |
| 5       | Organismo         | Organism             | FOOD        | Concrete-General |
| 6       | Prodotto agricolo | Agricultural product | FOOD        | Concrete-General |
| 7       | Verdura           | Vegetable            | FOOD        | Concrete-General |
| 8       | Pianta            | Plant                | FOOD        | Concrete-General |
| 9       | Leccornia         | Delicacy             | FOOD        | Concrete-General |
| 10      | Alimento          | Food                 | FOOD        | Concrete-General |
| 11      | Arredamento       | Furniture            | TOOL        | Concrete-General |
| 12      | Attrezzo          | Tool                 | TOOL        | Concrete-General |
| 13      | Arnese            | Instrument           | TOOL        | Concrete-General |
| 14      | Recipiente        | Container            | TOOL        | Concrete-General |
| 15      | arma da taglio    | Bladed weapon        | TOOL        | Concrete-General |
| 16      | Oggetto           | Object               | TOOL        | Concrete-General |
| 17      | Copertura         | Cover                | TOOL        | Concrete-General |
| 18      | Posata            | Cutlery              | TOOL        | Concrete-General |
| 19      | Utensile          | Utensil              | TOOL        | Concrete-General |
| 20      | Artefatto         | Artifact             | TOOL        | Concrete-General |
| 21      | Carnivoro         | Carnivore            | ANIMAL      | Concrete-General |
| 22      | Felino            | Feline               | ANIMAL      | Concrete-General |
| 23      | Suino             | Pork                 | ANIMAL      | Concrete-General |
| 24      | Mammifero         | Mammal               | ANIMAL      | Concrete-General |
| 25      | Essere vivente    | Living being         | ANIMAL      | Concrete-General |
| 26      | Bovino            | Cattle               | ANIMAL      | Concrete-General |
| 27      | Uccello           | Bird                 | ANIMAL      | Concrete-General |
| 28      | Pennuto           | Feathered animal     | ANIMAL      | Concrete-General |
| 29      | Volatile          | Flying animal        | ANIMAL      | Concrete-General |
| 30      | Insetto           | Insect               | ANIMAL      | Concrete-General |

|    |                 |                 |        |                   |
|----|-----------------|-----------------|--------|-------------------|
| 31 | Platano         | Plantain        | FOOD   | Concrete-Specific |
| 32 | Carota          | Carrot          | FOOD   | Concrete-Specific |
| 33 | Uvetta          | Raisins         | FOOD   | Concrete-Specific |
| 34 | Fragola         | Strawberry      | FOOD   | Concrete-Specific |
| 35 | Tartufo         | Truffle         | FOOD   | Concrete-Specific |
| 36 | Melanzana       | Eggplant        | FOOD   | Concrete-Specific |
| 37 | Peperoncino     | Chili pepper    | FOOD   | Concrete-Specific |
| 38 | Pomodoro        | Tomato          | FOOD   | Concrete-Specific |
| 39 | Crostata        | Tart            | FOOD   | Concrete-Specific |
| 40 | Zucca           | Pumpkin         | FOOD   | Concrete-Specific |
| 41 | Lampada         | Lamp            | TOOL   | Concrete-Specific |
| 42 | Martello        | Hammer          | TOOL   | Concrete-Specific |
| 43 | Scopa           | Broom           | TOOL   | Concrete-Specific |
| 44 | Bottiglia       | Bottle          | TOOL   | Concrete-Specific |
| 45 | Coltello        | Knife           | TOOL   | Concrete-Specific |
| 46 | Trapano         | Drill           | TOOL   | Concrete-Specific |
| 47 | Ombrellone      | Beach umbrella  | TOOL   | Concrete-Specific |
| 48 | Forchetta       | Fork            | TOOL   | Concrete-Specific |
| 49 | Pastello        | Crayon          | TOOL   | Concrete-Specific |
| 50 | Pennello        | Brush           | TOOL   | Concrete-Specific |
| 51 | Pastore tedesco | German Shepherd | ANIMAL | Concrete-Specific |
| 52 | Leone           | Lion            | ANIMAL | Concrete-Specific |
| 53 | Maiale          | Pig             | ANIMAL | Concrete-Specific |
| 54 | Dromedario      | Dromedary       | ANIMAL | Concrete-Specific |
| 55 | Agnello         | Lamb            | ANIMAL | Concrete-Specific |
| 56 | Mucca           | Cow             | ANIMAL | Concrete-Specific |
| 57 | Colomba         | Dove            | ANIMAL | Concrete-Specific |
| 58 | Gallina         | Hen             | ANIMAL | Concrete-Specific |
| 59 | Pappagallo      | Parrot          | ANIMAL | Concrete-Specific |
| 60 | Farfalla        | Butterfly       | ANIMAL | Concrete-Specific |
| 61 | Alterazione     | Alteration      | PSTQ   | Abstract-General  |
| 62 | Avvenimento     | Event           | PSTQ   | Abstract-General  |
| 63 | Assetto         | Arrangement     | PSTQ   | Abstract-General  |
| 64 | Grandezza       | Size            | PSTQ   | Abstract-General  |
| 65 | Astrazione      | Abstraction     | PSTQ   | Abstract-General  |
| 66 | Esito           | Result          | PSTQ   | Abstract-General  |
| 67 | Penalità        | Penalty         | PSTQ   | Abstract-General  |
| 68 | Rimedio         | Remedy          | PSTQ   | Abstract-General  |
| 69 | Fatica          | Fatigue         | PSTQ   | Abstract-General  |
| 70 | Valuta          | Currency        | PSTQ   | Abstract-General  |
| 71 | Ideale          | Ideal           | PS     | Abstract-General  |

|     |              |              |      |                   |
|-----|--------------|--------------|------|-------------------|
| 72  | Asserzione   | Assertion    | PS   | Abstract-General  |
| 73  | Stimolo      | Stimulus     | PS   | Abstract-General  |
| 74  | Attributo    | Attribute    | PS   | Abstract-General  |
| 75  | Causalità    | Causality    | PS   | Abstract-General  |
| 76  | Aldilà       | Afterlife    | PS   | Abstract-General  |
| 77  | Enigma       | Enigma       | PS   | Abstract-General  |
| 78  | Misfatto     | Misdeed      | PS   | Abstract-General  |
| 79  | Convinzione  | Belief       | PS   | Abstract-General  |
| 80  | Buonsenso    | Good sense   | PS   | Abstract-General  |
| 81  | Indole       | Nature       | EMS  | Abstract-General  |
| 82  | Allegria     | Cheerfulness | EMS  | Abstract-General  |
| 83  | Legame       | Bond         | EMS  | Abstract-General  |
| 84  | Contrarietà  | Setback      | EMS  | Abstract-General  |
| 85  | Gentilezza   | Kindness     | EMS  | Abstract-General  |
| 86  | Ritorsione   | Retaliation  | EMS  | Abstract-General  |
| 87  | Malessere    | Discomfort   | EMS  | Abstract-General  |
| 88  | Turbamento   | Turmoil      | EMS  | Abstract-General  |
| 89  | Affinità     | Affinity     | EMS  | Abstract-General  |
| 90  | Sentimento   | Feeling      | EMS  | Abstract-General  |
| 91  | Ripresa      | Pickup       | PSTQ | Abstract-Specific |
| 92  | Genesi       | Genesis      | PSTQ | Abstract-Specific |
| 93  | Diagramma    | Diagram      | PSTQ | Abstract-Specific |
| 94  | Sezione      | Section      | PSTQ | Abstract-Specific |
| 95  | Costante     | Constant     | PSTQ | Abstract-Specific |
| 96  | Paga         | Pay          | PSTQ | Abstract-Specific |
| 97  | Castigo      | Punishment   | PSTQ | Abstract-Specific |
| 98  | Cura         | Treatment    | PSTQ | Abstract-Specific |
| 99  | Sacrificio   | Sacrifice    | PSTQ | Abstract-Specific |
| 100 | Dollaro      | Dollar       | PSTQ | Abstract-Specific |
| 101 | Virtù        | Virtue       | PS   | Abstract-Specific |
| 102 | Epiteto      | Epithet      | PS   | Abstract-Specific |
| 103 | Provocazione | Provocation  | PS   | Abstract-Specific |
| 104 | Redenzione   | Redemption   | PS   | Abstract-Specific |
| 105 | Provvidenza  | Providence   | PS   | Abstract-Specific |
| 106 | Elisio       | Elysium      | PS   | Abstract-Specific |
| 107 | Indovinello  | Riddle       | PS   | Abstract-Specific |
| 108 | Idolatria    | Idolatry     | PS   | Abstract-Specific |
| 109 | Giudizio     | Judgment     | PS   | Abstract-Specific |
| 110 | Dialettica   | Dialectic    | PS   | Abstract-Specific |
| 111 | Pacatezza    | Calmness     | EMS  | Abstract-Specific |
| 112 | Euforia      | Euphoria     | EMS  | Abstract-Specific |

|     |              |             |     |                   |
|-----|--------------|-------------|-----|-------------------|
| 113 | Sodalizio    | Fellowship  | EMS | Abstract-Specific |
| 114 | Rivolta      | Revolt      | EMS | Abstract-Specific |
| 115 | Beneficenza  | Charity     | EMS | Abstract-Specific |
| 116 | Ripicca      | Spite       | EMS | Abstract-Specific |
| 117 | Angoscia     | Anguish     | EMS | Abstract-Specific |
| 118 | Umiliazione  | Humiliation | EMS | Abstract-Specific |
| 119 | Predilezione | Preference  | EMS | Abstract-Specific |
| 120 | Fobia        | Phobia      | EMS | Abstract-Specific |

**Table A.1.** Full list of Italian words varying in concreteness and specificity used in the behavioral tasks presented in Chapter 4.

| <b>N_nonwords</b> | <b>Nonwords</b>   |
|-------------------|-------------------|
| 1                 | Framba            |
| 2                 | Onsuccia          |
| 3                 | Gollo             |
| 4                 | Lute              |
| 5                 | Oclifupro         |
| 6                 | Blagassa ispodica |
| 7                 | Suprite           |
| 8                 | Guagge            |
| 9                 | Toscemboa         |
| 10                | Odalendo          |
| 11                | Ancosofulpa       |
| 12                | Effriggo          |
| 13                | Aggalo            |
| 14                | Dimiciampo        |
| 15                | Eclo do liglio    |
| 16                | Opponna           |
| 17                | Pubuccema         |
| 18                | Nadesa            |
| 19                | Asufrena          |
| 20                | Avremalbo         |
| 21                | Tinvucumo         |
| 22                | Sulito            |
| 23                | Dueta             |
| 24                | Tempefama         |
| 25                | Ebrede cagitre    |

|    |                 |
|----|-----------------|
| 26 | Dalina          |
| 27 | Arranfa         |
| 28 | Settafe         |
| 29 | Podanogo        |
| 30 | Anvento         |
| 31 | Spotuna         |
| 32 | Bovola          |
| 33 | Efolmira        |
| 34 | Brodiro         |
| 35 | Puclifo         |
| 36 | Nopettefa       |
| 37 | Ponemelzobo     |
| 38 | Lebunado        |
| 39 | Cregnalo        |
| 40 | Casba           |
| 41 | Gormisa         |
| 42 | Silcaldo        |
| 43 | Trela           |
| 44 | Ferloglio       |
| 45 | Dindesda        |
| 46 | Gredapa         |
| 47 | Addranvizo      |
| 48 | Cicchinsa       |
| 49 | Tespenco        |
| 50 | Busdetro        |
| 51 | putrole nituspe |
| 52 | Suese           |
| 53 | Loieme          |
| 54 | Branutiduo      |
| 55 | Ocheldo         |
| 56 | Gamme           |
| 57 | Tomagga         |
| 58 | Candana         |
| 59 | Licreterro      |
| 60 | Farrenva        |
| 61 | Avrozinauge     |
| 62 | Orroserindo     |
| 63 | Iltacro         |
| 64 | Spinseppe       |
| 65 | Esdraliosa      |
| 66 | Imafe           |

|     |              |
|-----|--------------|
| 67  | Dovoso       |
| 68  | Sumofio      |
| 69  | Visama       |
| 70  | Dideco       |
| 71  | Ecio         |
| 72  | Erbompio     |
| 73  | Trenada      |
| 74  | Attrimosa    |
| 75  | Liofelava    |
| 76  | Ottega       |
| 77  | Evundo       |
| 78  | Discelso     |
| 79  | Firdatriore  |
| 80  | Fionnonga    |
| 81  | Utruce       |
| 82  | Imbosbia     |
| 83  | Saruso       |
| 84  | Simproziana  |
| 85  | Dordococce   |
| 86  | Dozigruomo   |
| 87  | Lusattade    |
| 88  | Seplurecce   |
| 89  | Ettefama     |
| 90  | Forsalegga   |
| 91  | Rubrata      |
| 92  | Vocemi       |
| 93  | Ciaprenge    |
| 94  | Saviozo      |
| 95  | Geclante     |
| 96  | Seca         |
| 97  | Rugraro      |
| 98  | Lada         |
| 99  | Sechenizia   |
| 100 | Lastoco      |
| 101 | Surga        |
| 102 | Abafaco      |
| 103 | Pronofifiote |
| 104 | Remispiabo   |
| 105 | Stindindanse |
| 106 | Elipia       |
| 107 | Polirumago   |

|     |              |
|-----|--------------|
| 108 | Orufobria    |
| 109 | Giarituo     |
| 110 | Miabibrina   |
| 111 | Colavilva    |
| 112 | Eumozia      |
| 113 | Cecozicia    |
| 114 | Momesdo      |
| 115 | Ligafifilza  |
| 116 | Firasta      |
| 117 | Oltoprio     |
| 118 | Enequatuoce  |
| 119 | Frevironiena |
| 120 | Zoteo        |

---

**Table A.2.** Full list of nonwords used in the behavioral tasks presented in Chapter 4.

## APPENDIX B

The tables below list the full set of 60 Italian target sentences and 4 types of follow-ups used as experimental stimuli in the behavioral tasks discussed in Chapter 5 (Plausibility judgment and Rating Task).

| N_Sentence | Sentence                           | English_version                | Subcategory | Sentence_type     |
|------------|------------------------------------|--------------------------------|-------------|-------------------|
| 1          | Ho mangiato una banana baby        | I ate a baby banana            | FOOD        | Concrete-Specific |
| 2          | Ho mangiato un frutto              | I ate a fruit                  | FOOD        | Concrete-General  |
| 3          | Ho acquistato dell'uva sultanina   | I bought sultanas              | FOOD        | Concrete-Specific |
| 4          | Ho acquistato delle bacche         | I bought some berries          | FOOD        | Concrete-General  |
| 5          | Ho comprato delle fragole di bosco | I bought wild strawberries     | FOOD        | Concrete-Specific |
| 6          | Ho comprato del cibo               | I bought food                  | FOOD        | Concrete-General  |
| 7          | Ho tagliato la melanzana viola     | I cut the purple eggplant      | FOOD        | Concrete-Specific |
| 8          | Ho tagliato un ortaggio            | I cut a vegetable              | FOOD        | Concrete-General  |
| 9          | Ho cucinato un peperone giallo     | I cooked a yellow bell pepper  | FOOD        | Concrete-Specific |
| 10         | Ho cucinato una verdura            | I cooked a vegetable           | FOOD        | Concrete-General  |
| 11         | Ho preparato una crostata          | I prepared a tart              | FOOD        | Concrete-Specific |
| 12         | Ho preparato un dolce              | I prepared a dessert           | FOOD        | Concrete-General  |
| 13         | Ho usato un martello pneumatico    | I used a jackhammer            | FOOD        | Concrete-Specific |
| 14         | Ho usato un utensile               | I used a tool                  | TOOL        | Concrete-General  |
| 15         | Ho preso una caraffa               | I took a jug                   | TOOL        | Concrete-Specific |
| 16         | Ho preso un contenitore            | I took a container             | TOOL        | Concrete-General  |
| 17         | Ho affilato un coltellino svizzero | I sharpened a Swiss army knife | TOOL        | Concrete-Specific |
| 18         | Ho affilato un attrezzo            | I sharpened a tool             | TOOL        | Concrete-General  |
| 19         | Ho lavato il forchettone           | I washed the carving fork      | TOOL        | Concrete-Specific |
| 20         | Ho lavato la posata                | I washed the utensil           | TOOL        | Concrete-General  |
| 21         | Ho lavato uno spazzolino           | I washed a toothbrush          | TOOL        | Concrete-Specific |
| 22         | Ho lavato un oggetto               | I washed an object             | TOOL        | Concrete-General  |

|    |                                |                               |        |                   |
|----|--------------------------------|-------------------------------|--------|-------------------|
| 23 | Ho accarezzato il dalmata      | I petted Dalmatian            | ANIMAL | Concrete-Specific |
| 24 | Ho accarezzato un animale      | I petted an animal            | ANIMAL | Concrete-General  |
| 25 | Ho fotografato il leone bianco | I photographed the white lion | ANIMAL | Concrete-Specific |
| 26 | Ho fotografato il felino       | I photographed the feline     | ANIMAL | Concrete-General  |
| 27 | Ho visto un dromedario         | I saw a dromedary             | ANIMAL | Concrete-Specific |
| 28 | Ho visto un mammifero          | I saw a mammal                | ANIMAL | Concrete-General  |
| 29 | Ho catturato una farfalla      | I caught a butterfly          | ANIMAL | Concrete-Specific |
| 30 | Ho catturato un volatile       | I caught a flying animal      | ANIMAL | Concrete-General  |
| 31 | Ho creato un diagramma         | I created a diagram           | PSTQ   | Abstract-Specific |
| 32 | Ho creato una rappresentazione | I created a representation    | PSTQ   | Abstract-General  |
| 33 | Ho delimitato l'area urbana    | I delimited the urban area    | PSTQ   | Abstract-Specific |
| 34 | Ho delimitato l'estensione     | I delimited the extent        | PSTQ   | Abstract-General  |
| 35 | Ho calcolato la percentuale    | I calculated the percentage   | PSTQ   | Abstract-Specific |
| 36 | Ho calcolato la quantità       | I calculated the quantity     | PSTQ   | Abstract-General  |
| 37 | Ho fatto un ritratto           | I made a portrait             | PS     | Abstract-Specific |
| 38 | Ho fatto una dichiarazione     | I made a statement            | PS     | Abstract-General  |
| 39 | Ho compreso il movente         | I understood the motive       | PS     | Abstract-Specific |
| 40 | Ho compreso la ragione         | I understood the reason       | PS     | Abstract-General  |
| 41 | Ho pensato alla sorte avversa  | I thought about bade fate     | PS     | Abstract-Specific |
| 42 | Ho pensato al futuro           | I thought of the future       | PS     | Abstract-General  |
| 43 | Ho risolto un indovinello      | I solved a riddle             | PS     | Abstract-Specific |
| 44 | Ho risolto un problema         | I solved a problem            | PS     | Abstract-General  |
| 45 | Ho commesso un peccato mortale | I committed a mortal sin      | PS     | Abstract-Specific |
| 46 | Ho commesso una malefatta      | I committed a wrongdoing      | PS     | Abstract-General  |
| 47 | Ho espresso una delibera       | I expressed a resolution      | PS     | Abstract-Specific |
| 48 | Ho espresso un parere          | I expressed an opinion        | PS     | Abstract-General  |
| 49 | Ho percepito la serenità       | I perceived serenity          | EMS    | Abstract-Specific |
| 50 | Ho percepito l'attitudine      | I perceived the attitude      | EMS    | Abstract-General  |
| 51 | Ho stretto un sodalizio        | I formed a fellowship         | EMS    | Abstract-Specific |
| 52 | Ho stretto un rapporto         | I formed a relationship       | EMS    | Abstract-General  |
| 53 | Ho evitato una rivolta         | I avoided a revolt            | EMS    | Abstract-Specific |

|    |                                 |                                |     |                   |
|----|---------------------------------|--------------------------------|-----|-------------------|
| 54 | Ho evitato un contrasto         | I avoided a conflict           | EMS | Abstract-General  |
| 55 | Ho fatto un favore              | I did a favor                  | EMS | Abstract-Specific |
| 56 | Ho fatto una buona azione       | I did a good deed              | EMS | Abstract-General  |
| 57 | Ho controllato l'angoscia       | I controlled the anxiety       | EMS | Abstract-Specific |
| 58 | Ho controllato lo stato d'animo | I controlled the state of mind | EMS | Abstract-General  |
| 59 | Ho affrontato il terrore        | I faced the terror             | EMS | Abstract-Specific |
| 60 | Ho affrontato l'emozione        | I faced the emotion            | EMS | Abstract-General  |

**Table B.1.** Full list of Italian target sentences varying in concreteness and specificity used in the behavioral tasks presented in Chapter 5.

| Experiment | Follow_up      | English_version      | Follow-up_type |
|------------|----------------|----------------------|----------------|
| 1          | Non ho capito  | I did not understand | Uncertainty    |
| 1          | In che senso?  | What do you mean?    | Uncertainty    |
| 1          | Ottimo lavoro  | Good job             | Certainty      |
| 1          | Hai fatto bene | You did it well      | Certainty      |
| 2          | Raccontami     | Describe             | Curiosity      |
| 2          | Dimmi di più   | Tell me more         | Curiosity      |
| 2          | Ok             | Ok                   | End            |
| 2          | Grazie         | Thank you            | End            |

**Table B.2.** Full list of Italian follow-up responses expressing uncertainty, certainty, need for additional information (i.e., curiosity), and willingness to end the conversation, used in the behavioral tasks presented in Chapter 5.