

UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL
INSTITUTO DE INFORMÁTICA
PROGRAMA DE PÓS-GRADUAÇÃO EM COMPUTAÇÃO

SANDRO RAMA FIORINI

**Similarity, Structure and Spaces:
Representation of Part-Whole Relations in
Conceptual Spaces**

Thesis presented in partial fulfillment
of the requirements for the degree of
Doctor of Computer Science

Advisor: Prof. Dr. Mara Abel

Porto Alegre
Maio 2014

CIP – CATALOGING-IN-PUBLICATION

Fiorini, Sandro Rama

Similarity, Structure and Spaces: Representation of Part-Whole Relations in Conceptual Spaces / Sandro Rama Fiorini. – Porto Alegre: PPGC da UFRGS, 2014.

136 f.: il.

Thesis (Ph.D.) – Universidade Federal do Rio Grande do Sul. Programa de Pós-Graduação em Computação, Porto Alegre, BR–RS, 2014. Advisor: Mara Abel.

1. Part-whole relations. 2. Conceptual spaces. 3. Similarity. 4. Concept representation. I. Abel, Mara. II. Título.

UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL

Reitor: Carlos Alexandre Netto

Pró-Reitor de Coordenação Acadêmica: Rui Vicente Oppermann

Pró-Reitora de Pós-Graduação: Vladimir Pinheiro do Nascimento

Diretor do Instituto de Informática: Luís da Cunha Lamb

Coordenador do PPGC: Prof. Luigi Carro

Bibliotecária-chefe do Instituto de Informática: Beatriz Regina Bastos Haro

“— We demand rigidly defined areas of doubt and uncertainty!”
— DOUGLAS ADAMS, *The Hitchhiker’s Guide to the Galaxy*

ACKNOWLEDGEMENTS

As in any other kind of work, nobody can do anything without other people's help, directly or indirectly. Here I pay tribute to the institutions and people that helped me to carry the work that resulted in this thesis and in the doctor that is writing these words.

Firstly, I would like to thank UFRGS and the Informatics Institute for all help and support during the last 12 years since I started as an undergrad student. I am honoured to have their names imprinted in my history.

I also would like to thank the agencies that supported and funded my work continuously during the four years of my studies, namely CNPq and CAPES. In particular, I would like to thank ANP/Petrobras' PFRH-227 program for the increased support during the final year of this work.

No scientific work is done without a proper forum in which to discuss new ideas freely. The BDI Group is that kind of place, in which I learned how to do science and how to be professional. I am thankful to all colleagues and friends I made there throughout the last four years: Luan Fonseca Garcia, Jose Lozano, Gabriel Moser, Ricardo Werlang, Ricardo Linck, Guilherme Schievelbein and Alexandre Lorenzatti. Particular thanks to Joel Carbonera for the friendship and for his willingness to discuss the ideas I wrote in this volume (and also other, more tangent topics). I also thank Marnes Augusto Hoff, Vitor Jorge and Vitor Rey for their help and friendship.

One of the pillars of this and my previous work has been the support of my supervisor, Prof. Mara Abel. I am thankful for the confidence she has put on me, allowing myself to pursue my own research objectives since early on in this thesis. I am especially thankful to her for the guidance in all matters of academic and non-academic life since 2002. I would also like to thank Prof. Edson Prestes and Prof. Taisy Weber for the support in the latter stages of this work.

Part of this work has been carried out in Lund University, Sweden. I made some nice friends there, so I would like to thank all people at the LUCS lab, in particular Joel Parthemore and Frank Zenker. I would also like to thank Prof. Peter Gärdenfors, my supervisor at Lund, for suggesting the topic of this thesis and for the non-trivial amount of help in carrying out the research that led to the main chapter in this volume. I am thankful to Lund University for the hospitality and support.

Friends are important in all situations, particularly when the need for keeping mental sanity is in need during a phd. I would like to thank Gustavo Coelho, Bill Roesler, Leandro Heck, Francisco Scornavacca, Raquel Kieling, Laura Leão, Eliete Tiburski for their support. In Lund, I also made nice friends. First, I would like to thank my hosts Bálint and Lena Wagner for their kindness and their house's roof. I thank Guillermo Chaminade,

Maria Rodriguez-Rabadan, Fabio Cristiano and Tina Hannemann for the friendship and great support in all matters. I am also thankful to the friends at ENDEEPER for the cheerful lunch times during the last four years, particularly to Julia Zacouteguy for the other especially cheerful times.

I would like also to thank my family, especially Maria Rama and Fabian Rama for the company, encouragement and food. I am also in debt with my parents, Olir Fiorini and Marlene Rama Fiorini, for all the education and the values they provided me, which made me the person I am now.

Finally, thanks Universe for not throwing a coconut onto my head just yet. That would be a total waste of energy.

Similarity, Structure and Spaces: Representation of Part-Whole Relations in Conceptual Spaces

ABSTRACT

One of main remaining challenges in Artificial Intelligence is how to make intelligent systems to ground high-level abstractions in reality. At least part of this problem comes down to the question of what representation framework is best suited in a way that facilitates object recognition. Animal cognition, particularly in humans, can give a clue of how such representation framework looks like. Studies on the cognition of object recognition suggests that representation in cognition incorporates similarity and holistic-structural (i.e. whole-part) information about concepts. We hold the view that computer systems dealing with part-whole relations should be able to represent similarity and holistic-structural information. However, there exists no representation framework that allows such information to be represented in an integrated way. This thesis proposes a concept representation framework for representing information about similarity between part-whole structures. We base our proposal on the theory of conceptual spaces, which are mathematical spaces where regions and points represent concepts and objects respectively, such that their distance denotes their similarity. In our proposal, parts and wholes are represented in distinct conceptual spaces called holistic and structure spaces. Holistic spaces allow for systematic similarity judgements between wholes. On the other hand, structure spaces allow for systematic similarity judgement between structures of parts. A point in the structure space denotes a particular part structure; regions in the space represent different general types of part structures. By redefining conceptual spaces as a product of holistic and structure spaces, our goal is to allow one to represent similarity information between different wholes, taking into consideration the similarity between shared parts and their configurations. This thesis has three main contributions: a general theory about holistic and structure spaces; a formulation of the theory founded on products of metric spaces; and a generic algorithm for object recognition implementing holistic-structural processing.

Keywords: Part-whole relations, conceptual spaces, similarity, concept representation.

Similaridade, Estruturas e Espaços: Representação de Relações de Parte-todo em Espaços Conceituais

RESUMO

Um dos principais desafios remanescentes em Inteligência Artificial é como fazer sistemas inteligentes ancorar abstrações de alto nível na realidade. Pelo menos parte do problema vai em direção da questão de qual arcabouço de representação é mais apropriado de maneira que facilite o reconhecimento de objetos. A cognição em animais, particularmente em humanos, pode dar pistas de como tal arcabouço de representação se parece. Estudos na cognição do reconhecimento de objetos sugere que o problema da representação na cognição incorpora informações a respeito de similaridade e informação holística-estrutural (i.e. todo-parte) a respeito de conceitos. Temos a visão de que sistemas computacionais que lidam com relações de todo-parte deveriam representar informações holístico-estruturais e similaridade. No entanto, não existe arcabouço de representação que permite tais informações serem representadas de forma integrada. Esta tese propõe um arcabouço de representação de informação de similaridade entre estruturas de todo-parte. Nossa proposta é baseada na teoria dos espaços conceituais. Estes são espaços matemáticos onde regiões e pontos representam conceitos e objetos respectivamente, tal que a distância entre estas entidades denota a sua similaridade. Na nossa proposta, todos e partes são representados em espaços conceituais distintos, chamados espaços holísticos e estruturais. Espaços holísticos permitem o julgamento de similaridade sistemático entre todos. Por outro lado, espaços estruturais permitem o julgamento de similaridade sistemático entre estrutura de partes. Um ponto em um espaço estrutural denota uma estrutura particular de partes; regiões neste espaço representam diferentes tipos de estruturas de parte. Através da redefinição de espaços conceituais como um produto de espaços holísticos e estruturais, nosso objetivo é permitir a representação de informações de similaridade entre diferentes todos, levando em consideração a similaridade entre partes compartilhadas e suas configurações. Esta tese tem três contribuições principais: uma teoria geral sobre espaços holísticos e estruturais; uma formalização da teoria fundada em produto de espaços métricos; e um algoritmo genérico para reconhecimento de objetos, implementando processamento holístico-estrutural.

Palavras-chave: relações de parte-todo, espaços conceituais, similaridade, representação de conceitos.

LIST OF FIGURES

1.1	Relation between types of representation frameworks and representation levels.	19
1.2	Diagram representing an example of object recognition by holistic and structural processing.	21
2.1	Examples of targets and distractors used by Alexander and Zelinsky (2012) to test the impact of part similarity in visual search	30
2.2	The categories of wholes according to compositional structure.	33
2.3	Deformable worms and circles.	37
2.4	Conceptual spaces of shapes.	38
2.5	Parts as image patches.	39
2.6	Depiction of a multilevel representation of parts. Elements in each level correspond to a composition of elements in lower levels.	39
3.1	Colour spindle as the colour space.	44
3.2	Graphical schema of realist semantics: (a) extensional semantics; and (b) intensional semantics.	47
3.3	Graphical schema of cognitive semantics.	47
3.4	An image schema for “across”.	49
3.5	Simplified version of Chella’s cognitive architecture.	53
4.1	Example of diagrams depicting the conceptual space of apple	58
4.2	Structure space of apple concept	62
4.3	Structure spaces according to paronymy types.	64
4.4	Parts hierarchy model proposed by Marr and Nishihara.	67
4.5	Examples of different contexts in the apple structure space.	69
5.1	Example of diagrams depicting the conceptual space of apple.	83
5.2	The two stages of a dimensional filter applied to the concept Seed	93
5.3	Commutative diagram of <i>dimensional filter</i> (τ_i) and <i>part relation</i> (ρ_i)	94
5.4	Example of the relationship between Bike and Wheel	96
6.1	General structures of the object interpretation algorithm.	100
6.2	A small section of a gamma ray well log with the typical geological features identified in well log analysis.	110
6.3	Conceptual model of sequence and parasequence.	112
6.4	The two well logs used in the demonstration.	113

6.5	Visual depiction of how the structural processing algorithm checks for possible parts in an object.	115
6.6	Example of how holistic-structural similarity is measured between sequences.	116

LIST OF TABLES

2.1	Types of meronymic relation according the three differentiating criteria.	32
3.1	The concept Apple.	45

CONTENTS

1	INTRODUCTION	17
1.1	Background	17
1.2	Motivation	22
1.3	Objectives	24
1.4	Scope	24
1.5	Methodology and Structure	25
2	PERSPECTIVES ON PART-WHOLE RELATIONS	27
2.1	Part-Whole Relations in Cognition	28
2.2	Part-Whole Relations in Linguistics	31
2.3	Part-Whole Relations in Logic and Ontology	34
2.4	Part-Whole Relations in Computer Science	36
2.5	Summary and final remarks	40
3	CONCEPTUAL SPACES	43
3.1	Basic Notions	43
3.2	Cognitive semantics	46
3.3	Mathematical formulations	49
3.3.1	Proposals by Aisbett, Gibbon and Rickard	49
3.3.2	Proposals by Raubal and Adams	51
3.3.3	Formal Ontology	51
3.3.4	Other formulations	52
3.4	Applications	52
3.5	Summary and final remarks	54
4	PART-WHOLE RELATIONS AS CONCEPTUAL SPACES	55
4.1	Preliminaries: Bubble diagram	57
4.2	Representing Parts and Wholes in Conceptual Spaces	58
4.3	Types of Structure Domains	63
4.3.1	Complex and Collective	63
4.3.2	External Partitions	66
4.4	Marr's Hierarchical Model, Revisited	67
4.5	Same Whole, Different Parts	68
4.6	Partonomies and the problem of transitivity	69
4.7	Other Ontological Considerations	71
4.8	Computational Object Recognition with Structure Spaces	76
4.9	Summary and final remarks	78

5	PART-WHOLE RELATIONS AS PRODUCTS OF METRIC SPACES . . .	81
5.1	A short review of conceptual spaces and holistic-structure spaces	82
5.2	Metric Spaces for Conceptual Spaces	84
5.3	Product Space	87
5.4	Conceptual Spaces as Metric Spaces	87
5.4.1	Properties and Prototypes	90
5.4.2	Concepts	91
5.5	Part Relations and Dimensional Filters	91
5.6	Parts and wholes	95
5.7	Discussion: benefits to software applications	97
5.8	Summary and final remarks	98
6	AN ALGORITHM FOR HOLISTIC AND STRUCTURAL INTERPRETA-	
	TION	99
6.1	Representation framework	101
6.2	Interpretation algorithm	103
6.2.1	Specification	103
6.2.2	Implementation requirements	108
6.3	Example: well log interpretation task	108
6.4	Summary and final remarks	117
7	CONCLUSION	119
7.1	The Theory	120
7.2	The Formulation	122
7.3	The Algorithm	123
7.4	Open issues and future work	124
	REFERENCES	127

1 INTRODUCTION

This is a theoretical thesis about concept representation for computer systems. We aim at investigating what requirements similarity and holistic-structural processing in visual cognition impose on concept representation frameworks for visual object recognition by computers. The following sections we introduce its background, motivation, objectives and scope.

1.1 Background

We regard humans and some other animals as intelligent beings for their capability of perceiving, thinking and acting onto the world. In particular, humans seem to be capable of processing the stream of unstructured information incoming from their senses into impressive structures that can be later used in their advantage for survival. This process involves a series of mechanisms that characterize how human mind works, including recognition, memory, reasoning, language, planning, etc. An interesting issue that permeates all these mechanisms is how such those impressive structures are organized.

Since the early ages of civilization, humans have been trying to build intelligent machines. Machines that can perceive, think and act in the world in the very same way humans and other animals do – at least on a functional level. The importance of having such machines is unquestionable, as they could replace humans in tasks at which humans are not so good. Computers also employ mechanisms for representing information from the outside world. Presumably, there are differences and similarities in how humans and computers represent information. There are obvious differences in relation to the substrate where these processes occur; i.e. biological *versus* electronic. More similarities can be found in a functional level. Humans seem to abstract information in form of concepts, that is, structures generalizing sets of co-occurring patterns, or objects¹. The real nature of such mental structures is still open for debate (BARSALOU, 2010; PITT, 2013). Concepts are grounded in reality by means of a multitude of cognitive processes, such as action and perception processes (BARSALOU, 2008). We hold the position that recognition

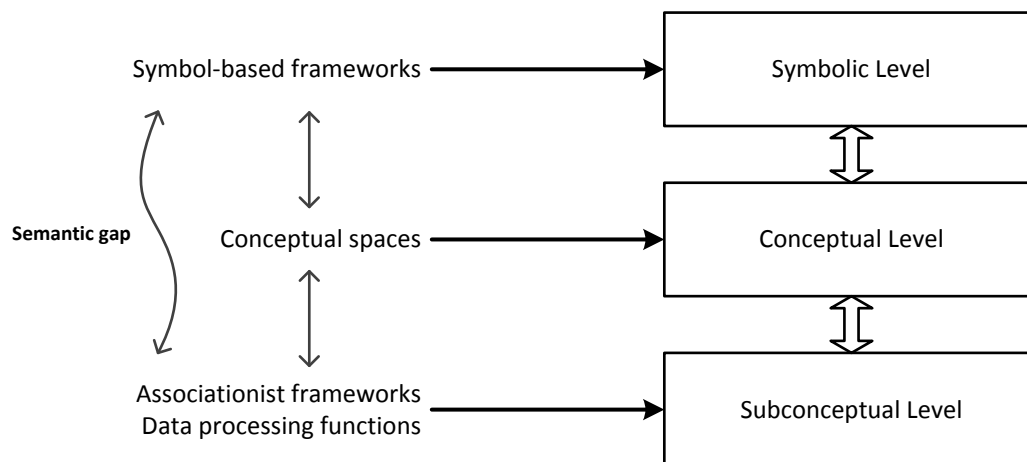
¹In this thesis, we use the term “object” as a synonym of instance or individual of a given concept.

plays a role in grounding though perception. Recognition is a process that can recognize the objects and concepts in perceived stimuli, effectively grounding them. In many ways, computers and other computational systems seem to have gone in a similar path. Algorithms process information represented in some way, such as data structures or databases, producing actions upon the world. The representations must reflect the perceived world structure in order to allow algorithms to be effective.

There are two main categories of frameworks for concept representation in computers. Concept representation can be based on associationist approaches, such as *artificial neural networks* (ALN). These techniques can aggregate representation, mechanism and grounding in the same structure, not needing the use symbols. They can associate recurrent perceptual patterns directly to output tokens denoting concepts and objects, effectively representing these entities. However, associationist approaches typically have problems to produce more elaborate abstractions, which are common in high-level human cognition. On the other hand, concept representation in computers can also be based on symbols systems. These systems are usually founded on logic or mathematical formalisms. Symbols carry low semantic content in themselves; their meaning is, by definition, ascribed by some sort of convention (GÄRDENFORS, 2000), defined outside the system. They must be ultimately grounded in reality to have meaning. This aspect is summarized by Harnad (1990) as the *symbol grounding problem*. It states that the meaning of a system of symbols cannot be fully defined in terms of a second system of symbols without incurring in an infinite regression. Ultimately, symbols must be grounded in a system that is not symbolic in nature. Evidently, the associationist and symbolic frameworks categories complement each other in their strengths and weaknesses. Not surprisingly, there are many computer systems that try to combine associationist (including signal processing algorithms) and symbolic approaches in order to implement full cognitive systems (FIORINI; ABEL, 2010, for a review on vision systems)

However, associationist and symbolic approaches do not capture the whole palette of features involved in representation. This problem is reflected in a common issue in hybrid computer systems. There is a *semantic gap* between symbolic representations and the output of low-level processing of perceptual stimuli (usually, resulting from connectionist or signal processing approaches). It is usually referred as a gap between stimuli information and the interpretation of this stimuli (cf SMEULDERS et al., 2000), which translates to a gap between high-level and low-level representation structures (Figure 1.1). This gap must be somehow bridged in order to allow computers to generalize perceptual patterns into symbols in a smooth way. The bridge over the semantic gap has the role of mapping patterns of sensorial input into the symbol that takes the place of these patterns, representing them. This thesis is motivated by the need to construct such *semantic bridge*. More specifically, we are interested in the features of the information represented by this bridge not contemplated by other representation frameworks.

Figure 1.1 – Relation between types of representation frameworks and representation levels.



Source: the authors.

The ways in which humans bridge the gap between linguistic constructions and motor-perceptual information can give us some clues about how it might work in computers. In particular, *we are interested in how the cognition of object recognition helps in building the semantic bridge*. Before that, however, there is a paradigmatic issue whether it is possible to use theories about human (or animal) cognition as inspiration for constructing intelligent machines. A common assumption in Artificial Intelligence is that, in order to build such machines, we should understand and reproduce the mechanisms of human intelligence. This thesis commit to this view.

At a cognitive level, some mechanisms are known to be involved in the object recognition, such as segmentation, memory, representation and so on. These mechanisms affect how cognition represents information. An important mechanism is *similarity*. Whilst it has been subject of debate throughout the years, some authors linked similarity to the core of concept formation in cognition (GOLDSTONE, 1994, for a discussion). The idea is that humans (and many animals) group together stimuli that look similar: a *dalmata* is seen as a dog because it is more similar to other dogs than to horses or airplanes. Rosch (1978) proposed that humans categorize stimuli essentially by measuring their similarity to concept *prototypes*. Prototypes are objects that are typical to a given concept. For instance, robins are examples of prototypes to the category of birds, contrasting to penguins.

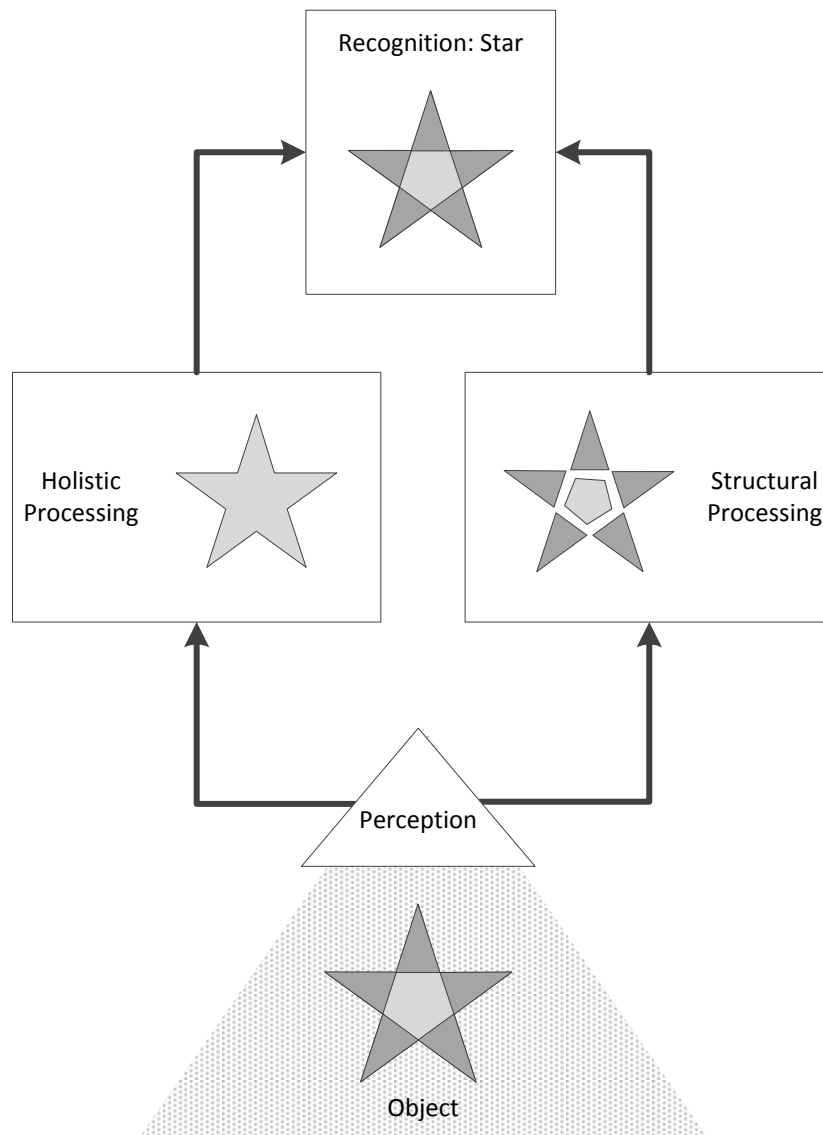
Similarity is the main construct of some frameworks for concept representation in cognition, especially in a class usually referred to as geometrical frameworks (SHEPARD, 1987; GÄRDENFORS, 2000). In general, objects are represented in some sort of feature space, where dimensions describe features that serve as criteria to compare objects. Concepts are entities in such spaces. Such frameworks distinguish themselves from

symbolic and associationist approaches. They are not symbolic, mainly because they do embed some of topological structure of reality. They are also not associative in the sense that they do impose some internal structure in how objects are categorized; i.e. they are not simply associations of inputs to outputs.

In particular, we adopt the *Theory of Conceptual Spaces*, proposed by (GÄRDENFORS, 2000). It defines concepts as regions and sets of convex regions in a mathematical space. In such spaces is possible to talk about similarity, prototypes and cognitive processes such as categorization, inductions and metaphors. One of its main features is that it acknowledges the existence of symbolic and associationist frameworks. They are included in what Gärdenfors calls *symbolic* and *subconceptual* levels to representation, respectively (Figure 1.1). We also include signal processing algorithms in the subconceptual level, for they also associate certain signal patterns to outputs. Gärdenfors argues for conceptual spaces to be the bridge between the symbolic and subconceptual levels, forming the *conceptual level*. Given this particular feature, conceptual spaces are particularly attractive to serve as a background theory for cognitive computer system. For instance, Chella, Frixione and Gaglio (1997) propose conceptual spaces to ground the meaning of logical formulas to perceptual stimuli in a computer vision application. However, conceptual spaces miss some important aspects of cognition.

Another aspect of object recognition is related to how cognition segment the world to categorize it. That is the difference between what we call *parts* and *wholes*. According to Palmer (1977), there is no formal difference between parts and wholes, except for the level of analysis. However, it has been shown that these entities have a very special and structuring role in human cognition (perhaps as much as the conceptual subsumption relation does). In an important work, Farah (1992) proposed that object recognition in vision includes two distinct processes. The *holistic process* recognize object as a whole entity. The *structural process* takes into consideration only the parts that form the object and how they are related. Farah suggests that they occur always in parallel during recognition, being suited for different categories of objects (Figure 1.2). For instance, subjects having impaired holistic processing have difficulties recognizing words and certain object. On the other hand, subjects with impaired structural processing have difficulties to recognize animals. Other studies corroborate with such separation (BEHRMANN; WILLIAMS, 2007; AUGUSTINE; SMITH; JONES, 2011). The role of whole and parts in object recognition has also been central to the model/view discussion on visual object representation (BIEDERMAN, 1987; EDELMAN, 1998), with recent approaches admitting that both parts and wholes are important aspects of visual object recognition (GRAF, 2006, for a review). This move corroborate to the view that object recognition is holistic *and* structural.

Figure 1.2 – Diagram representing an example of object recognition by holistic and structural processing.



Source: the authors.

Less studied is the interaction between holistic/structural processing and similarity effects. Nevertheless, research in cognitive science shows evidence of important interactions between both phenomena (ALEXANDER; ZELINSKY, 2012; FÖRSTER, 2009).

The question that poses itself at this point is how such body of empirical knowledge about the object recognition in humans can add to the state-of-the-art of concept representation in computers. If we assume that recognition requires information represented in some way, it is reasonable to expect that the mechanisms of the former affect the form of the later. This thesis argues to the view that similarity and holistic/structural information are fundamental aspects of concept representation and should be taken into consideration in computer approaches. In particular, we propose an extension to conceptual spaces that allows one to represent of holistic and structural information as geometric constructs in a similarity space.

1.2 Motivation

As computer systems develop towards ubiquity and autonomy, it becomes increasingly important for such systems to be able to perceive and interpret reality. One aspect of the challenge is to implement object recognition capabilities in these systems. We consider object recognition to be the problem of interpreting recurrent patterns of low-level sensory information into high-level abstractions that are meaningful and useful to cognitive agents, being these humans or intelligent systems.

Object recognition is both a *processing* and a *representation* challenge. It is a processing challenge in the sense that transforming low-level sensory information to high interpretation requires an active computational mechanism. At the same time, this mechanism requires some sort of representation structure to represent the information been processed. Such representation structure must support the processes necessary to achieve recognition.

A long standing issue in object recognition is the problem of *semantic gap* (SMEULDERS et al., 2000). It consists in effectively linking labels denoting concepts to low-level sensory information. The techniques to address the representation aspect of semantic gap vary, usually including a blend of techniques in the symbolic and subconceptual levels. As we have seen, these techniques have difficulties in dealing with some aspects that are inherent to object recognition, such as similarity between *conceptual* entities (i.e. such as concepts and objects, in contrast with pre-conceptual entities such as perceptual patterns) and prototypes.

We believe some of these difficulties can be addressed by the use of processing and/or representation frameworks that complement symbolic and subconceptual approaches. In this thesis, we commit to the theory of conceptual spaces as basic representation framework (GÄRDENFORS, 2000), which takes similarity and prototypes into account. Its

main feature is that concepts can be represented as regions in a multidimensional space. Similarity plays a central role: the dimensions of conceptual spaces provide the means for determining similarity between concepts and between objects. Conceptual spaces have been used as basis for representation frameworks in computer systems (AISBETT; GIBBON, 2001b; CHELLA; FRIXIONE; GAGLIO, 1997; ADAMS; RAUBAL, 2009b; KIRA, 2009). However, these approaches do not address one of the main aspects of object recognition in cognition: part-whole relations.

In our view, an adequate representation framework for object recognition should support similarity as well as representation of information required for holistic and structural processing. From a computational point of view, such framework would be highly beneficial. For instance, the modelling of part-whole relations is sometimes problematic in symbolic representation paradigms, such as ontologies (e.g., RECTOR et al., 2005). The algebraic approaches to part-whole relations, such as mereology, are not generally driven by the cognitive phenomena (SIMONS, 2006), such as similarity and holistic/structural processing. Some foundational ontologies already embed some notions of similarity (GANGEMI et al., 2002; GUIZZARDI, 2005), but it is not explained how similarity is related to part-whole representation.

We believe our approach is particularly beneficial when it comes to object recognition in computer systems, where notions of similarity and holistic/structural processing are more common. We find approaches employing either holistic processing (e.g., LOWE, 2004) or structural (part-based) processing (e.g., AGARWAL; AWAN; ROTH, 2004; FIDLER; BOBEN; LEONARDIS, 2008)). Some alternatives combine both strategies (e.g., MOTTAGHI; RANGANATHAN; YUILLE, 2011; GRAF, 2006), using global and local descriptors for a review. However, usually the focus greater on immediate issues related to shape recognition, and less on addressing broader questions associated to concept representation.

Furthermore, no representation framework based on conceptual spaces has directly addressed structural information. Aisbett and Gibbon (2001b) indirectly touch this point, but do not give it full attention. As a matter of fact, computer recognition proposals based on conceptual spaces have relied on symbolic level techniques to model part relations, such as in (CHELLA; FRIXIONE; GAGLIO, 2001). In this thesis, we extend conceptual spaces to be able to represent part-whole relations in integrated way, reflecting aspects of holistic and structural processing and being amenable to computer implementations.

1.3 Objectives

Our broad goal in this thesis is to investigate what requirements similarity and holistic-structural processing impose on concept representation frameworks for object recognition by computers. The expected result is a *theoretical* framework mainly grounded on evidence from cognitive sciences, which allows one to represent similarity between conceptual entities (i.e. concepts and objects) with regard to their holistic and structural aspects. In addition, the focus of this thesis is on representation of information regarding *physical* objects.

We adopt conceptual spaces as a starting point. Conceptual spaces already address the issue of similarity and integration with other representation paradigms. Thus, the more specific goal of this thesis is to investigate how holistic-structural information can be integrated into conceptual spaces and propose a theory of how this integration can take place.

Given we are interested in applications in computer science; the second goal is to construct a formal framework of the theory such that it specifies a mathematical scaffolding to computer implementations of the theory.

The last specific objective is to demonstrate the applicability of the results expected to be achieved by the previous goals. In particular, our goal is to show an interpretation algorithm based on holistic-structure spaces that can be applied in a domain-specific task.

The specific objectives can be summarized by:

Objective 1 (The Theory) : To investigate and propose a concept representation structure based on conceptual spaces that combines the notions of conceptual similarity and holistic/structural information;

Objective 2 (The Formulation) To define a formal concept representation framework based on the proposal in Objective 1. This formulation has to be amenable to computer implementation;

Objective 3 (The Algorithm) To specify an object-recognition algorithm that demonstrates the usefulness of formulation in Objective 2 compared to the state-of-the-art.

1.4 Scope

This thesis touches into a variety of domains, such as cognitive science, linguistics and computer science. As such, it is important for us to place it among other aspects of the issues it addresses.

Most of the motivation behind this thesis comes from the problem of symbol grounding and the semantic gap. As Objective 1 states, our proposal is based on conceptual spaces, which addresses these topics by introducing itself as an intermediate representation layer between the symbolic and the subconceptual levels. Whilst our proposal extends the capabilities of conceptual spaces, we do not directly address how these extensions are connected to constructs in the adjacent levels, as done by Aisbett and Gibbon (2001a), for instance.

Conceptual spaces propose itself as a model for cognitive phenomena. While we do use cognition as basis for our theory, we do not target the exact same objective of Gärdenfors (2000). We propose a theory that is intended as a framework for computation in computers (see Objective 2), not as a model for animal cognition.

This work is commonly related to the study of parts and wholes in logic and philosophy literature, which goes under the name of Mereology. These studies are frequently concerned with the ontological aspects of entities and their parts. Whereas we do contrast some of our ideas with topics commonly discussed in Mereology (such as modality and transitivity of part relation), we must emphasize that our concerns with concept representation are mainly epistemological. We are *not* concerned with *what* concepts and object are “out there” in reality. Instead, our concern here is how to represent concepts in computers such that they accommodate holistic and structural information according to *cognitive phenomena*. Most of all, we do not intend to replace or reform Mereology; nor we give a cognitivist interpretation to Mereology (even if such an idea is tempting).

Finally, the focus of this thesis is in representation of physical objects and their categories. As such, we cannot guarantee the framework presented here can be used to represent other entities, such as events and abstract entities.

1.5 Methodology and Structure

Our broad objective is to propose a concept representation framework that is capable of representing similarity and holistic-structural information in one conceptual structure that is useful for building computer systems.

To achieve Objective 1, we revise the literature in parts and wholes in the many different domains where it is studied, as well as conceptual spaces. The former aspect is addressed in Chapter 2 and the latter aspect is achieved in the first part of Chapter 3. Subsequently, we develop our theoretical framework in Chapter 4. The chapter present the theory’s core and discuss its ramifications in related aspects of part-whole relation. We mainly propose two conceptual spaces called *holistic space* and *structure space*. Holistic spaces coincide with the usual conceptual spaces theory and encode holistic aspects of concepts ad objects (i.e. the whole). On the other hand, structure spaces encode structural

aspects of concepts and objects, such what are the constituting parts, what are their properties and how they are related to the whole (i.e. structure). We then redefine concepts as a product of holistic and structure spaces. The resulting construct allows one to express holistic and partonomic similarity between entities.

Further in Chapter 4, we discuss other aspects of part relations in association with holistic and structure spaces, such as context, types part relations and some ontological characteristics, such as part transitivity and essentiality.

Having accomplished Objective 1, we turn to Objective 2 in two steps. In Chapter 3 we review the existing formulations for conceptual spaces. Based on this knowledge, in Chapter 5 we propose our own formulation for the main aspects of the theory introduced in Chapter 4. The framework is based on a metric space interpretation of conceptual spaces, branching on the work of Aisbett and Gibbon (2001b). In this context, we also define holistic and structure spaces as metric spaces with a particular inner structure, defining the part-relations as structure preserving morphisms between metric spaces.

Both theory and the representation framework resulting from achieving objectives 1 and 2 will allow us to accomplish Objective 3 in Chapter 6. We propose a general purpose algorithm for holistic-structural processing based on the ideas and constructs presented in previous chapters. In particular, this algorithm takes advantage of part similarity to achieve top-down interpretation of features in digital signals. The rest of Chapter 6 is intended to show the algorithm instantiated as an interpretation algorithm for geological data.

Finally, in Chapter 7 we close this thesis with further thoughts on its impact and future work.

2 PERSPECTIVES ON PART-WHOLE RELATIONS

The use of part-whole relations spans many areas of everyday life. It is prominent in cognitive processes and natural language. It is also used as main construct in the specification of science and engineering models. As such, the theory and use of part-whole relations are studied across a whole set of fields, under many different perspectives. This chapter review what we consider the main aspects of part-whole relations that are necessary to the subsequent discussion in this thesis.

Guarino, Pribbenow and Vieu (1996) suggest two perspectives from which part-whole relations can be studied. The *logico-philosophical approach* takes the perspective of formal ontology and algebraic theories of parts, such as mereology (the formal study about parts and wholes), as well as other derived theories (VARZI, 2011; SIMONS, 2003). The other approach is the *cognitive-linguistic approach*, which looks at the problem in language processing, perception and action planning.

Taking inspiration in that organization, we propose four broad fields in which the contributions can be separated: *cognition, language, logic and ontology* and *computer science*. This separation follows an axis that goes from human-based research to more formal, implementation-based research.

Research in *cognition* of part-relations shows that part structure and wholes are of primary relevance to important cognitive processes. In particular, we focus on object recognition, given its importance to the motivation of this thesis. Research in *linguistics* tries to capture the semantic of part-whole constructs in natural language. These often serve as basis to more formal proposals in *logic* and *ontology*, such as in mereology. Finally, we investigate how part are used in *computer science*. Many algorithms and software applications employ part-whole relations, at least in a naïve way and often based on logic frameworks. Whilst our focus is on *representing* concepts and their parts, such approaches present interesting practical solutions involving partonomical relations that are relevant for this discussion

2.1 Part-Whole Relations in Cognition

Given the importance of the cognition of parts and wholes for human reasoning, research on it is not as extensive as one would expect. For instance, early work by Barbara Tversky and her colleagues (TVERSKY, 1989; TVERSKY; HEMENWAY, 1984) showed that parts play a central role in differentiating between base-level concepts, and also suggested that parts form a bridge between perceptual and functional knowledge.

A good amount of the research in the cognition of part and wholes is centred on shape recognition. During the 1990s, the discussion concerning shape recognition gravitated around two general set of theories in which the importance given to part relations was a distinguishing feature. The *view-independent* theories, mainly influenced by Marr's computational models of vision (MARR, 1982) and Biederman's work on *geons* (BIEDERMAN, 1987), postulated that objects are represented and perceived based on configurations of visual primitives that are invariant to viewpoint changes. On the other hand, there are *view-dependent* theories, like the ones proposed by Edelman (1998) and Ullman (2000), which state that objects are represented by "snapshots" (i.e., images) of the object's different angles, dismissing the importance of structural information. View-independent theories have the tendency to give more relevance to part relations (between visual primitives), while this aspect is not so much emphasized in view-dependent theories. Recently, though, evidence from cognitive and neurosciences (FOSTER; GILSON, 2002; NEWELL et al., 2005) supports that both processes are needed in object recognition (GRAF, 2006). Reviewing the state of the art in object recognition, Peissig and Tarr (2007) argue that the discussion regarding dependence of view in object recognition is orthogonal to the actual importance of part structure in recognition of physical objects. In this thesis, we focus on the latter aspect.

There are several relevant streams of empirical research in part-whole reasoning and representation. The first stream comes from studies of patients with certain cognitive impairments. *Integrative agnosia* is a rare kind of impairment that makes recognition of wholes difficult, but which leaves relatively recognition of parts unaffected. In one experiment, Behrmann et al. (2006) asked a patient with integrative agnosia to compare objects formed by different parts. The patient could recognize dissimilarities between objects that did not share the same parts. However, the patient was unable to recognize dissimilarities when objects shared parts that were arranged in different ways. Their conclusion is that the brain seems to encode part arrangement (part structure) independently of part shape (part qualities). In other study, Behrmann and Williams (2007) suggest that patients with integrative agnosia also have difficulties to chunk parts together in bigger parts in order to recognize complex objects. Interestingly, patients tested by Behrmann and Williams (2007) also had difficulties forming prototypes in category learning tasks with complex

objects. On the other hand, patients with *simultagnosia* can recognize wholes, but have difficulties to recognize parts. After examining object identification in two simultanagnosic patients, Riddoch and Humphreys (2004) indicate that the impairment might be caused by an interplay of issues in attention and information coding.

The findings about integrative and simultagnosia suggest that object identification employ information about wholes, parts and structure in an independent fashion. Further still, they indicate that there are two distinct, but correlated representation systems in the brain: one based on parts and the other based on the whole object. Evidence for this distinction also comes from a meta-analysis carried out by Farah (1992) on research pertaining to patients with different types of agnosia. She suggests that the brain employs two parallel but distinct cognitive processes in object recognition. In the structural process, whole objects are recognized by recognizing its constituent parts. In the holistic process, the recognition rests on the whole object, independent of its parts. The recognition of certain categories of objects usually relies more on the one or the other. In support of this position, there is evidence of a double dissociation between impairments in word recognition (regarded as structurally based) and impairments in face recognition (regarded as holistically based). At the same time, object recognition also seems to be partially affected in both impaired conditions, suggesting that object recognition is dependent on both structural and holistic representations.

Developmental psychology also provides some insights into this topic. It has been shown that recognition of objects by children under two years of age is mostly part-based. However, children later acquire the ability to recognize objects by their full shapes. For instance, a series of experiments with 18- to 30-month-old children, conducted by Smith and her colleagues (AUGUSTINE; SMITH; JONES, 2011; SMITH, 2009), suggest that the representations of geometric structures of whole objects are built over time, and, more broadly, that shape and part relations are two distinct components of children's judgements of shape similarity. Rackison and colleagues (WU; MARESCHAL; RAKISON, 2010) showed a similar trend in the development of children's cognition, also suggesting that *salient parts* play a role in object categorization. The perception of parts also seems to affect generalization in learning. Son, Smith and Goldstone (2008) found that teaching children the names of simple, featureless versions of new objects (e.g., some inner parts of the objects) helps them to generalize the names to similar but more complex versions of the objects. This finding supports the idea that young children focus their attention on small details (parts) when learning words for new objects. Presenting them with simple objects steers their attention to more general geometrical structures, helping them learn and generalize words for basic-level concepts. However, parts are important when differentiating objects that are similar in overall shape (e.g., cows and horses). A child normally first notices high-level part similarities, but for some concept distinctions, more attention

must be given to lower levels. For instance, dogs and cows have quite similar overall parts, and children sometimes do not distinguish them in their naming. Then they learn to differentiate on lower levels of the hierarchy, such as by noticing that dogs and cows have differently structured noses and tails.

This thesis focus on representation of part (structural) and whole (holistic) information and how such information affects similarity judgements. Many of the studies mentioned previously include some sort of similarity judgement (or matching) between stimuli related to parts and wholes. Nevertheless, some studies focused directly on the interaction between part-whole processing and similarity effects. Alexander and Zelinsky (2012) showed that part similarity plays an important role in visual search of real-world object. In a target-distractor kind of experiment, different quantities of parts of photorealistic distractors were replaced with target parts, creating increasing levels of target-distractor similarity (Figure 2.1). They found that, indeed, as the distractors share more parts with targets, the time of visual search increases progressively. The same happens when distractors share parts. In particular, Alexander and Zelinsky (2012) state that “very few parts had to be transplanted from target to distractor in order for the two objects to be perceived as highly similar”. This indicates the high impact of part/structural similarity in visual similarity. Förster (2009) performed nine experiments on how people judge similar/dissimilar stimuli in global and local processing. They found that global (holistic) processing tends to focus on similarities, whilst local (structural) processing tends to focus on dissimilarities.

Figure 2.1 – Examples of targets and distractors used by Alexander and Zelinsky (2012) to test the impact of part similarity in visual search. The caption on the top of each bear indicates how many parts it has in common with the target.



Source: Alexander and Zelinsky (2012)

There also few works investigating the role of parts in object identity, which can be considered as a form of (total) similarity. The issue is particularly interesting in the problem of how cognition trace the identity of object through time and changes. Take the classic problem of the ship of Theseus. A wooden ship goes through repair works, where each of its planks is replaced for a new plank, one after the other, until all planks are replaced and none of the original pieces of wood remains. The removed pieces are then reassembled to form another wooden ship. The question is: which of the two ships is the original one? Hall (1998) carried out experiments to test people’s choices (across many

ages) in a problem similar to ship of Theseus. Furthermore, in a subset of these experiments, the replacement parts had different colours than the original parts. The results seem to point out to the fact that, *in some situations*, people prefer to use continuity as a criteria for identify. Rips, Blok and Newman (2006) show that there is correlation between identity and similarity in some tasks, even though high similarity does not always imply in identity. These results transpire that whilst perceptual similarity (or similarity in general) does not always influence identity, similarity plays a role in certain contexts. In an interesting review, Scholl (2007) discusses the three foundational aspects of object persistence usually addressed in experimental psychology and philosophy, namely spatiotemporal continuity, property change and cohesion (maintaining a boundary); and the experimental evidence supporting them (e.g., XU, 1997; SPELKE et al., 1995, etc.). While discounting property change as having a marginal role in object persistence, Scholl concludes that the usual procedures in experiments testing how people deal with object identity are usually too alien from day-to-day experience (such as the problem of the ship of Theseus and its parts). He argues that in some cases the experiments tells less about how people trace identity, and more about how people react to questioning in unusual situations. More importantly, Scholl argues that identity is probably a result of the subject's experience in the world, rather than a result of some metaphysical property about identity. While we do not discount the specific roles of spatiotemporal continuity, property change and cohesion (and other phenomena, e.g. WAXMAN; MARKOW, 1995) in identity and similarity, in this thesis we shall focus on perceptual similarity (and related mechanisms, such as holistic and structural processing) as a foundation for identity.

2.2 Part-Whole Relations in Linguistics

Studying language is perhaps the most direct way to understand how humans deal with parts and wholes. Linguists have studied the semantics of the different expression of part-whole relations in language. Such expressions are usually referred to as meronymic expressions. Studies in linguistics of part whole-relation distinguishes themselves from studies in pure cognition mainly by their focus: research in linguistics focuses in language and how it is used, while studies in cognition of part-whole relations focus in the underlying mechanisms supporting part-whole reasoning.

In an influential work, Winston, Chaffin and Herrmann (1987) argue for different kinds of meronymic (part-whole) relations in language, describing a collection of criteria for distinguishing among them. Meronymic relations differ with respect to three properties: a part can be functional/nonfunctional to its whole; parts can be homeomerous/nonhomeomerous; and parts can be separable/inseparable from their wholes. Parts are *functional* if, given their function, they have particular spatial or temporal displacements in the whole. This is the case of the *handle-cup* relation. Parts are *homeomerous* if

they are of the same kind of their wholes, such as in *slice-pie* relation. Parts are *separable* if they can, in principle, be separated from their wholes. For example, the relation *handle-cup* is separable, while *steel-bike* is not. The combination of these properties results in Table 2.1.

Table 2.1 – Types of meronymic relation according the three differentiating criteria.

Relation	Examples	Functional	Homeomorous	Separable
Component-Integral object	handle-cup, punchline-joke	+	-	+
Member-Collection	tree-forest, card deck	-	-	+
Portion-Mass	slice-pie, grain-salt	-	+	+
Stuff-Object	gin-martini, steel-bike	-	-	-
Feature-Activity	paying-shopping, dating-adolescence	+	-	-
Place-Area	Miami-Florida, oasis-desert	-	+	-


Source: the authors.

Component-integral object relations invoke the more prototypical sense of part relation. It involves wholes that have a given internal structure, wherein parts fill a certain specific (functional) role. These include physical and abstract objects. It is important to emphasize that, according to Winston, Chaffin and Herrmann (1987), not any piece of an integral object is a component. Hacking a computer to pieces would not produce components, but only pieces, which belong to a different category of part relation. *Member-collection* describes relations where the parts have no specific function or arrangement in the whole. *Portion-mass* relations are the typical example of part-whole homeomorosity, where the parts are similar to each other and to its whole. *Stuff-object* relations correspond to expressions such as “the bike is partly steel.” *Feature-activity* relations are similar to component-integral object relations, with the difference that it applies to object existing across time, such as events and processes. Finally, *place-area* relations are similar to portion-mass, differing with respect to their lack of separability from the whole. Interestingly, Winston, Chaffin and Herrmann (1987) classifies meronymic relations as subtypes of inclusion relation, which also includes class membership and spatial inclusion relation.

Gerstl and Pribbenow (1995) improved on the previous ideas of Winston, Chaffin and Herrmann (1987), proposing what they call a *common-sense* theory of part-whole relations. The authors drawn heavily from the main approaches to part-relation (cognitive, logic and linguistic), arguing that their proposal is language-independent. Nevertheless, we classify their work as essentially linguistic, for linguistic examples form the basis of their argumentation, which is also inspired by a critique on the six categories of Winston, Chaffin and Herrmann (1987). The theory divides two broad categories of part relations.

The main category consists of relations based on the inherent compositional structure of the whole. There are three relations (Figure 2.2). The *component-complex* relation comprises the case in which a heterogeneous collection composes a complex whole; parts are differentiated with respect to their spatio-temporal displacement and functional role in the whole. This is the typical notion of part, such as in *engine-car*. The *elements-collection* relations are defined between uniform collections of parts forming a whole; e.g. *ship-fleet*. The third category, *quantity-mass* correspond to cases where the whole is homogeneous. Interestingly, *masses* are wholes characterized for *not having* a compositional structure; parts are made by carving *quantities* out of the mass.

Figure 2.2 – The categories of wholes according to compositional structure.



Source: Gerstl and Pribbenow (1995).

The second broad category is based on external partitions on wholes, independent of the wholes' compositional structure. This category constitutes a different kind of part relations that is not usually found in other part-relation taxonomies. The category is further divided into *segments* and *portions*. A *segment* is a part resulting from the application of an *external scheme* on the whole. An example of external scheme is the *upper/lower scheme*, which can be applied to physical objects. For example, a house can be divided into the *upper part* and the *lower part*. *Portions*, on the other hand, are construed by selecting parts according a particular *property dimension* or feature of the whole. For instance, the utterances “the red parts of a painting” and “the beautiful parts of the song” are examples of portions.

Coming closer to ontological issues (see next section), Moltmann (1996) argues for the notion of *situated part structure* in the semantics of natural language. According to Moltmann, a whole does not consist only of parts and an ordering between them, but they are also constituted of conditions that specifies the entity as an *integrated whole* (SIMONS, 2003). Moltmann's argument is that different situations allow (or give rise) for distinct integration criteria that operate on entities, making their part structures situation-dependent.

2.3 Part-Whole Relations in Logic and Ontology

Parts and wholes have been also investigated from a formal standpoint, in logic and ontology. The motivation is usually the necessity of representing information about parts and wholes in a formal, well-characterized fashion. In this section, we review some of the research in this view, going from logic theories of parthood relation to conceptual modelling frameworks to knowledge representation.

One of the fundamental formal theory of parts is known as *extensional mereology*. It appears in two logical forms, namely Mereology by Leśniewski and The Calculus of Individuals by Leonard and Goodman. Both establish the same main definition of part relation, specifying it as an antisymmetric, reflexive and transitive binary relation between entities. Mereology and Calculus of Individuals were originally intended as a concurrent to set theory as a theoretical tool (SIMONS, 2003). Nevertheless, they have been used as basis for the ontological characterization of part relations because they allow one to construct logically well-formed theories about parts and wholes. Not surprisingly, a great portion of the subsequent developments on theories of parthood in ontology are based on extensional mereology (GUIZZARDI, 2005). This influence range from top-level ontologies, like SUMO (NILES; PEASE, 2001), to simple knowledge representation patterns (RECTOR et al., 2005).

However, the strict extensional interpretation of part relation of these theories becomes a hindrance when they are applied in representing complex domains. In general, some aspects of part relation in mereology are not intuitive from a cognitive and linguistic point of view, affecting their usefulness for concept representation. As discussed by Guizzardi (2005) and Simons (2003), there are three main issues with basic theories about mereology, summarized below:

Arbitrary fusion problem Stricter theories of mereology allow for any fusion (or sum) of arbitrary parts to be a whole. For instance, according to it, an entity having *Plato's head*, the planet *Venus* and *Bill Clinton* would be valid according to extensional mereology. While this might be a valid set, the existence of such entity does not seem make sense from a cognitive point of view.

Non-essentiality Some variations of mereology consider two objects with the same parts as being the same object. Now consider the case of a statue of a person and the clay that composes it. They are two (conceptually) different objects that happen to share the same extensional parts. Extensional mereology cannot account for such conceptualizations. Another version of the same issue is depicted in the following example. Consider the statue lose one arm. The statue in the past and present are the same entity, even having different parts. This raises the issue of *essentiality* of parts. Certain parts are essential to object, others are not. Mereology cannot account for such cases easily.

Transitivity problem Mereology defines part relation as transitive. This property might be true in some contexts, but fails in many others. The canonical counter-example is “*hand is part of person, person is part of company, therefore hand is part of company*”, which is not usually considered to be true.

Recent formal theories of parts consider these criticisms in order to propose new representation frameworks. One of the main arguments is to take into account a *unifying condition*. There is a difference between any sum of parts and the natural individuals we usually find in the world. The later ones are called *integral wholes*, entities in which parts are unified together by means of a special relation (MOLTMANN, 1996; SIMONS, 2003). This special relation is the *unifying condition*, and can range from functional contribution to spatial inclusion. For instance, the unifying condition that defines an organism as an integral whole would be “be connected to the body”. Thus, all parts connected to given organism *a* form a sum that constitutes an integral whole; i.e. the organism *a*. The sum of the parts of the organism *a* with a part of the organism *b* is not an integral whole, because, presumably, it does not fulfil any unifying condition. Thus, the notion of integral whole helps to address the arbitrary fusion problem. Interestingly, as we have seen, Linguistics corroborate with the existence of such condition (MOLTMANN, 1996).

The notion of integral whole also serves as basis for addressing the transitivity problem. Part relation in extensional mereology does not differentiate between its relata. However, linguists have shown that there exist different senses to the word part depending on the things being related. According to studies, the different “kinds” of part relation are not transitive between them (WINSTON; CHAFFIN; HERRMANN, 1987). For instance, the part relation in hand-person is different from the part relation in person-company. Therefore, the transitivity does not follow from one to the other, i.e. hand is not part of company. Guizzardi (2005) and Gangemi et al. (2001) have argued in the lines that the transitivity of part relation depends ultimately on the unification criteria; the transitivity should in general hold within sums with the same unification criteria.

The problem of the non-essentiality in some versions of mereology has been tackled by introducing modal operators to the basic mereological axioms. The details are not relevant here. It is just enough to say in *modal mereology* (SIMONS, 2003), an entity can have different parts in multiple *possible worlds*. In this view, an entity is not identified with the sum of its parts at a given time, but rather it has many part configurations in different worlds. In this setting, one can differentiate between different kinds of parts in relation to the whole, such as essential, optional and replaceable parts.

These notions have recently influenced many symbolic concept representation frameworks, in particular, ontologies. In Computer Science, ontologies are representation devices which explicitly specify the conceptualization shared by a group with respect of a give domain (STUDER; BENJAMINS; FENSEL, 1998), usually represented in a formal language. Representation of part relation in ontologies is treated sometimes in a

simple manner (e.g., RECTOR et al., 2005). However, some approaches are more elaborated. Neumann and Moller (2008) propose a representation construct called *aggregate*, which allows one to represent and reason about concepts and their parts, satisfying certain constraints. More fundamental work has been done in foundation ontologies and top-ontologies, usually intended to serve as basis for more domain-specific ontologies. For instance, the *Unified Foundation Ontology* (GUIZZARDI, 2005), or UFO, proposes four types of part relation based on what kinds of entities they relate. For instance, the *subCollectionOf* relation can only relate collectives, while *memberOf* can only relate singular entities to collectives. The SUMO (NILES; PEASE, 2001) and DOLCE (BORGO; MASOLO, 2010) provide a similar framework of different types of restricted part relations. ONTOCLEAN, a framework for evaluation ontologies (GUARINO; WELTY, 2009), defines that wholes must have *unity criteria*, which should allow one to tell at least what is not part of that object. Based on it, the framework imposes some rules, such that if a concept carries unity criteria, its subconcepts must carry the same criteria.

2.4 Part-Whole Relations in Computer Science

Algorithms and computer systems have often to deal with structure and partonomic relations. To review all possible alternatives would be a daunting task, if not impossible. However, in certain tasks, part-whole relations are essential to the problem solving. In this section, we review computer systems where part-whole processing is a central issue.


In general, we can list two aspects of partonomic processing in computer systems. The first aspect concerns how parts and wholes are represented in computer system. The second is how the relation itself is represented. Developers frequently employ symbolic formalisms, such as UML (RUMBAUGH; JACOBSON; BOOCH, 2004) and ER (CHEN, 1976), to model information systems. Such modelling languages represent part relations as symbolic links between entities, as any other relation, with little special treatment. Developers translate these representations to software in different ways, but still keeping the symbolic nature. Consider a part relation $A < B$ between two concepts. In object-oriented software, where A and B are represented as classes, the part-relation can be very well represented as an attribute of the class B (e.g., GUIZZARDI; FALBO; FILHO, 2002). In a relational databases, where A and B are tables, the relation can be represented as a foreign key in B , and so on and so forth (for a more detailed proposal, see HALPER; GELLER; PERL, 1998). In logic-based models, where A and B are concepts, one can represent part relation as a binary predicate between A and B (e.g., ARTALE; KEET, 2008; DONNELLY; BITTNER, 2005) or as a set of predicates (e.g., PADGHAM; LAMBRIX, 1994). The software and the user interface using the information usually give the meaning of these representations. Some systems give special treatment to the part relation (e.g.,

MOTSCHNIG-PITRIK, 1993), some incorporating notions of mereology (e.g., BENEVIDES; GUIZZARDI, 2009; GUIZZARDI, 2005), as seen in Section 2.3. Nevertheless, the picture is still the same: *symbolic* descriptions of parts and wholes linked by special *symbolic* constructs denoting different part relations.

The specifics of part representation seem to be more important to computer vision systems. This is not a surprise, given the importance of structural information to object recognition in humans, as we have seen in Section 2.1. Vision systems commonly employ geometrical descriptor and graph trees in order to represent part structure. Early proposals, such as the one by Marr and Nishihara (1978), employed generalized cylinders as part descriptor, joining them in a hierarchical manner to construct more complex structures.

Zhu and Yuille (1996) proposed an algorithm to extract and match the part structure of objects. Parts are represented by deformable geometrical entities called *worms* as well as deformable circles (Figure 2.3). Together, these can approximate parts of natural entities (such as animals), organized in a tree hierarchy. The matching between learned and new stimuli is carried out by comparing the hierarchy tree they generate. The matching algorithm has provisions to deal with the over-sensitivity of the skeletonization algorithm to slight variation in pose stimuli.

Figure 2.3 – Deformable worms and circles.

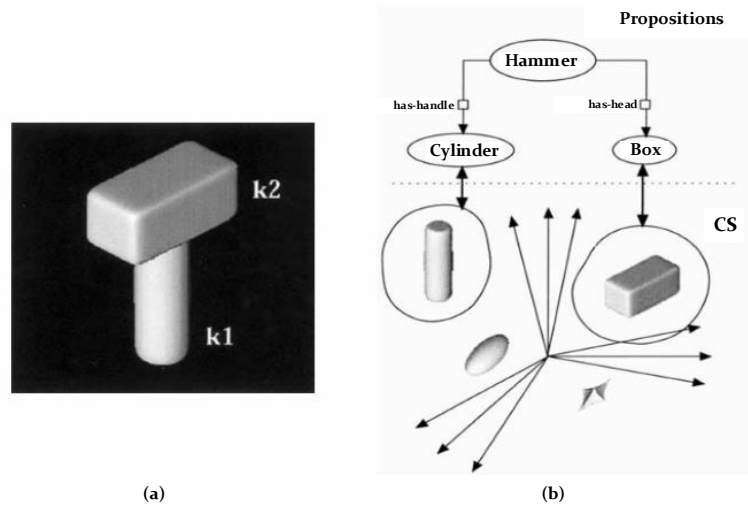


Source: Zhu and Yuille (1996).

Chella, Frixione and Gaglio (2001) took an approach to part representation that is more conceptual in nature. They propose a framework that represents parts as regions in a conceptual space of parameterized superquadrics (further discussed in Chapter 3). Such regions serve as conceptual representations of their possible properties, such as shape, position and orientation (Figure 2.4a). These predicates are mapped to predicates in a different level, where a first order-like language can be used to logically relate these parts (Figure 2.4b).

An existing technique for image classification and object recognition is the *bag-of-visual-words* (JIANG et al., 2010; DESELAERS; PIMENIDIS; NEY, 2008). Images are redescribed in terms of local descriptors encoding local features. There are many used types of local descriptors, such as SIFT (LOWE, 2004). A related technique is the *bag-of-features* (NOWAK; JURIE; TRIGGS, 2006; AGARWAL; AWAN; ROTH, 2004;

Figure 2.4 – Chella et al. represent parts as regions in a conceptual space. Points in this space correspond to specific volumes (called *knoxels*) and regions correspond to classes of volumes. Subfigure (a) depicts two *knoxels* forming a hammer. Part relation is specified in a symbolic level, such as depicted in Subfigure (b).




Source: adapted from Chella, Frixione and Gaglio (2001).

BAR-HILLEL; WEINSHALL, 2008), where local descriptors are image patches, closer to the actual notion of part. An example of this technique is the one by Agarwal, Awan and Roth (2004), which extracts relevant image patches from images of a particular object, such as car. These image patches are clustered, forming a vocabulary of part types (Figure 2.5a). Each image in a training set is represented as a feature vector describing the image by the occurrence of exemplars of part types and their relationship (orientation and distance)(Figure 2.5b). The authors then train a classifier on the set of feature vectors, such that it can take a new image redescribed as a feature vector and tell if a car is depicted. This classifier represents the concept of car. A similar approach by Murphy et al. (MURPHY et al., 2006) also includes holistic information to help object detection.


Some approaches represent wholes and parts in increasing levels of complexity. For instance, Fidler, Boben and Leonardis (2008) propose a multilevel representation where parts (segments of contour) in one complexity level is represented as a combination of the parts in the previous layer (Figure 2.6). The first layer is composed by simple Gabor filters encoding very simple contour primitives. High-order parts are spatial compositions of lower-order parts. Concepts are mapped to elements in all levels. The rationale is that direct parts of a concept exist in all levels. For instance, the wheels of a bike, having a simple shape, are already recognized in the lower levels of the hierarchy, while other parts are recognized only later given their more complex shapes. Thus, to avoid replicating simple parts to higher levels, a single concept can map to parts in all levels.

Figure 2.5 – Parts as image patches. (a) depicts clusters of image patches representing different parts of cars. (b) depicts the slicing of an image (left) into relevant image patches (middle) and its re-encoding using the vocabulary in (a).



Source: adapted from Agarwal, Awan and Roth (2004).

Figure 2.6 – Depiction of a multilevel representation of parts. Elements in each level correspond to a composition of elements in lower levels.



Source: Fidler, Boben and Leonardis (2008)

2.5 Summary and final remarks

This chapter reviews the main perspectives on representing parts, wholes and their relations. The perspectives are considerably diverse; nevertheless, we can highlight some conclusions regarding each perspective and the global view.

With respect to the evidence coming from cognition, it unveils two important aspects of part-whole representation by humans. First, information about parts and wholes are processed separately, in structural and holistic processes, respectively. Second, there is evidence to support the presence of similarity processing in holistic and structural processes. Consequently, one is to expect that a cognitive-inspired concept representation framework should account for such phenomena in some way.

Research in linguistics provides support for the idea that the semantics of “part of” term is varied and complex (as found out by Chaffin, Herrmann and Winston (1988) in empirical testing). Perhaps, such variability should indicate that natural language examples should not be taken at face value as motivation for proposing constructs in concept representation frameworks. Primarily, the way we deal with natural language can be seen as a reflection of cognitive phenomena. Nevertheless, this reflection might not be complete. As such, representation of part relations should take into account evidence of mechanisms both in Linguistics *and* in Cognitive Science. Accordingly, the relation between cognitive and linguistic phenomena regarding part-relation remains to be better explained.

Whilst logico-ontological approaches deal with a certain notion of part relation, it is not entirely clear whether Mereology reflects the way humans think and communicate about parts and wholes. Surely, that was not the case in *Classical* Mereology; but that might not be the case even for more recent approaches as well. There is a natural plasticity in the way humans deal with part relations (and other conceptual constructs, for that matter) that escapes the scope of such frameworks. The cause of this detachment might be as well be caused by differences in scope: Formal Mereology is usually concerned with what things *are*, rather than how they are reflected in the mind of the thinker.

Computational approaches are influenced by diverse theories, including the previous ones. Nevertheless, they also make a class of their own, mainly because they are pragmatic. Also, they impose pragmatic constraints in how other theories about parts and wholes can be put in actual use. It seems that hierarchy and part encoding are important aspects. However, few approaches try to give a more conceptual treatment to part-whole relations. In particular, it seems that similarity measurements commonly take place at subconceptual level, in lower levels of abstraction and before concepts can actually be used. We believe that such approaches would benefit from the added semantics if more conceptual information about wholes and parts could be used in similarity comparisons.

From the previous analysis, we conclude that a new framework for representing part-whole relations should have the following characteristics. It should represent part-whole relations taking into consideration how holistic and structural processing occurs in cognition. For instance, such framework should take into account the separation between holistic and structural processing and how similarity relates to both. Also, the framework should take into account the existence of different part relations evidenced by linguistics. Preferably, the framework should also consider the formal aspects of part-relations as presented by mereology. Perhaps, more importantly, the framework should complement the existing techniques for part-whole processing in computer systems.

3 CONCEPTUAL SPACES

The *theory of conceptual spaces* (GÄRDENFORS, 2000) describes a framework for representing concepts using *geometrical* and *topological* structures, in the tradition of other geometrical concept representation proposals, such as (SHEPARD, 1987). It has been employed in works ranging from computer science and robotics (CHELLA; FRIXIONE; GAGLIO, 2001; ADAMS; RAUBAL, 2009a; FIORINI; ABEL; SCHERER, 2013) to philosophy of science (GÄRDENFORS; ZENKER, 2011). A rationale for proposing conceptual spaces is that concept similarity is essential to understanding concept formation. The theory complements two other major approaches to concept representation: symbolic (logical) and associationist (connectionist), supplying an intermediate representation level.


3.1 Basic Notions

A conceptual space is a multidimensional space where concepts are projected and similarities represented. It can be understood as a space in the mathematical sense, such as a Euclidean space. Concepts correspond to regions in a conceptual space, whereas instances (objects) correspond to points (or, equivalently, vectors). If the space is provided with a metric, concepts and instances can be compared. Similarity between concepts and instances can then be defined as a function of their distance in the space. In the following discussion, we assume that the conceptual spaces are metric.

The representational power of conceptual spaces depends on the selection of the dimensions of the space for an application area. The *quality dimensions*, as they are called, represent different ways in which instances and classes in the space can be compared. A canonical example is the colour space that contains three dimensions: hue, saturation, and brightness (Figure 3.1). The combination of these three dimensions represent the space of perceivable colours. In this space, each point represents a particular colour. In everyday life, we do not refer to colours with such precision; we use general labels instead, like “red” and “yellow”, to refer to distinct sets of shades in which members look sufficiently similar to each other to be referred to by the same label. Geometrically speaking,

the concept of “yellow” corresponds to a particular convex region of the colour space. However, different languages carve up the colour space in different ways. Jäger (2008) has provided strong empirical support from more than 100 languages for the convexity of colour concepts. Once convexity of concept regions is required, it becomes natural to define prototypical instances as points that are central to a region. For example, the focal red that can be experimentally identified will be at the centre of the region representing the concept of red.

Figure 3.1 – Colour spindle as the colour space.



Source: National Instruments¹.

Conceptual spaces also introduce the notion of *quality domains*. A quality domain is a group of *integral* dimensions. Quality dimensions are integral when one cannot assign to an object a value in a given dimension without giving the same object a value in the other(s) (GARNER, 1974; MADDOX, 1992; MELARA, 1992). A coloured object cannot be given a hue value without also giving it brightness value; the pitch of a sound always goes with a certain loudness. Dimensions that are not integral are separable: for example, the size and hue dimensions. Using this distinction, we define a domain as a set of integral dimensions separable from all other dimensions. The three colour dimensions constitute a prime example of a domain in this sense: hue, saturation, and brightness are integral dimensions separable from all other quality dimensions. Other examples of domains are the taste domain and the shape domain.

Concepts defined exclusively as a region within a single domain are called properties. Gärdenfors define properties according to the *Criterion P*:

¹<http://zone.ni.com/reference/en-XX/help/372916P-01/nivisionconcepts/color_spectrum/>

Definition 3.1 (Criterion P (GÄRDENFORS, 2000)). A natural property is a convex region of a domain in a conceptual space.

The term natural means here properties that are natural for the purposes of cognitive processes, such as problem-solving, planning, memorizing, communication and so on. For example, “yellow” and “red” are properties, since they are single regions defined in a single domain, i.e., the colour space. These contrasts with “unnatural” properties. For instance, consider BlueYellow property, which is the result of the union of the Blue and Yellow regions in the colour domain. So, BlueYellow characterize objects that are already characterized as Blue or Yellow. The BlueYellow property does not make cognitive sense, as it is not useful for the general cognitive processes; e.g., why one would memorize a yellow object as having the BlueYellow property when it can be memorized simply as having the Yellow property. The same conclusion can be reached by evaluating the convexity of the region corresponding to BlueYellow in the colour domain; it is not convex, therefore it is not natural.

More complex concepts span regions in different domains. According to Gärdenfors, such concepts are defined according the *Criterion C*:

Definition 3.2 (Criterion C (GÄRDENFORS, 2000)). A natural concept is represented as a set of regions in a number of domains together with assignments of salience weights to the domains and information about how the regions in the different domains are correlated.

The Apple concept is a good example (Table 3.1): it comprises regions in domains like colour (red, green), taste, shape (cycloid), texture, smell, and nutrition.

Table 3.1 – The concept Apple.

Domain	Region
Colour	Red-Yellow-Green
Shape	Cycloid
Texture	Smooth
Taste	Regions of the sweet and sour dimensions
Fruit	Specification of seed structure, flesh and peel type, etc. according to the principles of pomology
Nutrition	Classes of values of sugar content, vitamins, fibers, etc

Source: adapted from Gärdenfors (2000).

The notion of salience weights in domains adds a new element to conceptual spaces. By giving weights to different dimensions/domains in similarity comparisons, it is possible to represent *context*. For instance, when comparing apples for eating, the taste domain might have more weight than the colour domain.

Gärdenfors (2000) has shown that conceptual spaces defined in terms of Criterion P and Criterion C explain cognitive phenomena such as concept combination, learning, inductive reasoning and metaphors. Since these aspects are not directly relevant to this thesis, we shall skip them.

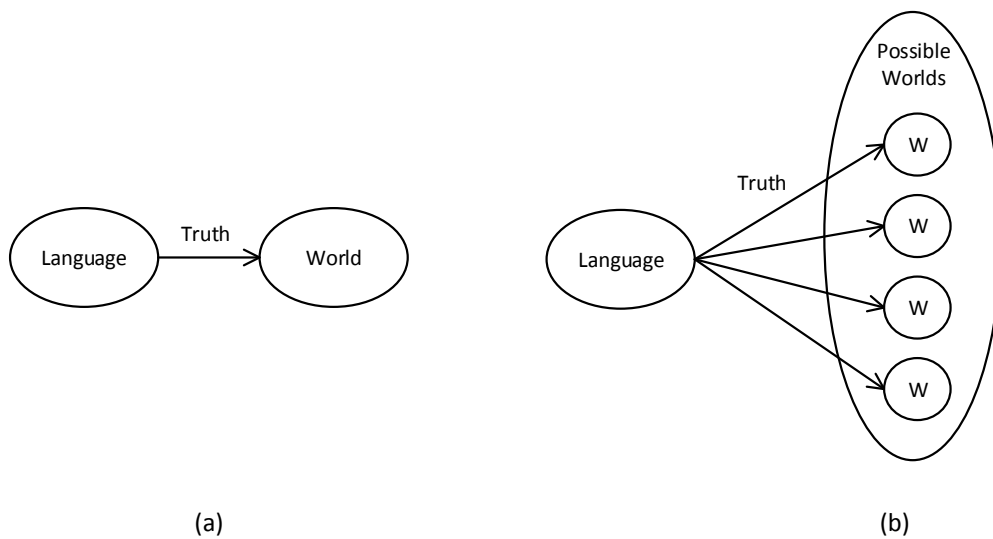
3.2 Cognitive semantics

Perhaps the theory of conceptual spaces can be better put in context by looking at it as a framework for language semantics; and then checking how it contrasts with other frameworks. The first question that arises when considering a theory of semantics is *what meaning is*. This is the ontological question for semantics (GÄRDENFORS, 2000). At first, semantics can be divided in *referential* and not *referential*. Referential semantics assume there is some kind of objects that constitute the meaning of linguistic expressions. It can give two answers to the ontological question. The *realist* answer states that the meaning of a word or expression is something out there *in the world*. On the other hand, the *cognitivist* (or *conceptualist*) answer states that the meanings are mental entities of some sort. Among the nonreferential semantics, there is also the functionalist approach, which would not be relevant here.

Realist semantics comes in two flavors: *extensional* and *intensional*. The extensional type of semantics defines it as mappings from linguistic sentences to objects in a “world” (Figure 3.2a). Constants are mapped onto objects; predicates are mapped onto classes of objects or relations between objects. Composing these mappings allows one to map sentences to truth values. The *truth conditions* for sentences determine how the world should look like for the sentence to be true. For instance, the semantics of the term “cat” is mapped to the class of all cats in the world and the semantics of the term “Felix” is mapped to a single cat. It allows one to verify the truth value of the sentence “Felix is a cat” by checking the membership of the object mapped by “Felix” to the class of objects mapped by “cat”. The idea of such schema is that it defines semantics of linguistic constructs independent of the actual speakers.

Extensional semantics has shortcomings, however. The main issue is that the semantics of sentences depend on the *present* state of affairs. For instance, the semantics of the term “cat” is dependent on the current existing cats. Therefore, its semantic would change as cats are born and die, changing the class to which the term refers. Since this is unacceptable in most cases, intensional semantics has been developed to account for these issues. In intensional semantics, linguistic expressions are mapped to objects and classes of objects in *possible worlds* (Figure 3.2b). Possible worlds can be related to “possible state of affairs”. Thus, the meaning of a sentence is given by all possible worlds where the


Figure 3.2 – Graphical schema of realist semantics: (a) extensional semantics; and (b) intensional semantics.



Source: adapted from Gärdenfors (2000).

sentence is made true. For instance, the term “cat” maps to different classes of objects that are cats in different worlds; some cats exist in some worlds, others do not. Such mappings are related to the notion of *ontological commitment* of a language to a conceptualization (see GUARINO, 1998).

Figure 3.3 – Graphical schema of cognitive semantics.



Source: adapted from Gärdenfors (2000)

Alternative to realist semantics, the *cognitive* approach to *semantics* states that the meaning of linguistic expressions is given by mappings from these entities to mental entities in the heads of the language users. The language itself is seen as part of the agent's cognitive structure (Figure 3.3). These mental entities form a *conceptual structure*, which, "via successful and less successful interactions with the world", become attuned to reality (GÄRDENFORS, 2000).

Gärdenfors (2000) argues that conceptual spaces are the appropriate framework for the cognitive structure. He justifies that by proposing *six tenets* of conceptual semantics and how conceptual spaces addresses them.

Tenet 3.1. Meaning is a conceptual structure in a cognitive system.

This tenet puts cognitive semantics in contrast with other semantic theories in that it does not require meaning to refer to reality; it requires only a cognitive structure in the agent's mind. As such, the slogan of cognitive semantics is *meanings are in the head*. The fact that cognitive semantics puts meaning before truth-conditions is the main contrast between it and realist semantics, as truth concerns just the relation between the conceptual structure and the world.

Tenet 3.2. Conceptual structures are embodied.

Cognitive structures are, at least partially, perceptually grounded. Gärdenfors (2000) argues even further, saying that conceptual structures are embodied; that is, grounded in bodily experiences and emotions. Conceptual spaces are appropriate for this purpose since quality domains and its basic dimensions are well suited to be directly grounded in perception and kinaesthetic.


Tenet 3.3. Semantic elements are constructed from geometrical or topological structures.

This third tenet contrasts with symbolic approaches to semantics. Conceptual spaces fulfils this tenet by defining conceptual constructs based on similarity measurements.

Tenet 3.4. Cognitive models are primarily image-schematic. Image schemas are transformed by metaphoric and metonymic operations.

This tenet contrasts with common propositional frameworks usual in realist semantics. On the other hand, image schemas are the most important structure in cognitive semantics. It is common to assume that image schemas are a form of representation usually linked to perception, memory and semantic meaning. For instance, Langacker (2002) define the meaning of "across" in accordance with the image schema depicted in Figure 3.4. This image schema depicts a "trajector" in different relation with a "landmark" (outside, inside, outside). It also has one time dimension and two spatial dimensions (represented by the sequence of snapshots along time).

Figure 3.4 – An image schema for “across”.



Source: Langacker (2002).

Conceptual spaces helps in giving a more precise definition to image schemas. For instance, Chella, Frixione and Gaglio (2003) suggests that types of physical movements can be represented in conceptual spaces by regions in the space generated by decomposing trajectories in its Fourier components.

Tenet 3.5. Semantics is primary to syntax and partly determines it.

In cognitive semantics, syntax cannot be described independently from semantics. Semantics exist before syntactic structures are fully developed.

Tenet 3.6. Concepts show prototype effects.

This contrasts with the Aristotelian paradigm based on necessary and sufficient conditions that define concepts. Prototypes shifts the focus of defining concepts from their conceptual borders (i.e. as in necessary and sufficient conditions) to their centre, which seems more plausible from a cognitive perspective (ROSCH, 1978).

3.3 Mathematical formulations

The theory of conceptual spaces is just a meta-theory. It lacks the necessary degree of formalization and completeness to be a fully-fledged concept representation framework. As such, there are many logic/mathematical formulations of conceptual spaces in the literature proposing different ways for filling the gaps. They mainly complement conceptual spaces with the necessary apparatus in order to make it usable in practical applications. In the following, we summarize the main approaches existent in the literature.

3.3.1 Proposals by Aisbett, Gibbon and Rickard

Aisbett and Gibbon (2001b) proposed a metric spaces-based formulation to conceptual spaces, encompassing both the symbolic and the conceptual levels. The core of Aisbett and Gibbon’s proposal is the idea that similarity induces a metric space of concepts and individuals. The authors also introduce a more general way of defining the topology of conceptual spaces based on the *betweenness relation*. Based on that, they propose two

complementary spaces. The first one is the actual *conceptual space*, which is unique and based on a definite set of dimensions. It can be arranged in levels to represent composite concepts, where each level is a copy of the primitive base space. For instance, the concept of “car pulling a trailer” is represented by a two-level space, where one-level represent the concept of car and the other the concept of trailer. The second space is a *symbols space* where symbols are mapped to properties in conceptual space by means of a symbol space. It resembles a feature space in which each property is mapped by more than one symbol in a fuzzy fashion. Interestingly, Aisbett and Gibbon (2001b) also put forward a dynamic system that employ their framework as representation formalism. This system can dynamically update its conceptual landscape based on new incoming stimuli until it reaches a state of equilibrium. We shall come back to Aisbett and Gibbon (2001b) framework in more detail in Chapter 5.

Rickard (2006) propose to represent concepts as instances of correlation matrices describing the connection between properties (as in Criterion P). These *connection matrices* represents how often two properties appear in the set of objects object. Thus, if a given application has N properties (in a given number of domains), a concept is a point in a N^2 -dimensional unit space representing a certain connection structure (e.g. graph) of properties, where each dimension is a degree of correlation between two properties. This definition corresponds neatly with the Criterion C. Regions in this space correspond to certain meta-structures in the domain knowledge, such as types of concepts. Also, Rickard interprets concept points as fuzzy sets, which allows for fuzzy tools to analyse conceptual similarity.

Rickard, Aisbett and Gibbon (2007) unifies both formulations of Rickard (2006) and Aisbett and Gibbon (2001b) to create a framework where domains are sets equipped with a measure and properties are measurable membership functions on domains. Concepts are points in a property correlation space. This scheme allows one to isolate the more discriminative properties in concept and focus attention to them in classification tasks.

In another proposal by Aisbett and Gibbon (2001a), the authors describe voltage maps as construct for conceptual spaces, motivated by topographic representations in cortical layers. A conceptual space is characterized by a set of voltage maps, or images, such that an object is single image and the distance is the energy required to change one image into the other.

3.3.2 Proposals by Raubal and Adams

Raubal (2004) gave an initial formulation to conceptual spaces based on multidimensional vector spaces. The correspondence is quite straightforward. Quality dimensions are dimensions of the vector space, quality domains are sets of dimensions and points are individual vectors. One interesting aspect of this formulation is that distance (i.e. similarity) is measured on *standardized* dimensions, which is obtained by *z-transformations* (BAHRENBURG; GIESE; NIPPER, 1999). In this way, dimensions with bigger magnitude units are less likely to dominate similarity calculations.

Later on, Adams and Raubal (2009b) presented a *metric conceptual spaces algebra*. It develops on the work by Raubal (2004) in a more complete framework. It supports the representation of basic conceptual space constructs (e.g., dimensions, domains, context, properties, concept and objects). Dimensions can exhibit different scales and limits, and can also be circular. Regions in domains and dimensions are construed by set of points that defines the convex hull of the region; or as convex *polytopes*, sets of linear inequalities defining regions. Based on this formulation, Adams and Raubal (2009b) propose a collection of algebraic operations, such as intersection, inclusion test, and metric distance; and query operations, such as similarity and concept combinations involving normal concepts and contrast classes.

Based on the metric conceptual space algebra, Adams and Raubal (2009a) proposed the *Conceptual Space Markup Language* (CSML), intended as web representation language for the constructs proposed in Raubal (2004). It is intended for information sharing among computer systems employing conceptual spaces. One of its main features is the use of *Uniform Resource Identifiers* (URI) to identify entities in the model. This allows the integration of conceptual spaces models to conceptual models represented in other XML-based knowledge representation languages, such as OWL (e.g. FIORINI; ABEL; SCHERER, 2013).

3.3.3 Formal Ontology

Some formal ontologies also provided a model for conceptual spaces. This has been pioneered by Gangemi et al. (2002) in the DOLCE top-ontology. Conceptual spaces have been used as basic constructs to describe the values of qualities attributed to object (e.g. the value of the colour of a car). However, DOLCE does not make it clear what is the relationship between its constructs and criterions C and P. We shall come back to these aspects of DOLCE in Section 4.7. The Unified Foundation Ontology (UFO) proposed by Guizzardi (2005) went a step further, extending DOLCE's proposal to account for quality dimensions and domains. However, as pointed out by Fiorini, Abel and Scherer (2013), UFO still does not account explicitly for the role of conceptual regions. Both DOLCE and UFO do not provide a complete explanation of how such qualities behave in mereological relations between concepts.

3.3.4 Other formulations

There are other, simple alternative formulations to conceptual spaces. Augello et al. (2013) argues for a *current conceptual space*, which encapsulates the notion of agent attention on its current mental representation. The current conceptual space is a combination of bits and pieces of the agent's other conceptual spaces. The authors then propose an algebra for manipulating conceptual spaces in order to derive the current space. It is similar to some of our own ideas presented in Chapter 5; however, we do not assume a vector space framework for representing spaces.

The *conceptual space logic* is an attempt by Nilsson (1999) to create a logic-based axiomatization to conceptual spaces. It resembles certain logics, such as Description Logics (BAADER et al., 2003), but it switches the usual set-theoretic model semantics to one based on conceptual space constructs. More importantly, it provides a symbolic language to talk about conceptual spaces and conceptual operations.

There also some formulations that are particular to applications. We review such applications in the next section.

3.4 Applications


Along the last fifteen years there has been many applications of conceptual spaces, in a variety of domains, ranging from philosophy of science (GÄRDENFORS; ZENKER, 2011) to artificial creativity systems (FORTH; WIGGINS; McLean, 2010). We are mainly concerned in this thesis with applications in computer-based systems.

Possibly the main application of conceptual spaces is in robotics. Conceptual spaces serve for two main purposes in robots: (a) providing the link between high-level and low-level information; and (b) providing an organized perspective on perceptual stimuli (i.e., sensor information).

Possibly the main author in the field, Chella (CHELLA; FRIXIONE; GAGLIO, 1997) proposed a cognitive architecture whereupon conceptual spaces occupy the centre spot (Figure 3.5). It serves as the link between sensor information and symbolic data, representing, for instance, shape information. Perception algorithms (including algorithms based on neural networks) in the robots are able to interpret certain aspects of world. The output of the interpretation is described by points in a conceptual space. These points are mapped to predicates in the symbolic level, which has the role of describing the objects as wholes.

Beyond its main proposal, Chella and his colleagues added extensions to their framework. In (CHELLA; DINDO; INFANTINO, 2006), the authors added dynamics to conceptual spaces in order to support imitation learning in robots, proposing a scheme for movement representation in conceptual spaces.

Figure 3.5 – Simplified version of Chella’s cognitive architecture.



Source: Chella, Frixione and Gaglio (2001).

Robots can align their conceptualization of the world through conceptual space transformations. There are many applications involving agent communication. For instance, Gärdenfors and Williams (2008) proposed a conceptual representation scheme as basis for robot communication in *Robocup*. Dimensions are football parameters, such as score, number of player defending or attacking, and so on. Region in this space represent situations in the games, such as attacking, defending, goal, etc. These can then be used for shared planning.

Similarly, Kira (2009) proposed a technique for concept sharing among heterogeneous robots. Such robots have distinct perceptual capabilities, what makes conceptual alignment a challenge. Each robot maintains a set symbols mapped to regions in conceptual spaces. These regions are represented by Gaussian Mixture Models, adding uncertainty information to the representation. The technique consists in showing the same scene to distinct robots and checking which property labels co-occur. This information is then used to find correspondence between properties in different robots in order to align representations. This allows, for instance, for a robot representing colour in the RGB space to align its conceptual representation of “red” with the one of a second robot representing colour in the HSV space.

There are related applications also involving sensor fusion. Paola et al. (2009) proposed a conceptual spaces application in Ambient Intelligence as a tool for abstracting data from different sensors around the environment. The system uses the data to act upon the environment, aiming to bring the state of the system closer to the desires of the user. LeBlanc and Saffiotti (2008) proposed a similar system;

Another type of application of conceptual spaces as a framework for information alignment is the work by (DIETZE; GUGLIOTTA; DOMINGUE, 2009). The problem consists in aligning the different semantic web services ascribe to distinct terms. The authors propose a *mediation space* in which the terms are grounded and then can be used to verify if their meanings intersect. Using notions such as similarity, the system can then find the best service match for a given request.

Conceptual spaces also find applications in geographic information systems. For instance, Adams and Raubal (2009b) propose conceptual spaces for doing classification and query of geographic landmarks and other entities. The authors propose the use of contrast classes as form of helping the system to differentiate between concepts such as Tall House and Tall Mountain.

3.5 Summary and final remarks

This chapter introduced the theory of conceptual spaces and reviewed the main mathematical formulations to the theory and its use in computer science applications.

The conclusions we can take from the reviewed literature is that both formulations and applications are still fragmented. Take the formulations presented in Section 3.3, for instance. As one would expect, they overlap in certain aspects, but are still incompatible. This situation probably arises as result of how extensive the theory is in its proposals (i.e., addressing aspects ranging from perceptions to metaphors and inductive reasoning) and the different needs imposed on conceptual spaces by the distinct authors. This situation brings forward some questions: what should be the precise core axioms of the conceptual spaces theory? Also is there an axiomatic core common to all *interpretations* of the theory? Assuming a core exists, is there a way of organizing the many additional features of the theory in different levels of complexity? While we do not answer to these questions in the thesis, similarly to how the different description logics are organized regarding their expressiveness, we envisage a set of conceptual spaces formulations organized in different “idioms”, having different conceptual capabilities.

While some works propose sets of *conceptual operations* in conceptual spaces, they still lack a common interpretation beyond the basic ones described by Gärdenfors (2000), such as intersection, product and so on. An interesting research venue would be to develop and standardize more types of cognitive operations in conceptual spaces.

Finally, an aspect that is lacking in the theory itself regards the representation of relations. Gärdenfors (2000) reserved just a couple of pages in his book to address the topic. He suggests relations are products of the conceptual spaces of the *relata*. While this is a good first approach to the issue, it is not satisfactory for a more comprehensive theory of concept representation. In particular, there is no indication of how to represent information about the relationship between parts and wholes of objects. We try to remedy this particular deficiency in the next chapters.

4 PART-WHOLE RELATIONS AS CONCEPTUAL SPACES

There are several types of cognitive phenomena involved in how we deal with concepts. Such processes are unsurprisingly reflected in theories about how we believe concepts should be *represented*. However, the interplay between two important aspects related to concepts, namely *conceptual similarity* and *part-whole relations*, is usually overlooked when it comes to concept representation.

The intuitive idea behind the role of similarity in categorization is that two objects belong to the same concept if they are “similar enough”. The exact definition of “similar enough” depends on the actual representation framework (EDELMAN, 1998; TVERSKY, 1977), but the general importance of similarity in categorization is well established (GOLDSTONE, 1994). In addition to similarity, concepts are thought to show prototype effects (ROSCH, 1978); concepts are defined in relation to one or more individuals that are judged to be typical exemplars of that concept (for example, a robin could be seen as a prototype for the concept of bird, in contrast to a penguin). The classification of a new object is determined by measuring its similarity to the concept prototypes.

On the other side, human cognition can also represent the relations between entities and their parts, for example, between a horse and its four legs. These relations play an important role in how humans perceive and think about concepts. One of the central questions of this thesis is how the notions of similarity and prototypes are reflected in an analysis of parts and wholes. The idea of a prototypical whole seems to be intuitive enough to make it significant: for example, it is easy to think about a prototypical *pen*, with its typical configuration of parts. The degree of typicality of other pens can be measured by their similarity to the prototype. Partonomical (part-whole) similarity between wholes takes into account which parts are actually similar and how parts are structured. While some experiments have indicated prototype effects in part-whole relations (CHAFFIN; HERRMANN; WINSTON, 1988), the specific role of prototypes in these relations is not yet well understood. Nevertheless, as we have seen in Section 2.1, there are direct and indirect evidences that partonomical similarity plays a role in object recognition and concept learning.

Our goal in this chapter is to describe a *computational* representation framework for part-whole relations, taking into consideration the cognitive mechanisms involved in representation and reasoning with parts and wholes, in particular the role of similarity and prototypes. We will not propose a specific formalism for knowledge representation; this shall be presented in the next chapter. The aim here is, instead, to present a theory that can guide the development of such formalism. Furthermore, we restrict our discussion to part relations involving physical objects only.

Our proposal follows the tradition of cognitive semantics (see Section 3.2). It differs fundamentally from the realist semantics usually employed in knowledge representation, specifically in ontologies (GUARINO, 1998). Whereas realist semantics defines meaning as mappings from language to one or more “worlds”, cognitive semantics defines meaning as mappings from language to *conceptual structures* within an agent’s mind (GÄRDENFORS, 2000). Cognitive semantics provides a more principled account of the influence of cognitive mechanisms, such as concept learning, perception, and symbol grounding. Additionally, the purely symbolic languages usually employed for representing concepts, such as the Ontology Web Language (OWL), lead to difficulties in representing part-whole relations (RECTOR et al., 2005). This is a good motivation for looking for other representational formats. In particular, the advantages of using cognitive semantics to explain meaning have been investigated in relation to the knowledge systems (e.g., the Semantic Web) (GÄRDENFORS, 2004; ADAMS; RAUBAL, 2009a). For instance, Adams and Raubal (2009a) proposed the Conceptual Space Markup Language (CSML). CSML is a XML-based representation proposed as a complement to the Semantic Web languages.

As it has been argued by Guarino, Pribbenow and Vieu (1996), there are two approaches to the problem of representing of part-whole relations. One is the *logico-philosophical approach*, which takes the perspective of formal ontology and algebraic theories of parts, such as classical mereology (the formal study about parts and wholes) and other derived theories (VARZI, 2011; SIMONS, 2003). This approach seems to be dominant among ontologists, particularly in computer science. However, Simons seems to recognize that algebra is not enough:

“When it comes to the honest toil of investigating the principles governing what objects are parts of others, and what collections of objects compose others, it appears that most ontologists have been following the paradigm of abstract algebra when it would have been better to take a lead from sciences such as geology, botany, anatomy, physiology, engineering, which deal with the real.” (SIMONS, 2006)

On the other hand, there is the *cognitive-linguistic approach*, which is the one we adopt. Here, we consider the cognitive phenomena related to concepts that are usually ignored in other approaches, like prototype effects and similarity. Furthermore, we account for the *context* effects of using concepts that frequently turn up. We submit that a semantic framework suitable for implementing intelligent computational systems should be aligned with human cognition as much as possible.


To that effect, we base our analysis on the theory of conceptual spaces (GÄRDENFORS, 2000), reviewed in Chapter 3. More precisely, we discuss the role of similarity in part-whole structures; in other words, what it means to say that two objects (or concepts) are *similar* because they have a similar part-whole structure. We present constructive proposals for modelling the conceptual structures of parts and wholes. In brief, we argue that the conceptual space of a whole can be seen as a product space, composed by a conceptual space describing its global (holistic) properties, as well as the conceptual spaces of its parts accompanied by part structural information. The notion of using product spaces to form more complex conceptual construction is not new (AISBETT; GIBBON, 2001b); however we innovate by adding more information to product spaces based on cognitive phenomena.

Part-whole relations can also take many forms, having different meanings. For instance, the part relation *engine-car* is of a different nature than *tree-forest*. Many authors have proposed diverse categorizations of these forms (CHAFFIN; HERRMANN; WINSTON, 1988; GERSTL; PRIBBENOW, 1995; GUIZZARDI, 2005; SIMONS, 2003). We analyse how some of these forms manifest themselves in the conceptual world, allowing us to account for the plasticity of part-whole relations. In addition to that, we discuss how prototype structures in a partonomical hierarchy affect object classification; and how the same whole can be seen in different ways, taking into account contextual information on the parts. We finally present a simplified model of an object recognition approach based on our conceptual framework.

4.1 Preliminaries: Bubble diagram

Usually, conceptual spaces are constructed out of many dimensions and domains. That can make their depiction very challenging. We have devised a simple diagram that emphasizes the multidimensional composition of conceptual spaces as a product of quality domains. Figure 4.1 exemplifies this diagram for representing the concept “apple”. The apple space is represented as a product space of properties (smaller ellipsoids) in the quality domains that form the conceptual space (bigger ellipsoids). This diagram is inspired on the intuitive notion that a concept in conceptual spaces can be seen as a product of

Figure 4.1 – Example of diagrams depicting the conceptual space of apple: (a) shows the inner form of the apple space as a product of properties (smaller ellipsoids) in different quality domains (bigger ellipsoids); and (b) shows a compact representation of the apple space as a set of points (smaller ellipsoid) in a multidimensional space formed by the product of its quality domains.



Source: the authors.

regions (or subspaces) in a series of quality domains (Figure 4.1a); or as a region in a multidomain space generated by the product of quality domains (Figure 4.1b). The ellipsoids and domains can be drawn in different colours and sizes to convey additional information.

4.2 Representing Parts and Wholes in Conceptual Spaces

The cognitive grounding of the relation existing between parts and wholes must be founded on a broader theory of concepts. Our aim is to show that conceptual spaces can provide the basis for such a theory. In the next sections, we describe how part relations can be founded in conceptual spaces and discuss the consequences for concept representation. The general idea is that the relation between a whole and its parts is represented in a *structure space*, where structural similarity between wholes can be measured, and prototypical wholes can be identified. We start by exploring the relation between the whole and its structure.

As we suggested in the final remarks of Chapter 2, the cognition of part-whole reasoning seems to require that the descriptions of concepts take into account holistic and structural information. Additionally, similarity effects present in categorization seem to suggest that the same descriptions should also support similarity comparisons. Bringing all this together, we can redefine concept similarity as a function of holistic and structural similarity. Intuitively,

$$\text{Concept Similarity} = \text{Holistic Similarity} \oplus \text{Structural Similarity} \quad (4.1)$$

The goal is to define a representation structure that allows for such similarity calculations.

We assume that wholes and parts have their own representational units; that is, certain properties are exclusive to wholes and some properties are exclusive to parts. More specifically, we assume that the whole and each of its parts are represented in their own, distinct¹ conceptual spaces. For example, the concept of bird is placed in a conceptual space with its own dimensions, while the concept of beak, wings and feet are placed in three other conceptual spaces, with independent dimensions and domains. There are of course correlations between properties of the whole and properties of the parts. Nonetheless, we do not assume them to be necessarily linked. For instance, the concept of black woodpecker (*Dryocopus martius*) certainly occupies the “black” region of the colour domain, as would both of its wings; its crown, however, is not correlated to the colour of the whole and is to be positioned in the “red” region. We return to this issue when discussing David Marr’s hierarchical model in Section 4.4.

The relation between conceptual spaces of wholes and parts is represented in the conceptual space of the whole. It is structured in such a way that it implements the conceptual similarity defined by the Eq. 4.1 . Accordingly, we propose a definition of the conceptual space of any whole as a *product space* of two *subspaces*: the *holistic (sub) space*, which represents the properties of the whole, allowing for holistic similarity comparisons; and a *structure (sub)space*, which represents the relation of the parts with the whole, and allow for structural similarity comparisons. Thus, we have that

$$\text{Conceptual space} = \text{Holistic space} \otimes \text{Structure space.} \quad (4.2)$$

The inner form of the holistic space is, for the present purposes, reasonably unproblematic. Holistic spaces are standard conceptual spaces with dimensions and domains describing properties about the whole. For instance, the conceptual space of apple represented in Figure 4.1 can be seen as a holistic space, for it mainly describes the properties of the whole apple. In this space, whole apples can be compared regarding their similarity.

We are, however, more interested in the inner form of the structure space, which is naturally more intricate. It has to implement structural similarity, relating parts to the whole. Therefore, we have to initially consider in more detail what structural similarity is, before unveiling its inner form. The intuitive notion behind structural similarity is that two wholes are structurally similar if they share a similar set of parts. This explanation is nevertheless incomplete. For instance, according to it, a pile of Lego bricks and the assembled Lego toy would be considered similar entities, since they share the same set of parts. Yet, the two entities are clearly structurally dissimilar: although the pile and the toy share the same types and number of parts, they do not share the same internal

¹The real meaning of the word “distinct” here depends on the actual way chosen to implement conceptual spaces. Some representations tend to aggregate all possible concepts in a unique, comprehensive conceptual space including all domains. Others may represent concepts in really distinct conceptual spaces, formed by the combination of the relevant quality domains.

structure, which is essential for distinguishing between the two. Other authors have drawn attention to this separation between part and structure (HUMMEL; BIEDERMAN, 1992). Thus, we claim that structural similarity consists of two elements: part and configuration similarity. Two individual wholes have high *part similarity* if the parts in one whole are similar to the parts in the other. It takes into consideration part categories and their number. Additionally, two individual wholes have high *configuration similarity* if parts in both wholes are also arranged in a similar way. The combination of these two kinds of similarity defines structural similarity between two wholes: similar parts placed in a similar configuration.

With the notion of structural similarity in mind, we can define in more detail what kind of construction a structure spaces have. A structural space is a conceptual space where one can represent and compare many possible part configurations. It is a high-dimensional space, where each point (or vector) corresponds to a particular part configuration and regions denote different kinds of part arrangements. It is formed by a product of a subspace of (the conceptual spaces) of each constituent part, as well as a set of properties in special quality domains called *structure domains*. Structure domains modulate how parts bear to the whole, representing, for instance, displacement information (e.g. a coordinate space centred on the whole). They are quality domains in the sense they also describe qualities of a concept. Thus, following from Eq. 4.2, we have that:

$$\text{Concept} = \text{Holistic properties} \otimes (\text{Part properties} \otimes \text{Structure properties}) \quad (4.3)$$

More formally, given a whole C and its parts P_1, P_2, \dots, P_n , a *structure space* is a subspace of C formed by the product of subspaces of P_1, P_2, \dots, P_n and supplemented by structure domains S_1, S_2, \dots, S_n to each P_i . The motivations for defining structure space essentially as a product of the parts is the necessity of finding a straightforward way of measuring structural similarity as a distance function between wholes. When we accept that a particular object must be represented by a single point in a concept, then we must commit to the notion that this single point must encode the information about the specific set of parts composing the whole (encoding part similarity) and also its structural information (encoding configuration similarity). A good way of accomplishing this is to transfer the information to the quality dimensions: information about the set of parts composing the whole is encoded by joining the quality dimensions of the parts, and the structure information is encoded by joining the structural domains for each part.

As an example, consider in the conceptual space of apple in Figure 4.2. Each part (stem, seed, skin and flesh) is defined as a set of properties in their own conceptual spaces (Figure 4.2a). In order to form the structure space of apple, subspaces of the parts are used to compose the product space that forms the structure space of apple (Figure 4.2b). Notice that the use of subspaces of the parts is due to the fact that the whole relates just to

a subset of the individuals described by the general part concept. For instance, the concept fruit seed describes all kinds of fruit seeds, given that only a fraction of them (a subspace) can be said to be apple seeds; i.e., oval dark seeds. Furthermore, for each part there is a property in a structure domain that defines the configuration information of that part (depicted as lune shapes in Figure 4.2b). In the context of this example, configuration information can be regarded as the displacement coordinates (e.g., position and orientation coordinates) of each part of the apple in relation to the whole apple.² For instance, Figure 4.2d shows a two-dimensional structure domain of coordinates centred at the apple, with the region at the top corresponding to the coordinates allowed for the position of the stem. The same kind of information is added as regions in structure domains for each part in the structure space. The end result is the complete apple structure space, which can be seen as a multidimensional conceptual space itself (Figure 4.2c). A given point p in this space corresponds to a particular structure of apple, with a specific set of parts displaced at specific places. Points in the neighbourhood of p correspond to similar apples, such as apples having a slightly different stem positioned at a slightly different place on the apple. It is important to note that structure domains allow the representation of structural information in terms of object-centred configurations. They can also be combined to describe more complex information, such as spatial relations³ (e.g., right of, back of, etc.).


The link between structure spaces and parts is governed by what we call *dimensional filters*. A filter is a higher-order structure that defines which subspace of the part that is actually used to compose the structure space of the whole. In this context, the role of a filter is two-fold. First, it can be used to filter out sections of the structure space of a part that is not relevant to the whole; i.e. as in the *stem-apple* example above. Second, it can also be used to filter quality domains of the part (holistic or even structural) that are not relevant to the whole. For instance, a combustion engine may have quality domains describing its characteristics as a car engine and as a power generator. However, the quality domains regarding power generation must be filtered out when the concept is imported in the structure space of a car. In some ways, the dimensional filter works like a context, screening out the unimportant dimensions of its parts.

The formation of concepts in the structure space is, to a large extent, determined by prototype effects. Some part structures can be seen as more typical than others. These typical individuals — not necessarily any that exist in reality — determine the focal points of the convex regions that form concepts in structure space. Take the concept of an apple as an example. Its part structure would be determined by a prototypical exemplar of its kind, denoted by a point in the structure space of apple. In turn, this prototype determines the focal points of the convex regions, which fully define the concept of apple structure.

²An object-centred coordinate system seems to be preferable. Nonetheless, one can think of structural information based on an egocentric coordinate space, or even a retinal coordinate space (NEWELL et al., 2005).

³More about conceptual spaces and spatial relations can be found in (LIGOZAT; CONDOTTA, 2005)

Figure 4.2 – Example of structure space for the apple concept: (a) the conceptual spaces of each part of apple, their inner form (dimensions and domains) is omitted; (b) the conceptual space of apple as a product of the holistic space and subregions of parts and also structure information; (c) a compact representation of the holistic and structure space of apple; and (d) an graphical depiction of the regions defining the displacement for the stem in the structure space of apple.



Source: the authors.

Furthermore, regarding the relation between the prototypical whole and the prototypical parts, it is tempting to say that the prototype of a whole is also composed by the prototypes of its parts. However, even if this assessment could be true in some cases, it fails in a high number of situations. For instance, our prototypical notion of grasshopper includes a subspace of the concept of wing that certainly does not include what we consider to be the common prototype of wing.

At this point is important to highlight that this framework does not address *what* is a part or *how* we separate parts from the whole (e.g. in perception). We are just interested in determining a way of representing the relation between parts and wholes so that holistic and structural similarity can be measured. We return to this point in the next sections.

4.3 Types of Structure Domains


Structure domains and properties are fundamental parts of our framework. They qualify the partonomical relations between parts and wholes. As such, it is natural to expect they themselves might have a complex form. As a matter of fact, the elementary notion of being part of something can be specialized into more specific types of part-whole relations. In recent years, several authors have proposed different taxonomies of part-whole relations, based on different criteria, from linguistic to ontological (WINSTON; CHAFFIN; HERRMANN, 1987; GERSTL; PRIBBENOW, 1995; GUIZZARDI, 2005; JOHANSSON, 2012). In this section, we show how common kinds of part-whole relations reflected in language can be explained by different ways of constructing structure domains.

4.3.1 Complex and Collective

Essentially, there are two clear-cut kinds of part relations. A part relation can be *collective*, when a class of instances of a certain type generically composes the whole. For example, the part relation between tree and forest just defines that an instance of forest includes an arbitrary number of trees; no particular tree is conceptualized or has a specific role or place in the forest. On the other hand, *complex* relations give more specific information about how the parts relate to the whole, such as with functional or displacement information. For example, the part relation between cat and tail specifically defines that an instance of cat has an instance of tail positioned at a particular place on the cat. These two types of relations can be found in some part-whole taxonomy proposals, such as those of Winston, Chaffin and Herrmann (1987) and Gerstl and Pribbenow (1995).

Regions in structure domains, called *structure properties*, have the exact function of defining how a part fits into a whole. Different structure domains allow for different part relations. In the previous section, we talked about structure properties that represent displacement information only. These are *complex* structure properties. Collective relations

Figure 4.3 – Structure spaces according to paronomy types: (a) fragment of the structure space of cat showing the complex relation of tail; and (b) fragment of the structure space of tree showing the collective relation of forest. Both “rear” and “many” are property regions on their respective structure domains.



Source: the authors.

can be represented in structure spaces by changing the type of structure domain used. Take the relation $R(A, B)$ as a part relation between a part concept A and its whole B . Let S be the structure space of B . Let D_r be the structure domain added by the relation R to S . Then, based on the type of R , we can establish which kind of information is carried by D_r :

- 1) R is a complex relation iff D_r is formed by complex structure domains;
- 2) R is a collective relation iff D_r is formed by collective structure domains.

In complex relations, the structure domain denotes the specific configuration or role of the part, like the allowed positions and orientations. For example, the relation $part(tail, cat)$ is complex because the relation between the cat and its tail is modulated by a complex structure property; i.e. a region in a complex structure space denoting the allowed set of positions and orientation of the tail in a cat (Figure 4.3a). In collective relations, the structure domain *quantifies* the part concept, such as how many instances of the part are components of the collective whole. For instance, our general conceptualization of a forest is that it is simply composed by many trees. As such, the relation $part(tree, forest)$ is a collective relation because the structure domain that modulates it in the structure space of forest is a *collective* structure domain. That is, it does specify only a region (i.e. interval) in a *quantification space* corresponding to how many trees a forest could have, and perhaps how they are packed (Figure 4.3b). This scheme allows the representation of individuals as “*there are thousands/many/few trees in that forest*”.

Complex and collective structure domains can have a variety of forms, depending on the implementation. Regarding *complex* structure domains, we already described an example in the last section where an object-centred position space could work as an implementation. In such a space, the displacement of parts of an object is described as a two-dimensional coordinate of each part within the whole. However, one can devise more sophisticated implementations for complex structure. Consider displacement now defined as the geometric volume in the spatial extension of the whole where a given part can be found (or seen). For instance, if a person is asked to point where the engine is usually placed in a car, she will probably point at the front of the car, drawing an ellipsoid with her finger, while saying “around there.” This ellipsoid captures the intuition of the displacement volume we are talking about and it is naturally a function of the overall shape, position and orientation of the part. This sort of construction can be neatly represented as a property region in a geometric volume domain, such as superquadrics (CHELLA; FRIXIONE; GAGLIO, 2001). A point in this region represents the specific placed volume where the part can be found in the whole (e.g. “where the engine is placed in this car”). For example, the displacement of the skin of an apple can be represented as a region encompassing the whole apple; a point in this region encodes a volume that coincides with the surface of the apple.

Collective structure domains on the other hand, represent how a *set* of parts of the same type generically related to the whole. For instance, it might be represented as simple one-dimensional spaces denoting, for instance, how many instances of that given part are expected to be found in the whole. More interestingly, we believe that collective structure domains might be able to represent the notion of *ensembles* of parts (ALVAREZ, 2011); which are related to the capacity humans have to summarize groups of similar objects into a cognitive compact average representation. In this context, we could see the relation between “tree” and “Amazon Forest” as conjunction of a conceptual space describing an average Amazonian tree and a region in a collective structure property denoting how many trees this forest has (e.g. say, “billions”, or “many”).

Perhaps not surprisingly, complex and collective relations are intrinsically correlated. Collective relations can be seen as a generalization of many complex relations to which no displacement information is necessary. This might account to what Gerstl and Pribbenow (1995) call the *plasticity* of part-whole relations. For instance, the relation *ship-fleet* can be classified into two ways. If one considers a fleet as a uniform set of ships, the relation is collective. On the other hand, if each ship has a special role in a fleet, one can consider the relation complex. We argue that there is a third case where both kinds of relations are mixed: certain ships could have special roles while others are referred to generically (e.g., as in a fleet having a carrier and many destroyers). In this case, the relation can be seen as a hybrid of collective and complex. These processes of change can be explained by processes of folding and unfolding of complex relations into collective relations.

4.3.2 External Partitions

In general, we talk about parts as the building blocks of objects, in the sense that an apple is formed by many parts. However, we can also talk about arbitrary partitionings that are imposed on the objects. For instance, if we say that “the upper part of the house is blue”, we are imposing a somewhat arbitrary partitioning of the house. In some cases, these *external partitions* can be seen as part of the definition of certain concepts, thereby composing their structure space. For instance, the concept “planet” might include the notion that the planet’s polar regions are cold; the notion of “polar region” can be seen as an external partition imposed on planet; there is no actual inner structural part that corresponds to the poles of a planet.


External partitions can be implemented in conceptual spaces by means of certain region operations (e.g., intersection) on the quality domains of wholes and their parts. For instance, the concept of “the equatorial region of a planet” translates to an intersection of a property in the shape domain of “planet” (a sphere) with the property in the shape domain of “equatorial region” (a section of a sphere). This new concept helps to compose the structure space of “planet”.

External partitions in conceptual spaces might explain a kind of parthood that recurs in the literature, namely, what Gerstl and Pribbenow (1995) call *external parthood*. They argue that certain part relations derive from the internal structure of the whole, whereas others can be said to derive from external partitions imposed on the whole. In their original proposal, Gerstl and Pribbenow (1995) defined two types of external part relations: portions and segments. A *segment* is a spatiotemporal part that results from the imposition of an external scheme on the whole. This scheme distinguishes different parts of the object, indifferent to its internal structure (“the upper part of the body”, “the beginning of a story”). On the other hand, a *portion* is construed by using a property dimension to select parts out of the whole; for example, the dimension of colour is used in phrases like “the red parts of a painting” or “the annoying parts of the evening television show”. We can explain portions and segments in conceptual spaces by relating them to intersection operations in particular kinds of domains. In brief, portions can be seen as the results of restrictions in spatiotemporal quality domains, whereas segments are the result of restrictions on other quality domains. For instance, the *pole-planet* example translates a typical case of segmentation: a restriction on the spatial domain of “planet”. On the other hand, a phrase like “the red brush strokes of the painting are original” translates the case of portioning. In a collection of all parts of a painting, this portion is where the colour property regions intersect with the colour property “red”.

4.4 Marr’s Hierarchical Model, Revisited

The basic utilization of structure spaces can be exemplified by reinterpreting influential hierarchical model of cylinders proposed by Marr and Nishihara (1978). In the literature, there are many attempts to model the shapes of objects in object recognition (MARR; NISHIHARA, 1978; BIEDERMAN, 1987; PENTLAND, 1986; ZHU; YUILLE, 1996; CHELLA; FRIXIONE; GAGLIO, 2001). Many of these attempts take into consideration the visual part-whole structure of the objects, associated with some sort of shape primitive, like cylinders or more complex parametric volumes. The model by Marr and Nishihara (1978) employs sets of cylinders to approximate biological forms, as illustrated in Figure 4.4. The cylinders are combined in a hierarchical manner, with the torso on the first level, and the head and limbs (arms) on the second, forearms on the third and so on. In the following, we demonstrate how Marr’s model can be described in our framework.

Figure 4.4 – Parts hierarchy model proposed by Marr and Nishihara. The arm is deconstructed as finer parts along the chain of part structures.



Source: adapted from Marr (1982).

Each level of the Marr’s hierarchy can be seen as a single concept (“body”, “arm”, “hand”, etc.). Each concept in this hierarchy has two types of descriptions. One represents the whole (e.g., the whole body); the other connects the parts of the hierarchy together (e.g., the limbs and head). The description of the whole is called a *model axis*: a generalized cylinder that captures the overall outline and orientation axis of a particular hierarchy level (e.g., the cylinder that represents the whole body in Figure 4.4). It abstracts away details that are usually supplied by parts. This approach is closely related to the model of holistic versus structural descriptions: the model axis represents a holistic take on the object shape at a given hierarchy level, while the same level includes more detailed structural (part) descriptions.

The recreation of Marr’s model using our framework is quite straightforward. Each level in Figure 4.4 (body, arm, forearm, etc.) can be modelled as a concept with properties in holistic and structural domains. The elementary quality domain here is *shape*. In this formulation, the shape space can be represented by two dimensions: cylinder length and radius. Different regions in this space denote different cylinder shapes. This shape domain is used to compose all concepts in this hierarchy. Each concept is then formed by a property (a region) in the shape domain, plus a structure space formed by its parts.


The structure space is formed by quality domains, imported from subparts of the body, and structure domains. In this case, the structure domains are of the *complex* kind and encode information about position and orientation of the parts in the coordinate space of the whole. For instance, the concept of “hand” in Figure 4.4 has a structure space formed by the shape domains and properties of each finger, plus structure domains defining the position and orientation allowed for each finger. A point in the conceptual space of “hand” describes a particular whole shape cylinder for the hand, plus particular shapes, orientations, and positions of the fingers. In this space, we can refer to prototypical hand shape and configurations, and categories of hand shape.

4.5 Same Whole, Different Parts

Structural similarity judgements are influenced by *context*: two apples can look very similar to a child in a supermarket, but very different to a botanist. Distinct parts have distinct importance (or “goodness”, according to Tversky (1989)) depending the context. In the theory of conceptual spaces, context is represented as different weights given to each quality dimension (or domain) in the conceptual distance measurement (GÄRDENFORS, 2000). The same method can be used to account for the influence of context in structural similarity judgements. Different weights are given to the domains of different parts in structure space: less relevant parts are given smaller weights and more relevant parts are given larger weights. In the apple example, a botanist comparing apples will give more weight to internal parts when trying to decide if a given apple is included in the “good apple” category, whereas a consumer will give more importance to external parts (Figure 4.5). As we shall argue in Section 4.7, context might be also related to the notion of essential and mandatory parts at the symbolic level.

If we define *context* as a vector of weights, it implies the existence of a *context space*. A context space is a higher-order space where points denote different combinations of weights of the quality domains in a conceptual space. Again, we believe that context space also includes prototype structures: some context situations are more typical than others. Prototypes denoting typical part-importance scenarios will complement the context space for a structure space. For instance, the same person can play the role of a botanist and consumer at different times; the present context “moves” through different categories in

Figure 4.5 – Examples of different contexts in the apple structure space, where ellipses with smaller sizes denote spaces with less weight in a given context: (a) structure space of apple for a botanist; and (b) structure space of apple for a consumer.



Source: the authors.

the context space. In an apple dissection situation, one can pay attention to internal parts of the apple; nevertheless, the prototypical apple consumer context is the one where just the more external parts of the fruit are relevant for comparisons. Situations that are close to the prototype situations define concept regions in context space; these regions can be interpreted as kinds of context, like “apple dissection” context or “apple consuming” context.

Another example situation in which context plays a role happens when characteristics of different parts influence the categorization of the whole. For instance, a boat with a black hull and white sail may be “the black boat” in the context of boats with white hulls, and “the white boat” in the context of boats with black sails. Context influences part saliency, which in turn affects categorization.

4.6 Partonomies and the problem of transitivity

People frequently describe part structures as hierarchies, where more simple objects compose more complex ones and so. These hierarchies are usually called partonomies. A partonomy is a simple tree-like representation that highlights the apparent transitive nature of part relations, allowing one to visualize and navigate different levels of deconstruction. It is frequently seen as the canonical way of representing part structures. Whilst these are useful as tools for reasoning about parts, there is no general way of measuring

similarity between graphs.⁴ Thus, it is not possible to explicitly represent part hierarchies in conceptual spaces in a simple way (e.g. as a point). At the same time, they are not necessary for comparing the similarity of objects at conceptual level, as we have seen in previous sections. Nevertheless, these hierarchies are implicit in a sequence of structure spaces represented as a chain of whole-part pairs, paronomies could possibly be derived as a symbolic construct from the chain of structure spaces if they are needed at the symbolic level.

However, we still must address the fact that parts also have parts and how it relates with structure spaces. Structure spaces import the domains of all parts. Given that parts can also be wholes, and that they have structure spaces of their own, the final structure space of a more complex whole can become a transitive closure of all its parts and sub-parts. This is not desirable from the standpoint of cognitive economy, and, at a first glance, can be seen as a fatal limitation of our framework. The mechanism that allows us to deal with this issue is related to the transitivity problem. Any discussion of representation of paronomies usually arrives at the *problem of transitivity* in part-whole relations (SIMONS, 2003; VARZI, 2006), which is mainly an ontological one. The classical formulation of mereology holds that transitivity is one of the basic properties of the part relation. However, transitivity often breaks down. A good example of this breakdown is the following: the eye is usually regarded as part of the face and the retina as part of eye; however, the retina is hardly regarded to be part of face. There are some solutions to this problem (VARZI, 2006); one of the accepted solutions is to consider ontologically distinct types of part-whole relations, which are naturally not transitive between themselves. For example, the relation between eye and face is not of the same kind as the relation between retina and eye. This general solution, however, does not help us to solve our representation problem; mainly because we do not assume any *a priori*, crisp ontological distinctions between types of wholes in the conceptual level.

In order to give a partial solution to this problem, we start by assuming that part relations are not transitive. We see few reasons to consider a *cognitive* interpretation of part relation as necessarily transitive. First, there are plenty of examples of intransitive part relations. As other authors have pointed out (e.g. JOHANSSON, 2006), part relations might have different interpretations in different contexts. For instance, the tendency to ascribe transitivity to simple part relations possibly comes from its association with the notion of spatial inclusion, which is usually transitive. Second, following the general notion of *situated part structures* (MOLTMANN, 1996), we see the formation of part relations and structure spaces as a process that is generally linked to context, experience, and perception, rather than to pure deductive reasoning. We suggest that the structure

⁴A main reason is that there is no general way of measuring similarity between three graphs. There are some attempts to do so by converting tree graphs into graph spectra (e.g., SHOKOUFANDEH et al., 2005), but the lack of isomorphism between the two kinds of representation prevents the use of graph spectra as a framework for implementing structure similarity.

space of a concept comprises all its *experienced* parts. One tenet of cognitive semantics is that conceptual structures are *embodied* (see Section 3.2). Thus, concepts are dependent on bodily experiences and emotions (GÄRDENFORS, 2000). In a broad sense, what dictates whether an object is a direct part of another is the experience (or perception) by the agent of a direct partonomical relation between the two. This experience may be influenced by many factors, like the experience of other relations, such as causality and spatial inclusion. For instance, consider a concept *B* that is usually experienced as direct part of *C*. Now, take a concept *A* that is perceived as a direct part of *B*. The perception of *A* as part of *C* is exclusively dependent on one's experience of the world. We argue that the conceptual level do not draw any a priori deductive conclusions about partonomic relationship between *A* and *C*.

However, even discounting transitivity, concepts in holistic and structure spaces can still suffer from inflation problems stemming from their legitimate direct parts. Remember that a structure space is a subspace of the conceptual space of the whole. So, the inclusion of the quality domains of “eye” in the structure space of “face” will also bring in the structure space of “eye” which includes the quality domains of “retina”. This may lead to a situation where inflated concepts represent themselves *and* all their possible parts. Dimensional filters act as a countermeasure for this issue. As seen in the Section 4.2, dimensional filters are able to select a subspace of the part's conceptual space to compose the whole. In this context, dimensional filters can filter out aspects of parts that are not relevant to the whole.⁵ For instance, the relation between face and eye is mediated by a dimensional filter that blocks the fraction of the conceptual space of eye referring to retina.

4.7 Other Ontological Considerations

Part relations and partonomies are also a common object of study in Ontology. However, in this work we are not much concerned with part relations that are held between objects in a world without a knower. We are rather concerned about how to represent concepts as they appear in cognitive processes of a knower. Nevertheless, certain issues regarding the ontology of concepts, parts and wholes are relevant to our discussion. In this section, we discuss how the notions of structure and holistic spaces contrast with those issues and how they can contribute to the field.

We touched some of the ontological issues of part relations when we addresses part transitivity in the previous section. In essence, we argued that part transitivity in the conceptual level is not predetermined, in the sense it is dependent more on experience than in predetermined concepts or deductive knowledge. However, there are other ontological

⁵Just as the botanist and the customer put different weights on different apple properties (Figure 4.5), a filter can be seen as an *attentional* mechanism that picks out the aspects of the structural complex that are of interest to the user in a particular context.

issues besides transitivity. For instance, there is the problem of diachronic identity in conceptual spaces (GAUKER, 2007). In brief, the issue is how one represents the fact that objects change in time (i.e. change their properties), but can still be categorized as the *same* object. According to Gauker, since conceptual spaces identifies objects with single points in a space, the same object in different moments in time would have to be represented as a set of points. This could be a problem, since there is no apparent way of linking all this points in order to say they denote the same object. Therefore, there is something missing in conceptual spaces that dictates what is necessary to an object to keep its identity through time. We will not fully address this issue here, as we believe it falls of the scope of this thesis. However, we briefly discuss how one could tackle it. The first possible solution is a no-solution: certain domains are stable enough to do without diachronic identity at the conceptual level in order to be represented. A similar approach is to assume that the sub-conceptual level compensate for diachronic change (e.g. using viewpoint invariant quality domains). The third solutions is to assume objects are in fact *regions* in a conceptual space. Gärdenfors (2000) himself hints at this possibility when suggesting that objects are *narrow concepts*, consisting of regions degenerated to single points. We suggest a similar solution in chapters 5 and 6 in order to tackle a similar issue that appears when one consider objects appearing to have different properties depending the perspective from which they are seen. Guizzardi (2014) presents a similar solution to the diachronic identity problem, proposing that points in a conceptual space represent object *states*, therefore making objects regions in conceptual space. Finally, it seems to us that in certain cases, we can treat objects that changed in time as different objects. Our cognitive processes seem to be able to rearrange their *conceptual* representations according necessity. Our opinion aligns with the one of Scholl (2007), who defends that the metaphysical issues that usually ground the discussion of identity may be themselves grounded in cognitive mechanisms that are still not well understood. More importantly, the very role of parts and wholes in tracking an object's identity is not well established from the perspective of cognition. In view of all that, we believe we should keep holistic-structure spaces transparent to this issue. The theory as it is now is malleable enough to fit with all these solutions. Furthermore, we believe that, at least for now, the identity problem should be tackled by particular implementations of holistic-structure spaces (as we do in Chapter 6), until a firm theory of objects in conceptual spaces is reached.

Another issue, related to the identity problem, is that people frequently consider wholes to have *mandatory* or *essential* parts (SIMONS, 2003). For instance, it is usually said that brain is an essential part of human, while heart is just a mandatory part; i.e. a particular human has to have a particular instance brain, but any instance of heart. Parts can also be *optional*; in the sense that they might or might not appear in the whole to which they relate to. This sort of construction has no parallel in structure spaces. Again, our position regarding these aspects are similar to our position in relation to transitivity.

We do not specify any *intensional* mechanism in the representation that allows us to tag parts as essential, mandatory, etc. The main reason for not having such mechanisms is that they would remove the plasticity of our representation scheme. For instance, it is not difficult to find counter-examples for many typical illustration cases of essential parts; there are tragic cases of living people without a brain, such as newborns. Alternatively, one can even conceive the invention of a successful lobotomy procedure, making brain just a mandatory part of human. Humans can easily adapt their conceptualizations for such situations in a non-monotonic way, where parts stop being necessary or mandatory. Thus, there must be a place in our representation where these adaptations are possible. We believe this place is the conceptual level. Structure spaces can capture at least a portion of that non-monotonicity in part reasoning by defining levels in which a given part is entrenched in a whole by means of context changes (see Section 4.5). Parts that bear more weight are more entrenched in a given context than others. For instance, the brain might be less important when comparing humans in a crowd.

Furthermore, concept similarity models have been linked to the notion of *family resemblance* (ROSCH; MERVIS, 1975). Family resemblance is an effect present in categorization, where instances of a given concept share a great number of similar properties, but not a single property is shared by them all. Conceptual spaces tend to be more aligned with this perspective. Humans have in common they all have a particular brain, but we can still categorize individuals with no brain as humans because they have many other properties in common. This plays against an essentialist view on concept representation. From a conceptualistic point of view, the idea of that some parts seem essential or mandatory might make more sense as symbolic level constructions *reflecting* common set of occurrences in the conceptual level. This might as well reveal itself as a point of connection between conceptual and symbolic frameworks in the future.

Here it is important to distinguish between Ontology as a field concerned with the study of being, and ontology as a knowledge representation technique in Computer Science. Our work has more in common with the representation frameworks in ontology engineering, particularly with foundation ontologies (see Section 3.3.3). DOLCE foundation ontology (and its derivations, such as UFO) holds our particular interest, since it employs conceptual spaces as a basic structure underlying qualities:

“Qualities can be seen as the basic entities we can perceive or measure: shapes, colors, sizes, sounds, smells, as well as masses, lengths, electrical charges... The term ‘Quality’ is often used as a synonymous of ‘property’, but this is not the case in DOLCE: qualities are particulars, properties are universals. Qualities *inhere* to entities: every entity (including qualities themselves) comes with certain qualities, which exist exactly as long as the entity exists. Within a certain ontology, we assume that these qualities belong to a finite set of *quality types* (like color, size, smell, etc.), and are characteristic for (*inhere in*) specific individuals:

no two particulars can have the same quality, and each quality is *specifically constantly dependent* on the entity it inheres in: at any time, a quality can't be present unless the entity it inheres in is also present. So we distinguish between a quality (e.g., the color of a specific rose), and its "value" (e.g., a particular shade of red). The latter is called *quale*, and describes the position of an individual quality within a certain conceptual space (called here *quality space*) [...]. So when we say that *two* roses have (exactly) the same color their two colors have the same position in the color space (they have the same *color quale*), but still the two roses have numerically distinct color qualities." (GANGEMI et al., 2002, pg 228)

Gangemi et al. (2002) imply that quality spaces are conceptual spaces, but do not develop much further in the relationship between DOLCE and conceptual spaces (particularly in the how properties and concepts in conceptual spaces map to DOLCE).

More recently, Guarino (2013) provided a preliminary investigation about the mereological behaviour of qualities in the context of DOLCE. There is a general distinction between *global* and *local qualities*. Global qualities inhere to the whole object, while local qualities inhere to a *part* of the object. For instance, when we say that the Adriatic Sea has a volume, we are referring more-or-less to a global quality. But since the sea has many depths, when we say that the Adriatic Sea has a depth, we are referring to a quality inhering in a specific part of the sea (e.g. the deepest part). More importantly, not all qualities of parts of an object are local qualities of that object. For instance, while the many widths of the main body of a vase are local qualities of the vase, the many widths of its handle are not local qualities of the vase. Guarino argues that there is probably a "simple cognitive mechanism" that divides the vase in *canonical parts*, "whose width count as local widths of the vase".

Guarino goes further, proposing the notions of *quality fields* and *quality patterns* as distinguished ontological categories. A quality field is the mereological sum of all local qualities of a given entity. Whereas local quality inheres in parts of an entity, a quality field inheres in the actual object. For example, the local qualities of the Adriatic Sea form the depth (quality) field of that sea. Guarino argues that quality fields also help to describe the distinction between *variation* and *change*. Variation would be explained as a process of change in the focus of attention along different parts of a quality field. For instance, there is a variation in the depth of the Adriatic Sea at the Italian part to the depth at the Croatian part. The variation occurs as attention shifts from one part to the other. On the other hand, change is rather a diachronic process. It occurs when some individual qualities in the quality field change their *quale* values (also known as *qualia*) in time. For instance, the depth of the Adriatic Sea has change since the Roman age.

Quality patterns refer to an individual distribution of qualia in a quality field. For instance, the depth quality field of the Adriatic Sea had a specific qualia distribution at the Roman ages, which is different of the distribution present today. Guarino describes quality patterns as constituted by a part of a quality field. The difference between quality fields and patterns is that the precise qualia distribution is essential to a quality pattern, whereas it is not essential to a quality field. Guarino (2013) does not make it clear whether quality fields inhere in other entities.

Contrary to early work in DOLCE on qualities, Guarino (2013) does not explain how quality fields and patterns are related to quality structures. We believe holistic-structure spaces provide conceptual grounding and justification to some of Guarino's decisions. We provide a sketch of how the relationship can be drawn.

There is a clear relation between quality patterns and structure spaces. Both consist of distributions of qualia representing specific quality values attributed to parts of a given entity. However, structure spaces are broader in scope. Together with dimensional filters and context change effects (see Section 4.5), they capture a more general notion than quality patterns. They represent all important qualities (or qualia) of parts of a given entity that are relevant to the specific conceptualization of the whole entity. A quality pattern is then a particular context on a structure space that selects the same quality domain in many parts composing the structure. If the Adriatic Sea correspond to a point in a holistic-structure space, giving greater weight to the depth domains of its parts (i.e., shifting the context), then we are effectively representing a depth quality pattern, i.e. we are isolating the qualia that refers to depths of parts of the sea.

However, the holistic-structures can go further. A single point in the structure space represent a "pattern" of local *and* global properties; i.e. it represents a distribution of qualia related to local and global attributes. Guarino (2013) recognizes that there is a correlation between values of certain global qualities and the same qualities in the parts. This correlation is reflected by the concept of point in the holistic-structure spaces.

If we assume diachronic quality fields, we can also readily explain the attention process involved in variation. Variation in conceptual spaces can partially be a change in the context of a given similarity measurement. If we are discussing the Adriatic Sea and the depth of its Italian part, we are bound to give more weight to the domain corresponding to the depth domain of the Italian part. A variation is just a special case of representing quality patterns.

The full notion of quality field can be represented in conceptual spaces depending in how one models objects in these spaces. In the diachronic case, quality fields and quality patterns seem to be indistinguishable; both map to a single point in a structure space. In the synchronic case, the complexity increases. If a certain quality field of a single object can have many patterns (such as in the example of the Adriatic Sea depth field now and in the Roman ages), then it implies that a single object might correspond to a *region* in

the structure space. Therefore, representing quality fields and also *change* require a more elaborated theory of objects as regions in conceptual spaces. We employ this strategy of representing objects in the last chapter of this thesis, but just as an implementation device. The full treatment would require further analysis. Nevertheless, the notion of regions in structure spaces seem to suggest the existence of *categories* of quality fields, such as the sea depth field, with all its possible variations.

4.8 Computational Object Recognition with Structure Spaces

Holistic and structure spaces possibly fit in alternative computational implementations of different cognitive tasks involving conceptual representations. Here, we present a sketch of how a system for object recognition inspired by holistic and structural processing could be built using structure spaces.

Object recognition can be seen on an abstract level as a cyclic process of perception-action. A bottom-up process interprets raw perception data into high-level structures; and a top-down process converts partial high-level interpretations into directed attention in order to clarify the missing information (i.e. visual search). We can say that an agent implementing this system achieves interpretation once its internal state stabilizes in a particular set of high-level structures abstracting the perceived stimuli.

In this context, conceptual spaces provide a way of describing concepts in terms of sets of possible observations. Consider an agent equipped with a visual perception system and a conceptual system described in terms of conceptual spaces. Simply put, visual interpretation consists in converting raw visual stimuli into vectors in these conceptual spaces and then checking if the vectors are similar (close) enough to the concept prototypes; the closest prototype indicates the concept describing the perceived object. For example, consider this agent is equipped with the conceptual space of the (holistic) apple as presented in 4.1, as well as a similar space for pear. When this agent is presented with visual image of an apple, the raw-visual stimuli is converted into a vector in a conceptual space formed by quality domains related to the visual system, such as the shape space, the colour space, the texture space and so on. This vector is a representation of the perceived object. In order to recognize this object as an instance of “apple”, it is sufficient to check whether the perceived object vector is *placed inside the properties regions* that define “apple” in each domain. In this case, high-level classification is reduced to a relatively simple verification of geometric inclusion of a point in a region. The process can be further simplified by reducing geometric inclusion calculation to a distance measurement from the concept prototype. However, more than one concept can be activated by the previous classification process; i.e. maybe it was not possible to say if the object was an apple or a pear. If more information is necessary in order to achieve classification, then it is possible to use the candidate concepts to redirect perceptual attention in order to gather better informa-

tion on the object; i.e. in a top-down process. For instance, the mismatch between the texture of apple and pear (i.e. empty intersection in the texture domain) might be used to guide attention to get closer view of the surface of the object being observed. The new perceptual information complements the previous observation by refining its vector in the conceptual space, restarting the bottom-up processing and closing the object recognition loop.

Structure spaces can improve this scheme by allowing the definition of independent holistic and structural processing strategies, which might bring improvements in how different stimuli are classified in certain situations. A possible way in which this can be implemented is to consider the previous model as the holistic strategy and join it to a parallel structural process. Consider an agent now is equipped with the conceptual space of apple (and pear) as in Figure 4.2; i.e. with holistic and structural descriptions. As we have seen in Section 2.1, holistic and structural processes occur in parallel in object recognition of humans, with one or other having some speed advantage depending on the context (LOVE; ROUDER; WISNIEWSKI, 1999). For the sake of the following example, consider the holistic processes have a slight speed advantage. Likewise, when this agent is presented with a *whole* apple, the process described previously is triggered, using the holistic space of apple (and other concepts) as basis for classification. As soon as a stimulus is recognized as being holistically an apple or a pear, then the structural processing can start in parallel. It tries to disambiguate between the candidate concepts. This process takes information of structural information encoded in the parts of apple and pear in order to refocus perceptual attention from the whole towards specific parts of the object being perceived in order to disambiguate them. Let us assume pears have slightly longer stems than apples. The structure domain of this part in both concepts can be used to shift attention to the appropriate locus of the stem on the perceived apple. When the perception is shifted to a part, all processing is primed to process a part, instead of a whole. The previous holistic vector is augmented, receiving values in other domains corresponding to the perceived part, such as part shape, texture and colour and configuration. In our example, the vector now encodes information about the whole apple being observed, as well its stem. This vector is then matched against structural fragment of the candidate concepts. If they are close enough to these concepts in the structure space (i.e. to their structural prototypes), they are kept as candidates. In this case, the concept of “pear” could be discarded as the stem part will now fail to fall in the correct property regions. The perception/action loop will keep running until no more distinguishing properties or parts can be found and concepts disambiguated; then classification is said to be achieved.

It is important to mention again that the stimuli segmentation is mainly governed by mechanisms outside structure spaces. From a cognitive point of view, these mechanisms define what a part is in fact. Some examples are Gestalt grouping principles (LOVE; ROUDER; WISNIEWSKI, 1999), part boundary rules (HOFFMAN; SINGH, 1997), part

saliency and so on and so forth. In the computational models exemplified here, these principles would be implemented in lower processing levels, such as perception; nevertheless, they still might use high-level information to tune perception (such as expected shape information).

Context effects can also play a role in this process through external systems. If we can keep a short-term memory with recently perceived objects, then these can prime the contextual space, influencing the weights of the vector components being compared during the high-level processing. Consider an agent is in a situation where there are many similar sailing ships with triangular sails. If a ship appears now with rectangular sails, the agent will tend to shift more weight to the “sail” part, in order to better discriminate between the objects being observed.

The mention to short-term memory highlights the possibility that more complex cognitive architectures could be combined with holistic and structure spaces in order to specify more complex behaviour. For instance, cognitive architectures such as the one proposed by Chella, Frixione and Gaglio (2001) could benefit from our framework, allowing them to deal with parts in the conceptual level. Furthermore, structural analogy models such as BRIDGES (TOMLINSON; LOVE, 2006) and DORA (DOUMAS; HUMMEL, 2010) could benefit of having structure spaces as an underlying conceptual structure. For instance, while BRIDGES provides a good processing strategy for measuring structural similarity between exemplars and new stimuli, it ignores holistic similarities; which is accounted by the holistic space in our framework. DORA, on the other hand, depends heavily on low-level symbolic feature descriptors. These could be replaced by regions in quality domains, allowing fine similarity comparisons between object characteristics (i.e. *geons*).

4.9 Summary and final remarks

In this chapter, we have presented a cognitive approach to representing part-whole relations, founded on the theory of conceptual spaces. Parts are associated with the whole in a structure space, where structural similarity can be measured between wholes and types of wholes. The structure space can capture many aspects of part relations. We have discussed different types of part deconstruction for the same whole, prototypical part deconstruction, variations in part structure caused by context, and part hierarchies. We also showed how different constructions of structure spaces can explain some types of part-whole relations

The model presented here contributes to the discussion of whether cognitive semantics is necessary for knowledge representation in computation. As it has been argued by Gärdenfors (2004), technologies based on symbolic theories, such as the Semantic Web, should also include representations that take into account cognitive phenomena. Indeed, the difficulties in representing part-whole relations using ontology representation languages (RECTOR et al., 2005) serve as good arguments for approaching the problem using a cognitive semantics framework like the one presented here.

Whilst the presented study focus on concept representation from a computational point of view, it still might be useful as a source of insights in the inner workings of the human cognition. Mainly, it calls the attention to the combination two important cognitive phenomena: part-whole processing and similarity effects. While there is already evidence indicating the existence of interplay between these two phenomena in cognition, it would be interesting to test more thoroughly if it exists also in other contexts, such as in events and in abstract entities, and how they relate to other aspects of cognition.

We showed that parts, wholes, and their relations can be represented in conceptual spaces, albeit the spaces are high-dimensional and more complicated than for other perceptual properties. However, a lot of empirical and mathematical work remains to turn these sketches into practical working models.

In next chapter, we give a step in that direction by providing a formal model to most of the notions presented in this chapter.

5 PART-WHOLE RELATIONS AS PRODUCTS OF METRIC SPACES

From a computational point of view, such framework would be highly beneficial. For instance, the modelling of part-whole relations is sometimes problematic in symbolic representation paradigms, such as ontologies. The algebraic approaches to part-whole relations are not generally driven by the cognitive phenomena (SIMONS, 2006), such as similarity and holistic/structural processing. Some foundational ontologies already embed notions of similarity (GUIZZARDI, 2005), but they are not usually associated to part-whole representation. On the other hand, when it comes to object recognition, notions of similarity and holistic/structural processing are more common. We find approaches employing either holistic processing (e.g., LOWE, 2004) or structural (part-based) processing (e.g., AGARWAL; AWAN; ROTH, 2004; FIDLER; BOBEN; LEONARDIS, 2008). Some alternatives combine both strategies (e.g., MOTTAGHI; RANGANATHAN; YUILLE, 2011; GRAF, 2006), using global and local descriptors for a review. However, these approaches focus more on immediate issues related to shape recognition, not addressing broader questions associated to concept representation.

Just as a recap, in the last chapter, as well as in some papers (FIORINI; ABEL; GÄRDENFORS, 2013; FIORINI; GÄRDENFORS; ABEL, 2013), we proposed a theoretical framework based on the *theory of conceptual spaces* (GÄRDENFORS, 2000), which embeds the notions of concept similarity and prototypes. We extended the theory by defining *holistic* and *structural spaces*. These are special types of conceptual spaces that are suitable for representing holistic (whole) and structural (part) information about concepts and objects in an integrated way. In such spaces, concepts and object are represented as regions denoting qualities of the whole and the parts structure. Similarity measurements between objects and concepts are function of the whole similarity and structure similarity.

In this chapter, we propose a mathematical formulation of holistic and structure spaces that can be more readily translated computational formalism. Gärdenfors (2000) suggests that conceptual spaces could be formulated as *metric spaces*, where the distance function implements the notion of similarity. This idea have been further explored by Aisbett and Gibbon (AISBETT; GIBBON, 2001b), where they lay down the underlying formal axiomatization of a metric-based formulation of conceptual spaces. We fork their work to propose an axiomatization of holistic and structure spaces as a *product* of metric spaces.

5.1 A short review of conceptual spaces and holistic-structure spaces

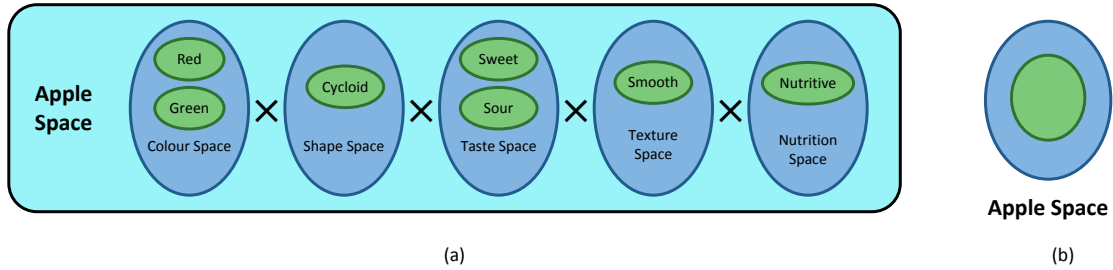
In this section, we give a short summary of the main notion of both conceptual spaces and holistic-structure spaces.

Gärdenfors' theory of conceptual spaces (GÄRDENFORS, 2000) puts forward a new way for representing concepts using geometrical and topological structures, which complements symbolic and connexionist approaches. Put it simple, a conceptual space is a space in the mathematical sense, where objects are points and concepts are regions or sets of regions. If this space has a well-defined metric, then it is possible to get the similarity between objects (and concepts) by measuring the distance between them: further objects are apart, less similar they are. The dimensions in a conceptual space have a special meaning: they denote the features — or qualities — through which entities can be compared and are frequently grounded in perception. Good examples are *hue*, *mass*, *height*, etc. Certain quality dimensions always co-occur, forming more complex spaces called *quality domains*. Examples of quality domains are *colour space*, *shape space*, *taste space*, etc.

The theory introduces the *Criterion P*, which states that *Convex* regions in quality domains define simple concepts, or *properties*. Properties such as Red, Circle, Sweet are examples of convex regions in colour, shape and taste domains respectively. The theory also introduces the *Criterion C*, which specifies that more *complex* concepts are defined as set of regions in many quality domains. For instance, the concept of apple can be defined as a product of colour properties in the colour space (Fig. 5.1a), cycloid regions in the shape domain, some regions in the taste domain and so on. An individual apple is represented by a single point in the multidimensional space formed by all quality domains of apple and contained in the property regions that form the concept of apple. In such space, a subtype of apple is a subspace of apple.

Our aim is to represent part-relations in conceptual spaces. However, conceptual spaces does not provide a complete solution for representing relations; it simply suggests that they could be represented by a product between the conceptual spaces of the related entities. Thus, in order to represent the relation between parts and wholes, we proposed a specialization of conceptual spaces in order to represent holistic and structural information about concepts and objects. Based on evidences from cognition, we first sug-

Figure 5.1 – Example of diagrams depicting the conceptual space of apple: (a) shows the inner form of the apple space as a product of properties (smaller ellipsoids) in different quality domains (bigger ellipsoids); and (b) shows a compact representation of the apple space as a set of points (smaller ellipsoid) in a multidimensional space formed by the product of its quality domains.



Source: the authors.

gested that the similarity between two concepts x and y defined as a function $\text{sim}(x, y)$ composed by two subfunctions $\text{sim}_h(x, y)$ and $\text{sim}_s(x, y)$, denoting holistic and structural similarity respectively; that is, it is a mixture (denoted by the binary operator \otimes) of the global similarity and part structure similarity. Informally,

$$\text{sim}(x, y) = \text{sim}_h(x, y) \otimes \text{sim}_s(x, y).$$

We can detail this function even further. From a structural point of view, we can compare object regarding the property of sharing the same set of parts; and also compare them regarding the property of having the parts arranged in a similar way. We refer to the former as *part similarity* ($\text{sim}_p(x, y)$) and the later as *configuration similarity* ($\text{sim}_c(x, y)$). By taking $\text{sim}_s(x, y)$ to be a mixture (denoted by the binary operator \oplus) of $\text{sim}_p(x, y)$ and $\text{sim}_c(x, y)$, we have that

$$\text{sim}(x, y) = \text{sim}_h(x, y) \otimes [\text{sim}_p(x, y) \oplus \text{sim}_c(x, y)].$$

This relationship between holistic and structural similarity induces further structure in how conceptual spaces are defined. We proposed that the conceptual space of a given entity is now formed by the *product* of two specific subspaces: an *holistic space* implementing the holistic similarity $\text{sim}_h(x, y)$; and a *structure space* implementing structural similarity $\text{sim}_s(x, y)$. A concept is then represented as a set of regions in the holistic space describing its overall qualities; and a set of regions in the structure space describing its inner structure of parts

More specifically, holistic spaces can be seen as conceptual spaces as the general definition proposed by Gärdenfors: a group of quality domains describing qualities about a whole entity, such its overall *shape*, *colour* or *mass*. On the other hand, given that structural similarity is more complex, structure spaces also have richer construction. A

point in a structure space denotes a particular configuration of parts of an individual. That is, a single multidimensional point encodes the information about what parts compose a whole and also about how these parts are configured. Similar configurations of parts are close together in this space. The information about how parts and structure can be compared is given by the dimensions and quality domains. Given a concept C and a set of concepts C_1, \dots, C_n that are parts of C , then we can generally define the structure space of C as the product space of the quality domains of C_1, \dots, C_n and other n quality domains denoting specific configuration information about each part P_i . The later quality domains are called *structure domains*; these spaces represent information about the configuration of each part in the whole; putting all together allow for part structure comparisons between different objects. An example of structure domain for a part is a whole-centred coordinate space denoting positions in relation to the whole.

A concept then is defined as a set of regions in both its holistic space and its structure space. For instance, consider again the conceptual space of the concept Apple. It consists of a holistic (such as the one in Fig. 5.1) and a structure space. The structure space of Apple could be formed by the product of Core, Flesh, Seed and Stem, plus regions denoting the general positioning each part in an object-centred coordinate system. A vector in the structure space of Apple denotes a particular apple-structure: a combination of individual apple parts, each with a specific value for colour, shape, taste and so on. Close points in this space represent similar apple-structures. The combination of regions of each part in the product restricts what are the valid individual components of an apple. More importantly, the structure space can be further divided into specific regions defining types of apple-structures; e.g., the concept of an apple with acid flesh and short stem.

5.2 Metric Spaces for Conceptual Spaces

As we have seen in Section 3.3.1, Aisbett and Gibbon (2001b) proposed a formulation for conceptual spaces based on metric spaces. We can divide their contribution into two parts. The first part sets the main geometric and topological properties that a metric space should have in order to support the main axioms of conceptual spaces. The second part is the formulation of conceptual spaces in terms of metric spaces. We shall reuse some elements of the first part as basis for our framework, but depart from what is proposed in the second part of their work. The main reason is that, while we share similar intuitions about how conceptual spaces relate to metric spaces, we believe it is possible to have a clearer conceptual formalism by employing the mathematical devices in a different way. In particular, whilst Aisbett and Gibbon (2001b) provide a basic mechanism to represent composite concepts (such as wholes), they do not address the distinction between holistic and structural similarity, as well as the existence of structural quality domains.

As we have mentioned earlier, the notion of metric spaces closely relate to how cognitive processes involving similarity measurements work. In this context, the metric generally corresponds to the similarity measure.

Definition 5.1 (*Metric space*). A *metric space* is an ordered pair (X, d) , where X is a non-empty set and d is metric on X (DEZA; DEZA, 2009). The metric d is a function $d : X \times X \rightarrow [0, \infty)$, which defines the distance between elements of X . For any $x, y, z \in X$, the following holds:

$$d(x, y) \geq 0 \text{ (non-negativity)} \quad (5.1)$$

$$d(x, y) = 0, \text{ iff } x = y \text{ (identity of indiscernibles)} \quad (5.2)$$

$$d(x, y) = d(y, x) \text{ (simmetry)} \quad (5.3)$$

$$d(x, y) \leq d(x, z) + d(z, y) \text{ (triangle inequality)} \quad (5.4)$$

Triangle inequality guarantees that d imposes the structure to X , which characterizes a metric space. For instance, taking $d(x, y)$ as a measure of how three persons x, y, z are in a conceptual space, the axiom guarantees that person z cannot be more similar to persons x and y than person x is similar to person y . The canonical example of metric space is the Euclidean 2-D space (\mathbb{R}^2, d) , where $d(x, y) = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$.

In general, similarity and distances can be related in different ways. Given a distance metric $d(x, y)$, we have that $sim(x, y) = f(d(x, y))$. According to Gärdenfors (2000), psychologists have argued for different types of function f , such as $f(d_{xy}) = e^{-cd_{xy}}$ or $f(d_{xy}) = e^{-cd_{xy}^2}$. The specifics of this relation are not important here; we shall assume that similarity is maximal at zero distance and that both covary monotonically.

Aisbett and Gibbon (2001b) employ *pointed metric spaces* as the basic formalism for conceptual spaces.

Definition 5.2 (*Pointed metric space* (AISBETT; GIBBON, 2001b)). A *pointed metric space* $(X \cup \{*\}, d)$ is a metric space that a distinguished point $*$, such that for any $x \in X$ in the space, $d(x, *) = \infty$.

The point $*$ represent similarity comparisons that are not defined or applicable in a certain context. This situation might arise, for instance, when we try to compare entities that have no defined value for certain dimensions. From now on, we will refer only to pointed metric spaces.

Another notion in conceptual spaces is *betweenness*. It is important to define connectivity and convexity, necessary properties for conceptual structures. From a conceptual point of view, the betweenness relation captures the intuitive notion that, given two objects in a conceptual space, there exists an intermediary object that relates them, e.g. a

ripe apple between a green one and a rotten one. The betweenness relation defines the structure of dimensions in conceptual spaces. Intuitively, for x, y, z in a space S , the relation $btw(x, y, z)$ holds when y is “between” x and z according the topology of S . For instance, in the Euclidean space we have that $btw(x, y, z)$ iff x, y and z are collinear.

Definition 5.3 (*Betweenness* (AISBETT; GIBBON, 2001b)). A betweenness relationship btw on a metric space (S, d) is a logical relation btw on $S \times S \times S$ such that for arbitrary $x_1, \dots, x_4 \in S$,

- (a) $btw(x_1, x_2, x_3) \rightarrow x_1 \neq x_2, x_2 \neq x_3, x_1 \neq x_3$;
- (b) $btw(x_1, x_2, x_3) \rightarrow btw(x_3, x_2, x_1)$,
 $btw(x_1, x_3, x_2) \rightarrow \neg btw(x_2, x_1, x_3)$;
- (c) $btw(x_1, x_2, x_3) \wedge btw(x_2, x_3, x_4) \rightarrow btw(x_1, x_2, x_4)$;
- (d) $btw(x_1, x_2, x_4) \wedge btw(x_2, x_3, x_4) \rightarrow btw(x_1, x_2, x_3)$;
- (e) $\neg btw(x_1, x_2, *) \wedge \neg btw(x_1, *, x_2) \wedge \neg btw(*, x_1, x_2)$;
- (f) $d(x_1, x_3) = d(x_1, x_2) + d(x_2, x_3) \rightarrow btw(x_1, x_2, x_3)$;

The axioms from (a-d) delineate the basic properties of the betweenness relations. Axiom (b) in particular avoids cycles. Axiom (e) forbids the distinguished point to participate in the betweenness. Axiom (f) decouples metric from the betweenness relation; in particular, it allows us to map spaces with different metric but that preserves the betweenness relation.

Recall that the Criterion P in the theory defines properties as *convex* regions. Taking into account the Definition 5.3, Aisbett and Gibbon state connectedness and convexity in a non-standard way by using betweenness, in order to avoid having to rely on a specific metric.

Definition 5.4 (*Connectedness, Region, Convexity* (AISBETT; GIBBON, 2001b)).

- (a) A subset A of (X, d) is called *r-convex* if:
 - (i) it is a singleton; or
 - (ii) for any pair $a, b \in A$, there exist $r - 1$ elements $x_1, x_2, \dots, x_{r-1} \in A$, where $a = x_0, b = x_r$ and such that $\forall c \in X btw(x_i, c, x_{i+1}) \rightarrow c \in A, 0 \leq i \leq r - 1$. A 1-convex space is called *convex*.
- (b) A subset A of a space (X, d) with a betweenness relation B is *connected* if:
 - (i) it is a singleton; or
 - (ii) for any pair $a, b \in A$, for some $n > 0$ there is a sequence $x_0, x_1, x_2, \dots, x_{n-1}, x_n$ in A where $x_0 = a, x_n = b$ and such that $\forall c \in X btw(x_i, c, x_{i+1}) \rightarrow c \in A, 0 \leq i \leq n - 1$.
- (c) A *region* of S is a connected closed subspace of S .

From now on, we depart considerably from the proposals of Aisbett and Gibbon.

5.3 Product Space

An important concept that is not introduced by Aisbett and Gibbon (AISBETT; GIBBON, 2001b) but that is important for the framework presented here is the notion of *product spaces* based on the betweenness relation. It differs from the standard definition of product of metric spaces in that it does not need to preserve local distances, but only betweenness.

Before defining product space, let us introduce a betweenness preserving morphism between metric spaces:

Definition 5.5 (*β -morphism*). Let X and Y be pointed metric spaces equipped with the betweenness relations btw_X and btw_Y respectively. A *morphism* $\beta : X \rightarrow Y$ is named a *β -morphism* if it is continuous, where $\beta(*) = *$; and preserves betweenness; i.e. for any $x_1, x_2, x_3 \in X \setminus \{*\}$, $btw_X(x_1, x_2, x_3) \leftrightarrow btw_Y(\beta(x_1), \beta(x_2), \beta(x_3))$

A product space is then a product of metric spaces where the projection morphisms are β -morphisms.

Definition 5.6 (*Product space*). A metric space $(\mathbf{X}_P, \mathbf{d}_P)$ is a metric *product space* of the finite family of metric spaces $\{(X_1, d_1), \dots, (X_n, d_n)\}$ if:

- (a) $\mathbf{X}_P = (X_1 \times \dots \times X_n) = \{(x_1, \dots, x_n) : x_1 \in X_1, \dots, x_n \in X_n\}^1$;
- (b) for each X_i , $1 \leq i \leq n$, there is a projection $\pi_i : \mathbf{X}_P \rightarrow X_i$, such that π_i is a β -morphism.

In many cases, the product metric \mathbf{d}_P is a function of all local metrics d_1, \dots, d_n , as in the usual interpretation of product of metric spaces.

5.4 Conceptual Spaces as Metric Spaces

The relation between similarity and metric spaces has been already established in the past. In his book, Peter Gärdenfors established that conceptual spaces can be seen as metric spaces. It is easy to see why metric spaces are a good tool for this. The main reason is that metric spaces naturally capture the notion of a similarity space with their distance metrics. More importantly, they allow us to abstract away issues about the *geometry* of conceptual spaces. Here we are concerned with the topological aspects involved in concept representation. More specifically, we show how certain mathematical notions, such as product space and projections, are particularly useful to formalize the main properties of holistic and structure spaces.

¹This is a typical case of abuse of notation in representing Cartesian products. Cartesian product is a binary operator, which is not associative in the rigorous sense, i.e. $X_1 \times (X_2 \times X_3) \neq (X_2 \times X_2) \times X_3$. However, $X_1 \times (X_2 \times X_3)$ and $(X_2 \times X_2) \times X_3$ can be taken to be isomorphic in most interpretations, which is the interpretation that we take in this thesis.

Conceptual spaces are about objects and the similarities between them. In the context of conceptual spaces, objects are usually represented as points. These points usually denote sets of stimuli (perceptual or not) or patterns, such as values of *colour*, *shape*, *weight*, etc. In certain contexts, they can be specific enough to denote real world objects, such as *John Lennon* or *Eiffel Tower*. These points are the primitives of our framework and we refer to them as *c-points*.

All conceptual space constructs can be seen as consisting of sets of c-points of the same sort. We can generalize this notion by introducing *c-spaces*².

Definition 5.7 (*c-space*). A *c-space* is a metric space (C, d) where C is a set of c-points equipped with a similarity metric d and a betweenness relation *btw* .

There are many different functions that can be used as d . Apart from the Euclidean metric, other kinds of metrics can be used depending on the application and the entities being represented, such as city-block metric (DEZA; DEZA, 2009).

We can specialize the definition of c-space in order to obtain metric space structures for conceptual constructs. We start by defining quality dimensions. Quality dimensions are the structuring constructs of conceptual spaces. They are one-dimensional structures usually related to elementary perceptual dimensions (e.g. *hue*, *pitch*, etc.)

Definition 5.8 (*Quality dimension*). A c-space Q is a *quality dimension* if there is a map from Q to subsets of \mathbb{R} such that $btw_Q(a, b, c)$ is preserved for any $a, b, c \in Q$.

Quality domains are c-spaces formed by products of integral quality dimensions. Quality dimensions are said to be *integral* when they cannot be valued independently of other dimensions.

Definition 5.9 (*Quality domain*). Given the quality dimensions Q_1, \dots, Q_n , a c-space \mathbf{D} is a *quality domain* if it is a product space $\mathbf{D} = Q_1 \times \dots \times Q_n$ for $n \geq 1$. If $n = 1$, \mathbf{D} is called *thin quality domain*.

The canonical example of quality dimensions and domain is the colour space, such as the HSL colour space. In this space, shades of different colours are represented as points in a three-dimensional space formed by the dimensions *hue*, *saturation* and *lightness*. These can be seen as quality dimensions forming a colour quality domain. For instance, the hue dimension can be seen as a quality domain (H, d_H) that is isomorphic to the real unit interval; its similarity metric d_H is the inverse of the one-dimensional Euclidean distance metric. Considering similar quality dimensions for saturation and lightness; i.e. (S, d_S) and (L, d_L) , we can define the HSL colour quality domain as $(\mathbf{Colour}, \mathbf{d}_{\mathbf{Colour}}) = (H, d_H) \times (S, d_S) \times (L, d_L)$, such that $\mathbf{Colour} = H \times S \times L$ and $\mathbf{d}_{\mathbf{Colour}}$ is a similarity metric in a cylindrical coordinate space

²The definition of c-spaces is a generalization of what Guizzardi (2005) calls *quality structures*.

There is an important aspect regarding the commutativity of Cartesian products when applied to conceptual spaces. Cartesian products are not commutative in general, i.e. $X_1 \times X_2 \neq X_2 \times X_1$. However, when considering the operands are metric spaces representing conceptual structures, the non-commutativity might look odd. For instance, there is no semantic difference between **Colour'** defined as $H \times S \times L$ and **Colour''** defined as $S \times H \times L$. Both product spaces are isomorphic. Therefore, commutativity does not have much importance at the present level of formalization of conceptual spaces. However, commutativity might become an issue when implementing the present framework in computers. For instance, certain implementations might rely on the ordering of dimensions in a quality domain in order to make the proper similarity calculations between points that domain. In the remainder of this thesis, we assume non-commutative product spaces as for simplicity, keeping in mind this property can be relaxed in actual implementations.

A *thin* domain has a single quality dimension, in which case the former inherits topology of the later (i.e. the metric). This case is not covered by the original definition of quality domain in conceptual spaces, but its inclusion reduces the complexity of the next definitions. Furthermore, any dimension to be used to compose conceptual spaces must be included in a domain.

In conceptual spaces, a great deal of emphasis is given to the quality dimensions as primitive building blocks. However, properties and concepts, the main constructs of the framework, are defined as being regions in *domains*, letting dimensions to a lesser important place. It is easy to see why. Quality domains can be easily linked to the idea of *perceptual modality* discussed in cognitive and neuroscience literature. They can be seen as high-level descriptions of perceptual modalities or portions of them, such as *colour*, *shape*, *taste* and so on. Defined as quality domains, these modalities can be treated as black boxes. In order to define a concept, we do not need to know *a priori* the exact inner dimensional structure of its quality domains; we need to know just if they obey certain principles, such as convexity, betweenness, etc. As a matter of fact, the dimensional structure of some domains such as shape are still open for debate (e.g., JOHANSSON, 2011). Nevertheless, this should not present any difficulty if we can still provide a betweenness relation and a (similarity) metric for the quality domain. Thus, from now on, we keep the focus on quality domains when describing c-spaces; in certain cases we might omit the dimensional structure of some quality domains.

Domains are banded together to form concepts and other structures. The space formed by a series of quality domains is a *support space*:

Definition 5.10 (*Support Space*). A c-space \mathbf{S} is a *support space* if it is a product space of $n > 0$ quality domains. If $n = 1$, we call \mathbf{S} a *thin* support space.

The support space of a given concept is what we usually refer to as a *conceptual space*. Nevertheless, any product of domains is a support space.

5.4.1 Properties and Prototypes

The theory of conceptual spaces states that properties are special types of concepts; properties are regions confined to quality domains and concepts are sets of regions spreading across many domains. According to the Criterion P, a property is a convex region in a quality domain. Focal (central) points of properties are the *prototypes* of these properties: the objects with high degree of typicality in relation to the property (ROSCH, 1978). This definition can be given more meaning if we ground it on Definition 5.4.

Definition 5.11 (*Property, prototype*). A property \mathbf{P} is a convex region in a quality domain \mathbf{D} and equipped with a single special point $\mathring{p} \in \mathbf{P}$ called *prototype*.

In general, domains contain many properties, grouped by domain *partitions*. For instance, the domain TASTE can be wholly separated into the partition $T_1 = \{\text{BITTERNESS, SALTINESS, SOURNESS, SWEETNESS}\}$ or just in the partition $T_2 = \{\text{TASTEFUL, UNTASTEFUL}\}$. Domain partitions can be generated in many ways. The main method is by defining properties as function of their prototypes. For instance, Aisbett and Gibbon (AISBETT; GIBBON, 2001b) have shown that it is possible to tessellate convex property regions from prototypes in a quality domain using Voronoi tessellation. For the purposes of this paper, it is enough to define the basic properties of a partition.

Definition 5.12 (*Domain partition*). Suppose a domain \mathbf{D} and a set of prototypes $\mathring{P} = \{\mathring{p}_1, \dots, \mathring{p}_n\}, \mathring{P} \subset \mathbf{D}$. A *domain partition* on D is a function $\Gamma : \mathring{P} \rightarrow 2^{\mathbf{D}}$ such that:

- (a) for any $\mathring{p}_i \in \mathring{P}$, $\Gamma(\mathring{p}_i) = \mathbf{P}_i$ is the property induced by the prototype \mathring{p}_i in \mathbf{D} ; and
- (b) given a set $T = \{\mathbf{P}_1, \dots, \mathbf{P}_n\}$ of all properties induced from a given partition function, then $\mathbf{P}_1, \dots, \mathbf{P}_n$ are pairwise disjoint and $\bigcup_{1 \leq i \leq n} \mathbf{P}_i$ is a connected region in \mathbf{D} .

Eventually, it is necessary to refer to sets of properties in a domain as single entities. For instance, consider that the concept Apple is defined within the properties Red and Green in the colour domain. The entity resulting of the union of these two properties is a *complex property*:

Definition 5.13 (*Complex property*). Consider a domain \mathbf{D} and partition Γ on \mathbf{D} . Let $T = \{\mathbf{P}_i\}_{i \in I}$ be the set of properties induced by Γ and indexed by I . A complex property is a set $\mathbf{PP} \subseteq \mathbf{D}$ such that $\mathbf{PP} = \bigcup_{i \in I'} \mathbf{P}_i, I' \subseteq I$.

Note that a complex property is not a property according to Criterion P; i.e., it is not necessarily convex or connected. Nevertheless, it shall be useful in the next definitions.

5.4.2 Concepts

Gärdenfors' original definition of concept is given by the Criterion C. According to it, a *natural concept* is represented as a set of regions in a number of quality domains together with an assignment of salience weights to the domains and information about how the regions in the different domains are correlated. At a first glance it might look unusual that Gärdenfors define concepts not in terms of *properties* in domains (i.e. convex regions), but rather in terms of generic regions in domains; one would expect Gärdenfors to take advantage of the representational power of properties, such as the presence of prototypes and convexity. This might have to do with the fact that certain concepts can be associated with complex properties, instead of just with properties. For instance, when we define that apples can be *red or green* only, the region formed by the union of RED and GREEN is not connected in the HSL colour space, let alone convex. This means that there is not an intermediary colour *between* red and green that can be attributed to apple.³ To that effect, our definition of concept based on metric spaces will rely on sets of complex qualities. This is a restriction on the view of concepts as simple set of regions (as it is defined in the Criterion C), to a compromised form based on sets of properties (convex regions).

Hence, a concept is a product space of complex properties, supported by the associated quality domains:

Definition 5.14 (Concept). A concept \mathbf{C} is a product space of a collection of complex properties $\mathbf{L}_1, \dots, \mathbf{L}_n$, such that $C = \{(p_1, \dots, p_n) : p_1 \in \mathbf{L}_1, \dots, p_n \in \mathbf{L}_n\}$.

The domains of each complex property $\mathbf{L}_1 \subseteq \mathbf{D}_1, \dots, \mathbf{L}_n \subseteq \mathbf{D}_n$ form the support space \mathbf{S}_C of \mathbf{C} , such that $\mathbf{S}_C = \mathbf{D}_1 \times \dots \times \mathbf{D}_n$ and $\mathbf{C} \subseteq \mathbf{S}_C$.

5.5 Part Relations and Dimensional Filters

As we have seen earlier, concepts (and objects) have holistic and structural fragments. The *holistic fragment* determines the properties of the whole and *structure fragment* determines what are the parts that compose the whole and how they are configured. The structure fragment is determined by dimensional filters on part concepts. In order to show how these notions can be formalized in metric spaces, we begin by considering how one can encode holistic and structural information about an object in a single multi-dimensional point.

Consider an object o having the parts o_1, \dots, o_n . For it to be represented as a single point in a conceptual space, we must encode it as a tuple defining its values in all relevant domains that characterize o . At the same time, this tuple must encode information about the parts that are present in o . Let h_k denote a value for the k -th domain of o describing

³A way of solving this is to assume the convex envelope of the complex property as the constituent of the main concept, then specializing the concept for each property.

the object as a whole; and $p_{(i,j)}$ denote a value for the j -th domain of the i -th part o_i that is *relevant* to o . So, the whole o can be defined written as:

$$o = \left(h_1, \dots, h_m, \overbrace{p_{(1,1)}, \dots, p_{(1,m_1)}}^{o_1}, \dots, \overbrace{p_{(n,1)}, \dots, p_{(n,m_n)}}^{o_n} \right), \quad (5.5)$$

In other words, the tuple is a combination of the holistic description of the object and the description of each of its parts. This definition of object as a point induces a specific notion of concept as space. Consider a concept \mathbf{C} abstracting o , such that $o \in \mathbf{C}$. Let $\mathbf{H}_1, \dots, \mathbf{H}_m$ be complex quality properties denoting holistic properties of \mathbf{C} such that $h_1 \in \mathbf{H}_1, \dots, h_m \in \mathbf{H}_m$. Similarly, consider the concepts $\mathbf{C}'_1, \dots, \mathbf{C}'_n$ abstracting parts of \mathbf{C} , such that $o_1 \in \mathbf{C}'_1, \dots, o_n \in \mathbf{C}'_n$. Let $\mathbf{P}_{(i,j)}$ denote the j -th complex property of the part \mathbf{C}'_i , such that $p_{(i,1)} \in \mathbf{P}_{(i,1)}, \dots, p_{(i,m_i)} \in \mathbf{P}_{(i,m_i)}$. Taking into consideration the definition of concept in (5.14), the concept \mathbf{C} can be seen as the product

$$\mathbf{C} = \prod_{k=1}^m \mathbf{H}_k \times \prod_{j=1}^{m_1} \mathbf{P}_{(1,j)} \times \dots \times \prod_{j=1}^{m_n} \mathbf{P}_{(n,j)}. \quad (5.6)$$

Intuitively, if we take $\mathbf{C}_h = \prod_{k=1}^m \mathbf{H}_k$ and $\mathbf{C}'_i = \prod_{j=1}^{m_i} \mathbf{P}_{(i,j)}$ then we can rewrite \mathbf{C} as

$$\mathbf{C} = \mathbf{H} \times \mathbf{C}'_1 \times \dots \times \mathbf{C}'_n, \quad (5.7)$$

where \mathbf{H} abstracts the holistic fragment of the concept \mathbf{C} . Additionally, if we take a space $\mathbf{C}_s = \mathbf{C}'_1 \times \dots \times \mathbf{C}'_n$, then we can write $\mathbf{C} = \mathbf{C}_h \times \mathbf{C}_s$ as the product space of its holistic (\mathbf{C}_h) and structural (\mathbf{C}_s) quality subspaces, where \mathbf{C}_h is referred as the *h-fragment* of \mathbf{C} and \mathbf{C}_s as the *s-fragment* of \mathbf{C} .


This provides a good intuition of how the relation between parts and wholes could be described in terms of metric spaces. However, the full notion of structure spaces requires two more constructs. Remind that the structure space of a concept is formed by the product of properties of the parts and properties in *structure domains*. The later kind defines how the parts interact with the whole, such as their relative position to the whole. We can rewrite (5.7) in order to accommodate structure domains. Take the concept $\mathbf{C} = \mathbf{C}_h \times \mathbf{C}_s$ as a product of its respective h-fragment and s-fragment. \mathbf{C}_s can be written as

$$\mathbf{C}_s = \mathbf{C}'_1 \times \mathbf{M}_1 \times \mathbf{C}'_2 \times \mathbf{M}_2 \times \dots \times \mathbf{C}'_n \times \mathbf{M}_n, \quad (5.8)$$

where \mathbf{M}_i is the structure property that modulates the structural relation of the part \mathbf{C}'_i in \mathbf{C} . The internal structure of \mathbf{M}_i is only relevant for the implementation.

In the majority of the cases, the concepts C'_1, \dots, C'_n correspond just to *projections* of the actual parts (denoted by C_1, \dots, C_n). These projections are given by what we call *dimensional filters* (see Section 4.2). For instance, a dimensional filter induces the concept Apple Seed from the concept Seed by means of a projection. The projection happens in two stages (see Figure 5.2). In the *first state*, the Seed space is projected to a lower-dimensionality support space, formed only by a subset of the domains in Seed. This is necessary to account for the cases where just selected characteristics of the part are relevant for the whole. For instance, Seed might have domains describing its own inner parts. These might not be relevant for forming the concept Apple, so they are filtered out by the first stage projection. In the *second stage*, a subset of the projected regions is selected. This corresponds to selecting the set of points that stand for actual apple seeds. The resulting projected concept is what is used to form the s-fragment of Apple. Therefore, there is a *part relation* between the whole and the projected part resulting of a dimensional filter applied to a part. In the following we formalize the main properties of dimensional filters and part relations interpreted in metric spaces. Then we discuss the real nature of these relations.

Figure 5.2 – The two stages of a dimensional filter applied to the concept Seed. This concept is represented as a cube in the Seed space (right). The first stage filters out domains (i.e. dimensions) that are not relevant to the concept Seed as part of apple (middle). The second stage produces a subset of the concept filtered in the first stage (right). This subset corresponds to apple seed as a subspace of the structure space of Apple.



Source: the authors.

Definition 5.15 (*Dimensional filter, part relation*). Consider the concept $C = C_h \times C_s$ and the concepts C_1, \dots, C_n such that each C_i is a part of C . Let $\{H_k\}_{k \in K}$ be a family of complex properties in the C_h and $\{P_j\}_{j \in J}$ be a family of complex properties, including structure properties, such that: (a) $J^s \subset J$ indexes properties in C_s ; (b) $J_i \subset J$ indexes properties in C_i . Consider a part C_i of C . Let $\{M_s\}_{s \in S} \subset \{P_j\}_{j \in J}$ be a family of structural complex properties. Based on this, we have that (Fig. 5.3):

- (a) a *dimensional filter* on a part \mathbf{C}_i is a β -morphism $\tau_i : \mathbf{C}_i \rightarrow \mathbf{C}'_i$ result of the composition $\tau_i = \zeta_i \circ \pi_i$, where: (i) $\pi_i : \mathbf{C}_i \rightarrow \mathbf{C}''_i$ is a β -morphism projection corresponding to the first stage projection, such that

$$\pi_i : \overbrace{\prod_{j \in J_i} \mathbf{P}_j}^{\mathbf{C}_i} \rightarrow \overbrace{\prod_{j'' \in J''_i} \mathbf{P}_{j''}}^{\mathbf{C}''_i}, \quad (5.9)$$

where $J''_i \subseteq J_i$ indexes complex properties of \mathbf{C}''_i ; and (ii) $\zeta_i : \mathbf{C}''_i \rightarrow \mathbf{C}'_i$ is a surjective β -morphism corresponding to the second stage projection, such that

$$\zeta_i : \overbrace{\prod_{j'' \in J''_i} \mathbf{P}_{j''}}^{\mathbf{C}''_i} \rightarrow \overbrace{\prod_{j' \in J'_i} \mathbf{P}_{j'}}^{\mathbf{C}'_i}, \quad (5.10)$$

where $J'_i \subset J^S$, $|J'_i| = |J''_i|$ and for each $\mathbf{P}_{j'}$ there is one and only one $\mathbf{P}_{j''}$ such that $\mathbf{P}_{j'} \subseteq \mathbf{P}_{j''}$. Alternatively, ζ_i^{-1} can be seen as an inclusion morphism (i.e. $\mathbf{C}'_i \hookrightarrow \mathbf{C}''_i$).

- (b) a *part relation* is a β -morphism $\rho_i : \mathbf{C} \rightarrow \mathbf{C}'_i$ that maps \mathbf{C} to its projection \mathbf{C}'_i , such that

$$\rho_i : \overbrace{\prod_{k \in K} \mathbf{H}_k \times \prod_{j' \in J'_1} \mathbf{P}_{j'} \times \mathbf{M}_1 \times \dots \times \prod_{j' \in J'_n} \mathbf{P}_{j'} \times \mathbf{M}_n}^{\mathbf{C}} \rightarrow \overbrace{\prod_{j' \in J'_i} \mathbf{P}_{j'}}^{\mathbf{C}'_i}. \quad (5.11)$$

By abuse of notation, the part relation ρ_i of the concept \mathbf{C} to the part \mathbf{C}_i can be also denoted as $\rho_i : \mathbf{C} \rightarrow \tau_i(\mathbf{C}_i)$. Let the complex properties \mathbf{C}'_i be called *part properties*.

Figure 5.3 – Commutative diagram of *dimensional filter* (τ_i) and *part relation* (ρ_i). The dimensional filter is a composition of two relations that we call *projections*. The *first stage projection* π_i removes dimensions of the part concept \mathbf{C}_i that are not relevant for the whole \mathbf{C} . The *second stage projection* ζ_i restricts the space resulting of the first stage projection to subset of it.

$$\begin{array}{ccccc} \mathbf{C} & \xrightarrow{\rho_i} & \mathbf{C}'_i & \xleftarrow{\zeta_i} & \mathbf{C}''_i & \xleftarrow{\pi_i} & \mathbf{C}_i \\ & & & & \text{---} & \text{---} & \\ & & & & \tau_i & & \end{array}$$

Source: the authors.

In some cases, we have that $J_i = J''_i$ for a part \mathbf{C}_i , meaning that all domains of \mathbf{C}_i are used to form the structure space (s-fragment) of \mathbf{C} .

The previous definitions just specify the basic formal properties that part relations and dimensional filters must have. Nevertheless, the real nature of dimensional filters is an open question, one that has an answer at the implementation level. If one implements conceptual spaces as mathematical representations in a programming language, then a

dimensional filter can be seen as functions transforming between different representations. For instance, Raubal (2004) represents properties in domains as polytopes, systems of equations defining regions in a conceptual space. Dimensional filters in this case could be seen as coded functions that transform polytopes to other polytopes. On the other hand, if ones implements conceptual spaces as sets of classifiers (in the sense of Edelman (1998)), then conceptual filters can be seen as activation links between classifiers. The metric definitions of dimensional filter proposed in this thesis are general enough to be implemented in different ways in actual computer systems.

5.6 Parts and wholes

Now it is possible to define what are parts and wholes in metric spaces. A whole a concept having at least a part property, accompanied by its structure property.

Definition 5.16 (*Whole, part*). Consider a concept \mathbf{C} as a product space of the family of complex properties $P = \{\mathbf{P}_j\}_{j \in J}$. Concept \mathbf{C} is a *whole* iff:

- (a) at least one $\mathbf{P}_j \in P$ is a part property; i.e., there is a concept \mathbf{C}_i , a dimensional filter τ_i and a part relation $\rho_i : \mathbf{C} \rightarrow \tau_i(\mathbf{C}_i)$ such that \mathbf{P}_j is mapped by ρ_i ; and
- (b) for each part property $\mathbf{P}_j \in P$, there is a structure property $\mathbf{P}_k \in P$.

A concept \mathbf{C}_i is a *part* if it is mapped by a pair of dimensional filter and part relation to a concept \mathbf{C} .

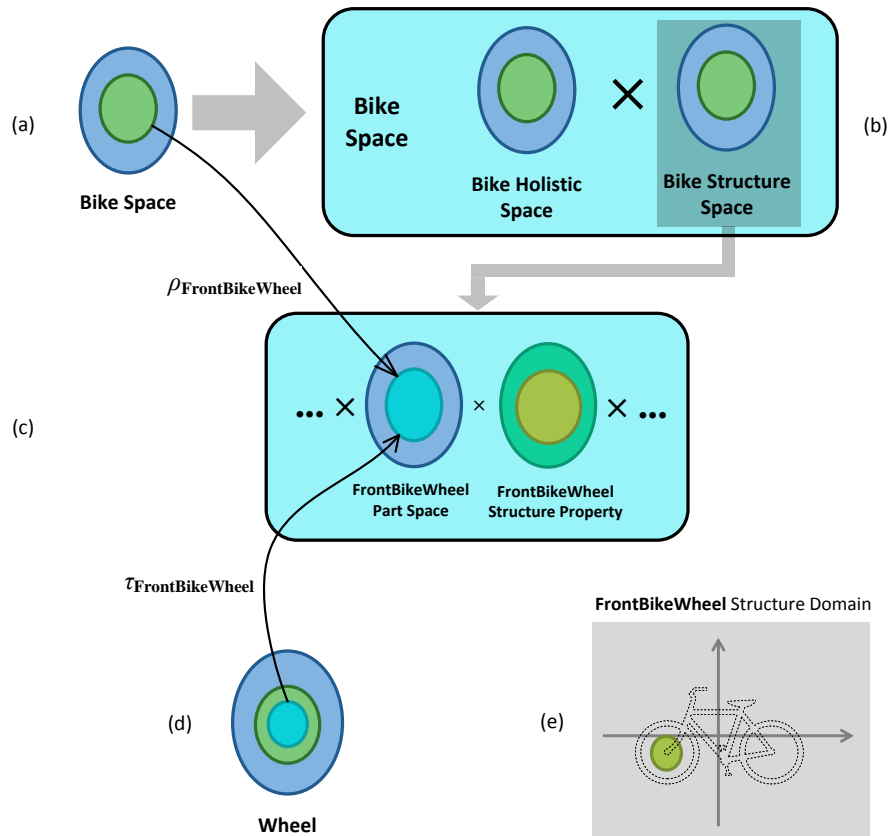
For example, consider the concept of bike with a part referring to its front wheel (Fig. 5.4). So, we have a concept **Wheel** related to the concept **Bike**, such that there is a dimensional filter $\tau_{\mathbf{FrontBikeWheel}} : \mathbf{Wheel} \rightarrow \mathbf{FrontBikeWheel}$ that defines which wheels can be the front wheel part of bikes; and a part relation $\rho_{\mathbf{FrontBikeWheel}} : \mathbf{Bike} \rightarrow \rho_{\mathbf{FrontBikeWheel}}(\mathbf{Wheel})$ that relates a projection of **Bike** to the concept **Wheel**.

Concepts that are not wholes, but still span more than one domain (i.e. concepts defined in just one domain are properties) are called *holistic concepts*. As the name suggests, these are basic concept having no structural fragment defined. These are concepts to which no part is conceptualized. For instance, rock might be seen as a holistic concept for a naïve observer, since it has no distinguishable inner parts.

The formalization of the notions of holistic and structure spaces follows from the definition of whole.

Definition 5.17 (*Holistic space, structure space*). Consider a whole \mathbf{C} with its support space (\mathbf{S}, \mathbf{d}) . The product of the domains containing part properties and structure properties form the *structure space* $(\mathbf{S}_s, \mathbf{d}_s)$ of \mathbf{C} ; and the product of the remaining domains form the *holistic space* $(\mathbf{S}_h, \mathbf{d}_h)$ of \mathbf{C} ; such that $\mathbf{S} = \mathbf{S}_h \times \mathbf{S}_s$.

Figure 5.4 – Example of the relationship between **Bike** and **Wheel**. Fig. (a) shows a compact representation of the **Bike** space, which is a product space of the holistic and structure spaces of bike (b). The structure space is exploded to show a fragment of its inner construction (c). The $\rho_{\text{FrontBikeWheel}}$ part relation projects the whole **Bike** space on its **FrontBikeWheel** space fragment, which in turn is a result of the dimensional filter projection $\tau_{\text{FrontBikeWheel}}$ on the concept **Wheel** (d). The **FrontBikeWheel** part space is modulated by a property in a structure domain denoting a bike-centered coordinate space (e), which defines the allowable positions for the front wheel in all instances of **Bike**.



Source: the authors.

From a cognitive point of view, the metric in \mathbf{S}_h should reflect the perceived holistic similarity in \mathbf{C} (i.e. $\text{sim}_h(x, y) = f'(\mathbf{d}_h(x, y))$) and the metric in \mathbf{S}_s should reflect the perceived structural similarity (i.e. $\text{sim}_s(x, y) = f''(\mathbf{d}_s(x, y))$), coordinating part and structure properties to account for part and configuration similarity respectively.

It is important to note that in certain interpretations of conceptual space, real-world objects (e.g. a person) might be represented also as regions in a conceptual space, in which the actual points represent specific *observations*, such as a specific measurement of colour, shape, taste and so on. For instance, consider the concept Tower. The individual *Eiffel Tower* is a subspace of Tower, which can be represented by a (convex) set of points in the shape domain denoting its many overall shapes from the many perspectives from which it can be looked at; plus a single point in the colour mass domain, representing its

unique mass. The framework presented here works in the same way for these entities. For instance, the structure space of Eiffel Tower maps to a specific subspace of Antenna representing the many possible aspects of the radio antenna attached to the top of the famous tower.

5.7 Discussion: benefits to software applications

We believe that holistic-structure spaces can be beneficial to software applications in two main ways. In general, we hope it provides the conceptual scaffolding for organizing different elements in a software application that has to deal with structural similarity and raw-data information.

For instance, the approach of Fidler, Boben and Leonardis (2008) is able to recognize objects by decomposing their contours into a hierarchy of shape parts. Small shapes aggregate to form figures with higher levels of complexity (Figure 2.6). In a given hierarchy layer, the similarity between two shapes is high if their subparts are similarly arranged. A part is basically characterized by its centre of mass and what are its subparts and where they are located. This similarity function corresponds roughly to our idea of structure similarity applied to the shape domain. Our work suggests how to enhance the semantic content of parts in Fidler, Boben and Leonardis (2008). For instance, quality domains other than shape could be added to the part descriptors, increasing the semantic content of those descriptions. These parts descriptors could be organized in quality domains. Our framework provide the blueprint how these different descriptors can be structured and used together. For instance, in order to increase performance of similarity comparisons between very complex parts, certain aspects of lower level parts could be filtered out of the comparison by dimensional filter, if the coarse topology of the similarity space is kept (i.e. the topology given by the *betweenness* in shape domain).

Applications based on conceptual spaces could also benefit from our work. For instance, Chella, Frixione and Gaglio (1997) propose a cognitive architecture that uses conceptual spaces to represent shapes of objects and their parts. However, the actual part relation is implemented in the symbolic level (Figure 2.4). So, similarity-based classification is solely based on individual part classification, not taking into account holistic and structural information. We could restructured this approach in terms of holistic and structure spaces. Instead of mapping entities from the symbolic level (i.e. a hammer) into multiple points denoting their parts, we could map these entities to single points having holistic and structural fragments. The previous points denoting parts become fragments of the same vector in the holistic-structure space. Thus, a single instance in the symbolic level becomes a single point in the conceptual level. This allows for simple holistic-structure similarity comparisons that were not possible in the original framework.

5.8 Summary and final remarks

In this chapter, we introduced a mathematical framework for representing the part-whole relations incorporating notions of cognition of object recognitions, namely holistic/structural processing and similarity. We hope it will serve as basis for computer algorithms implementing the same processes, such as computer vision programs. In the next chapter, we show an extensive example of how this can be accomplished.

The framework presented formalizes part of what is introduced in Chapter 4.2. Aspects such as context representation in part relations and specialization of different part relations still lack a formal account. The present formulation could be extended to address these aspects.

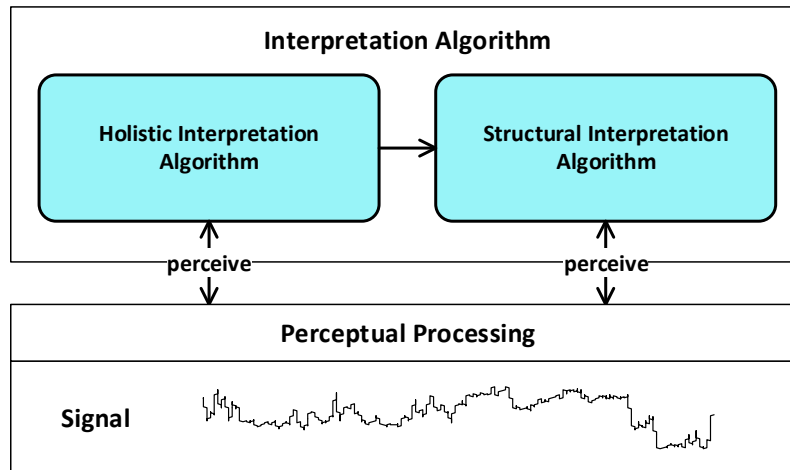
In a more speculative note, we believe that the framework presented here can also be useful to represent compositional relations between more abstract entities, such as actions and events.

6 AN ALGORITHM FOR HOLISTIC AND STRUCTURAL INTERPRETATION

There is a multitude of cognitive processes going on in an expert's mind when she is carrying out an interpretation task. As we argued in Chapter 1, any computer program trying to mimic such tasks would benefit of reproducing those cognitive processes. In this theoretical chapter, we focus on two particularly relevant processes involved in geological interpretation in general, namely *similarity* and *part-whole reasoning*, and examine an algorithm that can implement them in an effective way. We consider in particular how domain concepts can be specified taking into account those processes.

Similarity plays a significant role in psychological accounts for problem-solving, memory, prediction and categorization (GOLDSTONE; SON, 2005). A common example is in *object recognition*. Consider a situation where an object enters the field of view of a person. A person, in seeing this new object, compares it with mental representations of other objects she knows. Let us call these representations *mental objects*. If a given mental object is *similar* enough to the newly seen object, then the new object is recognized as been a new occurrence of that mental object. For instance, let us say Mary has a cat called Felix. Felix is represented as a complex mental object in Mary's mind. When a new object crosses in front of Mary, she compares the visual impressions of this new object with a series of mental objects she has in her mind. If the new object is similar enough to Felix, than she knows that the new object *is* Felix. Such mechanism works for both mental representations of *specific* objects as well as of more general *concepts*. In the case of concepts, similarity is measured between the new object and the known exemplars (or prototypes) of given concept. If they are similar enough, the new object is classified as being an instance of that concept. These exemplars are usually known as *prototypes* (ROSCH, 1978) of the concept. For instance, when seen a new object, Mary might compare it with some typical mental examples of the concept of cat. If the new object is similar enough to these mental prototypical cats, then she knows she has seen a cat.

Figure 6.1 – General structures of the object interpretation algorithm. The general algorithm is composed by two sub-algorithms. The holistic processing algorithm takes processed perceptual input (signal blobs) represented as points in a conceptual space and try to classify them holistically. Subsequently, the interpretation algorithm handles control to the structural processing algorithm, which tries to confirm the results of holistic processing by actively perceiving and classifying parts of the perceived objects. Both algorithms access signal (perceptual) processing algorithms through a generic interface represented by the function *perceive*.



Source: the authors.

As we have argued in previous chapters, part-whole processing also plays an important role in cognition. In particular, there is evidence that parts and wholes are important to object recognition in adult experts, being employed in different ways to solve tasks (BUKACH; GAUTHIER; TARR, 2006; HSIAO; COTTRELL, 2009). As a matter of fact, Richler, Wong and Gauthier (2011) argue that there is evidence of a continuous variation of processing strategy according to expertise. It generally goes from less holistic processing in novices to more holistic, automatic processing in experts.

In this chapter, we present *an object interpretation algorithm based on holistic and structural processing of perceptual input, using holistic and structure spaces* (Figure 6.1). We intended it to be abstract enough to be instantiated in actual object recognition and classification applications in computer systems. It takes as input pre-processed perceived objects represented as points in a conceptual spaces, and employ domain knowledge represented in holistic-structure spaces in order to classify the perceived objects. We do not assume any specific pre-processing algorithms (such as image processing software). Rather, we provide perception primitives defining the required interface that such algorithms must implement. While our algorithm in principle supports any kind of signal as perceptual data, we specified it with visual domains in mind, such as Medicine and Geology (CARBONERA et al., 2013; LORENZATTI et al., 2011). More specifically, we intend it to be used to automate object interpretation tasks where similarity between wholes and parts plays an important role in object recognition and classification.

Take, for instance, the task of stratigraphic interpretation in Geology. It is heavily based on object recognition in broad sense; geologists begin their work by recognizing and classifying geological features (i.e. objects) in the visual information items that they collect in various well logs, seismic blocks, outcrops, thin sections and so on. To our knowledge, there is not much research on the impact of the cognitive processes of object recognition in geological interpretation. However, considering the evidence issued from broader cognitive studies reviews in this thesis, there are reasons to believe that object recognition in geology does involve holistic and structural similarity. Shipley et al. (2013) argue that reasoning about parts and wholes is one of the skills needed by experts in structural geology. They hypothesize that expertise in geology is related to increased skill in reasoning about objects considered as single entities (wholes). The authors also emphasize the frequent necessity that geologists have to go across different spatial scales (such as *basin*, *reservoir*, *well*) in order to solve a particular interpretation problem. We believe this aspect is strongly related to the need of articulating holistic and structural processing for classifying together geological features that appear to be similar. In the second half of this chapter, we will notably show how our algorithm can be instantiated in order to implement stratigraphic interpretation.

6.1 Representation framework

The holistic-structural representation framework presented in the previous chapters does not establish all the aspects needed for specifying knowledge required for solving a task in a given domain. Thus, we introduce first the notion of conceptual system in order to do so. A conceptual system is a 4-tuple

$$SC = (S, C, K, P)$$

where S is a finite set of support spaces (see Definition 5.10); C is a finite set of concepts and objects, such that for all $\mathbf{C} \in C$, there is $\mathbf{S} \in S$ such that $\mathbf{C} \subseteq \mathbf{S}$; K is a set of *subsumption relations* between elements of C ; and P is a set of part relations (see Definition 5.15b) between elements of C .

Until this point in this thesis, we avoided any commitment to any specific theory about how taxonomic relations are represented in conceptual spaces. As we argued in Section 4.7, that would require a more elaborate theory about objects in conceptual spaces. However, an interpretation algorithm requires at least a simple treatment of taxonomic relations in order to be able to generalize input data. In conceptual spaces, points usually denote objects (such as the Eiffel Tower, my cat, Barack Obama, etc.) and regions denote concepts and properties (such as Building, Car, Person, Red, etc.) The concept membership relation is represented as a set inclusion relation between points and concept regions;

whereas the specialization relation between concepts is represented by subset relations between concept regions. However, in this work we have chosen to refer points in a conceptual space to patterns of percepts, rather than to actual whole objects. So, an object seen from two different sides becomes two different points in a conceptual space. That implies that objects, such as a car or a house, have to be considered as regions in these perceptual spaces. The points in these object regions denote different ways these objects can be perceived (i.e. a car seen from the front is perceptually different from a car seen from the back). Since concepts and objects are both regions, we eroded the usual epistemological distinction between objects and concepts. Concepts and objects have the same status in C , whereas both concept-concept (specialization) and concept-object (concept membership) relation collapse in a single type of relation in K , which we call *subsumption relations*. In other words, we take objects to be just special kinds of concepts. While this notion goes in an opposite direction to what is usually employed in concept representation (not without its exceptions; e.g. SCHNEIDER, 2010), it makes sense within the conceptual spaces framework. Gärdenfors (2000) already hints to the possibility of having objects as special types of concepts. He argues that objects in conceptual spaces are just like “narrow” concepts: concepts in which the regions are degenerated; corresponding to single points.

Representing object as regions gives us some practical advantages. This allows us to treat object and concepts homogeneously, which bringing simplicity to the interpretation algorithm when dealing with perceptual information. It allows one to process object recognition and classification in the same way. In object recognition, the algorithm checks how similar the perceptual input is to the typical view of a given object. In object classification, the algorithm checks the similarity between the perceptual input and the prototypes of a given concept. Also, that scheme allows us to represent the fact that an object might be linked to a property (such as red), rather than a specific quality value (e.g. a particular shade of red). This goes in the direction of increasing cognitive economy in cognitive-based representations (GÄRDENFORS, 2000).

More specifically, the subsumption relations in K capture the usual meaning of specialization relations between concepts as well as what is usually seen as the instantiation relation between objects and their classes. Thus, a subsumption relation $\kappa_i \in K$ is a total function, such that $\kappa_i : C_1 \rightarrow C_2$ for $C_1, C_2 \in C$. κ_i is also a β -morphism. The function κ_i can be seen as a betweenness-preserving inclusion function of a subconcept to its superconcept. Also, if $\kappa_i : C_1 \rightarrow C_2$ is in K , then $\kappa_i' : C_2 \rightarrow C_1$ is *not* in K , unless $C_2 = C_1$.

As with subsumption relations in K , part relations in P also apply both for concepts and objects in C . T

The support space $\hat{S} \in S$ is a distinguished support space that represents the immediate perceptual space of the agent having the system CS . Any object being perceived at a given moment is represented as a point in this support space. The domains $\mathbf{D}_1, \dots, \mathbf{D}_n \in S$ that form \hat{S} are perceptual domains grounded in vision, such as colour perception, shape perception, position perception, etc.

6.2 Interpretation algorithm

There are possible many ways in which holistic-structure spaces can be used to specify cognitive-like algorithms. In this section, we propose an algorithm that imitates some aspects of holistic-structural processing in order to classify objects present in the perceptual input (e.g. an image or any signal). This algorithm is to be used in tasks where holistic-structural processing is important, such as artificial object recognition and object classification. The general strategy of the interpretation algorithm is to apply holistic processing first (Algorithm 2), and then to use the result to restrict structural processing (Algorithm 3). When both holistic and structural information about objects are available, they are recognized (see Algorithm 1). In all processes, the algorithms shall measure similarity between objects in the perceptual input and entities (objects and concept) in a conceptual system CS .

6.2.1 Specification

Before introducing each algorithm, we must first introduce some basic primitive functions. Let $\text{perceive} : C \times X \rightarrow 2^C$ be a function that extracts objects from a signal $x \in X$ considering elements of C for attentional purposes and maps them to a set of *perceptual objects* embedded in the perceptual space \hat{S} . Perceptual objects are usually formed by regions degenerated to single points in the quality domains of \hat{S} . We denote perceptual objects in low case italics, as in $o \subset \hat{S}$. The application $\text{perceive}(\emptyset, x_i)$ should result in an unfocused attention processing of x_i (e.g. the whole signal x_i). We assume the function has default rules for segmentation

The function perceive encapsulates the connection from the conceptual system to the sub-conceptual level (see Figure 6.1). Given its output is a set of objects in the perceptual space \hat{S} , we can say that \hat{S} describes its output space. This has an interesting implication: *any sub-conceptual processing algorithm feeding the conceptual system has to obey the topology of \hat{S}* . In other words, \hat{S} defines the standard interface that processing algorithms have to implement in order to be able to create perceptual object in \hat{S} . So, for instance, consider a perceptual space \hat{S} consisting of two domain, the first being the HSL colour space (Figure 3.1), and the other being the superquadrics shape space (as in CHELLA; FRIXIONE; GAGLIO, 1997). So, any processing algorithm feeding objects to this space has to obey its established metrics and betweenness relation of both do-

mains. For instance, Chella, Frixione and Gaglio (1997) presents an algorithm based on neural networks that is able to extract superquadrics from images and obey that structure. Furthermore, this scheme implies that the output parameters space of any processing algorithm has a metric and a betweenness relation in order to be used with conceptual spaces. Also, if one ascribe different algorithms to different domains in $\hat{\mathbf{S}}$ (like the previous example), then both the perceptual space $\hat{\mathbf{S}}$ and the function perceive describe a principled scheme of the outputs of these processing algorithms can be combined.

We assume the function has default rules for segmentation. The perceive function can be implemented by a series of techniques, ranging from signal processing to neural networks.

Let $\text{conf} : \mathcal{O}^S \times \mathcal{O} \times \mathcal{S} \rightarrow \mathbf{M}$ be a function that determines the configuration of the perceptual object $o^S \in \mathcal{O}^S$ in relation to the object $o \in \mathcal{O}$ according to the configuration domain $\mathbf{M} \in \mathcal{S}$. The function returns a point $m \in \mathbf{M}$. For instance, if the configuration domain of a given object o is defined as a 2-D positioning space centred on the object (as it is usual in computer graphics, for instance), the function returns the position of the object o^S in relation to the coordinate space of o .

Algorithm 1 Interpretation algorithm

Require: A set $H \subset C$ of concepts to be checked

Require: A signal $x_i \in X$ to be interpreted

- 1: $\mathcal{O} \leftarrow \text{P-HOLISTIC}(H, x_i)$
 - 2: $\text{P-STRUCTURAL}(\mathcal{O}, x_i)$
 - 3: **return** \mathcal{O}
-

The holistic processing strategy is implemented by the function P-HOLISTIC specified in Algorithm 2. It takes a subset C' of concepts in C as input, representing the set of conceptual structures that the algorithm should use for classification. The algorithm starts by retrieving a set of perceptual objects \mathcal{O} from perception, with no modulated attention in perception (line 2). Each object in $o \in \mathcal{O}$ is a perceptual object having regions restricted to single points their spaces.

For example, consider an agent trained in screwdrivers classification running this algorithm (i.e. its conceptual system includes concepts about different types of screwdrivers, including information about their wholes and parts). Assume this agent is looking for Philips and flat screwdrivers. Consider we show this agent a box with different screwdrivers. At line 2, the agent would visually inspect the box and use its knowledge about screwdrivers to segment the visual input and creating a perceptual object representing holistic qualities of each segment, possibly corresponding to whole screwdrivers (assuming generic segmentation rules).

For each perceptual object O , the algorithm checks to which concepts in C' they are holistically similar (e.g. the agent tries to check if each segment is similar to the concept of Philips or flat screwdriver). More specifically, for each $o \in O$ and each $C \in C'$, the algorithm fetches its holistic fragment C^h (lines 5-6). It checks if the perceptual object o falls into C^h , by geometric inclusion in accordance with the metric in C^h . If that is the case, the algorithm adds to K a new subsumption map linking o (as a conceptual object) and C . This link serve as a hypothesis that o is a holistic instance, in a loose sense, of C . This link will be confirmed in the next phase. The algorithm outputs the recovered perceptual objects (line 11) to be used in the next phase.

Following in the same example, the agent running this algorithm would finish the execution of this algorithm by having in its conceptual system a set of perceptual objects (i.e. segments) *holistically* classified as being Philips or flat screwdriver (i.e. if the holistic qualities of the each seen object matches the holistic qualities defined in the holistic space of each concept). Note that a perceptual object might be subsumed by two concepts in certain contexts. For instance, a given perceptual object might look holistically similar to a Philips and flat screwdriver (since the difference usually lies at the tip of the instrument). In this case, the algorithm assumes that the perceptual object is subsumed by the two concepts. The disambiguation, if possible, would occur based on structural processing, which is the next algorithm.

Algorithm 2 Holistic interpretation algorithm

Require: A conceptual system $CS = (S, C, K, P)$

Require: A set $C' \subseteq C$ of concepts

Require: A signal x_i to be interpreted

```

1: function P-HOLISTIC( $C', x_i$ )
2:    $O \leftarrow \text{perceive}(\emptyset, x_i)$ 
3:   for all  $o \in O$  do
4:     for all  $C \in C'$  do
5:        $C^h \leftarrow \text{holi}(C)$ 
6:       if  $o \subset C^h$  then
7:          $K \leftarrow K \cup \{\kappa : o \rightarrow C\}$ 
8:       end if
9:     end for
10:  end for
11:  return  $O$ 
12: end function

```

The structural interpretation algorithm (Algorithm 3) takes as input the classified holistic objects and tries to confirm their classification also based on their parts. The general strategy is to try to perceive parts that are defined by the structural fragment of the entities subsuming those objects. For instance, if the holistic objects correspond to possibly types of screwdrivers, then the structural interpretation algorithm would try to check if each parts of the perceived objects match the parts defined by each type of screwdrivers.

The algorithm has as inputs the set of holistic objects O , a conceptual system CS and the actual signal x . It processes each holistic object in O separately. Given an object $o \in O$, the algorithm starts by creating a copy $CS' = (S', C', K', P')$ of the underlying conceptual system CS (line 3); it will accumulate the possible parts generated during the structural processing of o . In line 4, it retrieves a set C^{super} of super-entities of o (e.g., Philips or flat screwdrivers), which are associated to o through subsumption relations in K' (which is inherited from K). For each entity $\mathbf{C} \in C^{super}$, the algorithm gets all defined parts in the structural fragment of \mathbf{C} by retrieving all part relations involving \mathbf{C} (line 6). For instance, these would correspond to all parts that one should find in, say, an Philips screwdriver.

The algorithm then tries to find (in the perceptual input) the actual part referenced in each part relation of \mathbf{C} (lines 8-18). For each part relation $\rho \in P_{\mathbf{C}}$, it retrieves the actual projected part concept $\tau(\mathbf{C}^P)$ and its configuration domain $\mathbf{M}^{\tau(\mathbf{C}^P)}$ (line 8-10), forming a temporary concept \mathbf{C}_M^P . This entity and the object in focus o are then used to engage focused perception of the signal (line 11). We assume that the algorithm always retrieves one and only one perceptual object o^P associated with the conceptual part $\mathbf{M}^{\tau(\mathbf{C}^P)}$. This is akin to a person taking a picture of an arbitrary portion of an object in order to find a given part. The picture may or may not coincide with an actual part. Nevertheless, the “picture” is perceived and represented as a new perceptual object $o^P \subset \hat{\mathbf{S}}$; at this point it is considered as a possible part of o . For instance, if the concept Philips Screwdriver defines a particular type of tip, then the perceive function will move attention to the position of the perceived object where the tip might have been. The perceptual object would form even if the whole object in question was, say, a piece of pipe with no definable “tip”. In this case, this particular piece of perceptual object would form, but discarded in the following.

Next, the algorithm aggregates structural information to o^P , forming the object o_m^P (lines 12-13). The structural information m is calculated by the function conf , which defines a value for spatial configuration taking into account the whole o and the configuration domain $\mathbf{M}^{\tau(\mathbf{C}^P)}$ that is being tested. Then, in line 14, the algorithm tests whether o_m^P is classified by \mathbf{C}_M^P (i.e. if the perceived part object is indeed the tip of a Philips screwdriver). This test is based on the property regions and metric defined by \mathbf{C}_M^P . If the test results positive, then the object o^P is added to the set C' (line 15) and a new subsumption relation κ is added between the part object and the part concept (line 16). Moreover, a new part relation ρ between the object o and the considered part object o^P (line 17).

Algorithm 3 Structural interpretation algorithm

Require: A conceptual system $CS = (S, C, K, P)$
Require: A set O of quasi-objects, such that $O \subseteq \hat{S}$
Require: A signal x_i to be interpreted

```

1: function P-STRUCTURAL( $O, x_i$ )
2:   for all  $o \in O$  do
3:      $CS' = (S', C', K', P')$ , such that  $S' = S, C' = C, K' = K, P' = P$ 
4:      $C^{super} \leftarrow \{C \in \text{codom}(\kappa) : \kappa \in K' \wedge \text{dom}(\kappa) = o\}$ 
5:     for all  $C \in C^{super}$  do
6:        $P_C = \{\rho \in P' : \text{dom}(\rho) = C\}$ 
7:       for all  $\rho \in P_C$  do
8:          $\tau(C^P) \leftarrow \text{codom}(\rho)$ 
9:          $M^{\tau(C^P)} \leftarrow \text{struct}(\tau(C^P))$ 
10:         $C_M^P \leftarrow \tau(C^P) \times M^{\tau(C^P)}$ 
11:         $\{o^P\} \leftarrow \text{perceive}(\{o, C_M^P\}, x_i)$ 
12:         $m \leftarrow \text{conf}(o^P, o, M^{\tau(C^P)})$ 
13:         $o_m^P \leftarrow o^P \times m$ 
14:        if  $o_m^P \subset C_M^P$  then
15:           $C' \leftarrow C' \cup \{o^P\}$ 
16:           $K' \leftarrow K' \cup \{\kappa : o^P \rightarrow C^P\}$ 
17:           $P' \leftarrow P' \cup \{\rho : o \rightarrow \tau(o^P)\}$ 
18:        end if
19:      end for
20:      if  $o \subset C$  then
21:         $CS \leftarrow CS'$ 
22:         $K \leftarrow K \cup \{\kappa : o \rightarrow C\}$ 
23:      else
24:         $K \leftarrow K - \{\kappa : o \rightarrow C\}$ 
25:      end if
26:    end for
27:  end for
28: end function

```

At line 20, the algorithm already tried to detect all parts defined by C in the perceived object. So, at this point, the object o corresponds to a point in the whole holistic-structure space that supports C , with all detected and non-detected parts. Precisely at the line 20 the algorithm compares o , with all its classified parts, with the concept C . This is a full holistic-structural comparison, where both whole and parts are compared. For instance, at this point the perceived object in the screwdriver box is compared with the concept of Philips screwdriver, taking into account the whole and parts. If both entities are similar enough, then $o \subset C$. For instance, if o would have been a flat screwdriver, then it would fail to be similar to a Philips screwdriver regarding the tip. In the case where classification is successful, then the conceptual system CS' with the newly defined entities and relations becomes the new underlying conceptual system CS and a new subsumption relation be-

tween o and C is added to K (lines 21-22). Otherwise, any previous hypothesis that o is subsumed by C (i.e. made as result from the holistic processing phase) is removed from CS . For instance, if a piece of pipe was mistakenly classified as a screwdriver in holistic processing, then this classification would be removed by removing the appropriate κ relation from K .

The structural interpretation algorithm repeats the process for each quasi-object in O resulting from holistic processing. After both algorithms run, the quasi-objects become complete holistic-structural objects, having holistic properties and parts described in the conceptual system CS .

Eventually, the same object might still be subsumed by two different entities at the end of the whole process. This can be the case in a real situation, when a given object is too far in the field of view so that the agent cannot decide whether the object is, for instance, a Philips or a flat screwdriver.

6.2.2 Implementation requirements

There are many ways in which this algorithm can be implemented. However, there are some requirements and limitations. The interpretation algorithm can be implemented in any problem in which similarity and holistic-structure processing is involved. The first requirement is that the implementation should provide the domain knowledge in the form of concepts and objects in holistic-structure spaces. Also, these entities should be represented in a conceptual system.

Another requirement is that the implementation has to provide the actual processing algorithms that implement perception. There is no restriction in which kinds of algorithms they are; they can range from signal processing to neural networks. As we mentioned in previous section, the output space of these processing algorithms must comply with the topological constraints of the perceptual space \hat{S} in the conceptual system, particularly its metric. The algorithm also requires a static perceptual input; the processing signal cannot change during classification.

6.3 Example: well log interpretation task

The previous interpretation algorithm can be used to provide a basis for automation of tasks where holistic-structural processing of raw-data is involved. This kind of task is predominant in certain domain, such as Geology. In this section, we show how our algorithm provide a solution to the problem of interpretation of gamma ray well logs, a relevant task in Petroleum Geology domain.

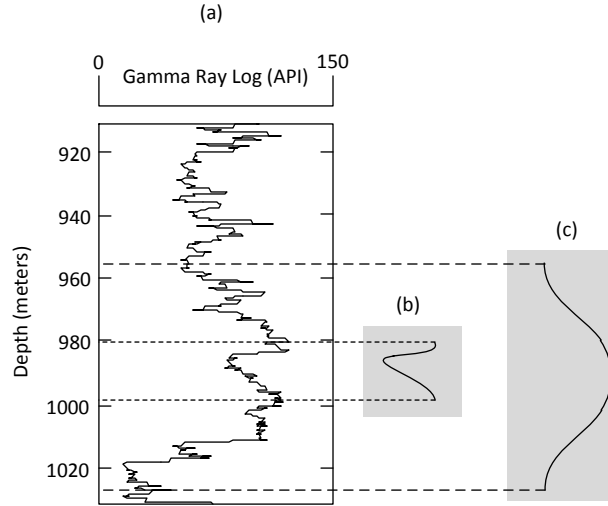
A gamma ray log is a one-dimensional signal originated from the capture of gamma radiations emitted by the rock strata crossed by the drilling. These radiations are captured by a probe along the depth of the well (Figure 6.2a). Since the features of the gamma ray signal are related to the rock geological properties, gamma ray logging is an alternative solution that is used in the many cases when direct rock samples (named *cores*) would be too difficult or too expensive to retrieve. The geologist's task then consists in plotting the well log signals as two-dimensional log charts and to visually inspect them, looking for curve patterns indicating relevant geological features. However, manually analysing gamma ray logs is a cumbersome task that requires a considerable amount of expertise and that could benefit from automation.

The stratigraphic interpretation in focus here consists in identifying two geological features: *sequences*, and *parasequences*. A sequence is a series of genetically related layers whose limits are defined as a response to a relative decline of the sea level (WAGONER et al., 1990). A sequence is usually identified in a gamma ray log as Gaussian-like curve having around 100 meters in length. A Gaussian curve in this context correspond to a log curve having a progressive variation followed by a regressive variation (Figure 6.2c), which correlates with relative changes in sea level. Parasequences, on the other hand, indicate small cycles of sea level variation within sequences. They are usually recognized in gamma ray logs as funnel-like curves having around 30 meters in length. A funnel curve is a well log curve characterized by a fast regressive variation followed by relatively smooth progressive one (Figure 6.2b). The main point here is that parasequences can be seen as *parts* of sequences. *Parasequences form the fragments of the well log curve of sequences.*

As argued about other similar tasks in Geology (SHIPLEY et al., 2013), it is reasonable to assume that geologists use parasequences as structural components to interpret sequences. Furthermore, geologists use previously learnt mental models for interpreting newly presented objects. For instance, the curve pattern associated with a sequence detected in a particular well log may be used as a template for detecting a similar sequence with a similar curve pattern in a neighbour well. This approach allows us to interpret the cognitive task of interpreting sequences and parasequences as a task of measuring the holistic-structural similarity between a learnt prototypical well and the new well.

Having this in mind, we propose an algorithmic solution to well log interpretation. We have addressed part of the problem already in the past using conceptual spaces (see FIORINI; ABEL; SCHERER, 2013). However, given the original theory of conceptual spaces is not clear in how to represent part relations, this aspect was not paid enough attention. Which is not ideal, given the importance of the part-relation between sequences and parasequences. Thus, our contribution with this work is that we demonstrate how the notion of similarity and holistic-structural processing casts a new way of looking at stratigraphic interpretation.

Figure 6.2 – A small section of a gamma ray well log with the typical geological features identified in well log analysis: (a) A 100m section of a gamma ray log; (b) The typical funnel-like curve pattern for parasequences. The dotted lines indicate an occurrence of a parasequence identified in the log; (c) The curve pattern for sequences showing a Gaussian curve shape. The dotted line also indicates an occurrence of a sequence.



Source: adapted from Fiorini, Abel and Scherer (2013).

In essence, we instantiate Algorithm 1 so that it tries to recognize elements of a set of sequences and parasequences in a new well log by means of holistic-structural processing. This instantiation has two general assumptions: (a) a conceptual system containing knowledge about sequences and their parasequences found in a given geological area, possibly learnt from a set of geographically close well logs; and (b) the new well log to be interpreted is geographically close and aligned to the ones used to create the conceptual system. The interpretation strategy has two steps. The first is to extract objects from the new well log and check whether these are *holistically* similar to the sequences we want to find in the new well log (by means of Algorithm 2). The second step is to check if the parts of the extracted objects are similar to the parasequences determined by each holistically recognized sequences (by means of 3). The result should be a set of recognized sequences and parasequences in the new signal.

We start by defining a conceptual system for this problem. Let

$$G = (S, C, K, P)$$

be a conceptual system for our well log interpretation task. We start by defining which domains compose the perceptual space $\hat{S} \in S$. For the sake of this example, we will assume the domains **Shape^D**, **Lenght^D** and **Position^D**; so that

$$\hat{S} = \mathbf{Shape}^{\mathbf{D}} \times \mathbf{Lenght}^{\mathbf{D}} \times \mathbf{Position}^{\mathbf{D}}.$$

The domain **Shape^D** is a log shape domain, in which the metric measures curve similarity: equal curves have zero distance. The domain **Lenght^D** is a one-dimensional space representing the curve length (in the depth dimension). The domain, **Position^D** is a four-dimensional space representing the global two-dimensional positions of points at both ends of the curve. Here, position is defined in relation to a global datum. Thus, points in $\hat{\mathbf{S}}$ represents different specific shapes of well log curve in different sizes, placed at different points in the global positioning system.

The set of entities C has some concepts representing the domain. The first concepts are **Parasequence** and **Sequence**. The concept **Parasequence** is taken here to be a holistic concept, with no parts. Thus, it is a simple subspace of $\hat{\mathbf{S}}$. Its properties regions are not important for this example; we just need to establish its prototype as a typical parasequence, which, as described previously, typically consists of a funnel-shaped curve with around 30 meters in length (e.g. Figure 6.2b).

A **Sequence** has parts. It consists of a product of its holistic and structure spaces, so that **Sequence** = **Sequence^h** \times **Sequence^s** has both a holistic and a structure fragment. The holistic fragment is a subspace of the perceptual space $\hat{\mathbf{S}}$. Its internal structure is not relevant for this example, only its prototype. Similar as we did with parasequences, let us assume that its prototype is a typical sequence, which corresponds to Gaussian-shaped curve with around 100 meters (e.g. Figure 6.2c). More importantly, sequences have parasequences as parts, so we have that the set of part relation P has a function

$$\rho : \mathbf{Sequence} \rightarrow \tau_{\mathbf{Parasequence}'}^{\mathbf{Sequence}}(\mathbf{Parasequence}).$$

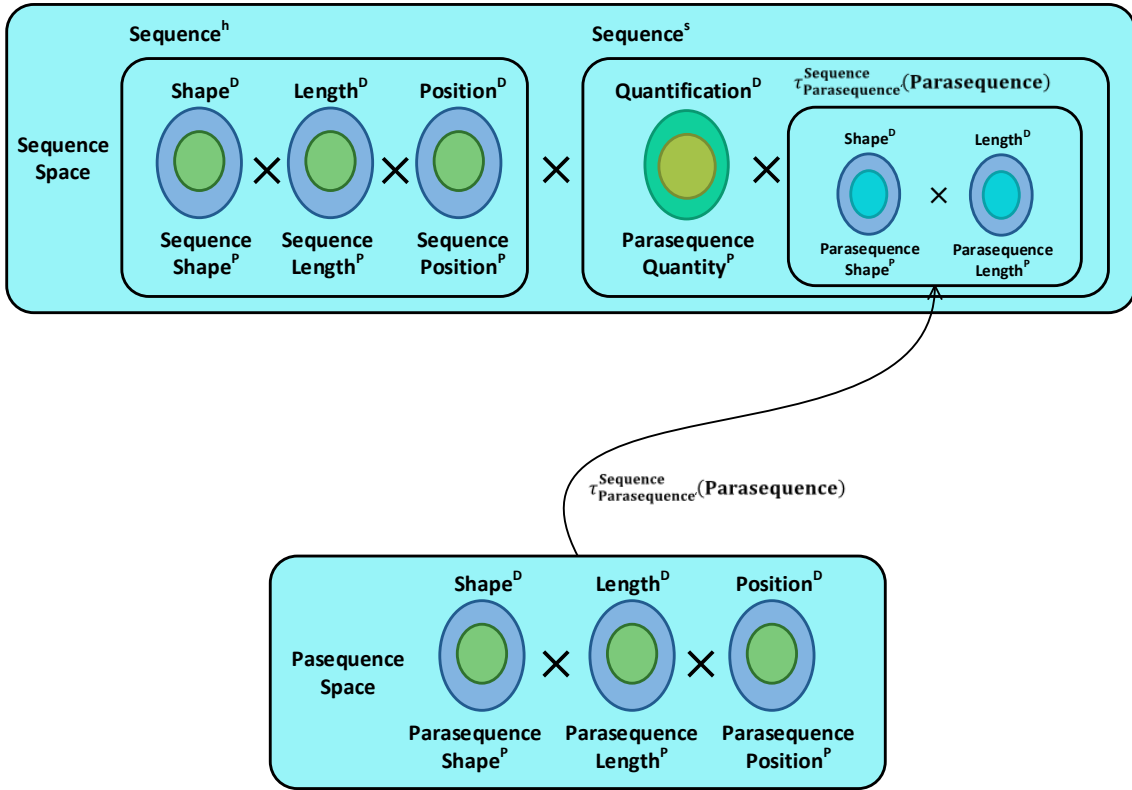
Based on the part relation, we can define the structure space **Sequence^s** as the product space

$$\mathbf{Sequence}^s = \mathbf{Quantification}^D \langle \mathbf{ParasequenceQuantity} \rangle \times \tau_{\mathbf{Parasequence}'}^{\mathbf{Sequence}}(\mathbf{Parasequence}),$$

where **Quantification^D** $\langle \mathbf{ParasequenceQuantity} \rangle$ denotes a property region within on a quantification structure domain representing the possible quantity of parasequences in a sequence. This property represents the possible quantity of parasequences a sequence might have. The relation $\tau_{\mathbf{Parasequence}'}^{\mathbf{Sequence}}(\mathbf{Parasequence})$ represents the actual part relation, which is materialized as the space **Parasequence'** resulting from a projection/filtering of the space **Parasequence**, removing information about the domain **Position^D**. This relation defines the concept of parasequence-as-a-part-of-sequence.

These concepts are not specific enough for stratigraphic interpretation, as they represent only general knowledge about sequences and parasequences. The interpretation task needs case-specific entities in order to function. Thus, given a geological area, our model presupposes ideal models of the sequences and parasequences present in that area. Such ideal models correspond to concepts mapping to subregions of **Sequence** and

Figure 6.3 – Conceptual model of sequence and parasequence. The position domain in **Parasequence** space has been “filtered out” by the dimensional filter (arrow), such that the filtered space takes the place of parasequence in the product space that composes **Sequence**.



Source: the authors.

Parasequence. In a given interpretation task, a new well log will be presented to the system and its sequences and parasequences will be extracted and compared to the ideal model through similarity comparison. The similarity between the ideal entities and the new entities will define the well interpretation.

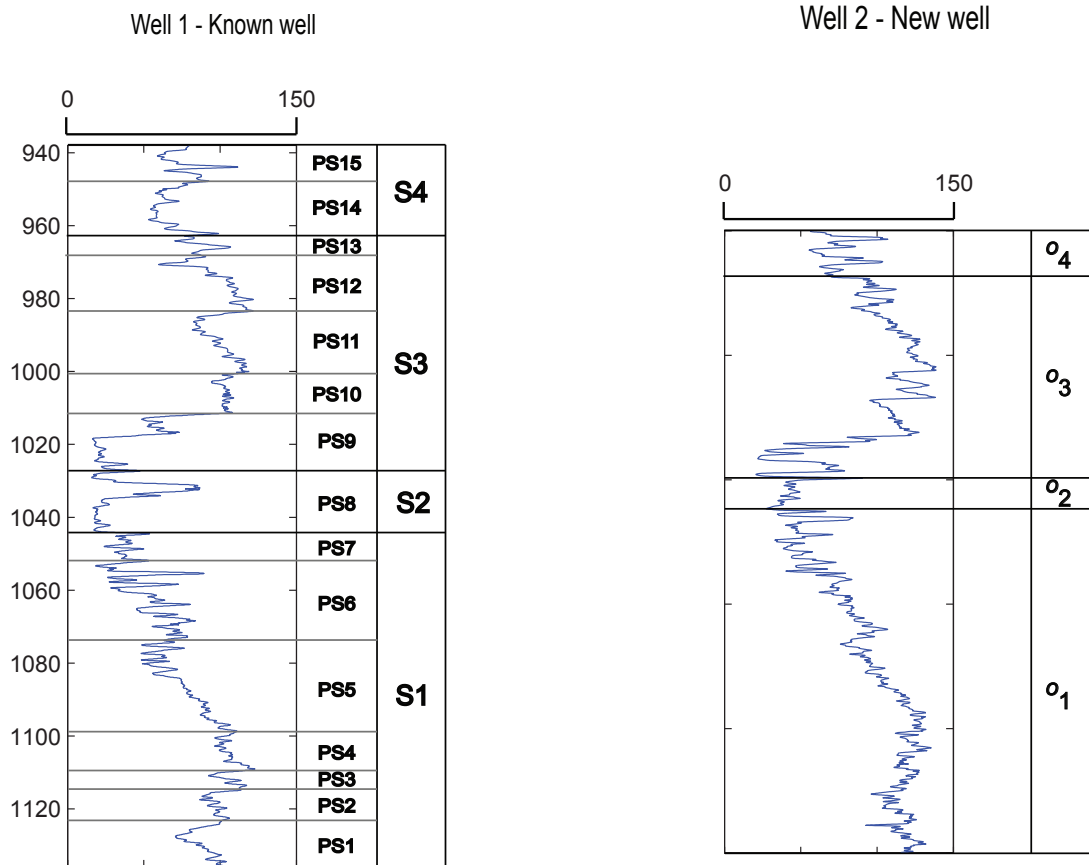
Let $x_i^1 \in X$ correspond to the gamma log of Well 1 in Figure 6.4. Let us assume an ideal conceptual model for sequences and parasequences where the prototypes correspond to the trained objects from x_i^1 . So,

$$\{S1, \dots, S4, PS1, \dots, PS15\} \subset C$$

and

$$\left\{ \begin{array}{l} \kappa_{S1} : S1 \rightarrow \mathbf{Sequence}, \dots, \kappa_{S4} : S4 \rightarrow \mathbf{Sequence}, \\ \kappa_{PS1} : PS1 \rightarrow \mathbf{Parasequence}, \dots, \kappa_{PS15} : PS15 \rightarrow \mathbf{Parasequence} \end{array} \right\} \subset K.$$

Figure 6.4 – The two well logs used in the demonstration used here. *Well 1*: Known gamma ray well log and interpretation used to guide interpretation. *Well 2*: new gamma ray log to be interpreted, with extracted holistic object. Interpretations are roughly based on (WAGONER et al., 1990).



Source: the authors.

Sequences have part relations with parasequences. For instance, the sequence **S3** has parasequences **PS9** to **PS13** as parts, so we have that

$$\{\rho_{\text{PS9}}^{\text{S3}} : \text{S3} \rightarrow \tau_{\text{PS9}'}^{\text{S3}}(\text{PS9}), \dots, \rho_{\text{PS13}}^{\text{S3}} : \text{S3} \rightarrow \tau_{\text{PS13}'}^{\text{S3}}(\text{PS13})\} \subset P.$$

For instance, sequence **S3** would be represented as a conceptual space

$$\begin{aligned} \text{S3} &= \text{S3}^{\text{h}} \\ &\times \text{Q}^{\text{S3}} \times \tau_{\text{Parasequence}'}^{\text{S3}}(\text{Parasequence}) \\ &\times \text{M}^{\text{PS9}} \times \tau_{\text{PS9}'}^{\text{S3}}(\text{PS9}) \times \text{M}^{\text{PS10}} \times \tau_{\text{PS10}'}^{\text{S3}}(\text{PS10}) \times \dots \times \text{M}^{\text{PS13}} \times \tau_{\text{PS13}'}^{\text{S3}}(\text{PS13}). \end{aligned}$$

The subspace $\mathbf{S3}^h \subset \mathbf{Sequence}^h$ is the holistic space of concept $\mathbf{S3}$. The subspace $\mathbf{Q}^{\mathbf{S3}} \times \tau_{\mathbf{Parasequence}}^{\mathbf{S3}}(\mathbf{Parasequence})$ is a constrain on the structure space of $\mathbf{Sequence}$ that was inherited through the subsumption relation $\kappa_{\mathbf{S1}} : \mathbf{S1} \rightarrow \mathbf{Sequence}$. It represents general information about the parasequences in $\mathbf{S3}$. Finally, the structure space includes configuration information (i.e. position) of each of its parasequences, through configuration properties and a set of filters on each parasequence. The general idea is that the prototype of $\hat{p} \in \mathbf{S3}$ corresponds to the curve in Well 1. It is important to note that we are taking $\mathbf{S3}$ to be a region in a conceptual space, where \hat{p} represent a typical occurrence of $\mathbf{S3}$ in a given well and other points correspond to other occurrences of $\mathbf{S3}$.

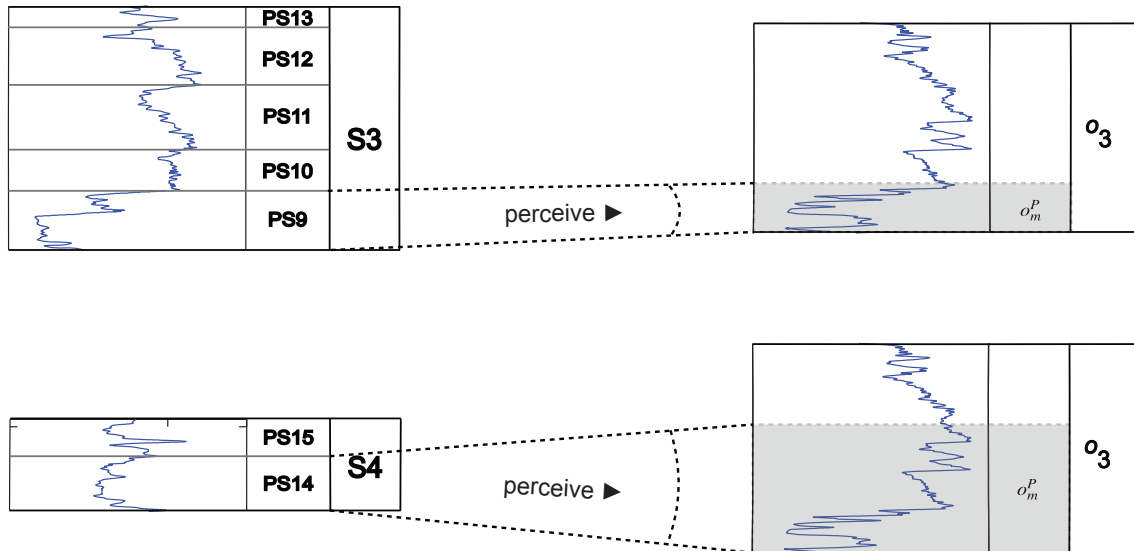
Having the appropriate conceptual system G in place, we can demonstrate an example of well log interpretation using our holistic-structural algorithm. Let $w_i^2 \in X$ denote the gamma ray curve in Well 2. We assume all signals in X to be described according to the same datum. Let $H \subset C$ be a set of sequences we want to find in w_i^2 , so that $H = \{\mathbf{S1}, \mathbf{S3}, \mathbf{S3}, \mathbf{S4}\}$. The algorithm starts by calling the holistic processing (P-HOLISTIC) algorithm, passing the set H and the signal w_i^2 .

Algorithm P-HOLISTIC starts by operating an unfocused perception on w_i^2 (line 2 in Algorithm 2). Considering purely perceptual features (such as abrupt signal changes), the signal is segmented and the segments are represented as objects in the perceptual space $\hat{\mathbf{S}}$. Let us assume that it extracts the blobs shown in Well 2. These are *holistic* sequences, or quasi-sequences, in the sense that they only represent the holistic aspects of possible sequences being extracted from the signal. All quasi-sequences are returned into the set O . Assume that the algorithms returns four (i.e., o_1, o_2, o_3, o_4) candidate holistic sequences in O (as seen in Well 2 in Figure 6.4). From this point on, the processing is straightforward. For each candidate holistic-sequence $o \in O$, the algorithm tests its similarity against the *holistic fragment* \mathbf{C}^h of each prototypical sequence in \mathbf{C} . If the similarity is high enough (line 6), a subsumption relation is created between them. For instance, let us take sequences $\mathbf{S3}$ and $\mathbf{S4}$ as example, together with o_3 extracted from Well 2. The distance of object o_3 to sequence $\mathbf{S3}$ in perceptual space is certainly small enough to place o_3 into the holistic fragment of $\mathbf{S3}$. This is not the case with o_3 and $\mathbf{S4}$, given they are quite far apart in the length and shape domain and thus dissimilar. Nevertheless, just for the sake of this example, let us assume that o_3 was finally categorized as subsumed by $\mathbf{S3}$ and $\mathbf{S4}$, so that P-HOLISTIC has changed the set H by including the new subsumptions relations:

$$\{\kappa_{o_3} : o_3 \rightarrow \mathbf{S3}, \kappa_{o_3} : o_3 \rightarrow \mathbf{S4}\} \subset K.$$

Finally, the algorithm returns the objects found (line 11), all of them O holistically classified. Some of them might fail classification, so that they will be ignored by the structural processing that follows.

Figure 6.5 – This figure shows how the structural processing algorithm checks for possible parts in an object. Object o_3 is assumed to be subsumed by **S3** and **S4** after holistic processing. During structural processing, the algorithm then tries to find a part of the whole object that corresponds to a part as described by the supposed subsuming concept. In this example, **S3** determines that the parasequence **PS9** would be found in the lowest part of o_3 . The algorithm tries to perceive a curve segment in the same position and then check for similarity. In this case, they are clearly similar, so that o_m^P would be categorized as an example of **PS9**. However, the same does not happen with **S4**, as the perceptual object o_m^P generated by the attempt of perceiving **PS14** is clearly dissimilar from the later.

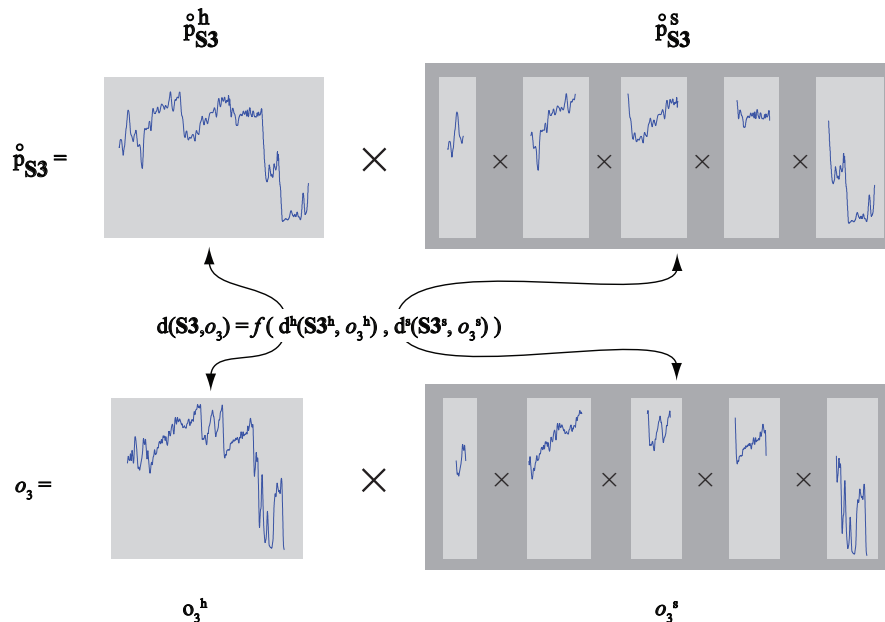


Source: the authors

The interpretation algorithm then calls the P-STRUCTURAL algorithm, passing the set of holistic sequences in O as well as the well $\log x_i^2$ that is being interpreted. This algorithm will try to confirm the existence of particular parasequences defined by each sequence subsuming the objects in O . At the end, we expect to have holistic-structural comparisons between the sequences that are being interpreted and the sequences present in the conceptual system. For the sake of this example, let us assume that at line 2 of P-STRUCTURAL, there only remains o_3 still left in O to be processed. So, let $o = o_3 \in O$ at line 2. The algorithm then instantiates a temporary copy of the conceptual structure G , called G' . In line 4, it follows the subsumption relation in K to check what are the super-concepts of o_3 . We established in the previous paragraph that o_3 is holistically subsumed **S4** and **S3**, so at line 4 we have that $C^{super} = \{\mathbf{S4}, \mathbf{S3}\}$.

First isolating $C = \mathbf{S3}$, the algorithm will in turn isolate its defined parts, which are $P_C = \{\mathbf{PS9}, \dots, \mathbf{PS13}\}$ according to the set of part relations P (and Figure 6.4). The algorithm then takes each parasequence, its related structure property, and tries to find each of them perceptually in the signal (lines 8-18). The idea is that having, for instance, $o = o_3$, $C = \mathbf{S3}$ and $C_M^P = M^{\mathbf{PS9}} \times \tau_{\mathbf{PS9}}^{\mathbf{S3}}(\mathbf{PS9})$, the algorithm will inspect the visual area within o_3 that would be similar to **PS9** if o_3 was really corresponding to sequence **S3** (see

Figure 6.6 – In order to decide whether, for instance, $o_3 \subset \mathbf{S3}$, they have to be compared in the holistic-structure space of $\mathbf{S3}$. This picture shows a visual depiction of how this can be done according to the theory of conceptual spaces. Assume Well 1 corresponds to the prototype point $\overset{\circ}{p}_{\mathbf{S3}} \in \mathbf{S3}$ and that o_3 also corresponds to a point. So, their similarity is given by a function of their metric distance in holistic and structure spaces (in the support space of $\mathbf{S3}$). If they are close enough so that o_3 falls within the property regions of $\mathbf{S3}$, then o_3 is said to be included in $\mathbf{S3}$



Source: the authors.

Figure 6.5). The inspection consists in projecting that part of the signal as a configured perceptual object o_m^P and in checking if it is included in $\mathbf{S3}$. A visual inspection in Well 1 (Figure 6.4) shows that the lower portion of $\mathbf{S3}$ corresponding to the parasequence $\mathbf{PS9}$ is similar to the well curve in the lower portion of o_3 . In this case, o_m^P is included as part of o with respect to the concept $\mathbf{S3}$. More specifically, the auxiliary conceptual system G' gets the new $\mathbf{SP9}$ -like parasequence denoted by o^P (line 15), which is subsumed by $\mathbf{SP9}$ (line 16) and mapped as part of o , in this case, o_3 (line 17).

After that all sequences defined by $\mathbf{C} = \mathbf{S3}$ have been searched in the curve and classified, the sequence o_3 has all possible detected parts according to $\mathbf{S3}$ represented as part relation in P' . Thus, at this point, the sequence o_3 has holistic and structural fragments (Figure 6.6). The algorithm then tests whether o_3 is embedded in the prototypical sequence $\mathbf{S3}$ (line 20). The metric used in this case here is a distance of the holistic and structural metrics (Figure 6.6). Since this is the case in this example, then the temporary conceptual system G' becomes the present conceptual system (with all new entities and relations), and the subsumption link between o_3 and $\mathbf{S3}$ is reinforced (line 22).

The algorithm advances for the next candidate concept to o_3 in C^{super} , which is **S4** in this example. As before, the algorithm isolates its defined parts, e.g. $P_C = \{\mathbf{PS14}, \mathbf{PS15}\}$, then takes each part, its related structure property, and tries to find each of them perceptually in the signal (lines 8 to 17). If $o = o_3$, $C = \mathbf{S4}$ and $C_M^P = M^{\mathbf{PS14}} \times \tau_{\mathbf{PS14}'}^{\mathbf{S4}}(\mathbf{PS14})$, then the algorithm inspects the visual area within o_3 that should be similar to **PS14** if o_3 was really corresponding to a sequence **S4** (see Figure 6.5). Again, the inspection consists in projecting that part of the signal as a configured perceptual object o_m^P and in checking its similarity to $C = \mathbf{S4}$. Another visual inspection in Well 1 (Figure 6.4) shows that the lower portion of **S4** *does not* correspond to the same well curve of the lower portion of o_3 . In this case, o_m^P is not included as part of o with respect to the concept **S4** (lines 14 to 18 are skipped). This implies that at least one part of o_3 is not similar enough to a part of **S4** (the parasequence **PS13**, in this case) to be included as part of o_3 as subconcept of **S4**.

If a part fails to be found, as in the previous case, the object o , in this case, object o_3 , would have a missing part when participating in the holistic-structural comparison with **S4**. In this example, it is expected that the metric and the threshold used are restrictive enough to avoid this classification.

When the structural processing runs on all holistic objects, we expect as a result, a set of holistic-structural objects corresponding to the sequences in Well 2, interpreted according to the entities learnt according to Well 1. The general idea is that any ambiguous classification generated in holistic processing of the wells is ironed out in structural processing.

6.4 Summary and final remarks

Our objective in this chapter was to demonstrate that the notions of holistic-structure spaces contribute to the solution of problems in specific domains. There are many ways in which to use holistic-structure spaces to specify problem-solving algorithms. In this chapter, we have shown a particular example of how these could be used effectively for stratigraphic interpretation.

As we suggested at the introduction of this chapter, the interpretation algorithm can find application in problems and situations where object recognition and classification might benefit from similarity information of parts and wholes. Apart from the application we already discussed in Section 5.7, there are other possible applications of this framework. For instance, take robot self-localization techniques based on particle filtering (e.g. THRUN et al., 2001). In brief, this technique includes a step where a series of “particles” are generated, representing possible positions in which the robots could be located in a known map; then testing how close they correspond to the actual location of the robot. This is done by simulating the perceptual input as if the robot was positioned where a particle is and then matching the simulation with the actual data perceived by the robot.

If they are similar enough, the position of the particle is elected as a possible position of the robot, since they are “seeing the same things”. Similarity is usually calculated by comparing a set of low-level features extracted from the robot sensors and from the simulation, or even directly comparing simulated and actual sensor values. Eventually, a great number of features and particles have to be matched, which is computationally costly. Our algorithm could speed up this process by bringing the similarity comparison a step up in the semantic ladder. For example, rather than simply comparing sensor data, the technique could include a step of object recognition in the simulated and actual sensor data and compare recognized objects in a conceptual space. If the simulated objects are similar to the actual seen object, the particle corresponds to the position of the robot. More importantly, each matching between a simulated particle and the robot could be done first in the holistic space, and only the more encouraging matches going to more detailed structural similarity comparison. The speedup could come from the need to use less particles to disambiguate the position of a robot (given the higher semantic content of the matching) and from the use of simpler holistic matching to discard particles that are outright implausible, saving more plausible particles for further processing.

Other possible applications include face recognition, where both holistic and structural information determinant, and recognition and classification of sequential data. For instance, recognizing if and how two audio files are similar include measuring similarity between the whole aspect of the audio and also comparing the individual parts to check how similar they overlap. This is particularly interesting to detect plagiarism in music, for instance.

The algorithm has a series of limitations. First, it is not able to process holistic and structure features in parallel; holistic processing occurs before structural processing and determines the results of the later. It would be interesting to have an algorithm without this restriction, in such a way that both holistic and structural processing could determine each other. For instance, it would be interesting to allow the recognition of a part of an object to trigger holistic processing. Another (related) limitation is that the user have to indicate what are the concepts and objects that will be used in classification. This input can be seen as a sort of *conceptual attention*, which isolates part of its conceptual system to be used for classification. However, the algorithm does not have any control mechanism to refocus this attention, if needed. We believe such algorithm would need more research in how attention and memory can be orchestrated with holistic-structural processing.

7 CONCLUSION

Connecting computers to reality is a challenging task. It involves linking data structures to recurrent patterns of sensory information. Perhaps the biggest obstacle to overcome is the different nature of the many levels of information representation, starting from raw sensory input, up to symbolic information. Creating the *bridge* over this semantic gap involves two separate problems: the processing challenge and the representation challenge. The former deals with the mechanisms necessary to transform information from sensory input up to conceptual structures. The later deals with how to represent the information needed for those mechanisms, in such a way that the link from lower to upper levels of representation is kept. For instance, computer vision systems frequently have to deal with the semantic gap when implementing object recognition or classification. Object recognition involves processing pixel data (e.g. images), extracting segments of the original data and interpreting them as being a member of a concept or a specific object. The interpretation mechanism involves a step where sets of pixels are mapped to character strings denoting the recognized object or concept. The semantic gap between these pixels and strings makes the processing difficult, *ad hoc* and cumbersome. *The main motivation of this thesis is to address this gap.*

There are many ways to address the semantic gap in computers (FIORINI; ABEL, 2010). However, when looking at how humans do it, there is a particular aspect not well addressed in computational approaches, namely, *similarity*. Processes involved in the cognition of object recognition and classification involve similarity comparisons between conceptual structures and the perceptual input. For instance, humans recognize that a newly seen object as an apple by measuring the similarity between the newly seen object and the mental representation of an apple; if they are similar enough, the object is classified as an apple. Peter Gärdenfors suggests a framework of how represent similarity between conceptual representations of such entities, namely, the theory of conceptual spaces (GÄRDENFORS, 2000), which is the background theory of this thesis. However, the cognition of object recognition involves another important set of mechanisms: holistic and structural processing (PEISSIG; TARR, 2007). These two processes articulate the interpretation of perceived entities from a holistic (whole) point of view and from a

structural (part) point of view. They essentially implement part-whole reasoning in recognition. As we have reviewed in this thesis, whilst computer systems do implement part-whole reasoning, they do not integrate these mechanisms with similarity. We argue that such integration would be beneficial in addressing the semantic gap in computers, since both similarity and holistic-structural processing provide a way of comparing similarity between parts and wholes, which cannot be easily done in symbolic and subconceptual levels. In this thesis, we leverage the representation power of conceptual spaces in order to provide the similarity-based representation framework that can bridge the semantic gap. However, conceptual spaces theory do not support representation of parts and wholes. *Providing this support constitutes the main specific motivation driving this work.*

This thesis has three main sets of contributions: the *theory* (Chapter 4), the *formulation* (Chapter 5) and the *algorithm* (Chapter 6). Since they are relatively independent from each other, we shall review the main contributions of each set separately.

7.1 The Theory

Chapter 4 presents the main contribution of this thesis: a *concept representation theory that provides a new view on representation of parts, wholes and their relations in computer systems*. The core proposal of the theory consists of representing the relationship between wholes and parts in conceptual spaces as products of their holistic properties and the relevant properties of their parts. Also, we define *structural properties* that qualify the relation between part and whole. This product lays in a special conceptual space formed by the product of *holistic* and *structure spaces*, which are the main epistemological constructs of our representation. The relation between the conceptual space of the whole and the part is also mediated by *dimensional filters*, which can filter out properties of the part that are not relevant for the whole.

This set of constructs allowed relating and contrasting our theory with other aspects of part relations. We argue that structure and holistic spaces can represent the two main types of part relation identified in linguistics, namely complex and collective relations. We argue that both types can be differentiated by how we define the structural properties governing specific part relations. We also show how the context representation mechanism of conceptual spaces (i.e. giving weights to domains in the similarity comparison) allows us to represent situations where there are contextual shifts in structural similarity, such as when certain parts are more relevant than others in similarity comparison. We also relate our theory with common representation issues in part relation representation. We argue that the part relation is not inherently transitive, such that our theory takes part as a perceived direct relation between the part and the whole. By this way and by the instrument of dimensional filters, we show we can avoid transitive closure problems in wholes containing many complex parts.

To our knowledge, no framework addresses the issue of representation of similarity between parts and wholes from a similar point of view. In one side, symbolic level approaches are too brittle to combine part-relations with similarity relations. On the other side, subconceptual approaches do not address the issue from the conceptual point-of-view; rather, they generally tend to be too specific about how to solve the issue. We believe the theory could complement symbolic approaches by providing a better grounding to them; and associationist approaches by providing a sort of “conceptual scaffolding” in which they can be organized. More specifically, *the theory provide some of the basic principles needed to understand and bridge the semantic gap in computer systems that have to interpret sensor data*, such as in robotics and intelligent systems in natural domain. Yet, the theory is agnostic to what technology is to be used to implement such principles as holistic and structural similarity. This is why, we believe, this theory is main contribution of this thesis.

In Chapter 4, we also delve even deeper in the recurrent ontological issues of part relation. While dealing with ontology is a secondary objective of this thesis, we believe that the theory might provide a background theory to connect the theory of qualities and the theory of parts. We informally demonstrate this by discussing how our theory would provide the background for Guarino’s (2013) mereological theory of qualities in formal ontology. In general, we believe the theory is a step in the direction of having more influence of cognitive phenomena in part-whole representation. Simons (2006) rises this issue by claiming that research in mereology should put begin to less emphasis in algebra and more emphasis on how the different domains deal with part relations. We believe that, by bringing up more cognition to this discussion, we help to achieve the goal of having more effective part-relation representation frameworks.

Finally, in Chapter 4, we provide a sketch of how holistic and structure spaces could be used in a cognition-inspired object interpretation algorithm. This is just a first step in developing a general algorithm that could implement full holistic and structural processing as they happen in cognition, using holistic and structure spaces. We come back to this issue in the following sections.

Finally, this theory is a general contribution to Conceptual Spaces. Conceptual spaces are showing promising applications in artificial systems, ranging from robotics to software integration. Until now, even if some works touch on the issue, the representation of part relations in conceptual spaces was still an open problem. We believe this might have happened mainly because, given the cognitive bias of conceptual spaces, any sounding extension to the theory also requires some grounding in cognition. This grounding was lacking in reference to part relations. By investigating how processes like similarity and holistic-structural processing interact in cognition, we were able to produce a sound theory of how the information supporting these processes can be represented in conceptual spaces. As such, we believe intelligent applications employing conceptual spaces will

benefit from our approach when there is a need to represent part-relations. The work of Chella, Frixione and Gaglio (1997) is a typical example. The authors represent knowledge about part relations uniquely as predicates in the symbolic level. We believe our approach could be a way of providing cognitive semantics to these predicates. It could allow for more processing about part relations to take place in the lower levels of the cognitive architecture, possibly lowering the chance of interpretation errors rising to the symbolic level.

7.2 The Formulation

In Chapter 5, we propose a *metric spaces formulation* of the theory. It is by no means definitive or unique, but we believe it captures the main notions of the theory in a plausible mathematical framework, making it more amenable to computer implementations. Thus, the main contribution of the formulation is this: *a formal mathematical representation framework for holistic and structure spaces*.

The formulation derives from early work by Aisbett and Gibbon (2001b), where they propose a formulation for conceptual spaces based on metric spaces and the betweenness relation. We fork their base axioms to propose that conceptual spaces are result of *product spaces* of other conceptual spaces. More specifically, we propose that conceptual spaces are products of holistic and structure spaces represented as metric spaces. The main twist in this formulation is that we define a specific product of metric spaces that can preserve the betweenness relation of the operating conceptual. So, for instance, if the topology of the colour domain dictates that green is between yellow and blue, then product space of the colour domain and other domains that form the conceptual spaces of *car* should preserve this topology, such that a green car is between a yellow car and a blue car.

We employ product spaces to formalize part relation. The driving idea is that part relation is a betweenness-preserving morphism from the product space representing the whole to the result of the dimensional filtering of the part concept. Thus, we first define what a *dimensional filter* is in terms of metric spaces. We specify dimensional filter as a composition of two projections that can generate a “subspace” of a concept in a conceptual space, such that filtered concept might have fewer dimensions and points than the original concept, but with betweenness preserved. Part relation is then a betweenness-preserving mapping from the whole to the filtered part concept. This structure of relations as basis to define holistic and structure spaces. A structure space is then the fragment of the conceptual spaces of a whole that is mapped to conceptual spaces of the parts through part relation morphisms, plus domains defining structure properties. A holistic space

is the complementary fragment of the conceptual space of the whole. The important contribution of this scheme of relations and space is to clarify how the topology of the part spaces behave when operated to form the whole. This is particularly important when developing computer system that implements the theory.

As we discussed in Chapter 5, the formulation could be used as a framework to organize the architecture of computer systems that need to interpret part relations in perceptual information, such as computer vision systems. The notion of quality domains as metric spaces essentially standardize the topological structure of the output space of subconceptual mechanisms, such as signal processing algorithms. Thus, our formulation could serve as a tool to combine and organize low-level perceptual system into a conceptual representation with more embedded meaning. Ultimately, this formulation gives the mathematical framework to articulate the outputs of these perceptual systems.

7.3 The Algorithm

Chapter 6 presents the last major contribution of this thesis: *an object interpretation algorithm that instantiate holistic and structural processing in holistic-structure spaces*. This algorithm is intended to serve as a blueprint to implement actual object interpretation systems in computer systems applied to specific domains.

Chapter 6 starts by using the constructs based on metric spaces of the formulation to define a *conceptual system*, a collection of sets containing concepts and objects defined as metric spaces to be used by an algorithm. The conceptual system also contains two types of relations between entities: part relations, as defined in Chapter 5; and subsumption relations. We define subsumption relations as betweenness-preserving mappings between entities in the conceptual system, such as concepts and objects. It implements the general notion of taxonomic relation present in conceptual modelling. This conceptual system, while simple, presents itself as a separate contribution of that chapter. It could be used with any other algorithm or representation framework implementing the ideas present in this thesis.

Later we specify the actual object interpretation algorithm. The input of this algorithm is a conceptual system and a signal. The output is the same conceptual system containing entities perceived in the signal. The algorithm is essentially the orchestration of two modules. The holistic interpretation module implements the holistic processing strategy by segmenting and extracting blobs of perceptual data (called perceptual objects) from the input signal, representing them as points in the underlying conceptual space and finally matching them against the holistic space of each entity in the conceptual system. These perceptual objects are expected to be objects perceived in the signal. Since holistic interpretation can generate ambiguous classifications for the perceptual objects, the object interpretation algorithm subsequently engages the structural interpretation algorithm on

the holistically-classified perceptual objects. Structural interpretation focus the perceptual attention within each object and tries to extract and interpret parts of the holistically-classified object *present in the signal*. When all possible parts have been found, the structural algorithm measures similarity of the complete perceptual object (with holistic and structural properties) to entities in the conceptual system. Classification (or recognition) is achieved if this complete perceptual object falls within the properties of a given entity in the conceptual model.

One of the main interesting aspects of the algorithm is that it establishes a clear interface between conceptual structures and subconceptual structures, mainly through primitive functions that can interpret low-level data and output them as points in a perceptual space, which is then used in the interpretation process.

We believe the algorithm could be instantiated in any computer system that needs to interpret sensory data, where similarity and part relations have an important role. To exemplify how this instantiation could be done, we take an important object recognition task of Petroleum Geology as a use case. We describe domain knowledge in terms of holistic and structural spaces and assume some perceptual primitives fit for the task. Finally, we also briefly describe how it could be applied in robotics and other domains.

We believe this algorithm is a first step in developing a more sophisticated holistic-structural processing algorithm. We discuss more about the next steps in the following section.

7.4 Open issues and future work

Whereas we believe this work contribute with answers to some questions, we also believe it bring about some new ones.

The first practical aspect is that the proposed formulation in Chapter 5 does not formalize every aspect of the theory in Chapter 4, such as context and types of part-relations. Some of these issues could easily be solved (such as context), but others require more work to be addressed properly (such as part types).

The second aspect relates to processing. In Chapter 1, we bring object recognition as our main motivation and argue that it is a processing and a representation challenge. That is, any representation, to be effective, must have an attached processing scheme. We only provide an outline of how such scheme could look like in relation to holistic and structure spaces in Section 4.8. More work is necessary in order to provide a more general, yet more formal, processing scheme for our theory. Such enterprise would certainly require further investigation on the role of attention and memory in holistic-structural processing.

Part relations and taxonomic relations are tightly related (TVERSKY; HEMENWAY, 1984). Yet we do not fully address how this relation is reflected in our framework. Part of the problems lays within conceptual spaces itself. For instance, it is still an open problem how to represent taxonomic relation between concepts in which the more specific concept have properties in many more specific domains than the generic concept. This problem is more acute when it is assumed, as we did, that certain concepts are represented in distinct conceptual spaces, with different sets of domains. For instance, a car has typically four distinct seats, but it can also have six distinct seats and still be a car. So, the latter category must have more domains corresponding to the added seats. In Chapter 6, we suggested an initial approach to the problem, which involved betweenness-preserving mappings between conceptual spaces of the related concepts. Nevertheless, more research is needed in order to address this issue.

In Section 1.4, we mention certain topics that are close, but not exactly in the scope of this thesis. One of them relate to the problem of mapping between the conceptual construct this thesis proposes (i.e. structure and holistic spaces) and structures in adjacent representation levels, namely the symbolic and the subconceptual level. Conceptual spaces already answers that question in a broad sense. Our theory should provide grounding for semantics of part-whole predicates in symbolic level and structure subconceptual processing of stimuli. However, it is still an open problem how the specifics of this mapping could work. We believe part of the answer could come by investigating how formal ontologies founded on conceptual spaces (such as DOLCE and UFO) deal with part-relations. Previously, it was not clear how they would their theories about part-relation could reflect to relation between quality structures in the conceptual level. Now, this thesis shed some light in how this could be done, taking a recent proposal by Guarino (2013) on quality fields as a connecting point between levels. However, more work is necessary, particularly regarding how reasoning dynamics in symbolic level affects conceptual structures.

Gärdenfors (2000) propose conceptual spaces as a model to cognition. We do not claim our theory can hold the same place. However, we hope it still can bring some new insights in the research of holistic and structural processing. For a start, it is clear more research is needed in how similarity interacts with those two processes in more complex objects. In particular, our work requires that variations in the qualities of the wholes should also affect the qualities of the parts, and vice-versa. Is there a correlation between the degree of similarity of two objects in their holistic space and in their structure space? There are questions that would need further research in cognition in order to be answered.

Finally, the holistic-structural processing algorithm proposed in Chapter 6 is just a first step in the development of a more complete processing scheme. More work is needed in order to fully evaluate the efficacy of holistic-structure spaces as tools for representing concepts in computer systems. As it stands now, the algorithm proposed in Chapter 6

is relatively simple. For instance, it would be interesting to examine if both holistic and structural processing could run independently, in parallel. As they stand now, their execution is serial, with structural processing serving as a filter for holistic processing results. In order to parallelize them, more research work is needed to understanding how experts articulate their attention, memory and recognition in order to solve a given task. On a more general note, it would be also interesting to evaluate the notion of having holistic-structure spaces as blueprints for creating neural networks that are able to process parts and wholes.

Finally, we believe this work adds relevant contributions to the state of the art in concept representation in computer systems. We also have confidence that this thesis opens interesting venues of research in the area.

REFERENCES

- ADAMS, B.; RAUBAL, M. Conceptual space markup language (csml): Towards the cognitive semantic web. In: IEEE INTERNATIONAL CONFERENCE ON SEMANTIC COMPUTING, 2009, Berkeley, USA. **Proceedings...** Washington, USA: IEEE, 2009. p. 253–260.
- ADAMS, B.; RAUBAL, M. A metric conceptual space algebra. In: 9TH INTERNATIONAL CONFERENCE ON SPATIAL INFORMATION THEORY, 2009, Aber Wrac'h, France. **Proceedings...** Berlin, Heidelberg: Springer-Verlag, 2009. p. 51–68.
- AGARWAL, S.; AWAN, A.; ROTH, D. Learning to detect objects in images via a sparse, part-based representation. **IEEE Transactions on Pattern Analysis and Machine Intelligence**, v. 26, n. 11, p. 1475–1490, 2004.
- AISBETT, J.; GIBBON, G. Conceptual spaces as voltage maps. In: MIRA, J.; PRIETO, A. (Ed.). **Connectionist Models of Neurons, Learning Processes, and Artificial Intelligence**. [S.l.]: Springer Berlin Heidelberg, 2001, (Lecture Notes in Computer Science, 2084). p. 783–790.
- AISBETT, J.; GIBBON, G. A general formulation of conceptual spaces as a meso level representation. **Artificial Intelligence**, v. 133, n. 1–2, p. 189–232, dec. 2001.
- ALEXANDER, R. G.; ZELINSKY, G. J. Effects of part-based similarity on visual search: the frankenbear experiment. **Vision research**, v. 54, p. 20–30, feb. 2012.
- ALVAREZ, G. A. Representing multiple objects as an ensemble enhances visual cognition. **Trends in Cognitive Sciences**, v. 15, n. 3, p. 122–131, Mar 2011.
- ARTALE, A.; KEET, C. M. Essential and mandatory part-whole relations in conceptual data models. In: INTERNATIONAL WORKSHOP ON DESCRIPTION LOGICS, 2008, Dresden, Germany. **Proceedings...** [S.l.]: CEUR-WS, 2008. v. 353.
- AUGELLO, A. et al. An algebra for the manipulation of conceptual spaces in cognitive agents. **Biologically Inspired Cognitive Architectures**, v. 6, p. 23–29, oct. 2013.
- AUGUSTINE, E.; SMITH, L. B.; JONES, S. S. Parts and relations in young children's shape-based object recognition. **Journal of Cognition and Development**, v. 12, n. 4, p. 556–572, 2011.
- BAADER, F. et al. **The Description Logic Handbook: Theory, Implementation and Applications**. New York, USA: Cambridge University Press, 2003.

BAHRENBERG, G.; GIESE, E.; NIPPER, J. **Statistische Methoden in der Geographie**. 2nd. ed. Stuttgart-Leipzig: BG Teubner, 1999.

BAR-HILLEL, A.; WEINSHALL, D. Efficient learning of relational object class models. **International Journal of Computer Vision**, v. 77, n. 1-3, p. 175–198, May 2008.

BARSALOU, L. W. Grounded cognition. **Annual Review of Psychology**, v. 59, n. 1, p. 617–645, 2008. PMID: 17705682.

BARSALOU, L. W. Grounded cognition: Past, present, and future. **Topics in Cognitive Science**, v. 2, n. 4, p. 716–724, 2010.

BEHRMANN, M. et al. Independent representation of parts and the relations between them: Evidence from integrative agnosia. **Journal of Experimental Psychology: Human Perception and Performance**, v. 32, n. 5, p. 1169–1184, 2006.

BEHRMANN, M.; WILLIAMS, P. Impairments in part–whole representations of objects in two cases of integrative visual agnosia. **Cognitive Neuropsychology**, v. 24, n. 7, p. 701–730, 2007.

BENEVIDES, A.; GUIZZARDI, G. A model-based tool for conceptual modeling and domain ontology engineering in ontouml. In: FILIPE, J.; CORDEIRO, J. (Ed.). **Enterprise Information Systems**. [S.l.]: Springer Berlin Heidelberg, 2009, (Lecture Notes in Business Information Processing, v. 24). p. 528–538.

BIEDERMAN, I. Recognition-by-components: A theory of human image understanding. **Psychological Review**, v. 94, p. 115–117, 1987.

BORGO, S.; MASOLO, C. Ontological foundations of DOLCE. In: **Theory and Applications of Ontology: Computer Applications**. [S.l.]: Springer Netherlands, 2010. p. 279–295.

BUKACH, C. M.; GAUTHIER, I.; TARR, M. J. Beyond faces and modularity: the power of an expertise framework. **Trends in Cognitive Sciences**, v. 10, n. 4, p. 159 – 166, 2006.

CARBONERA, J. et al. Visual interpretation of events in petroleum geology. In: IEEE INTERNATIONAL CONFERENCE ON TOOLS WITH ARTIFICIAL INTELLIGENCE (ICTAI), 2013, Dulles, USA. **Proceedings...** Washington, USA: IEEE, 2013. p. 189–194.

CHAFFIN, R.; HERRMANN, D. J.; WINSTON, M. An empirical taxonomy of part-whole relations: Effects of part-whole relation type on relation identification. **Language and Cognitive Processes**, v. 3, n. 1, p. 17–48, 1988.

CHELLA, A.; DINDO, H.; INFANTINO, I. A cognitive framework for imitation learning. **Robotics and Autonomous Systems**, v. 54, n. 5, p. 403–408, may 2006.

CHELLA, A.; FRIXIONE, M.; GAGLIO, S. A cognitive architecture for artificial vision. **Artificial Intelligence**, v. 89, n. 1-2, p. 73–111, jan. 1997.

CHELLA, A.; FRIXIONE, M.; GAGLIO, S. Conceptual spaces for computer vision representations. **Artificial Intelligence Review**, v. 16, n. 2, p. 137–152, 2001.

CHELLA, A.; FRIXIONE, M.; GAGLIO, S. Anchoring symbols to conceptual spaces: the case of dynamic scenarios. **Robotics and Autonomous Systems**, v. 43, n. 2-3, p. 175–188, 2003.

CHEN, P. P.-S. The entity-relationship model - toward a unified view of data. **ACM Trans. Database Syst.**, v. 1, n. 1, p. 9–36, mar. 1976.

DESELAERS, T.; PIMENIDIS, L.; NEY, H. Bag-of-visual-words models for adult image classification and filtering. In: IEEE INTERNATIONAL CONFERENCE ON PATTERN RECOGNITION, 2008, Tampa, USA. **Proceedings...** Washington, USA: IEEE, 2008. p. 1–4.

DEZA, E.; DEZA, M. **Encyclopedia of distances**. Dordrecht; New York: Springer Verlag, 2009.

DIETZE, S.; GUGLIOTTA, A.; DOMINGUE, J. Exploiting metrics for similarity-based semantic web service discovery. In: IEEE INTERNATIONAL CONFERENCE ON WEB SERVICES, 2009, Los Angeles, USA. **Proceedings...** Washington, USA: IEEE, 2009. p. 327–334.

DONNELLY, M.; BITTNER, T. Spatial relations between classes of individuals. In: COHN, A. G.; MARK, D. M. (Ed.). **Spatial Information Theory**. [S.l.]: Springer Berlin Heidelberg, 2005, (Lecture Notes in Computer Science, v. 3693). p. 182–199.

DOUMAS, L. A. A.; HUMMEL, J. E. A computational account of the development of the generalization of shape information. **Cognitive Science**, v. 34, n. 4, p. 698–712, 2010.

EDELMAN, S. Representation is representation of similarities. **Behavioral and Brain Sciences**, v. 21, n. 04, p. 449–467, 1998.

FARAH, M. J. Is an object an object an object? cognitive and neuropsychological investigations of domain specificity in visual object recognition. **Current Directions in Psychological Science**, v. 1, n. 5, p. 164–169, 1992.

FIDLER, S.; BOBEN, M.; LEONARDIS, A. Similarity-based cross-layered hierarchical representation for object categorization. In: IEEE COMPUTER SOCIETY CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, 2008, Anchorage, USA. **Proceedings...** Los Alamitos, USA: IEEE Computer Society, 2008. p. 1–8.

FIORINI, S. R.; ABEL, M. A review on knowledge-based computer vision. (manuscript). 2010. Available from Internet: <<http://www.inf.ufrgs.br/~srfiorini/wp-content/uploads/Sandro-Fiorini-Mara-Abel-Review-Knowledge-Based-Computer-Vision.pdf>>.

FIORINI, S. R.; ABEL, M.; GÄRDENFORS, P. Structure, similarity and spaces. In: THE SHAPES OF THINGS, 2013, Rio de Janeiro, Brazil. **Proceedings...** Germany: CEUR-WS, 2013. v. 1007, p. 71–74.

FIORINI, S. R.; ABEL, M.; SCHERER, C. M. An approach for grounding ontologies in raw data using foundational ontology. **Information Systems**, v. 38, n. 5, p. 784–799, jul. 2013.

FIORINI, S. R.; GÄRDENFORS, P.; ABEL, M. Representing part–whole relations in conceptual spaces. **Cognitive Processing**, Springer Berlin Heidelberg, p. 1–16, 2013.

FÖRSTER, J. Relations between perceptual and conceptual scope: How global versus local processing fits a focus on similarity versus dissimilarity. **Journal of Experimental Psychology: General**, v. 138, n. 1, p. 88–111, 2009.

FORTH, J.; WIGGINS, G. A.; McLean, A. Unifying conceptual spaces: Concept formation in musical creative systems. **Minds and Machines**, v. 20, n. 4, p. 503–532, nov. 2010.

FOSTER, D. H.; GILSON, S. J. Recognizing novel three–dimensional objects by summing signals from parts and views. **Proceedings of the Royal Society of London. Series B: Biological Sciences**, v. 269, n. 1503, p. 1939–1947, 2002.

GANGEMI, A. et al. Understanding top-level ontological distinctions. In: IJCAI-01 WORKSHOP ON ONTOLOGIES AND INFORMATION SHARING, 2001, Seattle, USA. **Proceedings...** Menlo Park, USA: AAAI Press, 2001. p. 26–33.

GANGEMI, A. et al. Sweetening ontologies with DOLCE. In: GÓMEZ-PÉREZ, A.; BENJAMINS, V. R. (Ed.). **Knowledge Engineering and Knowledge Management: Ontologies and the Semantic Web**. [S.l.]: Springer Berlin / Heidelberg, 2002, (LNCS, 2473). p. 223–233.

GÄRDENFORS, P. **Conceptual Spaces: The Geometry of Thought**. Cambridge, Massachusetts: The MIT Press, 2000.

GÄRDENFORS, P. How to make the semantic web more semantic. In: FORMAL ONTOLOGY IN INFORMATION SYSTEMS, 2004, Torino, Italy. **Proceedings...** Amsterdam, The Netherlands: IOS Press, 2004. p. 17–34.

GÄRDENFORS, P.; WILLIAMS, M.-A. Multi-agent communication, planning and collaboration based on perceptions, conceptions and simulations. In: SCHALLEY, A.; KHELENTHOS, D. (Ed.). **Mental States, Volume 1: Evolution, Function, Nature**. Amsterdam: Benjamins, 2008. p. 95–121.

GÄRDENFORS, P.; ZENKER, F. Using conceptual spaces to model the dynamics of empirical theories. In: OLSSON, E. J.; ENQVIST, S. (Ed.). **Belief Revision meets Philosophy of Science**. [S.l.]: Springer Netherlands, 2011, (Logic, Epistemology, and the Unity of Science, v. 21). p. 137–153.

GARNER, W. R. **The processing of information and structure**. [S.l.]: Lawrence Erlbaum, 1974.

GAUKER, C. A critique of the similarity space theory of concepts. **Mind & Language**, Blackwell Publishing Ltd, v. 22, n. 4, p. 317–345, 2007.

GERSTL, P.; PRIBBENOW, S. Midwinters, end games, and body parts: a classification of part-whole relations. **International Journal of Human-Computer Studies**, v. 43, n. 5–6, p. 865–889, Nov 1995.

GOLDSTONE, R. L. The role of similarity in categorization: Providing a groundwork. **Cognition**, v. 52, n. 2, p. 125–57, 1994.

GOLDSTONE, R. L.; SON, J. Y. Similarity. In: HOLYOAK, K. J.; MORRISON, R. G. (Ed.). **The Cambridge handbook of thinking and reasoning**. New York, NY, US: Cambridge University Press, 2005. p. 13–36.

GRAF, M. Coordinate transformations in object recognition. **Psychological Bulletin**, v. 132, p. 920–945, 2006.

GUARINO, N. Formal ontology and information systems. In: FORMAL ONTOLOGY IN INFORMATION SYSTEMS, 1998, Trento, Italy. **Proceedings...** Amsterdam, The Netherlands: IOS Press, 1998. p. 3–15.

GUARINO, N. Local qualities, quality fields, and quality patterns: A preliminary investigation. In: THE SHAPE OF THINGS, 2013, Rio de Janeiro, Brazil. **Proceedings...** Germany: CEUR-WS, 2013. v. 1007, p. 75–81.

GUARINO, N.; PRIBBENOW, S.; VIEU, L. Modeling parts and wholes. **Data & Knowledge Engineering**, v. 20, n. 3, p. 257–258, Nov 1996.

GUARINO, N.; WELTY, C. A. An overview of OntoClean. In: STAAB, S.; STUDER, R. (Ed.). **Handbook on Ontologies**. [S.l.]: Springer Berlin Heidelberg, 2009, (International Handbooks on Information Systems). p. 201–220. 4.

GUIZZARDI. Logical, ontological and cognitive aspects of objects types and cross-world identity with applications to the theory of conceptual spaces. 2014. Manuscript.

GUIZZARDI, G. **Ontological foundations for structural conceptual models**. Netherlands: CTIT, 2005. 416 p. (PhD Thesis Series, 05-74).

GUIZZARDI, G.; FALBO, R. d. A.; FILHO, J. A. G. A. P. Using objects and patterns to implement domain ontologies. **Journal of the Brazilian Computer Society**, sciELO, v. 8, p. 43 – 56, 07 2002. ISSN 0104-6500.

HALL, D. Continuity and the persistence of objects: When the whole is greater than the sum of the parts. **Cognitive Psychology**, v. 37, n. 1, p. 28 – 59, 1998. ISSN 0010-0285.

HALPER, M.; GELLER, J.; PERL, Y. An oodb part-whole model: Semantics, notation and implementation. **Data & Knowledge Engineering**, v. 27, n. 1, p. 59 – 95, 1998.

HARNAD, S. The symbol grounding problem. **Phys. D**, v. 42, n. 1-3, p. 335–346, 1990.

HOFFMAN, D. D.; SINGH, M. Saliency of visual parts. **Cognition**, v. 63, n. 1, p. 29–78, Apr 1997.

HSIAO, J. H.; COTTRELL, G. W. Not all visual expertise is holistic, but it may be leftist: The case of chinese character recognition. **Psychological Science**, v. 20, n. 4, p. 455–463, 2009.

HUMMEL, J. E.; BIEDERMAN, I. Dynamic binding in a neural network for shape recognition. **Psychological Review; Psychological Review**, v. 99, n. 3, p. 480–517, 1992.

JÄGER, G. The evolution of convex categories. **Linguistics and Philosophy**, v. 30, n. 5, p. 551–564, Mar 2008.

JIANG, Y.-G. et al. Representations of keypoint-based semantic concept detection: A comprehensive study. **IEEE Transactions on Multimedia**, v. 12, n. 1, p. 42–53, 2010.

JOHANSSON, I. Formal mereology and ordinary language - reply to varzi. **Appl. Ontol.**, v. 1, p. 157–161, Apr 2006.

JOHANSSON, I. Shape is a non-quantifiable physical dimension. In: **THE SHAPES OF THINGS**, 2011, Karlsruhe, Germany. **Proceedings...** Germany: CEUR-WS, 2011. v. 812.

JOHANSSON, I. Natural science and mereology. In: BURKHARDT, H.; SEIBT, J.; IMAGUIRE, G. (Ed.). **Handbook of Mereology**. München, Germany: Philosophia Verlag, 2012.

KIRA, Z. Transferring embodied concepts between perceptually heterogeneous robots. In: **IEEE/RSJ INTERNATIONAL CONFERENCE ON INTELLIGENT ROBOTS AND SYSTEMS**, 2009, St. Louis, USA. **Proceedings...** Washington, USA: IEEE, 2009. p. 4650–4656.

LANGACKER, R. W. **Concept, image, and symbol: the cognitive basis of grammar**. 2nd ed.. ed. Berlin; New York: Mouton de Gruyter, 2002. (Cognitive linguistics research, 1).

LeBlanc, K.; SAFFIOTTI, A. Cooperative anchoring in heterogeneous multi-robot systems. In: **IEEE INTERNATIONAL CONFERENCE ON ROBOTICS AND AUTOMATION**, 2008, Pasadena, USA. **Proceedings...** Washington, USA: IEEE, 2008. p. 3308–3314.

LIGOZAT, G.; CONDOTTA, J.-F. On the relevance of conceptual spaces for spatial and temporal reasoning. **Spatial Cognition & Computation**, v. 5, n. 1, p. 1–27, 2005.

LORENZATTI, A. et al. Ontological primitives for visual knowledge. In: COSTA, A. da R.; VICARI, R.; TONIDANDEL, F. (Ed.). **Advances in Artificial Intelligence – SBIA 2010**. [S.l.]: Springer Berlin / Heidelberg, 2011, (Lecture Notes in Computer Science, v. 6404). p. 1–10.

LOVE, B. C.; ROUDER, J. N.; WISNIEWSKI, E. J. A structural account of global and local processing. **Cognitive psychology**, v. 38, n. 2, p. 291–316, Mar 1999. PMID: 10090805.

LOWE, D. G. Distinctive image features from scale-invariant keypoints. **International Journal of Computer Vision**, v. 60, n. 2, p. 91–110, nov. 2004.

MADDOX, W. T. Perceptual and decisional separability. In: _____. **Multidimensional models of perception and cognition**. [S.l.]: Lawrence Erlbaum Associates, Inc, 1992. (Scientific psychology series.), p. 147–180.

MARR, D. **Vision: a computational investigation into the human representation and processing of visual information**. [S.l.]: W.H. Freeman, 1982.

MARR, D.; NISHIHARA, H. K. Representation and recognition of the spatial organization of three-dimensional shapes. **Proceedings of the Royal Society of London. Series B. Biological Sciences**, v. 200, n. 1140, p. 269–294, Feb 1978.

MELARA, R. D. The concept of perceptual similarity: From psychophysics to cognitive psychology. In: _____. **Psychophysical Approaches to Cognition**. [S.l.]: North-Holland, 1992. p. 303–388.

MOLTMANN, F. A new notion of part structure for natural language. **Data & Knowledge Engineering**, v. 20, n. 3, p. 323–345, Nov 1996.

MOTSCHNIG-PITRIK, R. The semantics of parts versus aggregates in data/knowledge modelling. In: ROLLAND, C.; BODART, F.; CAUVET, C. (Ed.). **Advanced Information Systems Engineering**. [S.l.]: Springer Berlin Heidelberg, 1993, (Lecture Notes in Computer Science, v. 685). p. 352–373.

MOTTAGHI, R.; RANGANATHAN, A.; YUILLE, A. A compositional approach to learning part-based models of objects. In: IEEE INTERNATIONAL CONFERENCE ON COMPUTER VISION WORKSHOPS, 2011, Barcelona, Spain. **Proceedings...** Washington, USA: IEEE, 2011. p. 561–568.

MURPHY, K. et al. Object detection and localization using local and global features. In: PONCE, J. et al. (Ed.). **Toward Category-Level Object Recognition**. [S.l.]: Springer Berlin Heidelberg, 2006, (Lecture Notes in Computer Science). p. 382–400.

NEUMANN, B.; MOLLER, R. On scene interpretation with description logics. **Image and Vision Computing**, v. 26, n. 1, p. 82–101, jan. 2008.

NEWELL, F. N. et al. The interaction of shape- and location-based priming in object categorisation: Evidence for a hybrid “what + where” representation stage. **Vision Research**, v. 45, n. 16, p. 2065–2080, jul. 2005.

NILES, I.; PEASE, A. Towards a standard upper ontology. In: INTERNATIONAL CONFERENCE ON FORMAL ONTOLOGY IN INFORMATION SYSTEMS, 2001, Ogunquit, USA. **Proceedings...** New York, USA: ACM, 2001. p. 2–9. 1.

NILSSON, J. F. A conceptual space logic. In: EUROPEAN-JAPANESE CONFERENCE ON INFORMATION MODELLING AND KNOWLEDGE BASES, 1999, Iwate, Japan. **Proceedings...** Amsterdam, The Netherlands: IOS Press, 1999. p. 39–53.

NOWAK, E.; JURIE, F.; TRIGGS, B. Sampling strategies for bag-of-features image classification. In: LEONARDIS, A.; BISCHOF, H.; PINZ, A. (Ed.). **Proceedings...** Heidelberg, Germany: Springer Berlin Heidelberg, 2006. (Lecture Notes in Computer Science), p. 490–503.

PADGHAM, L.; LAMBRIX, P. A framework for part-of hierarchies in terminological logics. In: INTERNATIONAL CONFERENCE ON PRINCIPLES OF KNOWLEDGE REPRESENTATION AND REASONING, 1994, Bonn, Germany. **Proceedings...** [S.l.]: Morgan Kaufmann, 1994. p. 485–496.

PALMER, S. E. Hierarchical structure in perceptual representation. **Cognitive Psychology**, v. 9, n. 4, p. 441–474, oct. 1977.

PAOLA, A. D. et al. Human-ambient interaction through wireless sensor networks. In: CONFERENCE ON HUMAN SYSTEM INTERACTIONS, 2009, Catania, Italy. **Proceedings...** [S.l.], 2009. p. 64–67.

PEISSIG, J. J.; TARR, M. J. Visual object recognition: Do we know more now than we did 20 years ago? **Annual Review of Psychology**, v. 58, p. 75–96, jan. 2007.

PENTLAND, A. P. Perceptual organization and the representation of natural form. **Artificial Intelligence**, v. 28, n. 3, p. 293–331, May 1986.

PITT, D. Mental representation. In: ZALTA, E. N. (Ed.). **The Stanford Encyclopedia of Philosophy**. Fall 2013. Stanford University, 2013. Available from Internet: <<http://plato.stanford.edu/archives/fall2013/entries/mental-representation/>>.

RAUBAL, M. Formalizing conceptual spaces. In: FORMAL ONTOLOGY IN INFORMATION SYSTEMS, 2004, Torino, Italy. **Proceedings...** Amsterdam, The Netherlands: IOS Press, 2004. v. 114, p. 153–164.

RECTOR, A. et al. **Simple part-whole relations in OWL Ontologies**. [S.l.], 2005. Available from Internet: <<http://www.w3.org/2001/sw/BestPractices/OEP/SimplePartWhole/index.html>>.

RICHLER, J. J.; WONG, Y. K.; GAUTHIER, I. Perceptual expertise as a shift from strategic interference to automatic holistic processing. **Current Directions in Psychological Science**, v. 20, n. 2, p. 129–134, 2011.

RICKARD, J. T. A concept geometry for conceptual spaces. **Fuzzy Optimization and Decision Making**, v. 5, p. 311–329, oct. 2006.

RICKARD, J. T.; AISBETT, J.; GIBBON, G. Reformulation of the theory of conceptual spaces. **Information Sciences**, v. 177, n. 21, p. 4539–4565, nov. 2007.

RIDDOCH, M. J.; HUMPHREYS, G. W. Object identification in simultanagnosia: When wholes are not the sum of their parts. **Cognitive Neuropsychology**, v. 21, n. 2-4, p. 423–441, 2004. PMID: 21038214.

RIPS, L. J.; BLOK, S.; NEWMAN, G. Tracing the identity of objects. **Psychological Review**, v. 113, n. 1, p. 1–30, 2006.

ROSCH, E. Principles of categorization. In: MARGOLIS, E.; LAURENCE, S. (Ed.). **Concepts: Core Readings**. [S.l.]: Cambridge, MA, US: The MIT Press, 1978. p. 189–206.

ROSCH, E.; MERVIS, C. B. Family resemblances: Studies in the internal structure of categories. **Cognitive Psychology**, v. 7, n. 4, p. 573–605, Oct 1975.

RUMBAUGH, J.; JACOBSON, I.; BOOCH, G. **Unified Modeling Language Reference Manual, The (2nd Edition)**. [S.l.]: Pearson Higher Education, 2004.

SCHNEIDER, L. Revisiting the ontological square. In: FORMAL ONTOLOGY IN INFORMATION SYSTEMS, 2010, Toronto, Canada. **Proceedings...** Amsterdam, The Netherlands: IOS, 2010. p. 73–86.

SCHOLL, B. J. Object persistence in philosophy and psychology. **Mind & Language**, Blackwell Publishing Ltd, v. 22, n. 5, p. 563–591, 2007.

SHEPARD, R. N. Toward a universal law of generalization for psychological science. **Science**, v. 237, n. 4820, p. 1317–1323, 1987.

SHIPLEY, T. F. et al. Structural geology practice and learning, from the perspective of cognitive science. **Journal of Structural Geology**, v. 54, n. 0, p. 72 – 84, 2013.

SHOKOUFANDEH, A. et al. Indexing hierarchical structures using graph spectra. **IEEE Transactions on Pattern Analysis and Machine Intelligence**, v. 27, n. 7, p. 1125–1140, Jul 2005.

SIMONS, P. **Parts : a study in ontology**. [repr.]. Oxford: Clarendon Press, 2003.

SIMONS, P. Real wholes, real parts: Mereology without algebra. **The Journal of Philosophy**, v. 103, n. 12, p. 597–613, dec. 2006.

SMEULDERS, A. W. M. et al. Content-based image retrieval at the end of the early years. **IEEE Trans. Pattern Anal. Mach. Intell.**, v. 22, n. 12, p. 1349–1380, 2000.

SMITH, L. B. From fragments to geometric shape: Changes in visual object recognition between 18 and 24 months. **Current Directions in Psychological Science**, v. 18, n. 5, p. 290 –294, oct. 2009.

SON, J. Y.; SMITH, L. B.; GOLDSTONE, R. L. Simplicity and generalization: Short-cutting abstraction in children's object categorizations. **Cognition**, v. 108, n. 3, p. 626–638, Sep 2008. PMID: 18565504PMCID: 2584368.

SPELKE, E. S. et al. Spatiotemporal continuity, smoothness of motion and object identity in infancy. **British Journal of Developmental Psychology**, Blackwell Publishing Ltd, v. 13, n. 2, p. 113–142, 1995.

STUDER, R.; BENJAMINS, V. R.; FENSEL, D. Knowledge engineering: principles and methods. **Data & Knowledge Engineering**, v. 25, n. 1-2, p. 161–197, 1998.

THRUN, S. et al. Robust monte carlo localization for mobile robots. **Artificial Intelligence**, v. 128, n. 1–2, p. 99 – 141, 2001.

TOMLINSON, M.; LOVE, B. C. Learning abstract relations through analogy to concrete exemplars. In: COGNITIVE SCIENCE SOCIETY, 2006, Vancouver, Canada. **Proceedings...** Mahwah, USA: Lawrence Erlbaum Associates, 2006. p. 2269–2274.

TVERSKY, A. Features of similarity. **Psychological Review**, v. 84, n. 4, p. 327–352, 1977.

TVERSKY, B. Parts, partonomies, and taxonomies. **Developmental Psychology**, v. 25, n. 6, p. 983–95, 1989.

TVERSKY, B.; HEMENWAY, K. Objects, parts, and categories. **Journal of Experimental Psychology: General**, v. 113, n. 2, p. 169–193, 1984.

ULLMAN, S. **High-level vision: object recognition and visual cognition**. [S.l.]: MIT Press, 2000.

VARZI, A. Mereology. In: ZALTA, E. N. (Ed.). **The Stanford Encyclopedia of Philosophy**. Spring 2011. Stanford University, 2011. Available from Internet: <<http://plato.stanford.edu/archives/spr2011/entries/mereology/>>.

VARZI, A. C. A note on the transitivity of parthood. **Appl. Ontol.**, v. 1, n. 2, p. 141–146, Apr 2006.

WAGONER, J. C. V. et al. **Siliciclastic sequence stratigraphy in well logs, cores, and outcrops**. Tulsa, US: AAPG, 1990. (AAPG methods in exploration series No. 7).

WAXMAN, S. R.; MARKOW, D. B. Words as invitations to form categories: Evidence from 12- to 13-month-old infants. **Cognitive Psychology**, v. 29, n. 3, p. 257 – 302, 1995. ISSN 0010-0285.

WINSTON, M. E.; CHAFFIN, R.; HERRMANN, D. A taxonomy of part-whole relations. **Cognitive Science**, v. 11, n. 4, p. 417–444, Oct 1987.

WU, R.; MARESCHAL, D.; RAKISON, D. H. Attention to multiple cues during spontaneous object labeling. **Infancy**, v. 16, n. 5, p. 545–556, 2010.

XU, F. From lot's wife to a pillar of salt: Evidence that physical object is a sortal concept. **Mind & Language**, Blackwell Publishing Ltd, v. 12, n. 3-4, p. 365–392, 1997.

ZHU, S. C.; YUILLE, A. L. Forms: A flexible object recognition and modelling system. **International Journal of Computer Vision**, v. 20, Dec 1996.