Standpoint Logic:

A Logic for Handling Semantic Variability, with Applications to Forestry Information

Lucía Gómez Álvarez

University of Leeds

School of Computing

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

November 2019

Intellectual Property Statement

The candidate confirms that the work submitted is her own, except where work which has formed part of jointly authored publications has been included. The contribution of the candidate and the other authors to this work is explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others. In particular,

Chapter 2 elaborates on work published in:

- L. Gómez Álvarez. *Ambiguity: What is it that needs representing and what needs resolving?* (short paper) Ambiguity: Perspectives on Representation and Resolution. Sofia, Bulgaria, 2018.

Chapter 4 elaborates on work published in:

- L. Gómez Álvarez and B. Bennett. *Dealing with conceptual indeterminacy: A framework based on supervaluation semantics.* L. Bozzato, M. Homola, T. Mossakowski and A. Zimmermann, editors, *the Joint Proceedings of MedRACER'18 and WOMoCoE'18*, pages 38-50, CEUR-WS, 2018.

Chapter 6 includes work published in:

- L. Gómez Álvarez, B. Bennett and A. Richard-Bollans. *Talking about Forests: an Example of Sharing Information Expressed with Vague Terms.* S. Borgo, O. Kutz, F. Loebe and F. Neuhaus, editors, *Proceedings of the Joint Ontology Workshops (JOWO'17)*, pages 3:1-3:13, CEUR-WS, 2017.
- L. Gómez Álvarez and B. Bennett. *Classification, Individuation and Demarcation of Forests: Formalising the Multi-Faceted Semantics of Geographic Terms.* E. Clementini, M. Donnelly, M. Yuan, C. Kray, P. Fobliarone and A. Ballatore, editors, *Proceedings of the 13th International Conference on Spatial Information Theory (COSIT'17)*, pages 8:1-8:-15, LIPIcs, 2017.

The candidate confirms that she is the first author of all the jointly-authored publications, that the work contained within these publications is directly attributable to her and that Bennett's and Richard-Bollans' contribution has been in terms of general advice and assistance.

Acknowledgements

In the first place, I would like to thank my supervisor, Brandon Bennett, for all the trust, support and freedom, as well as all the good times shared in these four years. Freedom, while sweet, can sometimes be uncomfortable. I really appreciate having always felt trust and encouragement when making decisions and producing my own work.

Thanks to Adam Richard-Bollans. Having a friend with whom you can have interesting and constructive conversations about your research is a gem, and I hope we can keep working together in the future.

Thanks to my homes in Leeds. I am lucky that since I arrived at Leeds five years ago, to a completely new city, I have felt at home. Thanks to 'the Turnways' and particularly to Gauthier, for a lot of inspiring and familial conversations during endless breakfasts. And thanks to Firelight, my current home, for being an amazing housing co-operative and family. Thanks to its members, Ryan, Ersilia and all the rest of you. It has been amazing to live and share a project with you. I take with me a lot of learning and good memories, and I couldn't be happier that we have seen the main historic burden for Firelight's survival go.

Most importantly, thanks to Marta, who's company has always inspired, soothed and brightened my life. Thanks for being my family and unconditional friend, for supporting me through the hardest bits of writing my thesis and for sharing my joy during the best ones. Having had you living close to me these five years in Leeds is invaluable, and I want to believe that there will be a lot more to go.

Abstract

It is widely accepted that most natural language expressions do not have precise universally agreed definitions that fix their meanings. Except in the case of certain technical terminology, humans use terms in a variety of ways that are adapted to different contexts and perspectives. Hence, even when conversation participants share the same vocabulary and agree on fundamental taxonomic relationships (such as subsumption and mutual exclusivity), their view on the specific meaning of terms may differ significantly. Moreover, even individuals themselves may not hold permanent points of view, but rather adopt different semantics depending on the particular features of the situation and what they wish to communicate.

In this thesis, we analyse logical and representational aspects of the semantic variability of natural language terms. In particular, we aim to provide a formal language adequate for reasoning in settings where different agents may adopt particular standpoints or perspectives, thereby narrowing the semantic variability of the vague language predicates in different ways.

For that purpose, we present *standpoint logic*, a framework for interpreting languages in the presence of *semantic variability*. We build on supervaluationist accounts of vagueness, which explain linguistic indeterminacy in terms of a collection of possible interpretations of the terms of the language (*precisifications*). This is extended by adding the notion of *standpoint*, which intuitively corresponds to a particular point of view on how to interpret vague terminology, and may be taken by a person or institution in a relevant context. A standpoint is modelled by sets of precisifications compatible with that point of view and does not need to be fully precise. In this way, *standpoint logic* allows one to articulate finely grained and structured stipulations of the varieties of interpretation that can be given to a vague concept or a set of related concepts and also provides means to express relationships between different systems of interpretation.

After the specification of precisifications and standpoints and the consideration of the relevant notions of truth and validity, a multi-modal logic language for describing standpoints is presented. The language includes a modal operator $\Box_s$ for each standpoint s, such that $\Box_s \phi$ means that a proposition ϕ is unequivocally true according to the standpoint s — i.e. ϕ

is true at all precisifications compatible with s. We provide the logic with a Kripke semantics and examine the characteristics of its intended models, of the class $\mathfrak{M}_{\mathbb{S}_0}$. Furthermore, we prove the soundness, completeness and decidability of *standpoint logic* with an underlying propositional language, and show that the satisfiability problem is **NP**-complete. We subsequently illustrate how this language can be used to represent logical properties and connections between alternative partial models of a domain and different accounts of the semantics of terms.

As proof of concept, we explore the application of our formal framework to the domain of forestry, and in particular, we focus on the semantic variability of 'forest'. In this scenario, the problematic arising of the assignation of different meanings has been repeatedly reported in the literature, and it is especially relevant in the context of the unprecedented scale of publicly available geographic data, where information and databases, even when ostensibly linked to ontologies, may present substantial semantic variation, which obstructs interoperability and confounds knowledge exchange.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background

The fact that natural language terms do not have precise universally agreed definitions that fix their meanings is widely accepted. Instead, their applicability is unclear and it may vary depending on the context and pragmatics of use; there are *borderline* cases where it is a matter of judgement whether a term should be applied or not.

Semantic variation occurs in different ways, such as shifts in the thresholds of applicability of terms like ‘tall’, and fluctuations in the sense with which words like ‘game’ are used. In practice, the particular semantics that a term acquires in a natural or artificial scenario can be bounded both by explicit compliance to definitions or conceptual models and by implicit commitments derived from the previous usage of the term. For example, by asserting that ‘Tim is tall’, an agent implicitly commits to his threshold for tallness to be lower than Tim’s height.

As a consequence of the variability, language users that share the same vocabulary may differ in the specific semantics that they give to terms in a particular situation, even if they agree on essential relationships (such as subsumption or mutual exclusivity) and on the truth or falsity of a wide range of predications. For instance, two agents may agree that ‘Tim is tall’ and yet they may have different standpoints on the threshold for tallness; hence they could disagree on other predications. Moreover, individuals do not tend to hold permanent and precise interpretations of the meaning of terms themselves [EdGR13], wherewith the problem is not only the interoperability between agents but also the semantic variability with which the same agent may use natural language terms in different scenarios.

Humans seem to cope with the Semantic Variability of Natural Language Terms (SVoNLT) by making use of context and other pragmatic information to narrow the range of possible interpretations, generally achieving successful (or at least acceptable) information exchange and reasoning. How we do this has been a long-lasting question, and there is as yet little consensus, even regarding the basic principles and mechanisms that are involved. Investigations into this issue are fuelled by several motivations, some linked to philosophical inquiry (for example, that raised by the sorites paradox[1]), and others to current challenges in domains like information processing. Despite the

[1]The sorites paradox or the heap paradox is discussed in chapter 2, section 2.2.1. In a nutshell, it starts with a heap and, by iterating the assumption that ‘if we remove one grain from a heap we still have a heap’, it reaches the paradoxical conclusion that a single grain is a heap.

advances in Artificial Intelligence (AI), current limitations in automated language understanding hinder the integration and synthesis of a large amount of the information and knowledge that humans make use of through natural language.

The topic of the Semantic Variability of Natural Language Terms (SVoNLT) has consequently been approached within a highly interdisciplinary area of research, going from philosophy to linguistics, cognitive science and AI. Philosophical interest dates well back to Ancient Greek philosophy [Ari55, Ath93] and contemporary authors often refer to it as the *vagueness* of natural language. One traditional view is that the imprecision of terms presents an obstruction to good philosophy [Fre48], and should be circumvented by establishing precise definitions. In contrast, views of human language becoming prominent in the 20th century (e.g. [Gri75, Wit09]) have accepted the SVoNLT as a fundamental feature of human communication[1], and numerous logical theories of vagueness have been proposed to model different aspects of it. To date, some of the most popular logico-philosophical approaches are based on *many-valued logics* [Tar80], mainly *fuzzy logic* [Zad75, Zad65], *epistemic frameworks* [Wil99] and *supervaluation semantics* [Fin75]. Meanwhile, in the field of cognitive science, conceptual models have been proposed (e.g. *conceptual spaces* [Gär04]), as well as accounts of the processes of concept formation, manipulation and understanding [GL05]. Finally, research on computational linguistics and AI has embraced the challenge of semantic heterogeneity as key to language understanding and have delivered substantial advances in strategies for representation and resolution of the meanings of terms in context [Nav09].

Given the variety of approaches and techniques, it is not surprising that there is, to date, no one way of representing natural language terms or concepts that performs well in all situations when the phenomenon arises. While some techniques do well at disambiguation within big corpora, others allow for formal reasoning and inference. Some are cognitively valid (supported by cognitive evidence) and applicable in research with humans while others are good for handling significant amounts of data.

[1]Some research suggests that not only it is not a deficit of natural language, but it may enable more efficient communication [Zip49, PTG12], being, therefore, an advantageous feature.

1.2 Overview of the problem

1.2.1 Statement of the problem

This project is concerned with elucidating the logical mechanisms that underlie communication involving the use of semantically heterogeneous terms. In particular, we aim to develop a logic formalism suitable for knowledge representation and reasoning with vocabularies whose terms display semantic variability.

We explore how different agents can establish their standpoints on the meaning of terms, that typically involve narrowing the admissible variability in their semantics while not necessarily fixing a sharp interpretation, and possibly setting general constraints that must hold on the domain. Natural reasoning tasks involving standpoints include gathering unequivocal or undisputed knowledge, knowledge that is relative to a standpoint or a set of them, and contrasting the knowledge that can be inferred from different standpoints.

Furthermore, our attention with regards to the semantic heterogeneity lays not only on variations in degree but also in qualitative aspects such as the cases where the set of attributes that determine the applicability of a predicate is contentious. Also, rather than seeking a strategy for the 'resolution' of the semantic variability, we are interested in a framework with which we can perform general-purpose representation and reasoning tasks *with* semantically variable terms.

1.2.2 Semantic variability notions

In line with the stated problem, we intend to establish explicit criteria of consistency for sets of statements involving the use of semantically variable terms. Such statements can be relative to the admissible variability in the whole language or with respect to specific standpoints. The following are the kind of statements that we are seeking to represent:

N1. "It is unequivocal that ϕ."

N2. "In some sense ϕ."

N3. "It is definite whether ϕ (or not ϕ)."

N4. "It is borderline that ϕ."

N5. "According to *s*, [it is unequivocal that] ϕ."

N6. "According to *s*, in some sense ϕ."

N7. "According to *s*, it is definite whether ϕ (or not ϕ)."

N8. "According to *s*, it is borderline that ϕ."

N9. "Whatever is unequivocal according to *s* is also unequivocal according to *s'*. In other words, standpoint *s'* is sharper or more precise than standpoint *s*."

In the former sentences, ϕ is an independent clause and *s* and *s'* are standpoints. The notions 'unequivocal' and 'in some sense' should be self-explanatory, by 'definite' we mean that ϕ is either unequivocally the case or unequivocally not the case, and 'borderline' that ϕ is in some sense the case and in some sense not the case. Finally, a standpoint reads as a particular interpretation or narrowing of the semantics of the language, which may be taken by an agent in a particular context. E.g. 'According to the standpoint that Anna took when she was in the classroom, Silvia is tall'. Note that a standpoint does not represent an agent or a context as such; it only represents a particular sharpening of the variability of the language, which can be specified and used by different agents in different scenarios.

1.2.3 Motivation

The SVoNLT, in different shapes like the sorites paradox and polysemy, causes problems for sound communication and reasoning, which has been observed by philosophers since ancient times. Far from being only problematic to solve deep philosophical questions, it creates difficulties in various research fields, among them Knowledge Representation.

While the applied research on the topic is broad, a large extent focuses on computing contextualised or prototypical meanings, on selecting the intended senses of terms in use and on establishing metrics to grade phenomena that are typically presumed to be binary. Examples of the latter include assigning degrees of truth or probability distributions to vague predications and graded membership to instances of vague categories.

In contrast, in this research we focus on applications where retaining classical reasoning is desired, and where the object of interest involves handling fine-grained symbolic perspectives on the semantics of terms, knowing what can be inferred from the consensual semantics of sets of agents or which standpoints are compliant with a partial truth among others. In this way, we hope that this framework has interesting applications in scenarios where binary or qualitative reasoning is preferred over approximations.

Consequently, we approach the problem of the SVoNLT from a different angle from most of the research on the topic, thereby providing a framework that can be thought of as complementary to other strategies for representation. We base our framework on a well established philosophical theory of vagueness, supervaluationism, for which, despite its reasonable success in the philosophical literature, there have been only mild attempts to turn it into an operational framework for automated knowledge representation and reasoning.

Finally, we put forward the idea that the logical mechanisms that we study in this thesis have interesting practical applications. For instance, in scientific domains were rigorous definitions and consistent data are necessary, scientists are often required to (partially) 'precisify' vague terms for the pursuit of their investigation. This process is non-trivial and depends on the domain and purpose of use, thereby leading to a multiplicity of characterisations of the terms. Unsurprisingly, problems arise when the interoperation of different standpoints is needed, as reported in a variety of domains such as engineering [CGV11] and forestry [CBL16, Gra08]. We will use the latter for our application scenario.

1.3 Application scenario in forestry

We illustrate the scenes that we want to model and reason about with an application scenario in the domain of forestry, in which a broad range of forest concepts and definitions have been specified for different purposes, leading to discrepancies of estimates ([Gra08, CBL16]) and to the consequent confusion over both the global forest extent and its spatial distribution [Ful06, GZR05].

This has been recognised to be one of the key challenges [BHHR18] preventing the integrated use of the more and more datasets (e.g. [HPM+13]), portals (e.g. Global Forest Watch) and ontologies (e.g. EnvO ontology) that have recently emerged for enabling the tracking, comparing and understanding of the available information on land use-cover and global forest extent.

Our application scenario involves the analysis and representation of the semantic variation of the term forest, including both matters of degree and conceptual divergencies that stem from the different domains and purposes for which the concept is defined. In particular, we focus on the representation of established definitions in the domain. Moreover, we discuss four use cases for a framework to be used in conjunction with existing datasets, portals and ontologies:

UC1. *Reasoning and querying data with varying standpoints.* This involves the capacity to interpret a dataset according to different standpoints on the relevant objects to which the data refers. In particular, we analyse how an agent can establish standpoints on the semantics of forest that comply with reference definitions in order to analyse the results in *forest cover* from a dataset, namely [HPM+13], according to different interpretations. Moreover, we explore how an agent can establish intermediate standpoints and relations to investigate the implications of different semantic commitments between those definitions.

UC2. *Reasoning with multiple underlying standpoints.* This involves the capacity to determine when a proposition is unequivocally true, regardless of which standpoint is taken on the semantics of the terms involved. In this use case, we analyse how an agent can establish definite or potential deforestation alerts depending on the (semantic) unequivocability of the *deforestation process* with respect to a set of reference definitions and data sources.

UC3. *Publishing knowledge relative to an interpretation.* This involves the capacity of an agent or a system to distribute knowledge that is dependent on a particular semantic interpretation, in the assumption that more than one interpretation may exist or be necessary to characterise the domain. We explore scenarios in the domain of forestry where such interpretation is typically provided in natural language, but partial aspects can be formalised with our framework enriching the representations. Or where we infer complex knowledge from observational data using a particular standpoint, which may be useful for other research.

One must remark that the issues that arise in this application scenario are somewhat different from those in informal conversational settings. In this case, standpoints concerning the meaning of forest belong to institutions or scientific communities that produce and distribute data. While in common-sense applications rough meaning is often sufficient, here a more precise analysis of the implications of choosing one or another interpretation is required, and potential data conflicts need to be resolved [BHHR18].

1.4 Achievements and contributions

The main achievements and contributions of this research can be summarised as follows:

1. **The development of a novel modal logic framework, *standpoint logic*, based on the well known logico-philosophical framework *supervaluation semantics*.** *Standpoint logic* is intended as a reasoning system for Knowledge Representation (KR), particularly in scenarios where agents or systems need to represent and reason about the different semantics that they give to natural language terms. We analyse the characteristics of the framework, its fundamental elements, namely precisifications and standpoints, and the relevant notions of truth and validity. While we take a modal approach to supervaluationism, in line with previous work on the subject, we separate ourselves from previous research on axiomatisations of supervaluational modalities in order to accommodate the notion of standpoint, which is central in our study. As such, not only standpoint logic is a multi-modal framework, but also its modalities respond to different schemas than those proposed in previous work. We provide detailed examples of the capacities of the framework to reason about different forms of semantic variability, such as graded predicates like 'tall' and qualitative variations as in 'game', and with different standpoint arrangements and the relations between them.

2. **The stipulation of its syntax and semantics, together with proofs of soundness, completeness, decidability and complexity,** going further than previous work on modal supervaluationist frameworks. We specify a formal language, syntax and semantics for Standpoint Logic with a propositional base language; we discuss the proof theory, where the most characteristic features are the interaction axioms **AS4** and **AS5**. These are stronger than the well known **A4** and **A5** and correspond to what we call *trans-transitive* and *trans-euclidean* relations in the Kripke models. Moreover, the multiple modalities are partially ordered under the subset relation, which corresponds to axiom **AP** and encodes the *sharper* relation that can hold between standpoints (notion **N9**). We prove the soundness, completeness and decidability of propositional standpoint logic,

and establish the complexity of the satisfiability problem as **NP**-complete. This is a 'good' result that is a consequence of the stronger interaction axioms **AS4** and **AS5** and contrasts with the higher complexity of other multi-modal logics such as epistemic logic, that is **PSPACE**-complete. Finally, we also provide the syntax and semantics for standpoint logic with more expressive underlying logics, namely description logic and first-order logic.

3. **The study of an application scenario in the domain of Forestry.** In order to illustrate the applicability of *standpoint logic*, we have explored an application scenario (as introduced in 1.3) that tackles the difficulties that the semantic variability of the term *forest* poses in the context of the scientific multidisciplinarity of the forestry domain. In addressing it, we proceed to analyse the literature on ontological issues of geographical objects and determine three main aspects that the different standpoints may represent, namely the classification, individuation and demarcation of *forests*. We then provide formal representations for standpoints corresponding to selected definitions and characterisations of *forest* from leading organisations that map into publicly available data. Finally, we illustrate how the framework can handle the four use cases through different reasoning tasks on the created standpoint knowledge base.

1.5 The structure of this thesis

The organisation of this thesis is as follows. Chapter 2 provides the background and literature review of this research. This begins with a broad historical introduction to the study of *vagueness* and the semantic variability of language (in section 2.2), followed by the specification of what is meant by the Semantic Variability of Natural Language Terms (SVoNLT) and which phenomena fall under its umbrella (in section 2.3). We continue with the literature review, which considers the main contemporary theories and approaches to the problem (in section 2.4), mostly focusing on logical theories of vagueness, and finally we look at relevant formal logic frameworks (in section 2.5). This includes a recollection of preliminary work for this thesis, a review of modal frameworks for supervaluationism and vagueness and finally a discussion of other related logic frameworks to this research.

Chapter 3 is devoted to the presentation of the standpoint framework, which initiates with the motivation (in section 3.2) and a general overview (in section 3.3). Subsequently we proceed to the specification and analysis of its key elements, namely *precisifications* (in section 3.4) and standpoints (in section 3.5). Then, we establish the relevant notions of truth and validity (in section 3.6) and finally we briefly consider the phenomenon of higher-order vagueness (in section 3.7), which we do not model in our framework.

Chapter 4 provides the formal framework for the multi-modal standpoint logic. After a brief introduction (in section 4.1), propositional standpoint logic ($\mathbb{S}_0$) is presented (in section 4.2) and its syntax, proof theory and semantics are given. In addition, the class $\mathfrak{M}_{\mathbb{S}_0}$ of simplified Kripke models is declared and its correspondence with the axiomatisation is shown. Subsequently, we proceed to prove soundness and completeness of $\mathbb{S}_0$ with respect to the class $\mathfrak{M}_{\mathbb{S}_0}$ of models (in section 4.3) and we present decidability and complexity results (in section 4.4). We then conclude the chapter by considering more expressive underlying logics (in section 4.5), namely first-order logic and description logics.

Chapter 5 explores how can *standpoint logic* be used in order to represent different forms of semantic variability, and illustrates it with examples and formalisations. We begin by considering an adaptation of a propositional syllogism (in section 5.2) and then proceed to illustrate representations of 'sorites' or graded predicates (in section 5.3) and of conceptual or non-numerical vagueness in (in section 5.4). Following, we consider penumbral connections between the representations (in section 5.5) and then we inspect other aspects that do not fall under the umbrella of the SVoNLT, namely context, generality and ambiguity (in section 5.6). We finish the chapter by introducing the relations that may hold between standpoints and some combinations that can be performed with them (in section 5.7) and lastly the formulation of an appropriate normal form for standpoint formulae (in section 5.8).

Chapter 6 explores an application scenario in the domain of forestry, addressing the reported challenges that the variability of terms like 'forest' pose for scientific and public knowledge acquisition. We first provide a background for the problem, as reported in the forestry research (in section 6.1) and then consider the representational challenges (in section 6.2), drawing on the literature of geographic objects and concluding with some formal representations. We then establish a set of use cases (in section 6.4) and we illustrate the reasoning tasks that can be done with the former representations. We conclude the chapter with a consideration of challenges, opportunities and limitations (in section 6.5).

Finally, chapter 7 concludes this thesis with a brief recapitulation of its structure and main contributions (in section 7.1), a summary of further work (in section 7.2) and the conclusion (in section 7.3).

Chapter 2

Background and Literature Review

2.1 Introduction

The Semantic Variability of Natural Language Terms (SVoNLT) has been studied from a range of perspectives, with different motivations and at various levels of detail. The problem is open and well known in the domains of AI and computer science, with active research in areas such as Word Sense Disambiguation (WSD) and knowledge integration. Preceding the contemporary interest, the literature on the topic has a long tradition dating back to Ancient Greek philosophy. While much of the philosophical research has focused on the study of natural language, often in relation to logic, the state of the art is mostly driven by the technical challenges posed by the phenomenon, particularly for the automated handling of natural language terms and discourse. As such, the literature on the topic is rich, historic and widely multidisciplinary.

In this chapter we provide the context for this research. We begin, in section 2.2, with an introduction of the study area and its progression, intended to provide the general background to the topic by revisiting the historical interest that the SVoNLT has awaken. We continue in section 2.3 with a clarification of the terminology and a broader explanation of the phenomena that we want to model. Subsequently, in section 2.4 we analyse the main theoretical frameworks for modelling the SVoNLT. Given the breadth of the domain, we will primarily focus on the developments in philosophical frameworks dealing with vagueness, but we will also briefly consider other formalisms and tools that have been developed to tackle different aspects of the SVoNLT. Finally, we discuss relevant formal frameworks and previous work in section 2.5.

2.2 A brief historical introduction

In the study of the SVoNLT, the philosophical analysis of vagueness has been central. This section draws from the historical accounts provided in [Wil94] and [KS96a] (to which the reader should refer for a more detailed discussion) on the philosophical study of vagueness. Moreover, as we will see in section 2.3, we consider the SVoNLT to be a broader phenomenon than some philosophical accounts of vagueness. Hence we discuss here other relevant issues, most notably a last period where emerging applications in AI prompt the loss of centrality of philosophical issues in the debate (such as the analysis of higher-order vagueness[1]) in favour of a broader and more multidisciplinary focus in

[1] Higher-order vagueness is discussed in chapter 3, section 3.7

the research on the SVoNLT.

2.2.1 Early history and introduction to the sorites paradox

Theoretical interest in language dates well back to ancient Greek philosophy. While philosophical reflections arose in the context of other kinds of inquiry, the analysis of language became abstract, focused and systematic with Aristotle and the Stoics, who also did extensive work in the area of logic [Ade13]. Indeed, since its origins, the history of the philosophy of language has often been closely related to that of logic, to the extent that the boundaries between them have often been elusive.

With regards to the SVoNLT, Aristotle focused on the notion of ambiguity [Hin59]. He distinguished two different kinds of phenomena that were denoted by the term and needed different theoretical treatment: on the one hand, what Hintikka calls 'the multiplicity of applications' [Hin59], here referred as polysemy, by which a word has multiple and correlated senses; On the other hand, 'mere ambiguity' or homonymy, by which words may have two or more distinct meanings like in the paradigmatic example of *bank*. Alternative distinctions and confusion between both phenomena have prevailed in subsequent philosophical and linguistic literature about the variability of natural language, and will be discussed in the following section (subsections 2.3.4 and 2.3.5.3) in order to avoid terminological confusion.

Even more influential to subsequent developments on the study of the SVoNLT was the formulation of the Sorites paradox (also known as *the slippery slope paradox*), attributed to a contemporary of Aristotle, Eubulides [KS96a], and which has played a major role in the contemporary analysis of *vagueness*. The sorites paradox uses the example of a heap (*soros* in Greek) and goes:

> 'If we have a heap of sand and we remove one grain, we still have a heap of sand. So take a heap and remove grains one by one, until there is a single grain left. You will absurdly infer that such grain is still a heap of sand.' [KS96a]

The paradox relies on the intuition that a single grain does not make a difference to whether something is a heap. Yet, even though a person would not even perceive the removal of one grain, after performing over and over the same operation we are left with a single grain[1]. Arguments with a sorites structure can be formalised in classical logic, so that premises that appear clearly true yield to a clearly false conclusion: that one grain of sand is a heap.

This formulation became paradigmatic in the analysis of *vagueness*, which has been described as the phenomenon characterised by the presence of borderline cases in the applicability of predicates. The paradox motivated not only philosophical inquiry among the stoics, but also in subsequent and contemporary research on vagueness, and will be discussed further in section 2.3.2.

2.2.2 Approaches within the analytic tradition

After a long period without major interest in the topic, the so-called 'Linguistic Turn' in Anglo-American philosophy took place in the mid nineteenth century and drew back a great attention to language, which came to be seen as a 'a focal point in understanding belief and representation of the world' and a 'medium of conceptualisation' [Wol10].

The early years of the analytic tradition were highly marked by the positivist agenda and a high regard to science and the ideal of precision. It was under such circumstances that the interest in the SVoNLT (or rather on the lack of precision of natural language) arouse back, and that the notion of *vagueness* became established as being the phenomenon exemplified in the sorites paradox. Russel, in 'Vagueness' [Rus23], made the technical sense of vagueness canonical, providing the first systematical analysis of the problem in something close to its current form [Wil94].

Around the mid twentieth century, the analytic tradition shows a notorious shift towards ordinary language philosophy, pushing the focus of philosophical inquiry to natural language in use, rather than towards its relation to formal logic. The SVoNLT began to be seen as a feature providing flexibility to language rather than a deficiency, and logic got questioned as an adequate tool for investigating it. Examples of specific accounts of the SVoNLT along these lines are the family resemblance concepts of Wittgenstein [Wit09] and Waismann's open-texture concepts [Wai45].

At the same time, the first formal accounts of vagueness started to be developed; Max Black provided the first formal framework for the treatment of vagueness in natural language [Bla37], according to which the correctness of the use of a vague predicate is

[1] In the antiquity it was formulated in the constructive order and with a series of questions: 'Does a grain of wheat make a heap? Do two grains? Do n grains?' [Wil94]

equated to the 'statistical conformity with the behaviour of a certain group of users'. In this way, his work aimed to integrate the idea of meaning as use while still providing a systematic framework for its treatment, yet it raised objections over conflating statistics about the contingent use of terms with the abstract notion of logical validity [1]. Later, in [Meh58], Mehlberg continues the investigation of the relation between vagueness and truth, and in particular of the logical status of indeterminate statements. For that purpose, he provides an account of vagueness that can be considered supervaluationistic in essence, according to which a vague term 'can be characterised tentatively as one the correct use of which is compatible with several distinct interpretations'.

The former discussions generated a subsequent explosion of philosophical interest in the 1970s that gave rise to the development of contemporary formal theories of vagueness[2], most of which are non-classical in nature. Those are commonly grouped in epistemic views of vagueness, many-valued logics and supervaluationist theories, and will be discussed in detail in section 2.4, as they constitute the main theoretical background of this research (in particular the supervaluationist approach section 2.4.1). Likewise, nihilist positions concluding that the sorites paradox is unsolvable were notably explored by Dummett [Dum75] and Wright (eg [Wri87]) among others. Finally, in recent years, the focus of philosophical investigations of vagueness moved towards the 'unsolved' problem of higher-order vagueness (section 3.7), into whether vagueness is exclusively linguistic or there may be vague objects ('de dicto' and 'de re' vagueness), and to the role of context in vagueness, which is overviewed in chapter 5, section 5.6.1.

Overall, this subsection illustrates the history of the development of the main formal approaches to vagueness in the philosophical tradition. Those (supervaluationism, many-valued theories of truth and the epistemic view) will be discussed further in section 2.4, where we will justify using supervaluationism as the basis for our framework. Moreover, one may notice that the theories introduced here arise from the philosophical investigation of truth with regards to propositions involving vague predicates, rather than as computational frameworks. As such, subsequent developments on formal meth-

[1] Black modified his position in [Bla63], dropping his revision of logic, yet his work was significant and probabilistic approaches to vagueness are today established.

[2] Vagueness is often understood as being both single-dimensional and multi-dimensional. Therefore the notion can be considered to be analogous to that of SVoNLT. In practice, however, the philosophical literature on vagueness has a very strong focus on single-dimensional vagueness and the sorites paradox, and often the phenomenon of multi-dimensionality is only briefly mentioned and not compared to other phenomena such as polysemy or ambiguity in a rigorous way.

ods and implementations will be discussed separately, in section 2.5.

2.2.3 Multidisciplinarity and shift of the focus towards applied work

Meanwhile, roughly simultaneously with the development of the first (proto) formalism for vagueness, major developments were happening in other domains of science. The first digital computers were created in the 1930s and soon later, in the 50s, the domain of Artificial Intelligence (AI) became established, giving rise to new motivations for the study of the SVoNLT. In terms of techniques, the advances in AI meant, on the one hand, that for the first time logical frameworks could be used for automated logical deduction. On the other, it provided the technical means for other kinds of non-binary representation, notably machine learning algorithms.

As AI, and notably the field of Knowledge Representation (KR)[1], became a playground for the application of existing logical theories and frameworks, it became evident that some problems that had not been previously considered (such as the semantic representation of whole domains of knowledge or the interoperation between different vocabularies and logic theories), raised challenging issues in applied scenarios. In the particular case of vagueness, most of the philosophical literature was focussed on the formal analysis of the sorites paradox, yet most instances of semantic variability can not be cleanly formalised in this way. Since then, the KR community has worked on modelling vagueness in a broader sense that is better aligned with the representational challenges that emerge in the computing domain, although the research in the domains of mathematics, philosophy and KR regarding logical frameworks is closely intertwined.

Moreover, the need for creating rich representations of knowledge and the simultaneous move in linguistics from behaviourism towards constructivism (with Chomsky's generative grammars) resulted in the emergence of the new interdisciplinary field of cognitive science emerge. In subsequent years, the multidisciplinary work led to an explosion of sub-fields, and, as such, the domains now involved in the area of modelling the semantics of natural language terms range from Philosophy and Linguistics to a broader academic arena also including Cognitive Science, Psychology, AI and all the intersections: computational linguistics, cognitive linguistics, psycho-linguistics, etc.

[1] KR is the subfield of AI that is concerned with representing information such that a computer system can perform automated reasoning with it in order to solve complex problems. It incorporates findings from logic as well as inspiration from psychology and cognitive science about how humans solve these problems.

For the analysis of the SVoNLT, this led to the wide-ranging scene of the research in the topic today, which we will overview in section 2.4. Moreover, it motivated a noticeable shift from studying vagueness and the SVoNLT in order to have a good understanding of language towards the urge to solve specific (and often mundane) tasks which, however, test established theories and their applicability. Since then, studies on the SVoNLT have broadened in scope: they have included psychological tests and computational simulations of behaviour, general theories of representation that deal with the SVoNLT (implicitly and explicitly) have been proposed and more methods for resolution have rested on empirical and behavioural evidence. Perhaps the most salient domain looking into the SVoNLT today is Natural Language Processing (Natural Language Processing (NLP)), in particular the sub-field of Word Sense Disambiguation (WSD).

2.3 The Semantic Variability of Natural Language Terms (SVoNLT)

Following the overview of the study of the SVoNLT, we proceed to provide an explicit account of what the 'Semantic Variability of Natural Language Terms' stands for, given that the specific phenomena under study varies across the literature. Hence, in this section we aim to clarify what the SVoNLT refers to in the context of this research, and we overview the different elements that we consider that fall under its umbrella.

2.3.1 Terminology

The terminology used for the notions related to the semantic variability is, perhaps unsurprisingly, as diverse as the different approaches from which this feature of natural language has been studied. Different terms are linked to different domains of study and emphasise different aspects of this phenomena.

In this work, we have decided to adopt the notion of 'Semantic Variability of Natural Language Terms (SVoNLT)'. This is an unconventional jargon, particularly taking into account that we mostly draw on philosophical and logical literature on *vagueness*. We make this move to highlight the shift from the philosophical study of the sorites paradox towards a more practical aim of supporting knowledge representation mechanisms for scenarios involving semantically heterogeneous languages.

In what follows, we will call the *Semantic Variability of Natural Language Terms (SVoNLT)* to the phenomenon by which the meanings of natural language terms do not have precise definitions that fix their meanings; that is, neither their intension or extension[1] are well defined. Consequently, there are scenarios in which it is inherently unsettled whether a proposition is true or false, which are often referred as borderline cases of application, and where underdeterminacy is not a consequence of the lack of information or knowledge about the world, but rather of the lack of semantic precision of the proposition. This area of borderline applicability of vague predicates makes room for speakers to use terms and make judgements from a variety of perspectives, making further semantic commitments that are relevant to the context of use and other factors.

With regards to the phenomena that we consider that fall under the umbrella of the SVoNLT, we include: all the variants of *vagueness* as discussed in philosophy, that is, *single-dimensional*, *multi-dimensional* and *conceptual* or *non-numerical* vagueness; and *polysemy*, which we understand as a similar phenomenon to that of conceptual vagueness. On the contrary, we do not consider *ambiguity*, when understood as *homonymy*, and *generality*. The rest of this section is devoted to the description of these notions.

2.3.2 Vagueness and the sorites paradox

> 'A proposition is vague when there are possible states of things concerning which it is intrinsically uncertain whether, had they been contemplated by the speaker, he would have regarded them as excluded or allowed by the proposition. By intrinsically uncertain we mean not uncertain in consequence of any ignorance of the interpreter, but because the speaker's habits of language were indeterminate.' (Peirce, [Bal02]).

The former is Charles Sander Peirce's entry for *vague* in the 1902 Dictionary of Philosophy and Psychology [Bal02]. While there are varying definitions across the literature [KS96b], the vagueness of a predicate is generally characterised as the possession of some or all of the following features: the potential for borderline cases of application, the lack of well defined extensions or boundary and the susceptibility to

[1] Intension and extension, in logic, are correlative words that indicate the reference of a term or concept: intension indicates the internal content of a term or concept that constitutes its formal definition; and extension indicates its range of applicability by naming the particular objects that it denotes.

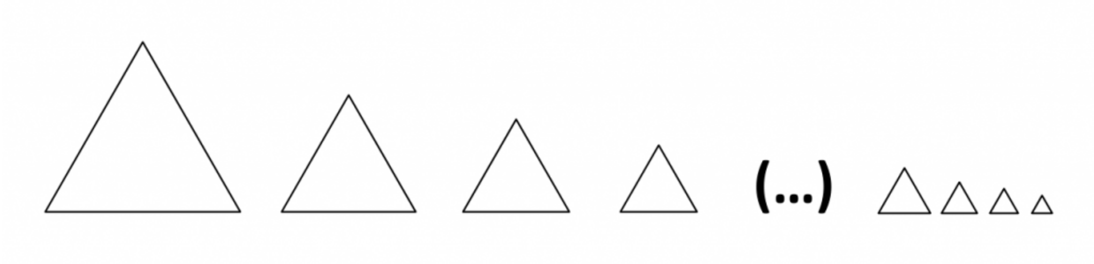

Figure 2.1: Illustration of the slippery slope effect.

the *sorites paradox*[1] (see Fig. 2.1). These features are strongly intertwined but not all apply in some scenarios. A longer discussion on the differences can be found in [KS96b]. In this thesis, we mostly focus on the possibility of borderline cases of application, given that not all vagueness is easily represented in the form of a sorites paradox and that smooth fuzzy boundaries are not necessarily easy to establish, as we will see for cases of multi-dimensional and conceptual or non-numerical vagueness.

Although definitions rarely specify it, most analysis of vagueness focus on the *single-dimensional* case, where the applicability of a predicate depends on a specific parameter and the thresholds that determine its truth or falsity are undetermined, thus potentially creating borderline cases. Archetypal examples of single-dimensional vagueness are the predicates 'tall', whose applicability depends on the property 'height', and 'heap', whose applicability depends on the number of grains (assuming an appropriate spatial arrangement).

Let us consider the case of 'tallness'. The predicate 'tall' lacks well-defined extensions, as there is no sharp boundary between tall people and people who are not tall. Instead, there may be instances such as 'Tina is tall' that are borderline cases of application. Moreover, with regards to the sorites paradox, let us imagine that it is clearly true that 'Nena is tall'. Then, it would be natural to assert that somebody who is $1mm$ shorter than Nena is also tall. It is easy to see how, by repeatedly iterating such claim, we would be forced to admit that any person is tall.

[1] Referred to a puzzle known as The Heap: Would you describe a single grain of wheat as a heap? No. Two grains? No. ... You must admit the presence of a heap sooner or later, so where do you draw the line? (Stanford Dictionary of Philosophy). The paradox is also formulated in the inverse way, so that we repeatedly assert: If this is a heap and I remove one grain, then this is still a heap.

2.3.3 Beyond single-dimensional vagueness

While *single-dimensional* examples of vagueness dominate the philosophical literature, the semantic variability of most natural language terms can not be expressed in this way, yet still fit the definition of vagueness that we provided the previous section.

Hence we consider *multi-dimensional vagueness*, which has been described as vagueness that manifests itself in multiple dimensions of the meaning of a predicate, in opposition to the single dimension of examples such as 'tall'. In [KS96a], Keefe gives the example of 'big', which depends on both height and volume (when used to describe people). Some recent literature studying multidimensional vagueness is [Ron11, Raf13] and some applied work can be found in [Gri12, Qiz06]. In previous work we have used the more general term *sorites vagueness* to describe the phenomena that occurs when the applicability of a predicate depends on specific measurable parameters but their thresholds are undetermined, hence combining both the notions of single and multi-dimensional vagueness [GÁB17].

In contrast, let us consider the example of 'Nice', which doesn't have a clear-cut set of associated dimensions. Examples of this kind are abundant and arise when there is a lack of clarity on which attributes or conditions are essential to the meaning of a given term, so that it is controversial how it should be defined. Thus, there is indeterminacy regarding to which property or logical combination of properties is relevant to determining whether a concept is applicable.

We call this *conceptual vagueness* [Ben05, GÁB17] (also refered as *non-numerical vagueness* [BC12]). In this case, the fact that there is not a total or even partial order with which we can represent the cases where a vague predicate applies, through borderline cases, to the cases where it does not apply, means that there is not a straightforward way to formalise this vagueness in the form of a sorites paradox (unless some sort of artificial similarity measure is computed).

For instance, this is the kind of vagueness that underlies the controversy about whether what determines a 'forest' is its land cover or its land use (or different combinations of both). As it is often the case, both are highly intertwined, since the use to which land can be put depends to a large extent on the material and ecological properties of its land cover, and conversely the land cover depends substantially on the use to which the land has been put.

Moreover, despite the fact that many (if not most) vague predicates display conceptual vagueness to some extent, this variety has been mostly neglected in the dominant literature on vagueness, perhaps because it is less attractive for logico-philosophical analysis than the single dimensional case illustrated by the sorites paradox. Nonetheless, the phenomenon has received more attention from other philosophers of language who have approached it in a more descriptive fashion. Indeed, family resemblance concepts [Wit09] and open-texture theories [Wai45] illustrate this kind of phenomena, although the authors did not consider logic to be an adequate tool for its investigation.

2.3.4 Polysemy

The phenomenon studied in linguistics and cognitive science under the name of *polisemy* denotes the capacity for a sign (such as a word, phrase, or symbol) to have multiple meanings or senses related by contiguity of meaning. Thus, if they where to be represented in a semantic field, polysemous terms would be closely clustered senses without well defined boundaries as opposed to homonymous ones, which would be disjoint and clearly separable [Tug93, VP06].

The accounts of the notion of polysemy in cognitive science and linguistic literature (e.g. [WPB05, VP06, RGMW02, KM02]) and that of conceptual or non-numerical vagueness can be considered somewhat parallel. In the case of polysemy there is often an underlying assumption that there is a set of senses, each of which can be listed and (at least roughly) separated, as opposed to a single concept who's semantics are highly variable. However, it is acknowledged in linguistic research that determining the number of senses of a word, defining them, and saying where one ends and another begins is non trivial and hence the idea of a well defined number of precise senses in polysemy should be relaxed. Some of the most recent work in Word Sense Disambiguation (WSD) goes further in this direction by completely avoiding the lists of senses and adopting 'a more Wittgensteinen approach' that fits better with cognitive research, pointing at the human ability to think flexibly [Dea88]. In this thesis we take a similar approach and hence include polysemy in the general notion of SVoNLT.

2.3.5 Related (yet distinct) phenomena

We finish this section by reviewing some phenomena that are often discussed (and sometimes confused) with vagueness in the literature, and that we do not consider to fall under the umbrella of the SVoNLT. Those are generality, uncertainty due to a lack of knowledge and ambiguity when defined as equivalent to homonymy.

2.3.5.1 Generality

While in some of the literature in cognitive linguistics no clear distinction is made between vagueness (and more generally the SVoNLT) and generality or underspecificity (e.g. [Lak70, Tug93, Zha98]), here we do differentiate them sharply, on the basis that generality does not necessarily display uncertainty in meaning.

For example, if we compare the statements 'I am in my twenties' and 'I am 29' we find that, although the first proposition is more general, it is not at all vague: It is true for me being any age exactly within the twenties and false otherwise [Ben98]. Moreover, 'I am in my twenties' is more general than 'I am 29' just in the same way that 'I am 29' is more general than 'I am 29 years and 3 months old', showing that there is no obvious criteria to decide what is general and what is specific. Instead, generality may be better understood as a relation that holds between propositions rather than a characteristic of a proposition itself. 'I am approaching 30', however, is vague because it doesn't have a clear range of applicability.

While generality and vagueness have often been conflated in the linguistic literature (e.g. [Tug93]), the distinction is generally recognised in the philosophical literature on vagueness, since it was proposed by Russell [Rus23]. According to him, a proposition is general when there is a well defined set of possible facts or states of the matter that would verify it, as opposed to unclarity about which states of the matter satisfy the proposition[1] (different accounts have been provided, such as those of Pierce [Pie05] and Burns [Bur95]). Consequently, we do not consider generality to be part of the SVoNLT, yet we will briefly consider the representational issues in our framework (section 5.6.2) when we tackle the expressive capabilities of the standpoint framework in chapter 5.

2.3.5.2 Uncertainty

We also differentiate vagueness and the SVoNLT from uncertainty due to lack of knowledge about the state of the matter: while the former refers to the variability with which language applies to a known state of affairs, the latter considers the lack of knowledge about the state of affairs itself.

[1] Curiously enough, it is reported in [Wil99] that Russell then contradicts himself confusing vagueness and generality when defining precision.

> "... when Tek is borderline tall, it does seem that the unclarity about whether he is tall is not merely epistemic (i.e. such that there is a fact of the matter, we just do not know it). For a start, no amount of further information about his exact height (and the heights of others) could help us decide whether he is tall. More controversially, it seems that there is no fact of the matter here about which we are ignorant: rather, it is indeterminate whether Tek is tall." [KS96a]

Given that one of the main philosophical theories of vagueness is the *epistemic view*[Wil94], the reader could think that this distinction is contentious in the literature. Yet, the epistemic ignorance of [Wil94] refers again to the lack of knowledge about the precise semantics of vague terms (which are presumed to exist), and hence of the criteria of applicability of a predicate to a given state of the matter. As such, despite knowing the exact height of Tek, we may be ignorant of the true threshold of applicability of the predicate *tall*. The epistemic view of vagueness is discussed in more detail in section 2.4.2.1.

2.3.5.3 Ambiguity

Another phenomenon related to the SVoNLT but not considered part of it in this thesis is ambiguity, which is generally defined as the existence of several meanings associated to a single term. As introduced in 2.2.1, two different kinds of multiplicity are usually distinguished: *polysemy*, as discussed previously, and *homonymy*.

Both polysemy and homonymy are studied in linguistics. While in the former those senses are related by contiguity, in the later they convey unrelated meanings, being often the result of mere linguistic coincidence. It can be interpreted that rather than there being a term with different senses, there are two different terms which are indistinguishable in their spelling [Tug93, VP06].

In the literature, ambiguity is usually described as either the conjunction of both polysemy and homonymy (e.g. [Tug93, WPB05]) or as homonymy alone (e.g. [IW07, Zha98]. Traditional strategies for ambiguity detection (e.g. conjunction reduction) similarly adopt the latter approach (implicitly).

Despite previous attempts to clarify what we understand as ambiguity and to analyse its nature, research on the field reveals a lack of uniformity on its theoretical treatment, accentuated by a tendency not to specify which are the commitments embodied in the proposed representation and resolution techniques. In this work we narrow the meaning of ambiguity to the linguistic notion of homonymy and, moreover, we do not consider it to be part of the SVoNLT. This is because, while ambiguity (i.e. homonymy) is akin to conceptual vagueness and polysemy in that the semantic variation is qualitative (rather than graded), in this case the meanings are entirely distinct. Consequently, rather than agents adopting different standpoints that narrow the semantics of terms in different ways, they face a binary choice, or rather they use the same word for two concepts. The underlying logic mechanisms are therefore different; further discussion on the limitations of our framework for modelling ambiguity will be provided in chapter 5, section 5.6, and, in contrast, an example of a logic intended for representing this phenomenon can be found in [WL16].

2.4 Theories of (and approaches to) the SVoNLT

After the clarification of the meaning of the SVoNLT and the discussion of the phenomena that the notion covers in the previous section (2.3), we move back to the point where our historical background ends. In this section, we overview the main contemporary formal theories of vagueness, introduced in section 2.2.2. We focus on supervaluationism, which is the underlying theory of *standpoint logic*, and we subsequently discuss the epistemic view on vagueness and many-valued approaches. Finally, we consider a set of alternative approaches that originate from a more multidisciplinary area of research (introduced in 2.2.3), highlighting the differences between them and our subject of interest.

2.4.1 Supervaluationism

According to the supervaluationist theory of vagueness, the failure of vague predicates to divide the domain into two sets (its positive and negative extensions), neatly and without a remainder, is explained by the fact that language can be precisely interpreted in many different yet acceptable ways [KS96a]. The contested area that may originate between interpretations is considered a truth-value gap [Fin75], and the possible interpretations are commonly referred to as *precisifications*.

Each precisification is an admissible yet arbitrary sharpening of the vague predicates (e.g. one precisification could determine that 4 grains of sand make a heap) and it is as good as any other sharpening. This results in predications having an undetermined truth value whenever they have borderline cases of application that can be sharpened in different ways. Overall, supervaluation semantics is involved with the representation and reasoning with such a set of possible sharpenings or classical interpretations of vague terms.

An early proposal that vagueness could be analysed in terms of multiple precise senses was made by Mehlberg in [Meh58] (as discussed in section 2.2.2). Lewis defended the position in [Lew70] and, in [Dum75], Dummett provided a description of the supervaluationistic approach and its attractions (which he rejected). Subsequently, drawing from the formal semantics used by Frassen [vF68], Fine applied a similar model to the analysis of vagueness in the influential [Fin75], and Kamp offered an analogous approach from a linguistic perspective in [Kam75]. After that, supervaluationism has become a popular theory of vagueness, which has been adopted by philosophers, logicians and linguists. Yet, there have been little attempts of applying it in computational scenarios.

A remarkable feature of the theory is that precisifications sharpen the semantics of vague terms but do not interfere with non-borderline scenarios or facts [Fin75]. Consequently, the precisification of a predicate must always validate what is unproblematically true and falsify what is unproblematically false, including both non-borderline cases of the predicate and *penumbral connections*. *Penumbral connections* are general statements that constrain the way vague predicates can be sharpened, helping to avoid implausible valuations of compound propositions . Examples of this are 'someone that is short is not tall' and 'anyone that is taller than a tall person is also tall'. In this way, penumbral connections ensure that any sharpening of a vague predicate respects certain laws or intuitions established in language use.

In order to support the former, supervaluationistic frameworks are not truth-functional. The truth-valuation of a proposition is, instead, determined by the collection of all complete precisifications of the language. In particular, a sentence is *super-true* iff *true* on all precisifications, and *super-false* iff *false* on all of them. Let us consider the valuation of some examples; Assuming that Tim is a borderline case of tallness, the proposition 'Tim is tall or Tim is not tall' is *super-true* because, wherever we set the threshold for tallness, 'Tim' will be either tall or not tall. Conversely, 'Tim is tall and Tim is not tall' is *super-false* because there is no way to make the language precise without it leading to a contradiction. Furthermore, 'Tim is tall and Tim is short' is similarly *super-false* if the appropriate penumbral connection is established.

One of the advantages of adopting supervaluationism is that it preserves classical theorems (they are super-true because they hold in all the complete — hence, classical — precisifications of the language) while the theory still provides the means for reasoning with vague predicates and borderline cases, as shown in the previous examples. The main criticisms to the theory are mostly directed at its handling of the philosophical problem of higher-order vagueness (see section 3.7) and at some proposed notions of validity, which we will discuss in section 3.6. However, these objections are not relevant for the aims of this thesis, given that the phenomenon of higher order vagueness does not occur naturally in language and that in our framework a standard notion of modal validity is used.

To conclude, we highlight some ways in which supervaluationism is an attractive theory with regards to the problem of interest of this research. In the first place, modelling vagueness through collections of admissible precise interpretations is not only useful to represent the variability of language as a whole, but it can also be applied to the notion of standpoint, which can be thought as a smaller collection of interpretations that results from the sharpening of the general semantics of terms, rending the framework suitable for multi-agent systems. Moreover, the preservation of classical theorems and penumbral connections is useful for a variety of applications like our forestry information use case, where different standpoints include a range of constraints that must be strictly satisfied for facts to hold according to that view-point. Finally, the acknowledgement that all admissible precisifications are true in some sense and are on an equal footing (because language is genuinely vague) is well suited for describing multi-agent scenarios where a different use of language does not imply that one of the agents is misusing it, in contrast to the epistemic view (following in section 2.4.2.1).

2.4.2 Other formal theories of vagueness

In addition to supervaluationism, there are two other prominent philosophical approaches to vagueness, namely the epistemic view and many-valued theories.

2.4.2.1 The epistemic view

According to the epistemic view of vagueness, attributed to the stoics [KS96a] and most influentially defended by Williamson [Wil94], vague predicates do, in fact, have sharp boundaries, and hence borderline cases are in reality either true or false. The lack of clarity portrayed in the sorites paradox is then interpreted as an unavoidable kind of ignorance: 'thresholds for vague predicates are not just unknown; they are unknowable' [Sor01] (yet they exist). The main consequence of such a stance is that, at least from a philosophical perspective, classical logic and semantics can be preserved.

Despite the apparent implausibility of such position, the epistemic view of vagueness has been significantly endorsed by authors such as Cargile [Car69], Sorensen [Sor88] and Williamson [Wil92, Wil94]. The defendants of this position often present it as a reasonable alternative after the exposition of the problems that non-classical alternatives face (mostly higher-order vagueness) and the drawbacks from abandoning classical logic and semantics. These authors have also explored different aspects of natural language in order to give accounts of how can vague terms have sharp boundaries [Wil92, Sor01], and provided a treatment of higher-order vagueness.

Relaxed interpretations of this approach include the pragmatic view that even if there is actually no such precise language in reality, individuals, when faced with decision problems about assertions, find it useful to behave as if the epistemic theory is correct. This has been refered to as the *epistemic stance*, whereby 'each individual agent in the population assumes the existence of a correct set of language conventions, governing what can appropriately (or truthfully) be asserted given a particular state of the world' [Kyb00, LT12].

It follows from the epistemic assumption of sharp boundaries that vagueness is a purely epistemic phenomenon: it is lack of knowledge, akin to beliefs and knowledge about the world but with regards to the language instead. In [Law06], Lawry names these two sources of uncertainty *possible worlds uncertainty* and *semantic uncertainty* respectively. Moreover, some frameworks for vagueness have considered both issues to play an important role on a characterisation of vagueness. This is the case of the bipolar belief framework [Law06], that models the phenomenon of vagueness as both the genuine lack of precision of the meaning of terms and also the agent's necessary lack of knowledge about the exact conventions that rule the use of language. In addition, the framework includes *possible worlds uncertainty* in the model. While we view this approach as realistic, we do not integrate epistemic issues in our theory. Instead, we focus our attention to a simple framework for representing the variability of language, and, rather than modelling the decision process that an agent faces to decide if a statement is assertible [LT12], we concentrate on the reasoning that can be done from her assertions once they have been made.

With regards to logic frameworks, the epistemic view is often formalised via epistemic modalities. For instance, Williamson, in [Wil94], proposes the 'logic of clarity', with the modality C (which reads as 'It is clearly the case that') built around the notion of a margin of error. Intuitively, if W is a set of possible worlds, d a measure of their similarity, α a margin for error and $[A]$ is the set of worlds at which A is true, then we say that CA is true ('It is clearly the case that A') at a world w if A is true at every world within the margin for error α of w [Wil94]. Subsequent research have often used, however, standard epistemic logic (eg. [FHMV95]).

Beyond the philosophical considerations on the merits and plausibility of this approach, we consider that the epistemic view on vagueness is not a natural framework for the representation of the problems of interest for this project, because the underlying assumption that vague terms are actually precise collides with our intended *standpoint* usage, where different agents can have equally acceptable standpoints on the meaning of terms, which can be conflicting between them.

However, in many instances frameworks implementing the epistemic view of vagueness and modal approaches to supervaluationism are formally very close. Hence, there is often room for the reinterpretation of the semantics given to a framework, such that they adapt to different theoretical interpretations.

2.4.2.2 Many-valued logics and fuzzy logic

Many-valued logics are non-classical systems that accept the principle of truth-functionality but have more than two truth values; in most cases, they are either three-valued or infinite-valued, with truth degrees ranging from '0' to '1'. They emerged as a subject in the early twentieth century with Łukasiewicz logics [Łuk20] and were conceived to use a third additional truth value for 'possible' in order to model modalities. Since, many systems have been proposed with different aims, such as Gödel [Göd32] or Jaśkowski logics [Jaś36].

Several many-valued theories of truth have used these logics for modelling vagueness on the basis that, if borderline case predications are not clearly true nor false, then additional non-classical truth values may be needed (e.g. see Tye [Tye94] for three-valued and Zadeh [Zad75] for infinite-valued). While in three-valued systems borderline cases are *indeterminate*, infinite-valued systems aim at reflecting the continuity of phenomena (such as the possible heights of people) by assigning a continuous range of truth values to vague predicates like *tall*. These non-classical systems have different semantics and notions of validity, often determined by either a certain threshold for truth (which could be *true* or '1', or a lower threshold such as '0.5') or by the preservation of degree of truth (e.g. the conclusion is at least as true as the premiss) – choices on the former impact on the appropriateness of the logic for modelling different phenomena.

Three-valued theories are the most simple representations of this sort, and consider a third value *borderline* that represents 'not true nor false' with Kleene logic connectives (e.g. [Tye94] and [LGR11] for more recent work). For instance, in [LGR11] vague predicates are handled by means of valuation pairs that capture the three values *true*, *borderline* and *false*, and that can express the basic three-valued connectives of conjunction, disjunction and negation. Moreover, these logics can also be used for representing uncertainty, when interpreting the third value as *unknown* instead of *borderline*.

However, the best known of the many-valued approaches to vagueness is fuzzy logic, an infinite-valued logic proposed by Zadeh in [Zad75] which develops fuzzy analogues of standard set-theoretical notions and has been remarkably successful in the domain of computing. In this system, functions (which correspond to vague predicates) map objects to numbers in the interval [0,1]. For example, the function for *tall* maps people to values between '0' and '1', where a person that is not tall will be assigned the value '0', a clearly tall person will be assigned the value '1' and borderline cases will have intermediate truth degrees such as '0.7'. Correspondingly, instances belong to the '*tall* fuzzy set' to the degree determined by the former function.

Since the 1970s, the expression 'fuzzy logic' started to be used in two different ways; not only as a particular logical theory developed for modelling partial truth but also, on a broader sense, to refer to logic systems, formalisms and techniques handling degrees, particularly when these are modelled with fuzzy sets: an extension of set theory in which the membership of elements to classes is a matter of degree. For example, in engineering contexts 'fuzzy logic' (also fuzzy control systems, fuzzy classification, ...) is aimed at efficient methods tolerant to suboptimality and imprecision [Ros04].

Contemporary research on mathematical fuzzy logic is mature, including detailed investigations of algebraic structures, graded notions of entailment, complexity issues, proof theory and automated theorem proving [CHN11]. Moreover, applied research in the area of computer science is widely recognised, most of it occurring in its broader sense, being remarkable its success in the engineering domain (see [Ros04]). In contrast, its applications for modelling vagueness in natural language are less prominent[1], and its suitability for its analysis has been sometimes questioned in the domains of philosophy and linguistic semantics [Sau11, Wol13].

Different semantics have been proposed for fuzzy logics modelling vagueness (see [DP97, Law06] for a review of some of them). For instance, some adopt a prototype theory approach, so that the degree of membership is interpreted in terms of similarity to the prototypical instances of a class, or probabilistic approaches, several of which have been put forward aiming to model the underlying uncertainty or variance of the vague terms, or the degree of preference or assertability among others. Examples range from the pioneer work of Black [Bla37] to contemporary frameworks such as [Law08], [Las11] and [LJ17], which suggest that semantic uncertainty is a likely consequence of the empirical way in which language is acquired and, in the latter, that stochastic assertion decisions can play a positive role in some communication scenarios.

Some general criticisms to fuzzy logic as a formalism for modelling vagueness were formulated in [Kam75] and [Fin75] (followed by [Wil94, KP95, Haa79] among others) and they generally question the appropriateness of truth-functional systems[2] for natural language vague terms. Despite the fact that truth functionality provides simplicity and intuitively generalises classical logic, critics point at the awkwardness of the valuation of certain composite expressions when applied to language. A fuzzy logic parallel of an example provided for supervaluationism is presented in [KS96a]:

> Suppose Tim is borderline tall (say, tall to degree 0.4) and Tek is taller (tall to degree 0.5); and assume negation flips values so that "Tek is not tall" is

[1] For example, if we analyse the results of searching for *fuzzy logic* in the last five years in Web of Science, we find that most of the results are in the branch of *engineering electrical electronic*. Overall, 9,287 out of 14,715 results were classified in an engineering/industrial category. This contrasts with the 701 results in *computer science information systems*, 67 in *logic*, 21 in *philosophy* and 4 in *linguistics*

[2] Note that some many-valued frameworks diverge from classical formalisations of fuzzy logic in that they are not truth-functional, particularly those based sets of classical valuations of the language (e.g. [LT12],[FK06])

> as true as "Tek is tall". Then consider (a) "Tim is tall and Tek is not tall" and (b) "Tim is tall and Tek is tall." Truth-functionality would imply that (a) and (b) must have the same value. But it seems that (a) must be false: if Tim is shorter than Tek, then it cannot be that Tim is tall and Tek is not. And (b) is surely not false for a degree-theorist, but is true to some positive degree.

Despite these criticisms, fuzzy logic has been successfully applied in different knowledge representation scenarios involving vagueness, such as ontologies and description logics [Str01, CC18]. The main appeal is the graded membership to the fuzzy set, which is used to capture not only vagueness but also uncertainty in a more general sense. Moreover, probabilistic approaches such as [Law06] avoid the odds of the aforementioned truth-functionality, and bridge towards the epistemic view and supervaluationism while preserving the infinite-valued system, have been proposed, for instance, in [FR09] and [LT12].

However, with regards to the problem under consideration in this thesis, infinite-valued systems are not directly suitable to represent the notion of truth relative to a context or view-point, hence we do not consider fuzzy logic a natural theory for the standpoint framework. One could conceive, however, an extension of the *logic of standpoints* supporting degrees of truth relative to a standpoint, perhaps by introducing some kind of degree of membership of each precisification in a standpoint. Yet, such a line of work is not explored in this thesis.

2.4.3 Other multidisciplinary approaches

The previously discussed approaches to vagueness are philosophical theories that arouse from the aim to solve the sorites paradox and were intended to dissect the underlying logical structure of vague terms; It was only subsequently that they were studied with regards to their capacity to solve problems arising in computational scenarios. In contrast, we now briefly overview a set of approaches and techniques that have been developed with a practical aim or cognitive foundation and that are relevant for addressing other aspects of the SVoNLT. In particular we will focus on three areas of research, namely sense approximation and resolution techniques and cognitive frameworks.

2.4.3.1 Sense approximation and resolution

In the first place, we consider the body of work on sense approximation and resolution of semantically variable terms (WSD). In a nutshell, the task involves the computation of the intended sense of a term in context. This can be done by selecting the relevant sense from a discrete set of options, or by computing a sense within a continuous semantic space. The field agglutinates a variety of techniques, mostly based on learning methods and many of which use vector representations [LT93, YCF01].

Most, but not all strategies for WSD use heuristic methods to acquire knowledge from big corpora and knowledge bases such as WordNet [Mil95], and they use it to predict the senses of words in context. There are supervised and unsupervised methods [GSD04], corpora may be annotated or not [AE07], and the resolution may consist of the selection of one sense from a set of possible ones or on the automatic generation of senses through clustering algorithms [ASPPBL06]. This area of research is mature; a comprehensive literature review on the topic with an overview of the current techniques and applications can be found in [Nav09]. In particular, the vector space model is a representation framework that supports approximate inference and performs well in this domain [LT93, YCF01].

To understand the relation of these frameworks to *standpoint logic*, we highlight the way in which they differ: The former focus on estimating the intended sense of a word in use (normally in a corpus), and generally consider senses that are either given or generated from the corpora; In the latter, we are interested in the way that semantic commitments enable agents to set the standpoints with respect to which they make assertions, and the logical reasoning that can be done with sets of them.

2.4.3.2 Cognitive frameworks

Secondly, work in cognitive science has stimulated the development of alternative models for understanding the semantics of concepts (some examples are *conceptual spaces* [GW01], *image schemas* [Joh87, Lak87] and *connectionist models* [BA91]). All the former are explanatory theories of language that account in some ways for the semantic variability of its terms and, while some of them have not been conceived with an intended application, they have all motivated or inspired subsequent applied research.

For instance, *conceptual spaces* [Gär04], based on prototype theory [KP95, Ros88], deal with concept formation and representation, where semantics are expressed within a geometrical space. For example, a colour concept such as ‘orange’ can be associated with a region in the space of all possible colours. This approach has been applied to concept comparison and conceptual adaptation to context (e.g. in the case of geographic information [AJ11]) and has been employed in accounts of prototypicality and borderline cases [DDDÉ13]. However, it is unclear whether conceptual spaces can deal with other forms of SVoNLT (particularly conceptual or non-numerical vagueness, see 2.3.3), as every concept is defined within a fixed set of quality dimensions; thus, a choice of the relevant ones in complex and ambiguous categories implies ruling out different and potentially close understandings of the meaning of such concepts that depend on different dimensions.

Cognitive frameworks are successful at modelling concepts in alternative ways other than a definition and they have a cognitive foundation. However, in the context of this research, they provide less flexibility than general purpose logic languages to encode semantic constraints and domain axioms. This imposes limitations on how can we characterise standpoints on the semantics of terms, particularly in scientific domains such as our applied scenario (in chapter 6). In contrast, these approaches are more focused on capturing the general ‘common-sense’ semantics of terms rather than the specific semantics with which words can be used in a certain context.

2.5 Formal logic frameworks

In this section we specifically overview relevant formal frameworks (rather than theories) for this research. We begin by reviewing previous work in establishing a logic of standpoints in subsection 2.5.1 and then we focus on the work on modal approaches 2.5.2, where we introduce modal logics, we review the literature formalising supervaluationism via a modal logics and we overview the work on multi-modal logics. Finally, we conclude by considering the non-modal frameworks (subsection 2.5.3) that have been used for modelling supervaluationism.

2.5.1 Preliminary work on a standpoint framework

In [Ben06], Bennett provides the first account of a logic of standpoints. There, a standpoint is associated with ‘a set of precisifications considered acceptable by some agent', and truth is primarily contemplated 'as a property of propositions that is relative to a particular precisification, rather than determined by the whole set'[1].

In Bennett's framework, two modes of vagueness, namely *sorites* and *conceptual vagueness* (see section 2.3.3), are modelled separately. Vague terms have various admissible definitions or predicate groundings that are specified in terms of primitive functions representing observable measurements, which models *conceptual vagueness*. Moreover, those are associated with thresholds of applicability, which encode the *sorites vagueness* of such groundings. Bennett then provides a two-dimensional semantics in which the interpretations of propositions are indexed by precisifications and possible worlds. Implementations of the framework in computer systems can be found in [BMT08] and [SBS05].

In this thesis, we revise the work in [Ben06] in different ways. We introduce a modal framework and we do not explicitly differentiate between predicate grounding and threshold assignment. Instead, we provide a uniform treatment of vagueness within which any formula expressible in the language can constrain the semantics of vague predicates. While we believe that the distinction made in the earlier version of the formalism is a useful way to organise semantic constraints, we do not consider that it is advantageous to separate them in the logic language.

Second, while Bennett's work was focused on the development of a grounding theory upon which the vague language was based, we take a somewhat different approach. We depart from the assumption that, in most scenarios, such grounding theories are not fully defined or that such information may not be available. For instance, a standpoint on the semantics of a vague predicate like *nice* is more naturally expressed or constrained in terms of other vague predicates such as *pleasant*, than in terms of a physically grounded theory relying on objective and measurable parameters. Hence, while the latter is expressible in our framework, softer commitments using vague terminology are also allowed.

Finally, the formalism that we now present goes significantly beyond [Ben11] in explicitly representing standpoints within the logic of the object language, which in our framework is done by employing modal operators (previously standpoints were handled by an auxiliary semantic apparatus). In what follows, we will provide our

[1]While [Ben06] is not a modal framework, it does adopt the notion of truth at a point/precisification characteristic of modal interpretations of supervaluationism.

language with standard Kripke semantics, and prove soundness and completeness of the logic with respect to a class of Kripke frames.

2.5.2 Modal frameworks

2.5.2.1 Introduction to modal logic

Modal logics are formalisms that extend classical logic by means of additional modal operators that qualify truth. For instance, a proposition P, standing for 'Phil is here', can be qualified with the traditional (alethic) modal operators of possibility and necessity, such that $\Diamond P$ reads as 'It is possible that Phil is here' and $\Box P$ as 'It is necessary that Phil is here'. Different systems of modal logics can be used broadly to formalise a variety of modal notions, such as those used for temporal, deontic, epistemic and doxastic modal logics.

The semantics of modal logics are usually given in the form of Kripke models, which will be formally defined in section 4.2.2.1. These are characterised by having a set of possible worlds or points of evaluation and an accessibility relation (or a set of them) on the set of worlds as well as the evaluation function. The intuitive idea is that 'It is necessary that Phil is here' if 'Phil is here' holds at all accessible worlds, and 'It is possible that Phil is here' if 'Phil is here' holds at some accessible worlds. Validity is then defined as truth in every possible world in every possible model.

A system of modal logic is characterised by a particular set of axioms and their corresponding restrictions on the relations of the Kripke models. Some well known systems of modal logic are determined by some or all of the following schemas:

$$\mathrm{K} : \Box(\phi \to \psi) \to (\Box \phi \to \Box_s \psi)$$

$$\mathrm{D} : \Box \phi \to \Diamond \phi \qquad \text{(serial)}$$

$$\mathrm{T} : \Box \phi \to \phi \qquad \text{(reflexive)}$$

$$4 : \Box \phi \to \Box\Box \phi \qquad \text{(transitive)}$$

$$5 : \Diamond \phi \to \Box\Diamond \phi \qquad \text{(euclidean)}$$

While **K** holds in standard Kripke models, the rest of the schemas impose the corresponding relation on the models. In this thesis, we will make special reference to the well known systems **KD45** and **S5**, which have the schemas K, D, 4, 5 and K, D, T, 4, 5 respectively. These axiomatisations restrict the relations on the corresponding models to be serial, transitive and euclidean in the case of **KD45**, and equivalence relations (as a result of adding reflexivity to the former) in the case of **S5**.

2.5.2.2 Modal approaches to supervaluationism

As in this thesis, there have been several proposals to treat vagueness as a modal phenomenon, and in particular to implement the supervaluationistic approach to vagueness with modal logic. Examples of that are [Var07, ADP09, Wil08, Cob08, Ben98] among others.

In a modal approach to vagueness, simple notions of truth are replaced by the notion of modes of truth with regards to vagueness [ADP09] in a similar way as it is done in other modal systems, such as epistemic or doxastic logic. In a supervaluationistic modal logic, a statement such as 'Phil is tall' can be qualified by saying that Phil is *unequivocally tall* (if Phil is a clear case for tallness) or that he is *in some sense tall and in some sense not tall* (if he is a borderline case for tallness). In the context of supervaluationism, in the former case we would be asserting that it is *supertrue* that Phil is tall; in the later, we are asserting that Phil is tall according to some but not all the precisifications of tallness.

When we provide the supervaluationistic modal logic with Kripke semantics, the structure of points (or worlds) of evaluation is used to model the precisifications of the language. In such a structure, the notion of truth-at-a-point replaces global notions of truth and reads as truth relative to one precisification [ADP09]. Modal truths are then derived from different configurations of points (or precisifications) and their valuation. A modality that is often provided in such logics is one that implements the notion of *super-truth* (truth at all points), but more operators have been proposed. Work exploring accounts along these lines include [Var07, ADP09, Wil08, Cob08, Cob11, Cob16], and a review on the different operators is provided in chapter 3, section 3.5.3.

It has been argued that some of the main philosophical objections to the supervaluationistic account of vagueness are not applicable if we interpret it as a modal phenomenon [ADP09]. Modal interpretations simplify Fine's presentation of the theory, which is somewhat obscure in its intensional and extensional accounts and has been attacked for its treatment of higher-order vagueness via the D (*determinacy*) operator.

Most of the referenced work on modal accounts of supervaluationism focus on analysing different forms of validity and logical consequence [Var07, Wil08, ADP09, Cob08, Cob11] and on responding to some of the main objections to supervaluationism [Far10, Wil18]. These, as well as the notion of validity in our system, will be discussed in chapter 3, section 3.6. Other research includes the combination of modal supervaluationistic approaches with other frameworks like fuzzy logic [FK06], some investigations on handling vague objects [Aki00, Swa14] and the formulation of responses to other objections to supervaluationism, such as the handling of indirect speech [GC10, Kee10].

However, established work providing fully-formalised supervaluationistic modal logics intended for operational frameworks for knowledge representation and reasoning is scarce. A preliminary example of this is [Ben98], where an early formalisation of a modal logic for vague concepts is given. Bennett's two-dimensional modal logic considers both possible worlds and possible interpretations, and provides a modal operator **U** for *unequivocality* (for modelling vagueness) as well as the standard alethic operators (necessity and possibility). In Bennett's work, precisifications are three-valued, hence they are not forced to be precise, in contrasts to the framework proposed in this thesis. While the potential applications of the framework in different computational scenarios are explored to a greater extent than in other related work, a formal semantics and proof theory of the logic are not provided. In [Ben01], a prototype of the framework is applied to the spatial reasoning domain and, along these lines, a modal logic for description logics is presented in [LC06]. Another implementation framework based on supervaluationism, in this case only for modelling the spatial terms near and far, can be found in [MM13].

Finally, more recently, [LYV16] presents a philosophical analysis of various previously suggested modal operators, together with an axiomatisation suitable for both supervaluationistic and epistemic logics of vagueness, which is contrasted with our analysis in chapter 3, section 3.5.3 and highlights the formal closeness of frameworks implementing the epistemic view of vagueness and supervaluationism, the latter particularly when based on local (modal) notions of truth. The framework, which is a three-dimensional modal logics with operators for actuality, 'supernecessity' and definiteness, explores and models philosophical aspects of epistemic, supervaluationistic and metaphysical [BW11] approaches to vagueness.

2.5.2.3 Relation to modal approaches to uncertainty

Modal frameworks are however much more commonly used to model uncertainty (as in lack of knowledge) rather than vagueness. This is particularly the case since the use of modal logic frameworks for modelling knowledge and belief [Hin62] became widespread. Epistemic logics [FHMV95] are modal frameworks for reasoning about incomplete knowledge and beliefs. They are well established in the domain of knowledge representation and they have been widely used to reason about epistemic uncertainty (e.g. [HP98, Auc10]). Formally, they are normal modal logics satisfying **S5**[1] for reasoning about knowledge and **KD45** for reasoning about beliefs [FHMV95]. The semantics is generally given in terms of Kripke structures, based on an accessibility relation between possible worlds or points. A modal proposition $\Box\alpha$ is true at a world w if and only if it is true at all accessible worlds w' under the relation R.

These logics are intended for modelling the knowledge and belief of an agent about the state of affairs, rather than the semantics of language. However, different variants have been used for reasoning about vagueness, particularly following an epistemic interpretation of vagueness, and on occasions due to the conflation of uncertainty (on the state of affairs) and vagueness (e.g. in [BD14] it is unclear for us whether they discuss possible worlds or possible interpretations).

Formally, *Standpoint logic* can be considered an extension of the formal frameworks that are commonly used in Epistemic logics, with a reinterpretation of its semantics (e.g. we model precisifications of the language rather than worlds) and additional schemata that capture the logics of standpoints.

One can also relate our work to MEL[BD14], a 'simple epistemic logic' to model uncertainty (on the state of affairs) that takes an unconventional approach in modal logics. In the first place, MEL uses only a fragment of the **KD** language, avoiding nesting of modalities and with no objective (modality-free) formulas. Moreover, instead of Kripke structures, its semantics is given in terms of epistemic states, understood as subsets of mutually exclusive propositional interpretations [BD14]. These epistemic states represent the 'incomplete knowledge (or belief[2]) about the real world possessed by an agent' as a disjunction of interpretations, 'one and only one of which is, according

[1] For more information on well known systems of modal logic such as **S5** and **KD45** the reader can refer to [Che80].

[2] It cannot distinguish belief from knowledge (true belief) since axiom T ($\Box\alpha \rightarrow \alpha$) cannot be expressed in the MEL language.

to this agent's beliefs, the actual state of the world' [BD14][1].

Indeed, the restricted language used in MEL could be, in a multi-modal flavour, interesting for *standpoint logic*. Like most modal logics, we use the general syntax extending propositional logic together with interaction axioms (in our case **AS4** and **AS5**, see chapter 4) to simplify the syntax. However, a restricted language may be interesting in some applied scenarios where the flexibility of the language is worth compromising, as we will consider in chapter 5, section 5.2.1.4.

Moreover, its semantics in terms of epistemic states relies on the results of [Pie09], that show that **K45** logics can be given simplified Kripke frames where the relations can be reduced to sets. While the semantics of *Standpoint logic* uses the same principle, extending it to the multi-modal case, we preserve the (simplified) Kripke style semantics (as in [Pie09] itself). In contrast, MEL's approach offers a logical grounding to other uncertainty theories like possibility theory and belief functions [BD14], and given that *Standpoint logic* could be given similar semantics despite its multi-agent nature, it highlights the links of our framework with other non-modal formal approaches, such as those introduced in section 2.5.3.

2.5.2.4 Multi-modal logics

Modal logics with a multiplicity of modal operators are common and have substantial applications in computer science. A prominent example is *epistemic modal logic* (logic of knowledge and belief) [Hin62, HM92], which uses the modal operators as a tool to represent the cognitive states of sets of agents (in particular their knowledge and/or beliefs) and provides the means for reasoning about them [Hin62]. This, in turn, provides a natural theoretical framework for reasoning with multi-agent systems [WJ94, dFL+97, FHMV95].

This multi-agent nature is achieved by introducing an accessibility relation for each agent in the system. Consequently, a multi-modal logic for n agents is obtained by joining n modal logics, usually described by the same logical system for simplicity [HS15]. In the case of epistemic logic, this provides expressivity to make statements such as 'agent α knows ϕ' or 'agent α knows that agent β knows ϕ'. From those further

[1] Note that in contrast to this, in supervaluationism the set of interpretations of language is conjunctive. It is in fact the whole set of admissible precisifications what characterises the meaning of a vague term.

notions can be defined, for instance *common knowledge*, that supports certain aspects of reasoning within groups of knowers.

In this research we are interested in taking a similar approach for the representation and reasoning with different perspectives that agents can take on the semantics of vague terms. Hence we aim to establish an adequate proof theory for language standpoints rather than mental states. This is a novel approach given that the literature reviewed so far on modal approaches to supervaluationism only discuss how modal operators can express the notions of *super-truth* and *definiteness* with respect to the whole language. We consider that modelling different perspectives of language has two fundamental benefits: on the one hand it provides a reasonable symbolic representational framework that characterises the semantics of vague terms in function of the uses that different agents make of them, which is particularly interesting for the modelling of conceptual vagueness; On the other hand, we hope to provide an interesting framework for reasoning with multi-agent systems in the presence of *semantic heterogeneity*.

2.5.3 Other logic frameworks for supervaluationism

Logic frameworks appropriate for supervaluationism must account for penumbral connections, which rends the logic non-truth functional. Beyond the modal approach, other alternatives have been proposed to model supervaluationist semantics, for instance through the use of valuation pairs (in this case particularly *supervaluation pairs*) as in [LD12]. Given a supervaluation pair, the lower truth valuation represents the criteria for absolute truth and the upper valuation represents the criteria for not absolute falsehood, which are evaluated over a set of classical valuations that encode the admissible precisifications.

This approach is, in form, strongly connected to *possibilistic logic*, a formalism related to fuzzy sets through possibility measures but which focuses on classical logic formulas pervaded with qualitative uncertainty [DP14], as discussed in [CDL14]. Moreover, through this connection, it also possible to relate this work to a variety of other formal frameworks for vagueness and uncertainty [CDL14].

In [LT12], *supervaluation pairs* are used in a framework that provides an integrated treatment of two modes of uncertainty, 'possible worlds uncertainty' and 'semantic uncertainty' (see section 2.3.5.2) with the assumption that interpretations of the language may admit truth-gaps, expressed with valuation pairs. Their model of epistemic uncertainty can be defined in terms of a probability measure on subsets of the cross product space of the precisifications (linguistic conventions) with the possible worlds, generating a probability distribution on valuation pairs.

In relation to the formal framework that is proposed in this thesis, there are close similarities. In particular, as will become obvious, standpoints are parallel, in Lawry and Tang's framework [LT12], exactly as sets of valuations (rather than individual valuations). This will be the case because standpoints capture partial accounts and commitments on the semantic variability of terms, rather than a well specified sharpening of the whole language. Hence, the framework proposed here allows for the unified representation of multiple standpoints, each of which is expressible as a set of supervaluation pairs.

Chapter 3

Introduction to Standpoint Logic

3.1 Introduction

In this chapter we present *standpoint logic*, our logical framework within the family of logics based on supervaluational semantics. The two key contributions to previous work on modal approaches to supervaluationism and vagueness are the introduction of the notion of *standpoint* (first proposed in [Ben11] in a non-modal framework) and the multiple modalities, intended to enable reasoning with vague languages in multi-agent settings.

Standpoint logic is a multi-modal logic with a particular characterisation of the notions of precisification, truth and validity, determined by a particular class of Kripke frames, and which intends to be an adequate system for multi-agent reasoning in the presence of a vague or semantically heterogeneous vocabulary.

This chapter is committed to the introduction and discussion of the main elements of the framework and to the justification of its particular interpretations within the supervaluationist approach to vagueness. In section 3.2 we introduce the motivations for the development of the framework with some very simple examples and in the following section, 3.3, we provide a general overview. Section 3.4 discusses *precisifications* in our logic, followed by *standpoints* in section 3.5. Finally, section 3.6 examines the notions of truth and validity in *standpoint semtantics* and section 3.7 briefly considers the issue of higher-order vagueness.

3.2 Motivation

We first introduce the kind of sentences and scenarios that we want to model. To begin with, let us consider an agent α stating the following assertions, (1) and (2).

1. (according to α) 'Nena is (unequivocally) tall.'

2. 'Pepa is taller than Nena.'

A competent English speaker should be able to draw the conclusion 'Pepa is tall' from the sentences (1) and (2). Yet, words like *tall* may mean different things to different agents, or even to the same agent in different contexts or occasions. However, while two agents may disagree on whether Nena is tall or not, they are quite unlikely to disagree on (2) and particularly on whether (3) is a sound conclusion from (1) and (2).

I.e. if we accept (2) and we are willing to accept (1) then the norms of language use require that we must also accept (3). This represents one of the paradigmatic examples of commonsense reasoning with vague terms. Consequently, our competent English speaker can safely draw the conclusion that, according to agent α, (3) is the case:

4 'Pepa is unequivocally tall according to α.'

Let now the sentence (4) be uttered by an agent β in the same context:

5 (according to β) 'Pepa is not tall.'

In that case, our competent English speaker can infer that:

6 'Nena is unequivocally not tall according to β'.

7 'Nena and Pepa are borderline tall'. Neither Pepa nor Nena are undisputed cases of tallness (they are not unequivocally *tall* or $\neg tall$).

Let us suppose that all agents know the height of both Nena and Pepa in the context of such conversation. Then it is highlighted that the two first inferences that our agent γ makes are not about α and β's beliefs about the state of affairs. Instead, they are inferences about their use of the term 'tall' facing the same state of affairs. Moreover, the last inference is valuable for our competent speaker as it highlights the fact that Pepa and Nena, as well as other instances of tallness within their height, are borderline cases of tallness and thus they can not be considered to be unequivocally *true* or *false*. Additionally, it is known that any person taller than Nena will be considered to be tall for α (if the context is constant) while any person shorter than Pepa will be considered to not be tall for β, and that any person considered tall by β will certainly be considered tall by α, who has a lower threshold for tallness, and that such person's height will be more than Pepa's height.

Hence we get information not only about the state of affairs but also we acquire information on the semantic commitments of the agents, which can be used in further communication.

3.3 Overview of the framework

The previous section highlights the multi-agent (or better multi-perspective) nature of the framework. In this section we overview its main elements and features, which will be subsequently discussed individually in the following sections.

In accordance with the general supervaluationist theory of vagueness, vague predicates such as *tall* fail to divide the domain into a positive and a negative extension, giving rise to borderline cases. Such cases can be interpreted as the discrepancy in the set of all possible precise interpretations of the language, its precisifications, and each of these precisifications is an admissible yet arbitrary sharpening of the vague predicates.

When an agent makes use of the vague language, she is unlikely to use a particular precisification, e.g. she is unlikely to commit to a fully precise language before referring to someone as *tall*. Instead, the agent is likely to narrow the variability of the relevant terms in order to communicate effectively, either explicitly or implicitly, for instance via semantic commitments. E.g. by stating that somebody is not tall she narrows the meaning of the term.

Our framework interprets the semantics of the vague language as a collection of admissible precisifications of the language. Agents using the language may refer to the whole set of precisifications or they may only use a subset of them, by adopting a standpoint that narrows or precisifies the meaning of the terms of the language. Standpoints therefore capture the intuition that an agent commits to the subset of precise interpretations of the language that are consistent with her own statements[1].

It must be highlighted that we do not consider epistemic limitations in our model, such as situations where different agents perceive different things. E.g.: 'The vision of α may be better than that of β and thus 'Pepa is taller than Nena' may hold for β but not for α. In fact, *standpoint logic* is not concerned with the perceptions or mental states of agents. Instead, it is concerned with possible interpretations of a language that agents can then use to describe the world in a non-precise manner. Particularly, with a language such that a proposition ϕ may hold or not for the same state of affairs, depending on the interpretation of the language (*standpoint*) to which the assertion ϕ is associated.

[1]This is presumed on the basis of general communication principles, for example supported by [Gri75]. The contrary would be regarded as a dishonest behaviour from the agent, which we do not model in this framework.

3.4 Precisifications

As discussed in section 2.4.1, the supervaluationistic approach to vagueness is built on the idea of precisifications. A *precisification* is identified with a precise interpretation of the language, such that every word has a unique, precise definition. The definitional theory may involve additional axioms and it must be consistent, such that for any state of the world there is a unique extension for every predicate.

A vague language is characterised by admitting several precisifications, and typically a statement using vague terms or predicates is said to be *super-true/super-false* iff it is *true/false* according to all of the admissible precisifications (see section 3.6 for the interpretation of *truth* in our framework).

There is however a substantial variation in the literature with respect to what precisifications are taken to be. In [Var07], Varzi outlines some of the key features present in the literature. The two first options relate to the conceptual interpretation of the framework and are:

(a) The distinction between precisifications being

  (a.1) classical languages in their own right, such that the vague language is a cluster of homophonic precise languages; or

  (a.2) precise interpretations or sharpenings of the vague language.

(b) The second distinction relates to whether

  (b.1) the vague term is literally defined by its precisifications; or

  (b.2) the vague term is analytically prior to the precisifications. Precisifications are then either the classical terms that have substituted the vague ones *(a.1+b.2)* or the interpretations obtained by sharpening them *(a.2+b.2)*.

It is argued that decisions on these aspects play an important role in both the criticisms to supervaluationist approaches to vagueness and their answers [Var07]. Both points are relevant to questions about the nature of natural language and the plausibility of the supervaluationist theory. Given that the present work is instead interested in supervaluation as a tool for modelling certain scenarios rather than as the definitive theory of vagueness we do not commit to any of those, and leave it open to the reader. Formally, our model is close to the interpretation *(a.1-b.1)*, yet it can be considered an

abstraction of any of the other. In fact, our reasoning with standpoints and commitments on the meaning of terms replicates the scenario in *(a.2)*, and *(b.2)* is certainly more compelling than *(b.1)*.
Varzi makes a last and more relevant distinction, this time focused on the scope.

(c) Precisifications can be

(c.1) relative to the whole language or

(c.2) we can speak of partial or limited precisifications, which only relate to the subset of the language that is relevant for the vague statement.

Precisifications in Standpoint Semantics are always relative to the whole language *(c.1)*. This feature facilitates the modal interpretation of precisifications as worlds or points in standard Kripke semantics.
Finally we do consider a last issue, on whether

(d) Precisifications

(d.1) ought to be *complete* (or classical) as in [Var07, Ben11], or

(d.2) can be incomplete, as suggested by Fine [Fin75] and others (e.g [ADP09, Ben98]), leaving certain propositions indeterminate.

In Standpoint semantics, like in most modal interpretations of supervaluationism (e.g. [Cob11, LYV16]), precisifications are classical/complete. Each precisification corresponds to a classical language with no borderline case scenarios or vagueness. Yet, partial orders of incomplete precisifications capture important intuitions on the process of sharpening terms and the relation between more and less precise interpretations of the semantics of sets of terms. Consequently, while we do not consider the possibility of incomplete precisifications, we do provide partially ordered structures for that purpose that we name *standpoints* and are discussed in the following section.

3.4.1 A note on admissible precisifications

The literature in supervaluation semantics discusses a further notion, that of *admissible* precisifications. In Varzi's words, 'precisifications should be admissible in the sense that some connections must be respected, such as analytic relations between expressions. These restrictions on the admissibility of a precisification enable the supervaluationist theory to endorse Fine's so-called penumbral connections, that is, connections that might hold among sentences even if these have a borderline status' [Var07].

How to account for such admissibility has been considered in the literature. In some modal frameworks it is suggested that admissibility could be formalised as the accessibility relation between precisifications/worlds [Cob11] and other work such as Fine's [Fin75] uses it as a primitive notion.

In the *standpoint* framework, the set of precisifications is the set of admissible precisifications. In the intended models, any precisification is complete and it ought to be consistent with the statements made in the language, including universal statements such as those corresponding to the notions **N1-4**. Given that analytic relations between expressions and other penumbral connections must be encoded as propositions of the kind 'It is unequivocal that ϕ', then any precisification in our semantics must be consistent with them and hence it is an admissible precisification.

Moreover our intuition is that, while an admissibility relation between precisifications has an intuitive role in a framework admitting incomplete precisifications (such that a precisification π' is admissible from another precisification π iff the former is more precise than the latter), there is not such a clear interpretation if we only consider complete precisifications.

3.5 Standpoints

Given that in Fine's terminology a precisification need not be completely precise, precisifications form a partial order where π_2 is more precise than π_1 if all propositions with a definite truth value in π_1 have the same truth value in π_2 but some propositions that are indeterminate in π_1 have a determinate truth value in π_2.

The partial order is useful for a variety of reasons, such as the fact that precisifications become comparable (one precisification can be more precise than another) and it models the intuition behind the sharpenings of the meaning of vague terms. For instance, we can sharpen the meaning of the vague predicate 'tall' by stating that a borderline individual is 'tall' (or 'not tall'), and then using this judgement as a precedent for subsequent uses of the word 'tall', thus narrowing the range of cases considered borderline.

In our framework, precisifications are akin to possible worlds and thus each corresponds to a complete, precise classical assignment. To capture the idea that the meaning of vague terminology is only partially defined, we propose and formalise the notion of *standpoint*, understood as a partial account of the semantics of the terms of a vocabulary, and we focus on how different standpoints interact.

Standpoints are modelled as collections of precisifications. They are typically organised by constraints (axioms and threshold limitations) that pick out a corresponding set of admissible precisifications that satisfy these constraints. In the example below, a standpoint might be constrained by a definition **D1**. Moreover, the value of the threshold t_tall linked to such definition might not be fixed but constrained to lie between certain values, as in **D2**:

(D1). $\forall x[\mathrm{Tall}(x) \leftrightarrow \mathrm{height}(x) > t\_tall]$

(D2). $t\_tall > 175cm \wedge t\_tall < 185cm$

Such values may not be explicitly given, since they could be inferred from assertions associated with the standpoint. For instance, if a person of height 175cm is asserted not to be tall, the threshold for tallness must be greater than 175cm.

Consequently, standpoint semantics uses a conceptually simple model that fits well with Kripke semantics, in which a partially determinate interpretation is called a *standpoint* and is modelled by a set of the (fully determinate) precisifications (or points) that are consistent with the (partially determinate) standpoint (relation).

Since they are modelled as sets, standpoints form a lattice under the subset relation; and, when one standpoint is a subset of another, we may also regard it as more precise, since it rules out certain interpretations of the vocabulary.

3.5.1 Standpoints as a tool for representation

Standpoints are straightforward tools for the representation of interpretations of the semantics of terms, which can be adopted by an agent or community, and be reused in the system. Moreover, we argue that the collection of standpoints is an interesting representation tool on its own right.

Beyond the expressivity of the underlying languages, with which individual standpoints can express different constraints such as **D1** and **D2**, the complete collection of standpoints of a system is a natural way to symbolically characterise the semantics of its vague terms as a whole.

Standard supervaluationistic frameworks model the abstract variation in meaning via the set of precisifications of the language, but such a set is unstructured. Whenever the whole semantic variation can be encoded via symbolic penumbral connections which apply to the global language (rather than individual standpoints), as in **D1** and **D2**, this is enough to represent the vague semantics of a term. However, conceptual or non-numerical vagueness resists this type of strategy, as we have seen in section 2.3.3, and hence previous supervaluationistic frameworks offer limited tools for the symbolic representation of concepts like 'clever'.

Alternatively, with standpoint logic we can provide a characterisation of the semantics of conceptually (or non-numerically) vague predicates by means of a (potentially big) collection of standpoints, which capture different constraints that hold in different scenarios and that encode different characteristic aspects of the meaning of such terms. This approach takes inspiration on the notions of *family resemblance* and *open-texture* concepts, yet it differs from *prototype theory* approaches in that, rather than learning the semantics from the instances of the concept, in *standpoint logic* they are characterised by the collection of different interpretations that have been specified and used in the system. The main upshot is that the characterisation is intensional rather than extensional, and that rich symbolic interpretations can be given, therefore providing a complementary framework to those based on similarity measures between instances.

3.5.2 Stanpoints in relation to Fine's theory

In Fine's presentation of supervaluation theory [Fin75], a specification space S is an arbitrary collection of partial models. Such a collection is, as shown, somewhat parallel to the set of standpoints and precisifications in our framework.

According to Fine, a rooted specification space is a specification space with a base point, one partial model identified as privileged. In our model, the analogy is the universal standpoint, which is referred to as $*$. The universal standpoint subsumes any other standpoint and it is characterised by the set of all admissible precisifications in the model; Anything that is unequivocally true in $*$ (that is, $\Box_* \phi$) is also unequivocally true in any other standpoint s (that is, $\Box_s \phi$).

Moreover, it is desirable (in Fine's proposal) for a supervaluationistic framework that the space satisfies the conditions of Fidelity (F), Completability (C), Stability (S) and Resolution (R).

FC *Completability* states that any point can be extended to a complete point within the same space. This can be formalised as follows, for t, w points in the specification space S and $\leq$ the extension relation :
$\forall t \in S, \exists w \in S(t \leq w \wedge w \textit{ complete})$

FF *Fidelity* states that truth-values at a complete point are classical.

FS *Stability* states that truth-values are preserved under extension points (in Fine's work it is modelled through monotonicity for precisifications).

FR The *Resolution Condition* for atomic sentences states that an indefinite atomic sentence can be resolved in either way (true or false) upon improvement in precision.

One obvious difficulty in order to draw parallels between the current framework and Fine's proposal is that, while Fine considers a single sort of entity (i.e. precisifications), we do consider both precisifications and standpoints. There are consequently different ways to interpret Fine's properties in our framework.

On the one hand, we may consider *(I1)* that any subset of precisifications could be a standpoint[1]. However, we may not want this to be the case, and indeed across this thesis it is presumed that, instead, *(I2)* only labelled subsets of precisifications are standpoints. While with interpretation *(I1)* we can draw the parallel exclusively between Fine's specification space and our standpoints, with interpretation *(I2)* we will consider both precisifications and standpoints.

For the case of **FS**, the condition holds unproblematically for standpoints given that they are monotonic (see chapter 4, section 4.2.2.5). Monotonicity in the sharpening relation between standpoints implies that one standpoint, s, is sharper than another, s', only if the former improves precision on the latter. That is, any proposition that is unequivocally true or false in s' has the same truth value in s. Additionally, s may precisify other propositions.

FF also trivially holds if we take interpretation *(I1)* and define *complete points* as standpoints containing a single precisification (we rename them as *complete standpoints*). Given that precisifications in our model are classical and complete, then the

[1]This is analogous to Fine's own specification if we consider partial points instead of standpoints.

truth-values of any proposition at the complete standpoint are also classical. In particular, for a *complete standpoint*, s_{cs}, and any proposition ϕ it is the case that either $\Box_{cs}\,\phi$ or $\Box_{cs}\,\neg\phi$. If we take interpretation *(I2)*, the condition holds trivially if we take complete points to be precisifications.

FC holds, for *(I1)*, in virtue of the partial order of standpoints under the subset relation. In particular, if the standpoints are modelled by all the possible subsets of precisifications of the language, then for any standpoint s there is another standpoint s' such that s' is complete (is modelled by a set consisting of one precisification) and $s' \preceq s$. Conversely, for *(I2)*, the condition should be rewritten so that *completability states that any standpoint can be sharpened to a precisification of the language.* In this way the intuition is preserved and the condition holds trivially, as any standpoint is non-empty.

Finally, the *Resolution Condition* **FR** holds in our system for *(I1)* given that for any standpoint s and proposition ϕ indefinite with respect to s, there are two sharper standpoints $s_p, s_n \preceq s$ such that they are identical to s except from the assertion of ϕ in s_p and its negation in s_n, thereby being a sharpenings of s and guaranteeing that ϕ is definite in both s_p and s_n. For the interpretation *(I2)*, again assuming that there is an atomic proposition ϕ that is indefinite with respect to a standpoint s, then by definition, both $\Diamond_s\,\phi$, i.e. there is necessarily at least a precisification $\pi \in s$ that validates it, and $\Diamond_s\,\neg\phi$, i.e. there is necessarily at least a precisification $\pi' \in s$ that validates its negation.

3.5.3 Modelling standpoints with modal operators

While precisifications are modelled as the points of our *standpoint logic*, standpoints are represented via the modal relations. Standpoint modal operators are denoted with the standard box and diamond. Each standpoint s is associated with the corresponding modal operators $\Box_s$ and $\Diamond_s$, and there is a universal standpoint $*$, corresponding to $\Box_*$ and $\Diamond_*$, which is the standpoint mapping to the set of all admissible precisifications. Differently from other standpoints, something that is necessarily true according to $*$ is also *super-true* (as in true in all precisifications. See section 3.6). For this reason, stronger axioms apply to $*$ (chapter 4, section 4.2.3).

$\Box_s\,\phi$ reads as 'ϕ is true in standpoint s', i.e. ϕ is true in all admissible precisifications for standpoint s. $\Box_*\,\phi$ means that ϕ is true in all precisifications. Consequently, the semantic heterogeneity notions **N1-9** in section 1.2.2 are formalised as follows:

NF1. $\Box_* \phi$ ———— "It is unequivocal that ϕ."

NF2. $\Diamond_* \phi$ ———— "In some sense ϕ."

NF3. $\Box_* \phi \lor \Box_* \neg\phi$ —"It is definite whether ϕ (or not ϕ)."

NF4. $\Diamond_* \phi \land \Diamond_* \neg\phi$ —"It is borderline that ϕ."

NF5. $\Box_s \phi$ ———— "According to s, [it is unequivocal that] ϕ."

NF6. $\Diamond_s \phi$ ———— "According to s, in some sense ϕ."

NF7. $\Box_s \phi \lor \Box_s \neg\phi$ — "According to s, it is definite whether ϕ (or not ϕ)."

NF8. $\Diamond_s \phi \land \Diamond_s \neg\phi$ — "According to s, it is borderline that ϕ."

NF9. $s' \preceq s$ ———— "s is sharper than s'."

In the semantics, the standpoints are interpreted as sets of admissible precisifications and are partially ordered under the subset relation. We say that a standpoint s is sharper than a standpoint s', $s \preceq s'$, if the set of precisifications compatible with s is a subset of the set of precisifications compatible with s'. Any standpoint s_i is sharper or equal than standpoint $*$, i.e. for all $s \in \mathcal{S}$ and for $* \in \mathcal{S}$ it is the case that $s \preceq *$. This ordering of interpretations was first proposed in [Sha06] and has been also referred as semantic precision in other work like [LD12].

Moreover, the embedding of standpoint operators is trivial. Although one could see standpoints to belong to agents (e.g. someone's standpoint), they are fundamentally relative to language. One standpoint can be shared among a community of agents and one same agent can adopt different standpoints in different circumstances. However, a standpoint s about another standpoint s' is no different than the latter itself (s'). This is discussed in more detail in chapter 4, section 4.2.3.4.

3.5.4 Other modal operators in the supervaluationist literature

We identify two main operators being discussed in the literature about modal approaches to supervaluation semantics, namely the definitely operator $\mathcal{D}$ and the supertruth operator $\mathcal{T}$ or $\mathcal{S}$, which we briefly review in relation to the standpoint operators.

3.5.4.1 $\mathcal{D}$ Definitely operator

The definitely operator $\mathcal{D}$ was first introduced by Fine in [Fin75], fundamentally to address the issue of higher-order vagueness, and became the source of a variety of arguments against supervaluationism. In particular, it has been argued that one of the principal benefits of supervaluationism, retaining classical consequence, was lost in the presence of the operator (see [Wil94, Far10] and [Var07, Cob11] for responses, as well as section 3.6 for the notion of consequence in *standpoint semantics*).

However, given the need for a mechanism to represent the definiteness or clarity of a predication in a vague language, most modal approaches to supervaluationism discuss some variation of the $\mathcal{D}$ operator (e.g. [Fin75, Var07, Cob11, LYV16, ADP09]). In [Var07], Varzi argues that there are two main readings of such an operator:

1. '$\mathcal{D}\phi$ is true on a precisification π if and only if ϕ is true on all precisifications accessible from π '.

2. '$\mathcal{D}\phi$ is true on a precisification π if and only if ϕ is true on precisification π.'

On the first and most common reading, $\mathcal{D}$ is said to be a necessity operator with a logic at least **KT** (see section 2.5.2.1 for the introduction of the main axioms and systems of modal logic). If, on the other hand, it is read according to the second option (2), then it could also be considered to be analogous to the actuality operator, hence at least **S5**. In our framework, according to the description (1), $\mathcal{D}$ would be analogous to a standpoint operator $\Box_s$ except in that $\Box_s$ does not satisfy axiom **T** (with the exemption of $\Box_*$), in contrast to part of the literature, where it is presumed that reflexivity is an obvious characteristic of $\mathcal{D}$ [Var07, Cob11].

The presumption of reflexivity is justified by the parallel to the necessity operator and the intuition that if "it is *definitely* the case that ϕ" then "it is the case that ϕ" ($\mathcal{D}\phi \rightarrow \phi$). This certainly holds in our system whenever "it is *definitely* the case (with respect to all standpoints)", i.e. when $\mathcal{D}$ is parallel only to the universal standpoint operator $\Box_*$ (which indeed satisfies **S5**). On the contrary, if $\mathcal{D}$ is set to be parallel to any standpoint's operator $\Box_s$, then it is not reflexive, as the fact that "it is *definitely* the case according to s that ϕ" does not guarantee that ϕ is the case in the 'actual' point (see chapter 4 section 4.2.2.5 for further analysis of the properties of $\Box_s$).

Last but not least, we do include in section 4.2.1 the defined operators $\mathcal{D}_s$ and its converse $\mathcal{I}_s$, which are simply shortcuts such that $\mathcal{D}_s\,\phi \;\equiv_{def}\; (\Box_s\,\phi \vee \Box_s\,\neg\phi)$ and

$\mathcal{I}_s\,\phi \;\equiv_{def}\; (\Diamond_s\,\phi \wedge \Diamond_s\,\neg\phi)$ and thus they are not to be confused with other accounts of $\mathcal{D}_s$ in the supervaluationistic literature such as those previously described in this section.

3.5.4.2 T True or $\mathbb{S}$ Super-true and Δ Locally-true operators

A truth operator T is discussed in [Var07, ADP09] (named $\mathbb{S}$ in the latter), with the reading "$T\phi$ is true on a precisification π if and only if ϕ is true on all precisifications" and satisfying **S5**. This operator is introduced to provide an account of global validity built upon the standard modal local validity, which will be discussed in the coming subsection. Given that in our framework all precisifications in the model are admissible and are accessible under $\Box_*$, this additional T operator would conflate with $\Box_*$.

Finally, the locally-true (Δ) operator is proposed in [Var07] with the reading "$\Delta\phi$ is true on a precisification π if and only if ϕ is true on precisification π " for the analysis and discussion of Williamson's objection to supervaluationism on the grounds of the failure of disquotationality[1] (see Williamson [Wil94] for the objection and Varzi [Var07] for the response). Although such an operator is definable in our logic we do not specify it because it is not deemed useful in applied scenarios.

3.6 Truth and Validity

Validity is generally defined in different ways that are equivalent in the case of a classical bivalent logic. An argument is said to be valid iff it is truth-preserving: 'whenever all the premises are true, one of the conclusions must be true'. Alternatively one may say that an argument is valid if and only if it is not possible for all the conclusions to be false when the premises are all true [Var07]. Moreover, it is usual to accept, in supervaluation semantics, that 'truth is super-truth'.

Discussions on appropriate notions of validity for supervaluationistic frameworks are abundant in the literature, examples of which are [Var07, Cob11, Cob08, Kee00] In the absence of a definitely (or any modal) operator and multiple conclusions, all the proposed notions of supervaluationist consequence coincide with the classical counterpart [Cob11]. This is not the case, however, in the aforementioned cases.

Typically, supervaluationistic frameworks rely heavily on the notion of *super-truth*,

[1] $\psi \leftrightarrow \mathsf{T}\psi$, which reads '$\psi$ is the case iff it is true that ψ is the case'

whereby a proposition ϕ is true iff it is true in all possible ways of precisifying ϕ. As commonly expressed, 'truth is super-truth', and any proposition that fails to be either *super-true* or *super-false* has no semantic value as such [Var07]. Along these lines, the so-called *global validity*, which has been put forward as the most natural notion in a supervaluationistic framework, evaluates the truth value of the premises and conclusions of an argument globally, and hence these can be *super-true*, *indeternimate* or *super-false*. However, global validity gives rise to counterexamples [Var07, Cob11, Cob08] that have been put forward as objections to the supervaluationistic framework overall [Wil94], as well as motivated reasonable interest on the right notion of validity in supervaluation semantics. Moreover, global validity fails to accommodate penumbral connections, which are arguably an important advantage of supervaluationism over other theories of vagueness [Var07].

This issue is discussed in depth in [Var07], where a variety of options are presented (see box 3.1). One of the main alternatives to the notion of *global validity* is that of *local validity*, which is the standard in modal logics [BdRV01] and hence particularly appropriate for modal approaches to supervaluation semantics. In a local validity approach, an argument is valid iff the argument is valid in all the precisifications of the language. Given that each precisification is classical in nature, there are no scenarios where the value *indeternimate* occurs, and hence classical consequence rules apply and behave well.

In standpoint semantics, along the lines of [ADP09], we consider truth to be a modal notion representing a way of being true rather than a fundamental notion of truth. Consequently, the fundamental role is played by the notion of truth at a precisification and we incorporate modal operators such that we can express not only truth relative to the whole language (which is akin to *super-truth*) but also truth relative to particular standpoints. This gives us a variety of applications that have not yet been explored, particularly for the process of concept negotiation and multi-agent systems, regarding the reasoning that can be done between standpoints.

In accordance to that, it is natural to adopt a local notion of validity, which is well studied in the domain of modal logics [BdRV01] and avoids the failure of well known rules of inference (contraposition, indirect proof, proof by cases and conditional proof), the failure of entailments with multiple arguments and the failure of entailments involving special operators such as $\mathcal{D}$ or $\Box_s$ in our logic.

Options of validity notions, extracted from [Var07]

- $\Sigma \models_{GA} \Gamma =_{df}$ If every $\phi \in \Sigma$ is: T on all $\pi \in \Pi$, then some $\psi \in \Gamma$ is: T on all $\pi \in \Pi$
- $\Sigma \models_{GB} \Gamma =_{df}$ If every $\psi \in \Gamma$ is: F on all $\pi \in \Pi$, then some $\phi \in \Sigma$ is: F on all $\pi \in \Pi$
- $\Sigma \models_{GC} \Gamma =_{df}$ If every $\phi \in \Sigma$ is: T on all $\pi \in \Pi$, then some $\psi \in \Gamma$ is not: F on all $\pi \in \Pi$
- $\Sigma \models_{GD} \Gamma =_{df}$ If every $\psi \in \Gamma$ is not : T on all $\pi \in \Pi$, then some $\phi \in \Sigma$ is: F on all $\pi \in \Pi$
- $\Sigma \models_{local} \Gamma =_{df}$ On all $\pi \in \Pi$: if every $\phi \in \Sigma$ is: T, then some $\psi \in \Gamma$ is: T.
- $\Sigma \models_{X} \Gamma =_{df}$ If on all $\pi \in \Pi$ every $\phi \in \Sigma$ is: T, then on all $\pi \in \Pi$ some $\psi \in \Gamma$ is: T.
- $\Sigma \models_{Y} \Gamma =_{df}$ If on all $\pi \in \Pi$ every $\psi \in \Gamma$ is: F, then on all $\pi \in \Pi$ some $\phi \in \Sigma$ is: F.

Table 3.1: Candidate notions of validity for a supervaluationistic framework, extracted from [Var07]. Local validity is the standard modal interpretation of ‘truth at a point’.

The main argument provided against *local validity* in a supervaluationistic framework is that, since 'for the supervaluationist that a sentence is true means that it is true in every precisification (that is, 'truth is supertruth')[...], local consequence does not preserve the supervaluationist-relevant notion of truth' [Cob11]. However, as put in [Var07], 'just as questions of truth may only be answered upon considering the precisifications of the language, so questions of validity may be answered only upon considering those precisifications. Just as a statement is rated true, supervaluationally, if and only if it is true on all admissible precisifications, so an argument may be rated valid if and only if, necessarily, its premises and conclusions stand in the appropriate relation on all admissible precisifications'. Given our modal approach and the nature of our intended applications, focused on supporting reasoning within multiple interpretations of the semantics of terms, it seems to us that the more reasonable option is to adopt *local validity*. Moreover, as discussed in [Cob08, Var07], reasoning with other notions of validity can be performed using a local notion as a base, and hence, wherever the intended use is gobal (or another) that is still supported in our logic by using the modal operators $\square_*$ (see the refered papers for more detail).

3.6.1 Super-truth or relative truth

Despite the fact that we focus on the modal notion of 'truth at a point' rather than on that of 'truth as super-truth', we do normally include the $\square_*$ operator, which is a special standpoint containing all the admissible precisifications of the language. This operator satisfies **S5**, hence being stronger than the rest of the standpoint operators $\square_s$.

There are, however, certain scenarios in which we may want to remove or dismiss $\square_*$, which we can do without prejudice to the rest of our logic. Such scenarios involve situations in which our system is used to model a variety of standpoints, say those relevant to a group of stakeholders for the domain of forestry, but were we want to acknowledge that there may be other admissible standpoints which are not currently formalised and could disagree with the axioms encoded by $\square_*$. In such cases, propositions that are true according to all the standpoints in the system are presumed to be super-true in the presence of $\square_*$, which can allow for further inferences and may not be intended by the users of the system. Hence, in these scenarios we can completely suppress the super-truth operator $\square_*$ and reason within *relative truth* (truth can only be relative to a standpoint).

3.7 Higher-order vagueness

The phenomenon of higher-order vagueness rests on the intuition that, as much as it is difficult to draw a sharp line between, say, a ‘heap’ and a ‘non heap’ (thereby there being borderline cases of heap), it is also difficult to draw a sharp line between a ‘heap’ and a ‘borderline-heap’. In other words, in a sorites sequence of heaps, each with one grain less than the previous one and ending in a single grain, it is not trivial to determine which is the last clear case of heap, followed by the first borderline case. This creates second-order vagueness and, by iteration, infinite orders of vagueness and, hence, unsharpeneability.

The issue of higher-order vagueness is prominent in the philosophical literature, where it is often considered that a logic of vagueness that does not provide an account of the phenomenon is deficient [Wil94, KS96a]. Consequently, important objections to the different theories of vagueness are made on these grounds [Wil94, Far10] In the case of supervaluation semantics, a notorious part of the literature ([Var07, Cob16, Fin75, Cob08, Var04] among many others) considers or deals with the phenomenon, presents counter-arguments to the objections and provides formalisations. In contrast, several philosophers reduce the phenomenon to an illusion [Wri10] which is deemed to be incoherent [Cri17] and which does not manifest in language use (‘speakers do not go around talking about borderline borderline cases and borderline borderline borderline cases and so forth’ [Raf05])

Given that our framework, standpoint semantics, is mostly intended as a formalism for knowledge representation and, on the grounds of the absence (or at least scarcity) of manifestations of higher-order vagueness in language use, we do not account for the phenomenon. In this way, a statement such as $\Box_* \Box_* \phi$ is semantically equivalent to $\Box_* \phi$. In fact, we have particularly strong axioms governing the iteration of standpoint operators, as can be seen in chapter 4, section 4.2.3.4. The intuition is that $\Box_* \phi$ holds iff ϕ holds in all the precisifications. Because precisifications are classical, then ϕ or $\neg\phi$ must hold in all $\pi \in \Pi$. Therefore, $\Box_* \phi$ can not be a vague expression itself and hence $\Box_* \Box_* \phi$ is redundant. In our framework, whenever an agent considers that the valuation of ϕ is not entirely definite (to any order), then their standpoint must be formalised such that $\mathcal{I}_s \phi$, hence the model accounts for at least a precisification in which ϕ does not hold.

On the other hand, *standpoint logic* does not force one to committ, for ϕ and a particular standpoint (or the whole language), to either $\Box_s \phi$, $\mathcal{I}_s \phi$ or $\Box_s \neg\phi$, nor to their universal counterparts. In fact, we may know that that $\Diamond_s \phi$ but we may not be able to infer neither $\Box_s \phi$ or $\mathcal{I}_s \phi$ because the standpoint s has not committed to either. This feature corresponds to the intuition that standpoints do not necessarily fix the truth conditions and gaps of the whole language and, moreover, they do not tend to fix sharp boundaries between unequivocal truth and borderline cases.

Finally, beyond the philosophical discussion and the adoption or rejection of theoretical positions disregarding higher-order vagueness as a meaningful phenomenon, we believe that its exclusion from our framework is unlikely to be a hindrance for its intended applications. On the contrary, the followed approach procures simpler semantic models and stronger axioms (**AS4** and **AS5**), hence resulting in a framework with better computational properties.

Chapter 4

A Formal Language for Standpoint Logic

4.1 Introduction to a logic of standpoints

In the previous chapter we have introduced the key elements of our framework, namely precisifications and standpoints, and we have justified our interpretation of truth and validity. In this chapter we proceed to characterise the framework formally, revisiting the notions discussed in its precedent and focusing on providing their formal definitions and discussing their logical properties.

We present *standpoint logic*, a logic with which we can assert that a proposition holds according to a particular standpoint, and that can express various relationships between standpoints. As it is conceived here, it is a multi-modal logic with Kripke semantics, in which modalities are used to model different interpretations of the semantics of the vague terms of the language.

From the technical point of view, the proposed multi-modal logic has a set of operators $\square_s$, one for each standpoint s, each of which satisfies the axioms of **KD45**. There is also a 'universal standpoint' operator $\square_*$, which is a stronger modality (whose accessibility relation is reflexive) that satisfies the axioms of **S5**.

Moreover, there are two main particularities of the standpoint logic with respect to other well known multi-agent systems of modal logic. On the one hand, the standpoint modal operators are partially ordered under the subset relation [AH10]. On the other hand, the relations of the Kripke frames are unary rather than binary, i.e. they are trivially binary (see section 4.2.2.4). This property is known to occur in **KD45** systems with a single modality [Pie09], but it does not normally extend to systems with multiple modalities. However, the stronger interaction axioms of *standpoint logic*, namely **AS4** and **AS5** (see section 4.2.3.4), extend this property to our multi-modal framework, consequently improving its computational properties as shown in subsection 4.4.2.

In this chapter we first describe in detail a *propositional standpoint logic* in section 4.2 by specifying its syntax (subsection 4.2.1), semantics (subsection 4.2.2) and proof theory (subsection 4.2.3). We then prove soundness and completeness in section 4.3, decidability and NP-complexity in section 4.4 and we conclude by introducing more expressive logics in section 4.5, namely a description standpoint logic (subsection 4.5.1), a predicate standpoint logic (subsection 4.5.2), and finally a two-dimensional standpoint logic (subsection 4.5.3).

4.2 Propositional Standpoint Logic, $\mathbb{S}_0$

In the current section we introduce a propositional logic for reasoning about standpoints, the system of modal logic $\mathbb{S}_0$.

4.2.1 The formal language $\mathcal{L}_{\mathbb{S}_0}$

Our *propositional standpoint logic* $\mathbb{S}_0$ consists of a non-empty finite set of propositional variables $\mathcal{P} = \{P_1, \ldots, P_n\}$, as well as connectives $\neg\phi$, $(\phi \wedge \psi)$ and a partially ordered set of standpoint operators $\square_{s_1}, \ldots, \square_{s_m}, \square_*$.

Definition 1. The language $\mathcal{L}_{\mathbb{S}_0}$ is the smallest set of formulas, such that $\mathcal{P} \subseteq \mathcal{L}_{\mathbb{S}_0}$ and, for all formulas $\alpha, \beta \in \mathcal{L}_{\mathbb{S}_0}$ and all standpoint operators $\square_{s_1}, \ldots, \square_{s_m}, \square_*$, we have:

$$\neg\alpha,\ (\alpha \wedge \beta),\ \square_s \alpha \quad \in\ \mathcal{L}_{\mathbb{S}_0}$$

The language also includes formulas constructed with definable operators, which have their usual definitions:

$$(\alpha \vee \beta),\ (\alpha \rightarrow \beta),\ (\alpha \leftrightarrow \beta),\ \lozenge_s \alpha \quad \in\ \mathcal{L}_{\mathbb{S}_0}$$

Beyond the standard definable operators ($\vee$, $\rightarrow$, $\leftrightarrow$, $\lozenge_s$), we can define other useful operators in $\mathcal{L}_{\mathbb{S}_0}$, namely:

- $\mathcal{I}_s\, \phi \ \equiv_{def}\ (\lozenge_s \phi \wedge \lozenge_s \neg\phi)$ — means that the truth of ϕ is indeterminate (i.e. borderline) according to standpoint s.
- $\mathcal{D}_s\, \phi \ \equiv_{def}\ (\square_s \phi \vee \square_s \neg\phi)$ — means ϕ has a determinate truth value according to standpoint s.

One can easily see that $\mathcal{I}_s \equiv \neg\mathcal{D}_s$ given that

$$\neg(\lozenge_s \phi \wedge \lozenge_s \neg\phi) \equiv (\neg \lozenge_s \phi \vee \neg \lozenge_s \neg\phi) \equiv (\square_s \neg\phi \vee \square_s \phi)$$

.

In the case where these operators are employed with respect to the universal standpoint, $*$, we have $\mathcal{I}_* \phi$, which can be read as 'ϕ is a borderline proposition' and $\mathcal{D}_* \phi$, which can be read as 'ϕ is a sharp proposition'[1]. Recall, from section 3.6.1, that $*$ is the set of all admissible precisifications, and hence *truth* with respect to $*$ can be regarded as *super-truth*.

Any further occurrences of the $\mathcal{D}_s$ operator in this work correspond to the definitions just given, unless otherwise indicated. (In chapter 3 section 3.5.4 we discussed some other interpretations of a determinacy operator, $\mathcal{D}$, which have been proposed in the literature on supervaluationism and vagueness.)

4.2.2 Semantics

In this section we present the formal semantics for standpoint logic. We begin by introducing Kripke semantics in 4.2.2.1, to then define the class of models $\mathfrak{M}_{\mathbb{S}_0}$ of standpoint logic in 4.2.2.2 and provide formal definitions for validity as in standard modal logics in 4.2.2.3. The two final subsections are devoted to the analysis of the particularities of the class $\mathfrak{M}_{\mathbb{S}_0}$ of Kripke models, namely the unary relations and their binary counterparts (4.2.2.4) and the properties of such counterparts (4.2.2.5).

4.2.2.1 Introduction to Kripke Semantics

Kripke models[2] (or variants thereof) are widely used to characterise the semantics of contemporary modal logics. They are, for example, the basis for many modern frameworks for reasoning about knowledge and belief [FHMV95], so that an agent's set of beliefs is characterised by a set of possible worlds that are interpreted as the states of affairs that are possible given what the agent believes. It is easy to see that a parallel can be drawn for a supervaluationist framework if we simply substitute the notion of possible worlds by possible precisifications, so that an agent's standpoint is the set of possible precisifications of the language that are compatible with her statements.

[1] Note that in other logics, for instance Ground S5, a formula like $(\Box_s \phi \vee \Box_s \neg\phi)$ is considered to be dishonest, because an agent 'cannot know one of without knowing one of them'. However, in our approach, standpoints are only structures that narrow the semantic variability of terms, and in that sense, the $\mathcal{D}_s$ expresses the non trivial commitment that the valuation of ϕ is classical.

[2] Kripke models are also referred as standard models [Che80]

A formal definition of a Kripke model reads as follows:

Definition 2. A Kripke model $\mathcal{M}$ is a triple $\langle W, R, V \rangle$, where W is a set of points or possible worlds, R is a binary relation on $W \times W$ known as the accessibility relation and V is the evaluation function $V : \mathcal{P} \to 2^W$, mapping each propositional constant $p \in \mathcal{P}$ to the set $V(p) \subseteq W$ of worlds at which that proposition is true. Formulas are interpreted as follows:

- $(\mathcal{M}, w) \models p$ if and only if $w \in V(p)$,
- $(\mathcal{M}, w) \models \neg\phi$ if and only if $(\mathcal{M}, w) \not\models \phi$,
- $(\mathcal{M}, w) \models \phi \to \psi$ if and only if $(\mathcal{M}, w) \not\models \phi$ or $(\mathcal{M}, w) \models \psi$,
- $(\mathcal{M}, w) \models \Box\phi$ if and only if $(\mathcal{M}, u) \models \phi$ for all u such that Rwu.

This kind of structure can be easily generalised for logics with multiple modalities, such that if the language contains a set of necessity operators $\{\Box_i \mid i \in I\}$ then the model has a set of accessibility relations R_i for each $i \in I$.

Definition 3. A Kripke model $\mathcal{M}$ for a multi-modal logic is a tuple $\langle W, R_1, \ldots, R_n, \models \rangle$, where W is a set of points or possible worlds, each R_i is a binary relation on W and V is the evaluation function $V : \mathcal{P} \to 2^W$, mapping each propositional constant $p \in \mathcal{P}$ to the set $V(p) \subseteq W$ of worlds at which that letter is true. Formulas are interpreted as in **Definition 2** except for the modal formulae, which are interpreted by:

- $(\mathcal{M}, w) \models \Box_i A$ if and only if $(\mathcal{M}, u) \models A$ for all u such that $R_i wu$.

Kripke models with no restrictions on the relations determine the modal logic K, which can be axiomatically specified as smallest modal logic containing axiom **AK** [Che80].

4.2.2.2 The class $\mathfrak{M}_{\mathbb{S}_0}$ of Kripke models

In order to characterise the semantics of *standpoint logic* we narrow our attention to a particular class of Kripke models for multi-modal languages, which we name $\mathfrak{M}_{\mathbb{S}_0}$.

Definition 4. A Kripke model of the class $\mathfrak{M}_{\mathbb{S}_0}$ is a tuple $\langle \Pi, (R, \preceq), \delta \rangle$ where:

- $\Pi = \{\ldots, \pi_i, \ldots\}$ is a set of points or precisifications.
- $(R, \preceq)$ is a partial order of $R_{s_1}, ..., R_{s_n}, R_* \in R$ under the sharper/subset relation $\preceq$, where each R_{s_i} is a unary accessibility relation on Π, i.e. a set of precisifications, and $s_1 \preceq s_2$ iff $s_1 \subseteq s_2$. In particular, $R_* = \Pi$.
- $\delta : \mathcal{P} \to 2^W$ is the evaluation function mapping each propositional constant $p \in \mathcal{P}$ to the set $\delta(p) \subseteq \Pi$ of worlds at which that letter is true.

Π is a set of points that we call precisifications and are analogous to possible worlds. $(R, \preceq)$ is a partial order of accessibility relations $R_{s_1}, ..., R_{s_n}, R_* \in R$, which are ordered under the subset relation ($\preceq \equiv \subseteq$)[1]. The relations in R are indexed by the variables s and s' or $s_1, \ldots s_n, *$, as they model the notion of standpoint (see chapter 3, section 3.5). Hence, $R_s \preceq R_{s'}$ is analogous to $s \preceq s'$.

R_* is the universal standpoint, such that $R_{s_i} \preceq R_*$ for all $R_{s_i} \in R$. Each R_{s_i} (including the special case R_*) is a unary relation on Π, hence a set of precisifications, and it is non empty. The relation between the unary R_{s_i} and its binary counterpart (standard in Kripke semantics) R'_{s_i} is analysed in section 4.2.2.4. Finally, δ is the interpretation function mapping each propositional constant p to the set $\delta(p) \subseteq \Pi$ of worlds at which that proposition is true.

4.2.2.3 Validity

The notion of validity in Standpoint Logic is discussed in section 3.6 and is the standard in modal logic. In this section we provide some formal definitions for subsequent use. For further reading on validity in modal logics see [Che80].

[1]The symbol $\preceq$ is used to denote the relation *sharper than* which holds between standpoints. In the semantics of the model, however, this is equivalent to the subset relation because standpoints are modelled as sets.

Definition 5. We write $(\mathcal{M}, \pi) \models \phi$ to mean that ϕ is true at precisification π in model $\mathcal{M}$, and it is defined recursively as follows,

- $(\mathcal{M}, \pi) \models p$ if and only if $\pi \in \delta(p)$,
- $(\mathcal{M}, \pi) \models \neg\phi$ if and only if $(\mathcal{M}, \pi) \not\models \phi$,
- $(\mathcal{M}, \pi) \models \phi \to \psi$ if and only if $(\mathcal{M}, \pi) \not\models \phi$ or $(\mathcal{M}, \pi) \models \psi$,
- $(\mathcal{M}, \pi) \models \Box_{s_i} \phi$ if and only if $(\mathcal{M}, \pi') \models \phi$ for all π' such that $\pi' \in R_{s_i}$.

Definition 6. A formula ϕ is valid in $\mathcal{M}$, written $\mathcal{M} \models \phi$, if it is true at every precisification in $\mathcal{M}$,

$$\mathcal{M} \models \phi \quad \text{iff} \quad \mathcal{M}, \pi \models \phi \;\; \text{for all} \;\; \pi \in \Pi \; .$$

Let $\mathfrak{M}_{\mathbb{S}_0}$ be the class of models $\mathbb{S}_0$,

Definition 7. We write $\models_{\mathbb{S}_0} \phi$ to mean that ϕ is valid in the class $\mathfrak{M}_{\mathbb{S}_0}$ of models and, hence, ϕ is a theorem of propositional standpoint logic. This is defined by:

$$\models_{\mathbb{S}_0} \phi \quad \text{iff} \quad \mathcal{M} \models \phi \;\; \text{for all} \;\; \mathcal{M} \in \mathfrak{M}_{\mathbb{S}_0}$$

4.2.2.4 The accessibility relation R and its counterpart R'

The accessibility relation R_{s_i} represents the set of precisifications that are admissible or accessible relative to some standpoint s_i. In particular, if $\pi \in R_{s_i}$, then we can think of π as being a possible precisification that is compatible with all that the standpoint s_i makes precise.

Following the former, R_{s_i} is a unary relation that holds (or not) for any precisification π in Π. But, although it is natural to specify our semantics by associating standpoints with sets of precisifications, it is sometimes convenient to recast this structure in terms of a standard (i.e. binary) Kripke relation over the precisifications. For each s_i, we define a relation R'_{s_i} over Π:

$$R'_{s_i} \equiv \Pi \times R_{s_i}$$

The standard accessibility relation R'_{s_i} determines, for a precisification π, the set of precisifications that are admissible (accessible) under a standpoint s_i. We can use this equivalent specification to relate properties of our models to properties that are well known with regard to Kripke semantics (e.g symmetry, transitivity, ...), as well as to proofs for the systems **S5** and **KD45** which our logic extends. However, in most of the presentation we will use R_{s_i} for simplicity and compactness.

Note that R being a unary relation reflects the fact that the framework assumes that any speaker of the heterogeneous language is able to understand what are the admissible interpretations of terms that are left open by any standpoint of the system, under the assumption that understanding how words can be used is a crucial part necessary for 'speaking' a language.

In practice, this means that assertions such as *'ϕ is the case according to standpoint s'* ($\Box_s \phi$) have the same truth conditions regardless of the current state or precisification (π). While this approach may seem unrealistic at first glance, it is due to the fact that standpoint operators do not model the epistemic state of agents. In fact, there is not a relation one to one from agents to precisifications or standpoints. This is because, instead of using fixed precisifications or standpoints themselves, the agents may change standpoints depending on the situation. Moreover, it is considered that a fundamental part of a speaker's knowledge of the language is to understand how words can be used differently in different settings, which justifies being able to access the precisifications of a standpoint from any state or precisification.

For instance, let us consider that *'according to the standpoint s Nena is tall'*, formalised as $\Box_s[\mathsf{NTall}]$. Then, regardless of her own standpoint, say s', and of the pointed precisification, say π, the agent understands that *'according to the standpoint s Nena is tall'*. Hence, any language user of a heterogeneous language can understand that another agent is using an interpretation of a term (for example a very slack use of *tall*) when that is made explicit, even though they would not consider using such standpoint in such situation themselves.

This feature of the precisification space implies that the models in $\mathfrak{M}_{\mathbb{S}_0}$ are a simplified subset of standard Kripke models, which in turn provides the logic with good computational properties as shown in the complexity proof in section 4.4.2. Moreover, the possibility to simplify normal modal logics satisfying **K45** has been already shown in [PKP19] for systems with a single modality, where a similar treatment is given.

4.2.2.5 Properties of the Accessibility Relations

Standpoint accessibility relations R_{s_i}, or the binary counterpart R'_{s_i}, represent the set of precisifications that are admissible relative to some standpoint s_i, i.e. the set of precisifications that are accessible through some standpoint s_i from each precisification π. As such, every R'_{s_i} contains a tuple $\langle \pi, \pi' \rangle$ for all $\pi \in \Pi$ to all $\pi' \in s_i$. In particular, if π' is such that $\langle \pi, \pi' \rangle \in R'_{s_i}$, then we can think of π' as being a possible precisification from the standpoint s_i.

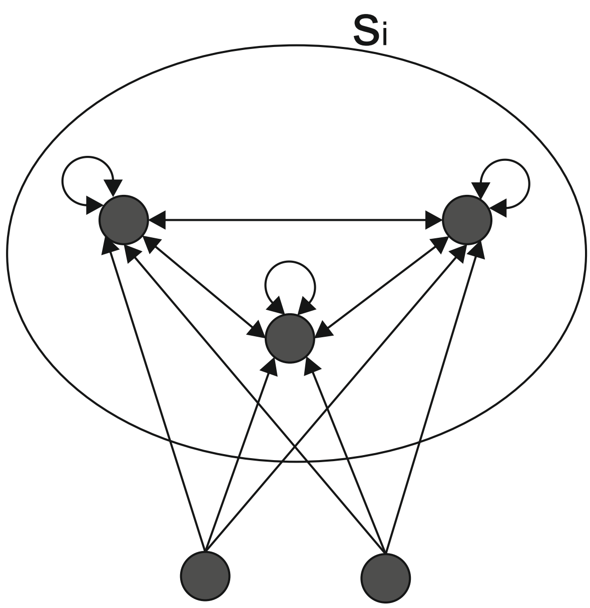

Figure 4.1: Accessibility relations for R_{s_1}.

Figure 4.1 shows the accessibility relations for standpoint s_1, which contains three of the five precisifications in the model. The grey dots are precisifications, the circle denotes the set of precisifications s_i and the arrows are the relations $\langle \pi, \pi' \rangle \in R'_{s_i}$, connecting all $\pi \in \Pi$ to all $\pi' \in R_{s_i}$.

In figure 4.1, R'_{s_i} consists of the tuples:

$$[\langle\pi_1,\pi_1\rangle,\langle\pi_1,\pi_2\rangle,\langle\pi_1,\pi_3\rangle,\langle\pi_2,\pi_1\rangle,\langle\pi_2,\pi_2\rangle,\langle\pi_2,\pi_3\rangle,\langle\pi_3,\pi_1\rangle,\langle\pi_3,\pi_2\rangle,\langle\pi_3,\pi_3\rangle,\\ \langle\pi_4,\pi_1\rangle,\langle\pi_4,\pi_2\rangle,\langle\pi_4,\pi_3\rangle,\langle\pi_5,\pi_1\rangle,\langle\pi_5,\pi_2\rangle,\langle\pi_5,\pi_3\rangle] \quad (4.1)$$

Every standpoint accessibility relation R'_{s_i} has the following properties:

1. *Serial*: For each $x \in \Pi$ there exists some $y \in \Pi$ such that $R'_s xy$; or alternatively, there exists some $y \in \Pi$ such that $R_s y$. This holds as a result of standpoints being non empty sets.

2. *Transitive*: $R'_s xy \wedge R'_s yz \rightarrow R'_s xz$. Moreover, since $R'_s xy$ and $R'_s yz$ we know that $R_s y$ and $R_s z$ from any $\pi \in \Pi$. Consequently, transitivity is a consequence of all $\pi \in s$ being accessible between them.

3. *Euclidean*: $R'_s xy \wedge R'_s xz$ implies that $R'_s yz$. This occurs again on the grounds that all the precisificatons in a standpoint s are related among them under R'_s. As a result, if $R'_s xy \wedge R'_s xz$ then $y, z \in s$ and consequently $R'_s yz$.

4. *Dense*: $R'_s xz \rightarrow (\exists y \in \Pi \mid [R'_s xy \wedge R'_s yz])$. $R'_s xz$ implies that there is also $R'_s zz$, hence the axiom holds easily for $y = z$.

In addition to the former, if we call R' to the assignment of a binary relation R'_s to each standpoint s, then the assignment is monotonic since if $s \preceq s'$ then we have $R_s \subseteq R_{s'}$ and hence $\Pi \times R_s \subseteq \Pi \times R_{s'}$ so that $R'_s \subseteq R'_{s'}$. Monotonicity ensures that, whenever a precisification is accessible from π under a standpoint s then it is also accessible from π under any other standpoint s' such that $s \preceq s'$. I.e. it is also accessible from any standpoint that is less sharp than s.

The standpoint with label $*$ is universal, hence for all $x, y \in \Pi$, $R'_* xy$. This is the case given that the standpoint $*$ encapsulates the notion of *supertruth* and consequently something that is necessarily true in $*$ must be true in all the precesifications in the language. Given that standpoint $*$ is the set of all admissible precisifications, it points to the whole set of precisifications Π.

Consequently, the relation R'_* has the following additional properties:

1. *Reflexive* $R'_* xx$.
2. *Symmetric* $R'_* xy$ implies $R'_* yx$

Finally, we discuss further restrictions or properties that hold in the models in $\mathfrak{M}_{\mathbb{S}_0}$, and that will be useful when proving completeness. These arise due to the interaction axioms that involve two different modalities. We shall say that the relations are *trans-transitive*, *trans-dense*[1] and *trans-euclidean*, where the properties hold across different relations in R'. This is shown in detail for two arbitrary relations R'_{s_1} and R'_{s_2}:

1. *Trans-transitive*: $R'_{s_2} xy$ and $R'_{s_1} yz$ implies $R'_{s_1} xz$. This holds in any $\mathcal{M} \in \mathfrak{M}_{\mathbb{S}_0}$ due to the fact that:
   (a) If $R'_{s_1} yz$, then $z \in s1$ and hence, for all $p \in \Pi$ it is also the case that $R'_{s_1} pz$:
   $R'_{s_1} yz \to \forall p \in \Pi[R'_{s_1} pz]$
   (b) For $p = x$ the property holds
2. *Trans-dense* $R'_{s_1} xz \to \exists (y \in \Pi)\ [R'_{s_2} xy \wedge R'_{s_1} yz]$: Whenever $R'_{s_1} xz$ then for all standpoints s_2 there exists a $y \in \Pi$ such that $R'_{s_2} xy \wedge R'_{s_1} yz$. This occurs on the grounds that:
   (a) Standpoints are non-empty, therefore for any standpoint s_2 there exists a precisification $q \in \Pi$ such that $q \in s_2$:
   $\exists (q \in \Pi)[q \in s_2]$
   (b) For all $r \in \Pi$ and $q \in s_2$, it is the case that $R'_{s_2} rq$:
   $q \in s_2 \to \forall (r \in \Pi)[R'_{s_2} rq]$
   (c) By 2(a) and 2(b) we have that for any $r \in \Pi$ and any standpoint s_2 there is a $q \in \Pi$ such that $R'_{s_2} rq$:
   $\forall (r \in \Pi), \exists (q \in \Pi)[R'_{s_2} rq]$
   (d) By 2(c) and the modified density axiom, let q be y and given that x is a precisification (i.e. $x \in \Pi$), we have that:
   $R'_{s_1} xz \to \exists (y \in \Pi)[R'_{s_2} xy]$ for any stanspoint s_2

[1] Density and its corresponding axiom $\Box\Box\phi \to \Box\phi$ hold in systems satisfying **KD45**, as they can be derived from the rest of axioms. See [Che80] for the proof.

(e) By 3(a) and 2(d): $R'_{s_1}xz \rightarrow \exists(y \in \Pi)[R'_{s_2}xy \wedge R'_{s_1}yz]$.

3. *Trans-euclidean* $R'_{s_2}xy \wedge R'_{s_1}xz \rightarrow R'_{s_1}yz$: This is identical to the *trans-transitive* case. Formally:

   (a) If $R'_{s_1}xz$, then $z \in s1$ and hence, for all $p \in \Pi$ it is also the case that $R'_{s_1}pz$: $R'_{s_1}xz \rightarrow \forall p \in \Pi[R'_{s_1}pz]$

   (b) For $p = y$ the property holds

The properties of the accessibility relations force the validity of the axioms that we discuss in the following section 4.2.3 in any model of the class $\mathfrak{M}_{\mathbb{S}_0}$, making them theorems of $\mathbb{S}_0$ by **Definition 7**, which will be proved in section 4.3.

4.2.3 Proof theory

We now specify a proof system for the logic $\mathbb{S}_0$. We write $\vdash_{\mathbb{S}_0} \phi$ to mean that ϕ is a derivable theorem of $\mathbb{S}_0$. We give a Hilbert-style axiomatic proof system for determining the theorems of $\mathbb{S}_0$.

We shall specify a set of axioms and four inference rules for the modal Standpoint logic with the partial order $(R, \preceq)$, which capture significant aspects of the semantics. The axioms are specified by schemas, where all occurrences of each of the meta-variables (ϕ and ψ) may be instantiated by any formula of $\mathcal{L}_{\mathbb{S}_0}$ (all occurrence of a given meta-variable must be substituted by the same formula, but where both ϕ and ψ occur in a schema they may be substituted by different formulas). Since standpoint logic is built upon an underlying classical logic, all classically valid formulas are theorems. We will not give a particular axiomatisation for classical logic but assume that we have a sound and complete method of determining which formulas are classically valid.

4.2.3.1 Axioms

$\mathbb{S}_0$ satisfies the following well-known modal axioms:

AK $\square_s(\phi \rightarrow \psi) \rightarrow (\square_s \phi \rightarrow \square_s \psi)$

AD $\square_s \phi \rightarrow \lozenge_s \phi$

AT* $\square_* \phi \rightarrow \phi$ (note that **AT*** only applies to the universal operator $\square_*$.)

Additionally, two modal axioms govern interactions between the standpoint modality operators:

AS4 $\square_s \phi \to \square_{s'} \square_s \phi$

AS5 $\lozenge_s \phi \to \square_{s'} \lozenge_s \phi$

And finally the partial order axiom:

AP $\square_s \phi \to \square_{s'} \phi$ for $s' \preceq s$

Axiom **AP** captures the meaning of $\preceq$, by ensuring that any proposition considered definite in a given standpoint is also considered definite in any sharper standpoint.

Moreover, we note that the well known modal axioms 4 and 5 are immediately derivable as special cases of **AS4** and **AS5**, where $s' = s$. Thus, each operator $\square_s$ obeys the axioms **AK**, **AD**, **4** and **5**, which are standard axioms for the normal modal logic **KD45** and the $\square_*$ operator satisfies an **S5** modality. This is discussed in more detail in section 4.2.3.4.

4.2.3.2 Inference Rules

The inference rules of $\mathbb{S}_0$ are: the rule that classical theorems are provable (**RC**), the rule that all instances of the axioms are provable (**RA**), the classical *modus ponens* rule (**MP**) and the rule of necessitation (**RN**):

RC. $\vdash_{\mathbb{S}_0} \phi$, if ϕ is a theorem of classical propositional logic (where we treat all modal sub-formulas of ϕ as atomic propositions).

RA. $\vdash_{\mathbb{S}_0} \phi$, if ϕ is an instance of one of the modal axioms (**AK**, **AD**, **AT***, **AS4**, **AS5**).

MP. If $\vdash_{\mathbb{S}_0} \phi$ and $\vdash_{\mathbb{S}_0} \phi \to \psi$, then $\vdash_{\mathbb{S}_0} \psi$.

RN. If $\vdash_{\mathbb{S}_0} \phi$, then $\vdash_{\mathbb{S}_0} \square_s \phi$, for all standpoints $s \in \mathcal{S}$.

A *proof* of ϕ in $\mathbb{S}_0$ is a sequence of formulas, ending in ϕ, such that each formula is either provable directly by **RC** or **RA**, or is derivable using **MP** or **RN** from formulas that occur earlier in the sequence.

4.2.3.3 Derivable theorems

We now list some theorems that are derivable using axioms $\Diamond$**Df** and **AK** and rules **RC**, **MP** and **RN** and will be used in establishing properties of the system.

AK1. $\Diamond_s(\phi \wedge \psi) \rightarrow \Diamond_s \phi$

AK$\Diamond$. $\Box_s(\phi \rightarrow \psi) \rightarrow (\Diamond_s \phi \rightarrow \Diamond_s \psi)$

AC. $\Box_s(\phi \wedge \psi) \leftrightarrow \Box_s \phi \wedge \Box_s \psi$

AC$\Diamond$. $\Diamond_s(\phi \vee \psi) \leftrightarrow \Diamond_s \phi \vee \Diamond_s \psi$

4.2.3.4 Axioms AS4 and AS5: Embedded standpoint operators

The standpoint operator $\Box_s$, its special case $\Box_*$ and the rest of the definable syntax ($\Diamond_s$, $\Diamond_*$, ...) flatten when embedding another standpoint operator. Consequently, complex propositions consisting of chains of two or more standpoint operators as well as nested expressions can be simplified into a single standpoint operator.

This property is what we call flattening standpoints, which consists of the fact that standpoint iteration is superfluous, not only for iterations of operators of the same standpoint but for any modal operator. Consequently, any string of box or diamond standpoint operators may be replaced by a single standpoint box or diamond, the last one in the string.

This feature is common to other systems of modal logic with a single modality satisfying the normal axioms **4** ($\Box \phi \rightarrow \Box \Box \phi$) and **5** ($\Diamond \phi \rightarrow \Box \Diamond \phi$), which make the iteration of modal operators redundant. An alethic interpretation of this feature would be to say that a statement such as 'It is necessary that it is necessary that π' says nothing more than 'It is necessary that π'. Moreover, saying that 'It is *possible* that it is *necessary* that π' is the same as saying that 'It is *necessary* that π'.

Yet, our **AS4** and **AS5** are stronger than **4** and **5**. Rather than being restricted to a single modal operator (and its dual) as in epistemic logic, where positive introspection reads as 'if I know ϕ then I know that I know ϕ' ($K_i\phi \rightarrow K_iK_i\phi$) but not necessarily 'if I know ϕ then x knows that I know ϕ' ($K_i\phi \rightarrow K_xK_i\phi$), in our system if ϕ is the case according to standpoint s, then it follows that according to any other standpoint s' it is the case that ϕ is the case according to standpoint s: $\Box_s \phi \rightarrow \Box_{s'} \Box_s \phi$.

Axioms		Relations	
Name	**Schema**	**Property**	**Schema**
AK	$\Box_s(\phi \to \psi) \to (\Box_s \phi \to \Box_s \psi)$	All	
AD	$\Box_s \phi \to \Diamond_s \phi$	Serial	$\forall x, \exists y\ R'_s xy$
A4	$\Box_s \phi \to \Box_s \Box_s \phi$	Transitive	$R'_s xy \wedge R'_s yz \to R'_s xz$
A5	$\Diamond_s \phi \to \Box_s \Diamond_s \phi$	Euclidean	$R'_s xy \wedge R'_s xz \to R'_s yz$
AS4	$\Box_s \phi \to \Box_{s'} \Box_s \phi$	Trans-transitive	$R'_{s_2} xy \wedge R'_{s_1} yz \to R'_{s_1} xz$
AS5	$\Diamond_s \phi \to \Box_{s'} \Diamond_s \phi$	Trans-euclidean	$R'_{s_2} xy \wedge R'_{s_1} xz \to R'_{s_1} yz$
AT*	$\Box_* \phi \to \phi$	Reflexive	$\forall x\ R'_* xx$
AP	$\Box_s \phi \to \Box_{s'} \phi \ \ (s' \preceq s)$	Monotonic	$R'_s xy \wedge s \preceq s' \to R'_{s'} xy$

Table 4.1: Correspondence between the axioms in $\mathbb{S}_0$ and the properties of the relations of the models $\mathfrak{M}_{\mathbb{S}_0}$.

An example of how this feature reads goes as follows: if it is the case that *'This is a forest according to the FAO[1]'*, which we formalise as $\Box_{FAO}$ IsForest, then according to any agent's standpoint α, it can only be the case that *'according to his standpoint (s_α), this is a forest according to the FAO[2]'*, which can be formalised as $\Box_\alpha \Box_{FAO}$ IsForest.

4.2.4 Correspondence between model properties and axioms

Correspondence theory is well studied for normal modal logics [Che80], and states that the restrictions in the Kripke or standard models (such as forcing relations to be reflexive, serial, ...) correspond to the validity of certain axioms in the syntactic counterpart (e.g. **AT*** and **AD** respectively). In this section we quickly overview the correspondence between our semantics and syntax (displayed in table 4.1) before proceeding to prove that the logic $\mathbb{S}_0$ is determined by the class $\mathfrak{M}_{\mathbb{S}_0}$ of Kripke models.

It is known that, given that standpoint logic is a normal modal logic, **AK** holds for all standpoints, being determined by the general class of Kripke models. Further, the correspondence between relation properties and the theoremhood of the axioms in $\mathbb{S}_0$ is

[1] Food and Agriculture Organisation of the United Nations (FAO)
[2] FAO

as follows. Serial standpoint relations s correspond to the axiom **AD** being a theorem of $\mathbb{S}_0$, the reflexive relation R'_* to the axiom **AT*** holding for the modal $\square_*$, transitive and euclidean relations to the axioms **A4** and **A5** respectively, monotonicity corresponds to the partial order axiom **AP**. Finally, as we prove in the following section, trans-transitive and trans-euclidean relations correspond to our stronger interaction axioms **AS4** and **AS5**. The set of correspondences can be seen in detail in table 4.2.3.4. More properties and axioms hold, such as density corresponding to $\square\square\phi \rightarrow \square\phi$, but they are derivable from the axioms and properties presented in this table[1].

4.3 Soundness and Completeness of $\mathbb{S}_0$

Following [Che80], a system of modal logic is said to be *sound* with respect to a class of models just in case every theorem of the system is valid in the class of models. Conversely, a system is *complete* with respect to a class of models if and only if every sentence valid in the class of models is a theorem of the system. Finally, a system of modal logic is *determined* by a class of models when it is both sound and complete, such that for every ϕ,

$$\vdash_{\mathbb{S}_0} \phi \quad \text{iff} \quad \models_{\mathbb{S}_0} \phi$$

In this section we shall show soundness and completeness, hence determination, of the system of modal logic $\mathbb{S}_0$ with respect to the class of models $\mathfrak{M}_{\mathbb{S}_0}$. Note that it is possible that a system of modal logic be determined by more than one class of models [Che80].

4.3.1 Soundness

We first record some preliminary definitions,

Definition 8. A *proof* in $\mathbb{S}_0$ is a finite sequence of sequents:

$$\Gamma_1|\theta_1, ..., \Gamma_m|\theta_m$$

where $\theta_i \in \mathcal{L}_{\mathbb{S}}$ and $\Gamma_i \subseteq \mathcal{L}_{\mathbb{S}}$ are finite, such that for each $i = 1, ..., m$ either the formula θ_i is an instance of an axiom of $\mathbb{S}_0$ or for some $j_1, ..., j_s < i$

[1] Axioms **A4** and **A5** are also derivable but we include them as they help illustrating our stronger interaction axioms **AS4** and **AS5**

$$\frac{\Gamma_{j1}|\theta_{j1}, \quad \dots \quad ,\Gamma_{js}|\theta_{js}}{\Gamma_i|\theta_i}$$

is an instance of one of the rules of proof.

Definition 9. For $\theta \in \mathcal{L}_{\mathbb{S}}$ and $\Gamma \subseteq \mathcal{L}_{\mathbb{S}}$, possibly infinite, we define:

$$\Gamma \vdash_{\mathbb{S}_0} \theta \iff \text{there is a proof } \Gamma_1|\theta_1, ..., \Gamma_m|\theta_m \text{ in } \mathbb{S}_0 \text{ such that}$$

$$\Gamma_m \subseteq \Gamma \text{ and } \theta_m = \theta$$

In this case, $\Gamma_1|\theta_1, ..., \Gamma_m|\theta_m$ is called a proof of θ from Γ in $\mathbb{S}_0$.

Theorem 1 (Soundness)**.** *If $\Gamma \subseteq \mathcal{L}_{\mathbb{S}}$ and $\theta \in \mathcal{L}_{\mathbb{S}}$, then,*

$$\textit{whenever } \Gamma \vdash_{\mathbb{S}_0} \theta \textit{ we must have } \Gamma \models_{\mathbb{S}_0} \theta \ .$$

Proof.

Let $\Gamma_1|\theta_1, ..., \Gamma_m|\theta_m$ be a proof of $\Gamma \vdash_{\mathbb{S}_0} \theta$. Hence, $\Gamma_m \subseteq \Gamma$ and $\theta_m = \theta$, and we show by induction on k that $\Gamma_k \models_{\mathbb{S}_0} \theta_k$ for $k = 1, 2, ..., m$.

With $\mathcal{M} = \langle \Pi, (R, \preceq), \delta \rangle$ an interpretation structure in $\mathfrak{M}_{\mathbb{S}_0}$, we show that $\Gamma_k \vdash_{\mathbb{S}_0} \theta_k \to \Gamma_k \models_{\mathbb{S}_0} \theta_k$ in the case (1) where $\Gamma_k|\theta_k$ is an instance of an axiom of $\mathbb{S}_0$ and in the case (2) where $\Gamma_k|\theta_k$ is obtained by applying a rule of $\mathbb{S}_0$ to some $\Gamma_1|\theta_1, ..., \Gamma_m|\theta_m$ where $1, ..., m < k$

4.3.1.1 Case 1. $\Gamma_k|\theta_k$ is an axiom of $\mathbb{S}_0$

AK. $\Box_s(\phi \to \psi) \to (\Box_s \phi \to \Box_s \psi)$

Let $\pi \in \Pi$ and $(\mathcal{M}, \pi) \models \Box_s(\phi \to \psi)$,

$$\Rightarrow (\mathcal{M}, \pi') \models \phi \to \psi \text{ for all } \pi' \in R_s \qquad (1)$$

Let $(\mathcal{M}, \pi) \models \Box_s \phi$,

$$\Rightarrow (\mathcal{M}, \pi') \models \phi \text{ for all } \pi' \in R_s$$

$$\Rightarrow (\mathcal{M}, \pi') \models \psi \text{ for all } \pi' \in R_s \quad \text{(by 1)}$$

$$\Rightarrow (\mathcal{M}, \pi) \models \Box_s \psi$$

$$\Rightarrow (\mathcal{M}, \pi) \models \Box_s \phi \to \Box_s \psi$$

$$\Rightarrow (\mathcal{M}, \pi) \models \Box_s(\phi \to \psi) \to (\Box_s \phi \to \Box_s \psi)$$

$$\Rightarrow \models \Box_s(\phi \to \psi) \to (\Box_s \phi \to \Box_s \psi) \text{ as } \pi \text{ is arbitrary.}$$

.

The case of $\Box_*$ follows in the same way.

AD. $\Box_s \phi \to \Diamond_s \phi$

We must show that for all $\pi \in \Pi$, $(\mathcal{M}, \pi) \models \Box_s \phi \to \Diamond_s \phi$. Let π and π' be arbitrary precisifications $\pi, \pi' \in \Pi$,

$$(\mathcal{M}, \pi) \models \Box_s \phi \Rightarrow (\mathcal{M}, \pi') \models \phi \text{ for all } \pi' \in R_s$$

$$\Rightarrow (\mathcal{M}, \pi') \models \phi \text{ for some } \pi' \in R_s \text{ as } R_s \neq \emptyset \Rightarrow (\mathcal{M}, \pi) \models \Diamond_s \phi$$

$$\Rightarrow (\mathcal{M}, \pi) \models \Box_s \phi \to \Diamond_s \phi$$

$$\Rightarrow \models \Box_s \phi \to \Diamond_s \phi \text{ as } \pi \text{ is arbitrary.}$$

AT. $\Box_* \phi \to \phi$

Let $\pi \in \Pi$, and assume $(\mathcal{M}, \pi) \models \Box_* \phi$,

$$(\mathcal{M}, \pi) \models \Box_* \phi$$

$$(\mathcal{M}, \pi') \models \phi \text{ for all } \pi' \in \Pi$$

$$(\mathcal{M}, \pi) \models \phi \text{ since } \pi \in \Pi$$

AP. $\Box_s \phi \to \Box_{s'} \phi$ **iff** $(s' \preceq s)$

Let $\pi \in \Pi$, and assume:

$$(\mathcal{M}, \pi) \models (\Box_s \phi)$$

then,

$$(\mathcal{M}, \pi) \models (\Box_s \phi) \text{ (*)}$$

And since $s' \preceq s$,

$$\Rightarrow R_{s'} \subseteq R_s \text{ and therefore for all } p \in R_s \text{ then } p \in R_{s'} \text{ (**)}$$

$$\Rightarrow (\mathcal{M}, \pi') \models \phi \text{ for all } \pi' \in R_s \text{ by *}$$

$$\Rightarrow (\mathcal{M}, \pi') \models \phi \text{ for all } \pi' \in R_{s'} \text{ by **}$$

$$\Rightarrow (\mathcal{M}, \pi) \models (\Box_{s'} \phi)$$

AS4. $\Box_s \phi \rightarrow \Box_{s'} \Box_s \phi$

Let $\pi \in \Pi$, and assume $(\mathcal{M}, \pi) \models \Box_s \phi$

$$(\mathcal{M}, \pi') \models \phi \text{ for all } \pi' \in R_s \ (1)$$

Let $\pi'' \in R_{s'}$

$$\langle \mathcal{M}, \pi'' \rangle \models \Box_s \phi \text{ by } 1$$

$$(\mathcal{M}, \pi) \models \Box_{s'} \Box_s \phi$$

AS5. $\Diamond_s \phi \rightarrow \Box_{s'} \Diamond_s \phi$

Let $\pi \in \Pi$ and assume $(\mathcal{M}, \pi) \models \Diamond_s \phi$,

$$(\mathcal{M}, \pi') \models \phi \text{ for some } \pi' \in R_s \ (1)$$

Let $\pi'' \in R_{s'}$,

$$\langle \mathcal{M}, \pi'' \rangle \models \Diamond_s \phi \text{ by } 1$$

$$(\mathcal{M}, \pi) \models \Box_{s'} \Diamond_s \phi$$

4.3.1.2 Case 2. $\Gamma_k|\theta_k$ is obtained by applying a rule

$\Gamma_k|\theta_k$ is obtained by applying a rule of $\mathbb{S}_0$ to some $\Gamma_1|\theta_1, ..., \Gamma_m|\theta_m$ where $1, ..., m < k$

By the induction hypothesis we know that $\Gamma_n \models_{\mathbb{S}_0} \theta_n$ for $n = 1, ..., m$.

Our only rules are that classical theorems are probable (**RC**), the rule that all instances of the axioms are provable (**RA**), Modus Ponens (**MP**) and the rule of necessitation (**RN**). Modus Ponens and **RC** follow as in the propositional calculus, since the verification always remains within one precisification. We have shown that all instances of the axioms are provable (**RA**). It remains to show that $\models_{\mathbb{S}_0}$ respects the necessity rule:

RN $\Box_s \phi$ for any theorem ϕ (i.e. if ϕ is provable from no premisses).

In other words, we may assume $\models_{\mathbb{S}_0} \phi$ and we must show. $\models \Box_s \phi$

Let $\pi \in \Pi$. Since $\models_{\mathbb{S}_0} \phi$ we know,

$$(\mathcal{M}, \pi') \models \phi \text{ for all } \pi' \in \Pi$$

$$\Rightarrow (\mathcal{M}, \pi) \models \Box_s \phi \text{ as any } R_s \text{ is } R_s \subseteq \Pi$$

$\Box$

This finishes our soundness proof of $\mathbb{S}_0$ with respect to the class $\mathfrak{M}_{\mathbb{S}_0}$ of Kripke frames. Note that the same proof could be written with the analogous binary relations $R'_s \in R'$ without major differences.

4.3.2 Completeness

We show that every formula ϕ that is valid according to the semantics of standpoint logic (i.e. $\models_{\mathbb{S}_0} \phi$) is provable in the above proof system.

Theorem 2 (Completeness)**.** *If $\phi \in \mathcal{L}_{\mathbb{S}_0}$, then*

$$\models_{\mathbb{S}_0} \phi \Rightarrow \vdash_{\mathbb{S}_0} \phi$$

Completeness of multimodal **KD45** systems with respect to transitive, serial and euclidean Kripke frames has been proven. We first introduce some preliminary ideas and methods used in the proofs. We then outline the completeness proof for multimodal KD45 logics provided in [FHMV95] and proceed to prove completeness of the extension of **KD45** presented here, namely *standpoint logic* $\mathbb{S}_0$ with respect to the class of models $\mathfrak{M}_{\mathbb{S}_0}$.

4.3.2.1 Maximal Consistent Formula Sets

Definition 10. A set of sentences Λ is consistent in $\mathbb{S}_0$ or $\mathbb{S}_0$-consistent ($Con_{\mathbb{S}_0}\Lambda$) just in case a constradiction ($\phi \wedge \neg\phi$) is not deducible from Λ.

We write $C\not{o}n_{\mathbb{S}_0}\Lambda$ just in case Λ is not consistent.

Definition 11. A set of sentences Γ is maximal consistent in $\mathbb{S}_0$ just in case it is $\mathbb{S}_0$-consistent and has only $\mathbb{S}_0$-inconsistent proper extensions, i.e. for any $\phi \in \mathcal{L}_{\mathbb{S}_0}$, either $\phi \in \Gamma$ or $\neg\phi \in \Gamma$.

All theorems of $\mathbb{S}_0$ are in every maximal $\mathbb{S}_0$-consistent set of formulas. Moreover, maximal consistent sets are closed under entailment, meaning that deducibility and membership coincide. We recapitulate theorem 2.12 from [Che80], where a proof is provided:

Theorem 3. *Let Γ be a $\mathbb{S}_0$-maximal set of sentences and let T be the set of theorems of $\mathbb{S}_0$ (i.e. $T = \{\phi \mid\ \vdash_{\mathbb{S}_0} \phi\}$). Then:*

1. $\phi \in \Gamma$ *iff* $\Gamma \vdash_{\mathbb{S}_0} \phi$.
2. $T \subseteq \Gamma$.
3. $\phi \in \Gamma$ *for any formula* ϕ *that is a theorem of classical propositional logic.*
4. $\phi \notin \Gamma$ *if* $\neg\phi$ *is a theorem of classical propositional logic.*
5. $\neg\phi \in \Gamma$ *iff* $\phi \notin \Gamma$.
6. $(\phi \wedge \psi) \in \Gamma$ *iff both* $\phi \in \Gamma$ *and* $\psi \in \Gamma$.
7. $(\phi \vee \psi) \in \Gamma$ *iff either* $\phi \in \Gamma$ *or* $\psi \in \Gamma$.
8. $(\phi \rightarrow \psi) \in \Gamma$ *iff either* $\phi \notin \Gamma$ *or* $\psi \in \Gamma$.
9. $(\phi \leftrightarrow \psi) \in \Gamma$ *iff either* $\phi, \psi \in \Gamma$ *or* $\psi, \phi \notin \Gamma$.

The next theorem is known as Lindenbaum's lemma. It is the proposition that every consistent set of sentences has a maximal extension.

Theorem 4 (Lindenbaum's lemma)**.** *If γ is $\mathbb{S}_0$-consistent, then there exists a Γ such that (i) $\gamma \subseteq \Gamma$ and (ii) Γ is a $\mathbb{S}_0$-maximal set of formulas.*

A proof for Lindenbaum's lemma is provided in [Che80].

4.3.2.2 The canonical model

One of the key ideas of this completeness proof is that of a *canonical model* for a system of modal logic. A canonical model is a particular type of model where the set of points is the set of all the maximal consistent sets of the language.

Definition 12. A canonical model of the system of modal logic $\mathbb{S}_0$ is a model $\mathcal{M}^C = \langle \Pi, (R', \preceq), \delta \rangle$ in which:

(a) Π is the set of all maximal $\mathbb{S}_0$-consistent sets of sentences. Hence, each precisification $\pi_\Gamma \in \Pi$ is identified with a maximal $\mathbb{S}_0$-consistent set of sentences, Γ.

(b) Each relation $R'_s \in (R', \preceq)$ is such that $(\pi_A, \pi_B) \in R'_s$ iff $\{\phi : \Box_s\, \phi \in A\} \subseteq B$.

(c) The evaluation function $\delta : \mathcal{P} \to 2^W$ maps each propositional constant $p \in \mathcal{P}$ to the set $\delta(p) \subseteq \Pi$ of precisifications $\{\pi_\Gamma \mid p \in \Gamma\}$.

Remark 1. Each relation $R'_s \in R'$ can alternatively be defined in terms of the diamond operator such that $(\pi_A, \pi_B) \in R'_s$ iff $\{\Diamond_s \phi : \phi \in \pi_B\} \subseteq \pi_A$. Hence, $\{\phi : \Box_s \phi \in \pi_A\} \subseteq \pi_B$ iff $\{\Diamond_s \phi : \phi \in \pi_B\} \subseteq \pi_A$ (see [Che80] p.174).

Because in the canonical model each of the worlds or precisifications is an $\mathbb{S}_0$-maximal set of sentences, then the sentences that are true at a world are those that are contained by it, such that π_Γ is the precisification determined by a $\mathbb{S}_0$-maximal sets of sentences Γ.

Lemma 1. *For ϕ an $\mathbb{S}_0$-consistent formula and Γ a $\mathbb{S}_0$-maximal consistent formula set,*

$$\mathcal{M}^C, \pi_\Gamma \models \phi \quad \textit{iff} \quad \phi \in \Gamma$$

The first part of our completeness proof consists of showing that any consistent formula in the $\mathbb{S}_0$-language, $\phi \in \mathcal{L}_{\mathbb{S}_0}$, has a model: the canonical model. That is, we prove that any consistent formula ϕ is true at some precisification of a model $\mathcal{M}^C$.

Subsequently, one must typically show that such a model is of the class of intended models. As we will prove, the relations in $\mathcal{M}^C$ are serial, trans-transitive and trans-euclidean, but R'_* is an equivalence relation rather than universal and relations can not be simplified into sets. Hence, $\mathcal{M}^C$ is in fact not of the class $\mathfrak{M}_{\mathbb{S}_0}$ of intended models[1].

However, we will then show that we can generate a submodel $\mathcal{M}^C_\Gamma$ of $\mathcal{M}^C$ such that R'_* is universal and all relations can be simplified into sets.

Lemma 2. *For each ϕ such that $\mathcal{M}^C, \pi_\Gamma \models \phi$, there is a model $\mathcal{M}^C_\Gamma$ equivalent to $\mathcal{M}^C$ with respect to Γ such that*

$$\mathcal{M}^C_\Gamma \in \mathfrak{M}_{\mathbb{S}_0}$$

[1] Note that it is possible that a system of modal logic be determined by more than one class of models (see [Che80] p.60). Until this point, the proof of completeness shows that $\mathbb{S}_0$ is complete with respect to the class of models that are serial, trans-transitive and trans-euclidean, where R'_* is an equivalence relation, and without further restrictions. Yet, because we are interested in simpler models, we will go further in proving that $\mathbb{S}_0$ is also determined by the more restricted class $\mathfrak{M}_{\mathbb{S}_0}$.

Note that the model $\mathcal{M}^C_\Gamma$ is equivalent to $\mathcal{M}^C$ with respect to Γ if for any formula ϕ then,

$$\mathcal{M}^C_\Gamma, \pi_\Gamma \models \phi \quad \text{iff} \quad \mathcal{M}^C, \pi_\Gamma \models \phi$$

Before proceeding to the detailed proofs, let us visualise an example of a canonical model $\mathcal{M}^C$ for a small $\mathbb{S}_0$-language consisting of a single propositional variable. Let us also visualise how a particular proposition ϕ in that language is satisfied by both $\mathcal{M}^C$ and a submodel $\mathcal{M}^C_{\Gamma_5}$ such that $\phi \in \Gamma_5$, and how the latter is of the class $\mathfrak{M}_{\mathbb{S}_0}$.

Example 4.3.1. Let us consider a $\mathcal{L}_{\mathbb{S}_0}$ consisting of the set of propositional variables $\mathcal{P} = \{p\}$ and the standpoints s and $*$. Now, for the sake of example, let us imagine that we want to see if the consistent sentence $\phi = p \wedge \Box_s \neg p$ is satisfiable.

We first proceed to build the canonical model. For the language $\mathcal{L}_{\mathbb{S}_0}$ the set of distinct maximal consistent sets is $\{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4, \Gamma_5, \Gamma_6, \Gamma_7, \Gamma_8\}$, such that each maximal consistent set is the closure under entailment of the following sets:

- $\{p,\ \Box_s p,\ \Box_* p\} \subseteq \Gamma_1$
- $\{\neg p,\ \Box_s \neg p,\ \Box_* \neg p\} \subseteq \Gamma_2$
- $\{p,\ \Box_s p,\ \mathcal{I}_* p\} \subseteq \Gamma_3$
- $\{\neg p,\ \Box_s p,\ \mathcal{I}_* p\} \subseteq \Gamma_4$
- $\{p,\ \mathcal{I}_s p,\ \mathcal{I}_* p\} \subseteq \Gamma_5$
- $\{\neg p,\ \mathcal{I}_s p,\ \mathcal{I}_* p\} \subseteq \Gamma_6$
- $\{p,\ \Box_s \neg p,\ \mathcal{I}_* p\} \subseteq \Gamma_7$
- $\{\neg p,\ \Box_s \neg p,\ \mathcal{I}_* p\} \subseteq \Gamma_8$

In figure 4.2, we an see the canonical model $\mathcal{M}^C$ generated from the maximal consistent sets, where each point π_i is associated with the respective set Γ_i and where the relations are established following the definition **12**. Moreover, note that the canonical model satisfies ϕ at point π_5,

$$\phi \in \Gamma_5 \quad \text{and} \quad \mathcal{M}^C, \pi_5 \models \phi$$

It should be easy to see how $\mathcal{M}^C$ does not belong to the class of intended models, as $*$ is clearly not universal. However, the proof of **Lemma 2** shows how for $\phi \in \Gamma_5$, the submodel $\mathcal{M}^C_{\Gamma_5}$, which includes Γ_5 and is closed under the $*$ relation, belongs to the class of intended models, $\mathfrak{M}_{\mathbb{S}_0}$, and equally satisfies ϕ.

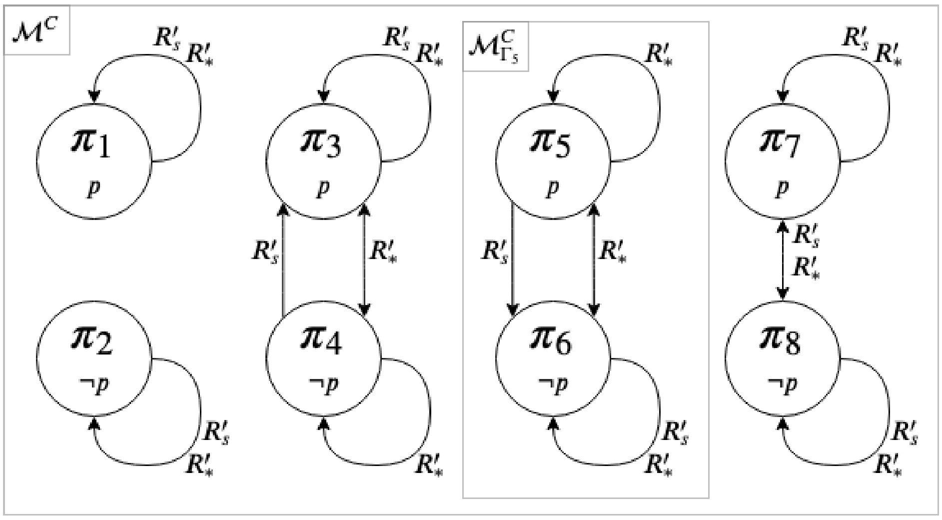

Figure 4.2: Representation of the canonical model $\mathcal{M}^C$ for the language $\mathcal{L}_{\mathbb{S}_0}$ with $\mathcal{P} = \{p\}$ and the standpoints s and $*$, as well as the model $\mathcal{M}^C_{\Gamma_5}$ which is equivalent to $\mathcal{M}^C$ with respect to the precisification π_5, such that $\phi \in \pi_5$.

4.3.2.3 Completeness proof

We are now ready to prove completeness of the system of standpoint logic $\mathbb{S}_0$ with respect to the class $\mathfrak{M}_{\mathbb{S}_0}$ of Kripke models.

The proof is by a Henkin-style method. We show that every consistent formula has a model; or, more precisely: for every formula ϕ that cannot be proved to be false, there is a model in which ϕ is true at at least one precisification:

Lemma 3. *For some $\mathcal{M} \in \mathfrak{M}_{\mathbb{S}_0}$ and $\pi \in \Pi_{\mathcal{M}}$,*

$$\textit{if} \quad \nvdash_{\mathbb{S}_0} \neg\phi \quad \textit{then} \quad \mathcal{M}, \pi \models \phi \ .$$

Proof of Theorem 2 (Completeness).

To see that **Lemma 3** entails **Theorem 2** suppose, for the sake of contradiction, that $\models_{\mathbb{S}_0} \phi$ and $\nvdash_{\mathbb{S}_0} \phi$. From **Lemma 3**, if $\nvdash_{\mathbb{S}_0} \phi$ then $\mathcal{M}, \pi \models \neg\phi$ for some $\mathcal{M} \in \mathfrak{M}_{\mathbb{S}_0}$ and $\pi \in \Pi_{\mathcal{M}}$. This is however a contradiction with $\models_{\mathbb{S}_0} \phi$, since, for ϕ to be valid, $\neg\phi$ must be false at every precisification of every model. Hence, we conclude that **Lemma 3** entails **Theorem 2**.

□

Proof of Lemma 3.

In order to prove **Lemma 3**, we use Lindenbaum's lemma (**Theorem 4**) to build a maximal $\mathbb{S}_0$-consistent formula set Γ from ϕ. Then, by direct application of **Lemma 1**, from $\phi \in \Gamma$ we get that $\mathcal{M}^C, \pi_\Gamma \models \phi$, where $\mathcal{M}^C$ is the canonical model. Finally, by **Lemma 2** we know that there is an equivalent model $\mathcal{M}^C_\Gamma$ with respect to ϕ such that $\mathcal{M}^C_\Gamma \in \mathfrak{M}_{\mathbb{S}_0}$. This suffices to prove **Lemma 3** given that we show that for all $\phi \in \mathcal{L}_{\mathbb{S}_0}$ there is a model $\mathcal{M}^C_\Gamma$ in which ϕ is true at least at one precisification, and that belongs to the class $\mathfrak{M}_{\mathbb{S}_0}$ of Kripke models, in correspondence with our proof theory. $\square$

Proof of Lemma 1.

In order to prove **Lemma 1** we build the $\mathcal{M}^C$ as follows. Given a set of formulas Γ, and a standpoint s_i, we define

$$\Gamma/\Box_{s_i} = \{\phi \mid \Box_{s_i}\phi \in \Gamma\} \ .$$

We then let $\mathcal{M}^C = \langle \Pi, (R', \preceq), \delta \rangle$ be such that

- Π is the set of maximal $\mathcal{L}_{\mathbb{S}_0}$-consistent sets of sentences,
- each R'_{s_i} in the partial order $(R', \preceq)$ is $R'_{s_i} = \{(\pi_\Gamma, \pi_\Sigma) \mid \Gamma/\Box_{s_i} \subseteq \Sigma\}$,
- $\delta(\pi_\Gamma, p) = \mathbf{t}$ iff $p \in \Gamma$,
- $\delta(\pi_\Gamma, p) = \mathbf{f}$ iff $p \notin \Gamma$.

Each precisification in $\mathcal{M}^C$ corresponds to a maximal consistent set of formulas, which are the formulas that are true according to that precisification. The accessibility relation R'_{s_i} relates each precisification π_Γ to all those precisifications $\pi_{\Gamma'}$, such that for every formula $\Box_{s_i}\phi$ that is true in π_Γ, the formula ϕ is true in $\pi_{\Gamma'}$.

We now proceed to show that every consistent formula ϕ has a model, in particular the canonical model $\mathcal{M}^C$. Recall that every consistent formula must be a member of at least one maximal consistent set, and the canonical model contains a precisification for every maximal consistent formula set. We will now show that the formulae true at each precisification in the canonical model are exactly the formulae in the associated maximal consistent formula set. Thus, ϕ will be true at at least one precisification in the canonical model.

We prove this by induction on the structure of ϕ, so that for all $\pi_\Gamma \in \Pi$ we have that $(\mathcal{M}^C, \pi_\Gamma) \models \phi$ iff $\phi \in \Gamma$. We thus assume that **Lemma 1** holds for all subformulas of ϕ and we show that it also holds for ϕ.

(a) Base Case, ϕ is an atomic formula. Suppose ϕ is an atomic formula P_i. If $P_i \in \Gamma$ then $\mathcal{M}^C, \pi_\Gamma \models P_i$ and if $P_i \notin \Gamma$ then $\mathcal{M}^C, \pi_\Gamma \not\models P_i$, both by definition of δ.

(b) ϕ is a conjunction $\alpha \wedge \beta$: If $\alpha \wedge \beta \in \Gamma$ then, by **Theorem 3(6)**, $\alpha \in \Gamma$ and $\beta \in \Gamma$. Assuming that the theorem has been established for α and β, it must be the case that $\mathcal{M}^C, \pi_\Gamma \models \alpha$ and $\mathcal{M}^C, \pi_\Gamma \models \beta$. Therefore, according to the semantics, $\mathcal{M}^C, \pi_\Gamma \models \alpha \wedge \beta$.

In the case where $\alpha \wedge \beta \notin \Gamma$, the closure of π_Γ under entailment ensures that at least one of α and β is not in π_Γ. Hence, either $\mathcal{M}^C, \pi_\Gamma \not\models \alpha$ or $\mathcal{M}^C, \pi_\Gamma \not\models \beta$ or both. Consequently, according to the semantics, $\mathcal{M}^C, \pi_\Gamma \not\models \alpha \wedge \beta$.

(c) ϕ is a negation $\neg\alpha$: If $\neg\alpha \in \Gamma$ then, by **Theorem 3(5)**, we must have $\alpha \notin \Gamma$. So, if the theorem has been proved for α, then $\mathcal{M}^C, \pi_\Gamma \not\models \alpha$. Hence, according to the semantics, $\mathcal{M}^C, \pi_\Gamma \models \neg\alpha$.

In the case where $\neg\alpha \notin \Gamma$ then, because π_Γ is maximal, we must have $\alpha \in \Gamma$. Thus, if the theorem has been proved for α, then $\mathcal{M}^C, \pi_\Gamma \models \alpha$. Hence $\mathcal{M}^C, \pi_\Gamma \not\models \neg\alpha$.

(d) ϕ is a modal formula $\Box_{s_i} \alpha$: We adapt and complete the proof provided in [FHMV95]. We first show that if $\Box_{s_i} \alpha \in \Gamma$ then $\mathcal{M}^C, \pi_\Gamma \models \Box_{s_i} \alpha$. First, we know that if $\Box_{s_i} \alpha \in \Gamma$ then $\alpha \in \Gamma / \Box_{s_i}$. By the definition of R'_{s_i}, if $(\pi_\Gamma, \pi_\Sigma) \in R'_{s_i}$ then $\Gamma / \Box_{s_i} \subseteq \Sigma$, and hence $\alpha \in \Sigma$. We now assume, by induction, that the theorem has been proved for $\alpha \in \Sigma$, so that $\mathcal{M}^C, \pi_\Sigma \models \alpha$ for all Σ such that $(\pi_\Gamma, \pi_\Sigma) \in R'_{s_i}$. Then, according to the semantics, it follows that $\mathcal{M}^C, \pi_\Gamma \models \Box_{s_i} \alpha$.

For the other direction, we show that if $\mathcal{M}^C, \pi_\Gamma \models \Box_{s_i} \alpha$ then $\Box_{s_i} \alpha \in \Gamma$. Assume $\mathcal{M}^C, \pi_\Gamma \models \Box_{s_i} \alpha$. Then, for the sake of contradiction, assume that the set $(\Gamma / \Box_{s_i}) \cup \{\neg\alpha\}$ is $\mathbb{S}_0$-consistent, such that $\alpha \notin (\Gamma / \Box_{s_i})$ and therefore $\Box_{s_i} \alpha \notin \Gamma$. By Lindenbaum's lemma (**Theorem 4**) $(\Gamma / \Box_{s_i}) \cup \{\neg\alpha\}$ would have a maximal $\mathbb{S}_0$-consistent extension Σ and, by construction, we would have $(\pi_\Gamma, \pi_\Sigma) \in R'_{s_i}$. By the induction hypothesis we have $\mathcal{M}^C, \pi_\Sigma \models \neg\alpha$, and so $\mathcal{M}^C, \pi_\Gamma \models \neg \Box_{s_i} \alpha$, hence reaching a contradiction.

Since $(\Gamma/\Box_{s_i}) \cup \{\neg\alpha\}$ is not $\mathbb{S}_0$-consistent then there must be some finite subset $\{\psi_1, ..., \psi_{n-1}\} \cup \{\neg\alpha\}$ that is not $\mathbb{S}_0$-consistent.[1] By the properties of deducibility and consistency (see [Che80], p.49 for the formal proof), we have that $\{\psi_1, ..., \psi_{n-1}\} \vdash_{\mathbb{S}_0} \alpha$.

$$C\o n_{\mathbb{S}_0}(\{\psi_1, ..., \psi_{n-1}\} \cup \{\neg\alpha\}) \text{ iff } \{\psi_1, ..., \psi_{n-1}\} \vdash_{\mathbb{S}_0} \alpha$$

Then again by the properties of deducibility and consistency (see [Che80], p.48 for the formal proof), we have that

$$\{\psi_1, ..., \psi_{n-1}\} \vdash_{\mathbb{S}_0} \psi_n \to \alpha \text{ iff } \{\psi_1, ..., \psi_{n-1}\} \cup \{\psi_n\} \vdash_{\mathbb{S}_0} \alpha$$

Hence, by iteration and propositional logic we can easily get that

$$\vdash_{\mathbb{S}_0} \psi_1 \to (... \to (\psi_n \to \alpha)...).$$

And by the necessitation rule ,

$$\vdash_{\mathbb{S}_0} \Box_{s_i}(\psi_1 \to (... \to (\psi_n \to \alpha)...)). \qquad \text{(d.1)}$$

By induction on n, together with axiom **AK** and propositional reasoning, we can show

$$\vdash_{\mathbb{S}_0} \Box_{s_i}(\psi_1 \to (... \to (\psi_n \to \alpha)...))$$

$$\to (\Box_{s_i} \psi_1 \to (... \to (\Box_{s_i} \psi_n \to \Box_{s_i} \alpha)...)). \qquad \text{(d.2)}$$

And from (d.1) and (d.2) we get

$$\vdash_{\mathbb{S}_0} \Box_{s_i} \psi_1 \to (... \to (\Box_{s_i} \psi_n \to \Box_{s_i} \alpha)...).$$

By **Theorem 3(2)** it follows that

$$(\Box_{s_i} \psi_1 \to (... \to (\Box_{s_i} \psi_n \to \Box_{s_i} \alpha)...)) \in \Gamma. \qquad \text{(d.3)}$$

Given that $\{\psi_1, ..., \psi_n\} \in (\Gamma/\Box_{s_i})$, then, by definition, $\{\Box_{s_i} \psi_1, ..., \Box_{s_i} \psi_n\} \in \Gamma$.

Finally, by repeatedly applying **Theorem 3(8)** we get $\Box_{s_i} \alpha \in \Gamma$ as desired. $\Box$

[1] Here, we are relying on a compactness property of the logic: a formula set is inconsistent iff it has a finite subset that is inconsistent. This is true for first-order logic and most modal logics. In the case of modal logics, compactness can be established by specifying a translation of modal formulae into first-order formulae, which explicitly represent the association of propositions to possible worlds and the accessibility relation between worlds. The properties of the accessibility relation also need to be specified by first-order. Since our axioms correspond to simple local constraints on the accessibility relation, it is clear that this can be done for $\mathbb{S}_0$.

We have now shown that for every consistent formula ϕ there is a Kripke model $\mathcal{M}^C$ in which ϕ is true at at least one precisification. We now must prove **Lemma 2**, which we used to show completeness with respect to the particular class $\mathfrak{M}_{\mathbb{S}_0}$ of Kripke models.

Proof of Lemma 2.

We show that for each ϕ such that $\mathcal{M}^C, \pi_\Gamma \models \phi$, there is a model $\mathcal{M}^C_\Gamma$ equivalent to $\mathcal{M}^C$ with respect to ϕ such that $\mathcal{M}^C_\Gamma \in \mathfrak{M}_{\mathbb{S}_0}$.

To do this, we first show, in **Part 1**, that the set of relations R' of the canonical model $\mathcal{M}^C$ is such that:

1. All the relations $R'_s \in R'$ are serial, transitive and euclidean.

2. R'_* is reflexive.

3. The relations $R'_s \in R'$ are trans-transitive and trans-euclidean.

4. The relations $R'_s \in R'$ are partially ordered under the subset relation.

Then, we proceed to show, in **Part 2**, that for each ϕ such that $\mathcal{M}^C \models \phi$, there is a model $\mathcal{M}^C_\Gamma$ equivalent to $\mathcal{M}^C$ with respect to ϕ where R'_* is the universal relation and all the relations $R'_s \in R'$ are such that, if $(\pi_A, \pi_B) \in R'_s$ then for any $\pi_\Gamma \in \Pi$, $(\pi_\Gamma, \pi_B) \in R'_s$ and for any $\pi_\Gamma \in s$ $(\pi_A, \pi_\Gamma) \in R'_s$. Hence the binary relations can be simplified into their unary counterparts as specified in the semantics (section 4.2.2.2). Consequently we show that $\mathcal{M}^C_\Gamma$ belongs to the class $\mathfrak{M}_{\mathbb{S}_0}$, hence finishing our completeness proof.

Part 1.

(1.1) Serial relations We first show that axiom **AD** forces the possibility relations in the canonical model $\mathcal{M}^C$ to be serial, i.e. for every $\pi_A \in \Pi$ there is a $\pi_B \in \Pi$ such that $(\pi_A, \pi_B) \in R'_s$.

To see this, recall that all instances of **AD** are true at all $\pi \in \Pi$. Then if $\Box_{s_i} \phi \in A$, by **AD** we have $\Diamond_{s_i} \phi \in A$. Moreover, by **RN**, for ϕ a propositional tautology we have $\Box_{s_i} \phi \in A$. Finally, by the construction of $\mathcal{M}^C$ we have $\{\Diamond_{s_i} \phi : \phi \in B\} \subseteq A$ iff $(\pi_A, \pi_B) \in R'_s$ (see **remark 1**). Consequently,

$$\Diamond_{s_i}(p \vee \neg p) \in A \text{ iff for some } \pi_B \text{ such that } (\pi_A, \pi_B) \in R'_s,\ (p \vee \neg p) \in B$$

Given that $(p \vee \neg p)$ is a theorem of propositional logic, then by and **AD** we get $\Diamond_{s_i}(p \vee \neg p)$, and hence it must be the case that $\Diamond_{s_i}(p \vee \neg p) \in A$. Therefore there must be some B such that $(\pi_A, \pi_B) \in R'_s$ and $(p \vee \neg p) \in B$.

(1.2) Transitive relations We show that axiom **A4** forces the possibility relations in the canonical model $\mathcal{M}^C$ to be transitive. To see this, suppose that $(\pi_A, \pi_B), (\pi_B, \pi_\Gamma) \in R'_{s_i}$ and recall that all instances of **A4** are true at all $\pi \in \Pi$.

It follows that, for any ϕ, if $\Box_{s_i} \phi \in A$ then also $\Box_{s_i} \Box_{s_i} \phi \in A$ by **A4**. Given $(\pi_A, \pi_B), (\pi_B, \pi_\Gamma) \in R'_{s_i}$ and $\Box_{s_i} \Box_{s_i} \phi \in A$, then by the construction of $\mathcal{M}^C$ we have $\Box_{s_i} \phi \in B$ and $\phi \in \Gamma$. Therefore we have $\phi \in \Gamma$ for $\Box_{s_i} \phi \in A$, i.e. $(A/\Box_{s_i}) \subseteq \Gamma$, and hence by the definition of the canonical relations we have $(\pi_A, \pi_\Gamma) \in R'_{s_i}$ as desired.

(1.3) Euclidean relations We now show that axiom **A5** forces the possibility relations in the canonical model $\mathcal{M}^C$ to be euclidean, i.e. if $(\pi_A, \pi_B), (\pi_A, \pi_\Gamma) \in R'_{s_i}$ then $(\pi_B, \pi_\Gamma) \in R'_{s_i}$.

To see this, we recall that all instances of **A5** are true at π_A, and we suppose that

(i) $(\pi_A, \pi_B) \in R'_{s_i}$, and hence $\{\phi : \Box_{s_i} \phi \in A\} \subseteq B$,

(ii) $(\pi_A, \pi_\Gamma) \in R'_{s_i}$, and hence $\{\Diamond_{s_i} \phi : \phi \in \Gamma\} \subseteq A$,

(iii) $\Diamond_{s_i} \phi \in A$.

If (iii) then $\Box_{s_i} \Diamond_{s_i} \phi \in A$ (by **A5**), and, given (i) we then have $\Diamond_{s_i} \phi \in B$. In particular we have $\{\Diamond_{s_i} \phi : \Box_{s_i} \Diamond_{s_i} \phi \in A\} \subseteq B$ which is he same as $\{\Diamond_{s_i} \phi : \Diamond_{s_i} \phi \in A\} \subseteq B$.

We recall that, by the construction of $\mathcal{M}^C$, we have that $\{\Diamond_{s_i} \phi : \phi \in \Gamma\} \subseteq B$ iff $(\pi_B, \pi_G) \in R'_s$. And, given that because of (ii) we have that $\{\Diamond_{s_i} \phi : \phi \in \Gamma\} \subseteq A$, it follows that $\{\Diamond_{s_i} \phi : \phi \in \Gamma\} \subseteq B$ and therefore $(\pi_B, \pi_G) \in R'_s$ as desired.

(1.4) Reflexive * relations We show that axiom **AT*** forces the * relations in the canonical model $\mathcal{M}^C$ to be reflexive, i.e. $\{\phi : \Box_* \phi \in A\} \subseteq A$. To see this, recall that all instances of **AT*** are true at all $\pi \in \Pi$.

By the construction of $\mathcal{M}^C$, we have that $\{\phi : \Box_* \phi \in A\} \subseteq B$ iff $(\pi_A, \pi_B) \in R'_*$. Then, suppose that $\Box_* \phi \in A$. It immediately follows from **AT*** that $\phi \in A$. Hence it is easy to see that $\{\phi : \Box_* \phi \in A\} \subseteq A$ and therefore we get $(\pi_A, \pi_B) \in R'_*$ as desired.

(1.5) Partial orders We show that axiom **AP** [$\Box_s \phi \to \Box_{s'} \phi$ iff $(s' \preceq s)$] forces the relations in the canonical model $\mathcal{M}^C$ such that $(s' \preceq s)$ to be partially ordered under subsumption, i.e. if $(\pi_A, \pi_B) \in R'_{s'}$ then $(\pi_A, \pi_B) \in R'_s$.

To see this, we recall that all instances of **AP** satisfying the partial order are true at any $\pi \in \Pi$. We now suppose that there are two standpoints, s_i and s_j, such that $(s_j \preceq s_i)$.

In order to show that for any pair (π_A, π_B) and the relation $(s_j \preceq s_i)$, if $(\pi_A, \pi_B) \in R'_{s_j}$ then $(\pi_A, \pi_B) \in R'_{s_i}$, we show that $(\pi_A, \pi_B) \notin R'_{s_i}$ and $(\pi_A, \pi_B) \in R'_{s_j}$ entails a contradiction in the presence of **AP**.

We suppose then, for the sake of contradiction, that $(\pi_A, \pi_B) \notin R'_{s_i}$ and $(\pi_A, \pi_B) \in R'_{s_j}$. Hence we get $\{\phi : \Box_{s_j} \phi \in A\} \subseteq B$ and $\{\phi : \Box_{s_i} \phi \in A\} \nsubseteq B$. But from **AP** we get $\Box_{s_j} \phi \in A$ for all ϕ such that $\Box_{s_i} \phi \in A$, and hence $\{\phi : \Box_{s_i} \phi \in A\} \subseteq \{\phi : \Box_{s_j} \phi \in A\}$. It is now easy to see that this entails a contradiction because it can not be that $A/\Box_{s_i} \subseteq B$, $A/\Box_{s_j} \subseteq A/\Box_{s_i}$ and $A/\Box_{s_j} \nsubseteq B$. Hence we conclude that the relations $R'_s \in R'$ satisfy the partial order under subsumption.

(1.6) Further restrictions of AS4 and AS5 We show that our axioms **AS4** and **AS5**, from which we can derive **A4** and **A5**, further restrict the possibility relations in the canonical model $\mathcal{M}^C$ to be trans-transitive and trans-euclidean. Precisely, such that

(I) if $(\pi_A, \pi_B) \in R'_{s_j}$ and $(\pi_B, \pi_\Gamma) \in R'_{s_i}$, then $(\pi_A, \pi_\Gamma) \in R'_{s_i}$.

(II) if $(\pi_A, \pi_B) \in R'_{s_j}$ and $(\pi_A, \pi_\Gamma) \in R'_{s_i}$, then $(\pi_B, \pi_\Gamma) \in R'_{s_i}$.

Both (I) and (II) can be easily proven in the same way as regular transitive and euclidean relations. We first show that axiom **AS4** forces the relations in the canonical model $\mathcal{M}^C$ to be *trans-transitive*. To see this, suppose that $(\pi_A, \pi_B) \in R'_{s_j}$ and $(\pi_B, \pi_\Gamma) \in R'_{s_i}$ and recall that all instances of **AS4** (and it's inverse) are true at all $\pi \in \Pi$.

Suppose $\phi \in \Gamma$. Then, by the construction of $\mathcal{M}^C$ we have $\Diamond_{s_i} \phi \in B$ and $\Diamond_{s_j} \Diamond_{s_i} \phi \in A$. Also, by the inverse of **AS4**, $\Diamond_{s_i} \phi \in A$. We recall that $\{\Diamond_{s_i} \phi : \phi \in \Gamma\} \subseteq A$ iff $(\pi_A, \pi_\Gamma) \in R'_s$. Therefore it straightforwardly follows that $(\pi_A, \pi_\Gamma) \in R'_{s_i}$ as desired.

Now we show that axiom **AS5** forces the relations in the canonical model $\mathcal{M}^C$ to be *trans-euclidean*. To see this, suppose that

(i) $(\pi_A, \pi_B) \in R'_{s_j}$, and hence $\{\phi : \Box_{s_j} \phi \in A\} \subseteq B$,

(ii) $(\pi_A, \pi_\Gamma) \in R'_{s_i}$, and hence $\{\Diamond_{s_i} \phi : \phi \in \Gamma\} \subseteq A$,

(iii) $\Diamond_{s_i} \phi \in A$.

If (iii) then $\Box_{s_j} \Diamond_{s_i} \phi \in A$ (by **AS5**), and, given (i) we then have $\Diamond_{s_i} \phi \in B$. In particular we have $\{\Diamond_{s_i} \phi : \Box_{s_j} \Diamond_{s_i} \phi \in A\} \subseteq B$ which is he same as $\{\Diamond_{s_i} \phi : \Diamond_{s_i} \phi \in A\} \subseteq B$.

Because of (ii), we have that $\{\Diamond_{s_i} \phi : \phi \in \Gamma\} \subseteq A$, and with $\{\Diamond_{s_i} \phi : \Diamond_{s_i} \phi \in A\} \subseteq B$ it follows that $\{\Diamond_{s_i} \phi : \phi \in \Gamma\} \subseteq B$. By the construction of $\mathcal{M}^C$, we have that $\{\Diamond_{s_i} \phi : \phi \in \Gamma\} \subseteq B$ iff $(\pi_B, \pi_G) \in R'_{s_i}$. Therefore it follows that $(\pi_B, \pi_G) \in R'_{s_i}$ as desired.

Part 2.

We now want to show that there is a model $\mathcal{M}^C_\Gamma$ such that for $\phi \in \Gamma$, $\mathcal{M}^C_\Gamma \models \phi$ iff $\mathcal{M}^C \models \phi$ and $\mathcal{M}^C_\Gamma \in \mathfrak{M}_{\mathbb{S}_0}$. We first introduce some concepts,

Definition 13. (Submodel induced by Π_{sub}) Let $\mathcal{M} = \langle \Pi, (R', \preceq), \delta \rangle$ be a $\mathbb{S}_0$-model and Π_{sub} be a subset of Π. The submodel $\mathcal{M}_{sub} = \langle \Pi_{sub}, (R'_{sub}, \preceq), \delta_{sub} \rangle$ induced by Π_{sub} is defined by:

- $R'_{sub_i} := R'_{s_i} \cap \Pi_{sub}$ for each $i \in S$
- $\delta_{sub}(p) := \delta(p) \cap \Pi_{sub}$

Definition 14. (Submodel generated by Π_{sub}) Let $\langle \Pi_{sub} \rangle$ be the smallest superset of Π_{sub} that is closed with respect to all relations R'_{s_i} in $\mathcal{M}$, so that for all $(\pi, \pi') \in R'_{s_i}$ and $\pi \in \Pi_{sub}$, then $\pi' \in \Pi_{sub}$. Then $\mathcal{M}_{sub}$ is the submodel generated by Π_{sub} if it is the submodel induced by $\langle \Pi_{sub} \rangle$.

Lemma 4. *Let $\mathcal{M} = \langle \Pi, (R', \preceq), \delta \rangle$ be a $\mathbb{S}_0$-model and $\mathcal{M}_{sub} = \langle \Pi_{sub}, (R'_{sub}, \preceq), \delta_{sub} \rangle$ be the submodel generated by the non-empty subset Π_{sub}. Then for each $\pi \in \Pi_{sub}$ and each formula ϕ,*

$$\mathcal{M}, \pi \models \phi \quad \textit{iff} \quad \mathcal{M}_{sub}, \pi \models \phi$$

The proof for **Lemma 4** can be found in [Wöl15].

(2.1) Universal * relations We have shown that in the canonical model $\mathcal{M}^C$ the relation R'_* is an equivalence relation, as it is reflexive, symmetric and transitive, and $\Box_*$ is **S5**-consistent. A proof that there is an equivalent model where R'_* is universal goes along the lines of [Wöl15], where they prove it for the system **S5** with a single modality, and it goes as follows.

Let us consider a maximal $\mathcal{L}_{\mathbb{S}_0}$-consistent set Γ built from a proposition ϕ and a canonical model $\mathcal{M}^C$. We have previously shown that R'_* is an equivalence relation, as it is reflexive, symmetric and transitive. Moreover, we have shown that any other relation R'_s is a subset of R'_* by definition of * and **AP**.

Let $\mathcal{M}^C_\Gamma = \langle \Pi_{sub}, (R'_{sub}, \preceq), \delta_{sub} \rangle$ be the submodel generated by $\{\pi_\Gamma\}$ on $\mathcal{M}^C$. We know that in $\mathcal{M}^C$, any pair $(\pi, \pi') \in R'_{s_i}$ is also $(\pi, \pi') \in R'_*$ because of the subsumption of all relations under R'_*. Moreover, in the generated model from π_Γ, Π_{sub} is the smallest superset of $\{\pi_\Gamma\}$ that is closed with respect to all relations, or what is the same, with respect to R'_*. It is straightforward that this model is connected under R'_{sub*} and, given that R'_* is an equivalence relation, then R'_{sub*} is universal.

Finally, it directly follows from **Lemma 4** that $\mathcal{M}^C_\Gamma, \pi_\Gamma \models \phi$ iff $\mathcal{M}^C, \pi_\Gamma \models \phi$ and hence we show that the system is complete with respect to models $\mathcal{M}^C_\Gamma$ with R'_{sub*} universal.

(2.2) Unary relations We finally show that in $\mathcal{M}^C_\Gamma$ not only R'_{sub*} is universal but also the following holds for all R'_{sub_i}

$$(\pi_\Gamma, \pi_A) \in R'_{sub_i} \quad \text{iff} \quad \text{for all } \pi_B \in \Pi_{sub}\ ,\ (\pi_B, \pi_A) \in R'_{sub_i}$$

This restriction together with the rest of properties of the relations means that we can generalise R'_{sub_i} into a unary relation R_{sub_i} containing the subset of accessible relations under such standpoint, as specified in the semantics (section 4.2.2.2) and hence $\mathcal{M}^C_\Gamma \in \mathfrak{M}_{\mathbb{S}_0}$.

Let $\mathcal{M}^C_\Gamma$ be the submodel generated by $\{\pi_\Gamma\}$ on all the relations $R'_{sub_i} \in R'_{sub}$. For one direction, suppose that $(\pi_\Gamma, \pi_A) \in R'_{sub_i}$. Let us consider another precisification $\pi_B \in \Pi_{sub}$. By the construction of $\mathcal{M}^C_\Gamma$ we know that there is a path from π_Γ to π_B, because Π_{sub} is closed with respect to all relations from π_Γ. Let us say, thus, that there is a series of relations $R'_{sub_1}, \ldots, R'_{sub_j}$ and a series precisifications $\pi_1, \ldots, \pi_n$ such that $(\pi_\Gamma, \pi_1) \in R'_{sub_1}, \ldots, (\pi_n, \pi_B) \in R'_{sub_j}$. By successive applications of trans-

transitivity, we know that $(\pi_\Gamma, \pi_B) \in R'_{sub_j}$. Moreover, if we know that $(\pi_\Gamma, \pi_B) \in R'_{sub_j}$ and $(\pi_\Gamma, \pi_A) \in R'_{sub_i}$, by the trans-euclidean property we get that $(\pi_B, \pi_A) \in R'_{sub_i}$ as desired. The other direction is trivial.

We have now shown that the model $\mathcal{M}^C_\Gamma$ is in the class of models $\mathfrak{M}_{\mathbb{S}_0}$, determined by **Definition 4**. We now assume that ϕ is an $\mathbb{S}_0$-consistent formula (**Definition 10**), hence it is consistent with the axioms of $\mathbb{S}_0$. By Lindenbaum's lemma (**Theorem 4**), we can create a maximal consistent set Γ such that $\phi \subseteq \Gamma$. Then it follows that ϕ is satisfiable in $\mathcal{M}^C_\Gamma$. $\square$

Corollary 1. *$\mathbb{S}_0$ is a sound and complete axiomatisation with respect to $\mathfrak{M}_{\mathbb{S}_0}$ for formulas in the language $\mathcal{L}_{\mathbb{S}_0}$*

Corollary 2. *$\mathbb{S}_0$ is determined by the class of models $\mathfrak{M}_{\mathbb{S}_0}$.*

4.4 Decidability and Complexity of $\mathbb{S}_0$

In this section, we present results showing that in $\mathbb{S}_0$ whenever a formula is valid it is decidable; that is, there is an algorithm that, when given a formula ϕ as input, will decide whether ϕ is valid. Moreover, such algotithm is NP-complete.

We introduce some definitions,

Definition 15. The size of a formula ϕ in $\mathcal{L}_{\mathbb{S}_0}$, denoted $|\phi|$, is its length over the alphabet $P \cup \{\neg, \wedge, (,), \Box_{s_1}, \ldots, \Box_{s_n}, \Box_*\}$, where P is the set of atomic propositions.

Definition 16. ψ is a subformula of ϕ if either $\psi = \phi$ or ϕ is of the form $\neg\phi'$ or $\phi' \wedge \phi''$ or $\Box_{s_i} \phi'$ and ψ is a subformula of either ϕ' or ϕ''. Let $Sub(\phi)$ be the set of all subformulas of ϕ. Then $|Sub(\phi)| \leq |\phi|$.

4.4.1 Decidability

We say that a system of modal logic is decidable if there is an effective finitary method for determining the theoremhood of a formula. This is shown to be the case if the system is axiomatisable by a finite number of schemas or axioms and it has the finite model property [Che80]. We have shown before that our system is axiomatisable by a finite number of axioms, presented in section 4.2.3, and hence the positive test for theoremhood can be done by a proof with the axioms and rules of inference. We therefore must show not that it has the finite model property so that we have a finite negative test for theoremhood.

Definition 17. A modal logic has the *finite model property* if and only if each non-theorem is false in some finite model of the logic [Che80].

We first note that if we have a finite model then there is a finite algorithm to determine whether a formula is satisfiable in that model or not. Specifically,

Lemma 5. *There is an algorithm that, given a finite Kripke structure $\mathcal{M}^{FIN}$ and a formula ϕ, determines, in time $O(||\mathcal{M}^{FIN}|| \times |\phi|)$, whether there is a state π of $\mathcal{M}^{FIN}$ such that $(\mathcal{M}^{FIN}, \pi) \models \phi$, for $|\phi|$ be the size of ϕ and $||\mathcal{M}^{FIN}||$ be the sum of the number of points π in Π and the number of pairs in the relations $R'_s \in R'$. I.e. there is an algorithm for checking if ϕ is satisfied in $\mathcal{M}^{FIN}$.*

See [FHMV95] for a proof of **Lemma 5**. The strategy to prove the finite model property for $\mathbb{S}_0$ follows closely the proof on [FHMV95] and consists in using the result that $\mathbb{S}_0$ is determined by the class of models $\mathfrak{M}_{\mathbb{S}_0}$ in section 4.3 to show that $\mathbb{S}_0$ is also determined by the class of $\mathfrak{M}^{FIN}_{\mathbb{S}_0}$ of finite models, so that

$$\vdash_{\mathbb{S}_0} \phi \quad \text{iff} \quad \models^{FIN}_{\mathbb{S}_0} \phi$$

Soundness is trivial because $\mathfrak{M}^{FIN}_{\mathbb{S}_0} \subseteq \mathfrak{M}_{\mathbb{S}_0}$, and hence it remains to show completeness with respect to $\mathfrak{M}^{FIN}_{\mathbb{S}_0}$. We first show that if a formula is $\mathbb{S}_0$-consistent, not only is it satisfiable in some structure (such as the canonical structure constructed in the proof of Theorem 2), but it is also satisfiable in a finite structure $\mathcal{M}^{FIN}$ in the class $\mathfrak{M}_{\mathbb{S}_0}$. The proof is actually just a variant of the proof of **Theorem 2 (*Completeness*)**.

Theorem 5. *If ϕ is $\mathbb{S}_0$-consistent, then ϕ is satisfiable in a structure in $\mathfrak{M}_{\mathbb{S}_0}$ with at most $2^{|\phi|}$ points or precisifications.*

Proof.

Let $Sub^+(\phi)$ be the set consisting of the subformulas of ϕ and their negations, that is, $Sub^+(\phi) = Sub(\phi) \cup \{\neg\psi \mid \psi \in Sub(\phi)\}$. Let $Con(\phi)$ be the set of maximal $\mathbb{S}_0$-consistent subsets of $Sub^+(\phi)$. Theorem 3 can be used to show that every $\mathbb{S}_0$-consistent subset of $Sub^+(\phi)$ can be extended to an element of $Con(\phi)$. Moreover, a member of $Con(\phi)$ contains either ψ or $\neg\psi$ for every formula $\psi \in Sub(\phi)$ (but not both, for otherwise it would not be $\mathbb{S}_0$-consistent), so the cardinality of $Con(\phi)$ is at most $2^{|Sub(\phi)|}$, which is at most $2^{|\phi|}$, since $|Sub(\phi)| \leq |\phi|$.

We can now construct a structure $\mathcal{M}^{FIN} = \langle \Pi^{FIN}, (R^{FIN}, \preceq), \delta^{FIN} \rangle$ in the class $\mathfrak{M}_{\mathbb{S}_0}$ identical to the canonical model (**Definition 12**) except that $\Pi = \{\pi_\Gamma \mid \Gamma \in Con(\phi)\}$. Note that the sets in $Con(\phi)$ are $\mathbb{S}_0$-consistent and that the axioms of $\mathbb{S}_0$ guarantee the properties of the relations in $\mathfrak{M}_{\mathbb{S}_0}$ and hence in $\mathcal{M}^{FIN}$ (as we have proven in section 4.3.2). Therefore the following restrictions hold for the relations in $(R^{FIN}, \preceq)$:

- $R'_s = \{(\pi_\Gamma,\ \pi_\Sigma) \mid \Gamma/\Box_s = \Sigma/\Box_s\,,\ \ \Gamma/\Box_s \subseteq \Sigma\}$
- $R'_* = \{(\pi_\Gamma,\, \pi_\Sigma) \mid \Gamma/\Box_* = \Sigma/\Box_*\,, \Gamma/\Box_s \subseteq \Gamma, \Gamma/\Box_s \subseteq \Sigma\}$

Finally, with a proof identical to that of **Theorem 2(a)-(d)**[1], we can show that for all $\psi \in Sub^+(\phi)$ and $\Gamma \in Con(\phi)$, iff $\psi \in \Gamma$ then $(\mathcal{M}^{FIN}, \pi_\Gamma) \models \psi$. Hence, ϕ is satisfiable in a structure in $\mathfrak{M}_{\mathbb{S}_0}$ with as many points as the cardinality of $Con(\phi)$, which is at most $2^{|\phi|}$ and, therefore we conclude that the provability problem is decidable.

$\square$

4.4.2 Complexity

After showing that the validity and probablility problems are decidable, we now show that they are in fact NP-Complete. In this section we partially use the proofs provided in [FHMV95, HM92]

Theorem 6. *The satisfiability problem for $\mathbb{S}_0$ is NP-complete, and thus so is the validity problem for $\mathbb{S}_0$.*

We first show that,

Lemma 6. *An $\mathbb{S}_0$ formula ϕ is satisfiable if and only if it is satisfiable in a structure $\mathcal{M}^{|\phi|} \in \mathfrak{M}_{\mathbb{S}_0}$ with at most $|\phi|$ precisifications.*

Proof.

Suppose that ϕ is satisfiable in a structure $\mathcal{M}^\phi = \langle \Pi, (R, \preceq), \delta \rangle$ in $\mathfrak{M}_{\mathbb{S}_0}$. Thus there must be at least one precisification $\pi \in \Pi$ such that $(\mathcal{M}^\phi, \pi^\phi) \models \phi$. From $\mathcal{M}^\phi$, let us

[1]See the proof by induction on the structure of ϕ with respect to the canonical model

define another structure $\mathcal{M}^m = \langle \Pi^m, (R^m, \preceq), \delta^m \rangle$. This model $\mathcal{M}^m$ will be a small submodel of $\mathcal{M}^\phi$ that is sufficient to satisfy the truth-conditions of ϕ and contains less than $|\phi|$ precisifications. It is constructed as follows:

R^m is the set of standpoints of ϕ such that $R^m_{s_i} \in R^m$ iff there is a subformula of the form $\Box_{s_i} \psi$ or $\Diamond_{s_i} \psi$ in ϕ. (We need only consider standpoints that are actually referenced in ϕ.) We now re-label the elements of the set R^m with the indexes $1, \ldots, n$ so that for any h, i, if $s_h \preceq s_i$ then $h < i$. So the order starts with the sharpest standpoints and proceeds to coarser standpoints until reaching standpoint $*$ if present. Since $\preceq$ need only be a partial order, we may have neither $s_h \preceq s_i$ nor $s_i \preceq s_h$, in which case they can be numbered in any order. Thus, the re-indexed set of standpoints can be in any strict order that is consistent with the original partial ordering induced by the $\preceq$ relation.

For every standpoint s_i, we specify the formula set F_{S_i}, containing all subformulae of ϕ that are not necessary according to s_i. Formally we define:

$$F_{S_i} = \{\neg\psi \;\mid\; \neg\Box_{s_i} \psi \in Sub^+(\phi) \;\text{ and }\; (\mathcal{M}, \pi) \models \neg\Box_{s_i} \psi \;\}$$

We now proceed to build the sets of precisifications for each standpoint.

Let us take F_{S_i} for every $R^m_{s_i} \in R^m$ in order of the newly assigned indices. Then, for every formula $\neg\psi \in F_{S_i}$, there is a precisification $\pi_\psi \in R_s i$ such that $(\mathcal{M}^\phi, \pi_\psi) \models \neg\psi$.

We now specify the set of precisifications for each standpoint, as follows. There are two cases:

$$R^m_{s_i} \;=\; \{\pi_\psi \mid \neg\psi \in F_{S_i}\} \;\;\cup\;\; \{\pi \mid \pi \in R^m_{s_h}, h \preceq i\}$$

provided this set is non empty;
otherwise, if the above set is empty, we set

$$R^m_{s_i} = \{\pi'\} \quad \text{for some arbitrary } \pi' \in R_s i$$

In the first case, the set of precisifications of standpoint s_i has one precisification π_ψ for each $(\neg\Box_{s_i} \psi) \in Sub^+(\phi)$ such that $(\mathcal{M}, \pi_\psi) \models \neg\psi$, which is sufficient to satisfy the formula. (Thus, for each subformula that is not necessary in the original model, there must be a precisificaiton where it is false in the reduced model.) In addition, it also includes all the precisifications from all sharper standpoints $s_h \preceq s_i$. The order of creation and the subsumption of previous standpoints preserve the partial order, such that whenever we calculate the set $R^m_{s_i}$, then any other standpoint subsumed by it, say $R^m_{s_h}$ such that $h \preceq i$, has already been calculated.

In the second case we ensure, by adding a precisification π' from $R^s_{s_i}$ in $R^m_{s_i}$ that, whenever the first case is an empty set, the standpoint s_i remains non empty and hence we preserve seriality. (In other words, we ensure that all standpoints contain at least one precisification, as required by the semantics.)

Now that we have specified the sets of precisifications R^m associated with each standpoint, the rest of elements of the model $\mathcal{M}^m = \langle \Pi^m, (R^m, \preceq), \delta^m \rangle$ can be defined straightforwardly, as follows:

- the set of precisifications $\Pi^m = \{\pi^\phi\} \cup \bigcup_{i=1}^{n} R^m_{s_i}$;
- δ^m is the restriction of δ to Π^m — i.e. $\delta^m(p) = \delta(p) \cap \Pi^m$;
- and the partial order $\preceq$ is the same as in $\mathcal{M}^\phi$ but restricted to R^m.

It is easy to see that the model $\mathcal{M}^m$ is such that $|\Pi^m| \leq |\phi|$, given that the sum of formulas in the sets F_s (or one for the empty F_s) is necessarily smaller than $|Sub(\phi)|$, which is in turn smaller than $|\phi|$.

We now show that the reduced model $\mathcal{M}^m$ reproduces the truth valuation of the original $\mathcal{M}^\phi$, with respect to all subformulae of ϕ. More precisely, we show that for all precisifications $\pi \in \Pi^m$ and all subformulas $\phi_{sub} \in Sub^+(\phi)$, then $(\mathcal{M}^m, \pi) \models \phi_{sub}$ iff $(\mathcal{M}^\phi, \pi) \models \phi_{sub}$. We can show this by induction on the structure of ϕ_{sub}. If ϕ_{sub} is atomic, then the specification of δ^m means that for any $\pi \in \Pi^m$, ϕ_{sub} is true at π in $\mathcal{M}^m$ if and only if it is true at π in $\mathcal{M}^\phi$. Given that both models have simple truth-functional semantics for evaluating formulae $\neg\alpha$ and $\alpha \wedge \beta$ at any given possible world, it must hold that, if the models have the same valuations for α and β at π, they must also give the same valuations for $\neg\alpha$ and $\alpha \wedge \beta$ at π. Thus, the only non-trivial case is when ϕ_{sub} is of the form $\square_{s_i} \psi$.

If $(\mathcal{M}^\phi, \pi) \models \square_{s_i} \psi$, then for all $\pi' \in R_{s_i}$ it is the case that $(\mathcal{M}^\phi, \pi') \models \psi$. Given that $R^m_{s_i} \subseteq R_{s_i}$ and that $R^m_{s_i}$ is non empty, then for any $\pi'' \in R^m_{s_i}$ it is also the case that $(\mathcal{M}^\phi, \pi'') \models \psi$, and by the induction hypothesis, $(\mathcal{M}^m, \pi'') \models \psi$. Hence, by definition it is the case that $(\mathcal{M}^m, \pi) \models \square_{s_i} \psi$.

On the contrary, suppose that $(\mathcal{M}^\phi, \pi) \not\models \square_{s_i} \psi$. This means that $(\mathcal{M}^\phi, \pi) \models \neg \square_{s_i} \psi$, and hence $\neg\psi \in F_{s_i}$. Therefore there is some $\pi' \in R_{s_i}$ such that $(\mathcal{M}^\phi, \pi') \models \neg\psi$ and, by the construction of $\mathcal{M}^m$, $\pi' \in R^m_{s_i}$. By induction we get that $(\mathcal{M}^m, \pi') \models \neg\psi$. So, again by the construction of $\mathcal{M}^m$, we know that if $\pi' \in R^m_{s_i}$ and $(\mathcal{M}^m, \pi') \models \neg\psi$ then $(\mathcal{M}^m, \pi^m) \models \neg \square_{s_i} \psi$ and hence $(\mathcal{M}^m, \pi^m) \not\models \square_{s_i} \psi$.

Finally, since $\pi^{\phi} \in \Pi^m$, we have $(\mathcal{M}^m, \pi^{\phi}) \models \phi$ iff $(\mathcal{M}^{\phi}, \pi^{\phi}) \models \phi$. Consequently, $\mathcal{M}^m$ satisfies ϕ.

$\square$

Proof of Theorem 6.

We now proceed to prove that the satisfiability problem is NP-complete giving an NP-algorithm for deciding whether a formula ϕ is $\mathbb{S}_0$-satisfiable.

Such an algorithm involves, for a formula ϕ, nondeterministically guessing a structure $\mathcal{M} \in \mathfrak{M}_{\mathbb{S}_0}$ with at most $|\phi|$ precisifications and exactly $|R^{\phi}|$ standpoints **(Part 1)**, and deterministically verifying that ϕ is satisfied in such $\mathcal{M}$ **(Part 2)**. **Part 1.**

In particular, given a formula ϕ, we guess a structure $\mathcal{M} = \langle \Pi, (R, \preceq), \delta \rangle$ where (a) Π is a set of n precisifications such that $n \leq |\phi|$, (b) R is a set of s subsets of Π such that $s = |\{i \mid \Box_{s_i} \psi \in Sub(\phi)\}|$ and (c) $\delta(\pi)(p) = \mathbf{f}$ for all $\pi \in \Pi$ and all primitive propositions not appearing in ϕ. Consequently, the nondeterministic guessing involves (1) deciding the size of Π, n, (2) the valuation $\mathbf{t}$ or $\mathbf{f}$ of the primitive propositions p appearing in ϕ at every point $\pi_{i \in \{1,\ldots,n\}}$, and (3) the membership or not of every point in each standpoint in ϕ, i.e $\forall \pi_i \in \Pi$, $\forall R_{s_j} \in R$ $\{\pi_i \in R_{s_j} \vee \pi_i \notin R_{s_j}\}$.

Such structure can be guessed in nondeterministic time $O(m^2)$ for $m = |\phi|$. In order to show this we must note that the number of primitive propositions in ϕ, p^{ϕ}, plus the number of distinct modal standpoint operators in ϕ, s^{ϕ} is clearly less than the size of ϕ (it easily follows from **Definition 15**). Therefore, given that $(p^{\phi} + s^{\phi}) \leq m$, and there is $n \leq m$ precisification, the structure can be guessed at $O(m^2)$ time.

Part 2.

We must now verify that ϕ is satisfied in such $\mathcal{M}$, hence there is a π such that $(\mathcal{M}, \pi) \models \phi$. We know by **Lemma 5** that this can be done in time $O(||\mathcal{M}|| \times |\phi|)$, for $||\mathcal{M}^{FIN}||$ the sum of the number of points π in Π and the number of pairs in the relations $R'_s \in R'$. In the worst case we have $n \times n$ relations in each R'_s and we have at most $s < m$ relations or standpoints. Consequently we know that we can verify it deterministically in at least $O(m^3)$ time.

Moreover, we can show that, because of the particularities of the models in $\mathfrak{M}_{\mathbb{S}_0}$, we can check whether ϕ is satisfied at some state $\pi \in \Pi$ in time $O(m^2)$ (see **Lemma 7** below).

Finally, by **Lemma 6**, we know that if ϕ is satisfiable then it will be satisfied by one of the models created in this way. If ϕ is not satisfiable, on the other hand, no guess will be right. Thus, we have a nondeterministic $O(m^2)$ algorithm for deciding if ϕ is satisfiable.

$\square$

Lemma 7. *There is an algorithm that, given a finite Kripke structure $\mathcal{M}$ of the class $\mathfrak{M}_{\mathbb{S}_0}$ and a formula ϕ, determines, in time $O(|\phi|^2)$, whether $(\mathcal{M}, \pi) \models \phi$ for some precisification π of $\mathcal{M}$, $|\phi|$ the size of ϕ and the upper bound of the number of standpoint relations $R_s \in R$.*

Proof.

Let $Sub(\phi) = \{\phi_1, \ldots, \phi_m\}$ so that the subformulas of ϕ are listed in order of length, with ties broken arbitrarily. Thus we have $\phi_m = \phi$, and if ϕ_i is a subformula of ϕ_j , then $i < j$. There are at most $|\phi|$ subformulas of ϕ (**Definition 16**), so we must have $m \leq |\phi|$.

We can label each precisification π in $\mathcal{M}$ with ϕ_j or $\neg\phi_j$ depending on whether ϕ_j is satisfied in $\mathcal{M}$ at π, i.e. $(\mathcal{M}, \pi) \models \phi_j$. Given that there are m subformulas and at most m states, we are at least at complexity $O(m^2)$. Now, let us see a sketch of the algorithm.

Let us take the subformulas of ϕ by order. For each subformula $\phi_j \in Sub(\phi)$, if ϕ_j is:

- An atomic proposition: for each $\pi \in \Pi$ we check if $(\mathcal{M}, \pi) \models P$ $(O(m))$.
- A formula of the form $\neg\phi_i$ or $\phi_i \wedge \phi_h$: for each $\pi \in \Pi$ we check if π is labelled with $\neg\phi_i$ or both ϕ_i and ϕ_h respectively. $(O(m))$.
- A formula of the form $\Box_s \phi_i$. We first check if all $\pi \in R_s$ (at most m) are labelled with ϕ_i. Then, if they are, we label all $\pi \in \Pi$ with $\Box_s \phi_i$. If not, we label all $\pi \in \Pi$ with $\neg \Box_s \phi_i$. $(O(2m) => O(m))$.

(Note that checking $\Box_s \phi$ formulae is more efficient than one might expect because although they depend on all precisifications within the set associated with standpoint s, the value of $\Box_s \phi$ must be the same at all precisifications, so only has to be computed once.)

Given that the size of $Sub(\phi)$ is at most m, we can easily see that the algorithm has a complexity $O(m^2)$. This is thanks to the fact that formulas of the kind $\Box_s \phi_i$ have the same valuation for all precisifications by the construction of $\mathfrak{M}_{\mathbb{S}_0}$.

$\square$

4.5 More expressive logics

We now briefly consider the syntax and semantics of more expressive standpoint logics, namely *description standpoint logic* $\mathbb{S}_{ALC}$ and *predicate standpoint logic* $\mathbb{S}_1$, as well as the semantics for a two-dimensional modal logic for standpoints and possible worlds, allowing us to express intensional and extensional meanings.

4.5.1 Description standpoint logic, $\mathbb{S}_{ALC}$

In the following, we provide a formalisation for *standpoint logic* with an underlying description logic language. Description logics are a family of languages that are (generally) more expressive that propositional logic but less than predicate logic, and that are normally interesting for their computational properties: the decidability and complexity of the reasoning problems in relation to the expressivity drives a substantial part of the research, and various languages are characterised by different sets of mathematical constructors (e.g. $\mathcal{ALC}$, $\mathcal{SHOIN}$ and $\mathcal{SROIQ}$ are common description logics). Moreover, description logics are widely used in AI, and in particular the 'de facto' language for representing formal ontologies (OWL) is based on description logics [Baa03].

Modal frameworks with an underlying description language have been proposed. The reader can refer to [BKW10] for a summary and, for related work, Lutz [LWZ08] provides a survey on modal description frameworks focusing on temporal extensions, Krieger [Kri16] integrates graded knowledge and temporal change in a modal fragment of owl (however with non-standard modal logics) and in [DNR97] a framework for autoepistemic description logics is developed.

In this section we specify the syntax and semantics of *standpoint logic* with an underlying $\mathcal{ALC}$ logic, following the lines of [BKW10].

4.5.1.1 The formal language and syntax of $\mathbb{S}_{ALC}$

The language of the description standpoint logic $\mathcal{L}_{\mathbb{S}_{ALC}}$ with $n+1$ unary modal operators $\Box_{s_1}, \dots, \Box_{s_n}, \Box_*$[1] consists of three sets:

- A set N_C of concept names: $C_0,\ C_1,\ \dots$
- A set N_R of role names: $R_0,\ R_1,\ \dots$
- A set N_O of object names: $a_0,\ a_1,\ \dots$

The concepts and roles of $\mathbb{S}_{ALC}$ can be inductively defined in the following way:

- The concept names $C_i \in N_C$ and $\top$ are concepts.
- If C and D are concepts and R is a relation, then the following are concepts: $C \sqcap D,\ C \sqcup D,\ \neg C,\ \forall R.C,\ \exists R.C,\ \Box_s C$.
- The role names R_i are roles.
- If R is a role, then $\Box_s R$ is a role.

Finally, a formula or axiom is defined as follows:

- If C and D are concepts, R is a role and a and b are object names, then $C \sqsubseteq D$, aRb and $a : C$ are atomic formulas or axioms.
- If ϕ and ψ are formulas, then the following are formulas: $\phi \wedge \psi$, $\neg\phi$, $\Box_s \phi$.

4.5.1.2 Semantics

A Kripke model for the modal description logic $\mathbb{S}_{ALC}$ is a tuple $\mathcal{M} = \langle \Pi, \langle \triangleleft, \preceq \rangle, \Delta^I, .^I \rangle$, where

- Π is a set of points or precisifications,
- $\langle \triangleleft, \preceq \rangle = \triangleleft_1, \dots, \triangleleft_n, \triangleleft_*$ is a partial order of unary accessibility relations on $\Pi \times \Pi$, ordered under the subset relation and where $\triangleleft_*$ is universal,

[1]Note that we omit the diamond operators $\Diamond_{s_i}$ here and in the rest of the section because they can be defined in terms of the box operators.

- Δ^I is a non empty set called the domain of the interpretation I, and
- $.^I$ is a function associating every precisification with an ALC-interpretation, such that it maps every object name $a \in N_O$ to an element of Δ^I and, for every $\pi \in \Pi$, every concept name $C_i \in N_C$ to a subset of the domain $C_i^{I,\pi} \subseteq \Delta^I$ and every role name $R_i \in N_R$ to a subset of relations $R_i^{I,\pi} \subseteq (\Delta^I \times \Delta^I)$, such that:
  - $\top^I = \Delta^I$ and $\bot^I = \emptyset$.
  - $(C \sqcap D)^{I,\pi} = C^{I,\pi} \cap D^{I,\pi}$,
  - $(C \sqcup D)^{I,\pi} = C^{I,\pi} \cup D^{I,\pi}$,
  - $\neg C = \Delta^I \backslash C^{I,\pi}$,
  - $(\exists R.C)^{I,\pi} = \{x \in \Delta^I |$ There is some $y \in \Delta^I$ such that $(x, y) \in R^{I,\pi}$ and $y \in C^{I,\pi}\}$,
  - $(\forall R.C)^{I,\pi} = \{x \in \Delta^I |$ For all $y \in \Delta^I$ if $(x, y) \in R^{I,\pi}$, then $y \in C^{I,\pi}\}$,
  - $(\Box_{s_i} C)^{I,\pi} = \{x \mid$ for all $\pi' \in \lhd_i$, $x \in C^{I,\pi'}\}$.
  - $(\Box_{s_i} R)^{I,\pi} = \{(x, y) \mid$ for all $\pi' \in \lhd_i$, $(x, y) \in R^{I,\pi'}\}$.

A pair $\langle \Delta^I, .^{I,\pi} \rangle$ for a $\pi \in \Pi$ is what is called in standard $\mathcal{ALC}$ semantics an interpretation I. In our case, we need different interpretation functions for the different points (precisifications) of our models, but given that the standpoint framework presented here assumes a constant domain of quantification we only need one set Δ^I.

We write $(\mathcal{M}, \pi) \models \phi$ to mean that ϕ is true at precisification π in model $\mathcal{M}$ and is defined as in standard $\mathcal{ALC}$ description logic except for the sentences of the form $\Box_{s_i} \phi$, for which:

$$(\mathcal{M}, \pi) \models \Box_{s_i} \phi \quad \text{iff} \quad (\mathcal{M}, \pi') \models A \text{ for all } \pi' \text{ such that } \pi' \in \lhd_i$$

To summarise, the former syntax and semantics for description standpoint logic are the natural extension of the logic presented so far to the description logic $\mathcal{ALC}$. For more material on this topic we refer the reader to [BHS08] for an introduction to description logics and [BKW10] for a general overview on modal extensions to description logics.

4.5.2 Predicate standpoint logic, $\mathbb{S}_1$

In this section we provide a specification of standpoint logic with an underlying first-order language, which allows us to reason about the semantic variability of terms with a more expressive language, desirable to model a variety of scenarios.

4.5.2.1 The formal language $\mathcal{L}_{\mathbb{S}_1}$

Our formal language $\mathcal{L}_{\mathbb{S}_1}$ is an extension of classical first-order calculus including numerical symbols and comparison relations ($=$ and $<$) as well as the usual boolean operators and quantifiers.

The non-logical symbols of the language are specified by a *vocabulary*, which is a tuple of the form:

$$\mathbb{V} = \langle \mathcal{N}, \mathcal{X}, \mathcal{P}, \mathcal{F} \rangle \ ,$$

where:

- $\mathcal{N}$ is a set of nominal constants,
- $\mathcal{X}$ is a set of nominal variables,
- $\mathcal{P} = (\mathcal{P}_1, \cup \ldots \cup \mathcal{P}_n \cup \ldots)$ is the set of predicate symbols, whose subsets $\mathcal{P}_n$, are the sets of n-ary predicate symbols,[1]
- $\mathcal{F} = (\mathcal{F}_1, \cup \ldots \cup \mathcal{F}_n \cup \ldots)$ is the set of function symbols, subsets $\mathcal{F}_n$, being the sets of n-ary function symbols,

4.5.2.2 Terms.

The language has two types of terms: one type refer to individual entities and the other refer to numerical magnitudes:

- $\mathcal{T}_n = \mathcal{N} \cup \mathcal{X}$ is the set of *nominal terms* of the language.
- $\mathcal{T}_m = \mathcal{T}_{\mathbb{D}} \cup \{f(\tau_1, \ldots, \tau_n) \mid \ f \in \mathcal{F}_n \wedge \tau_1, \ldots \tau_n \in \mathcal{T}_n\}$ is the set of *magnitude terms*.

[1] 0-ary predicates (propositional constants) have been omitted to simplify the presentation but could easily be added.

The set $\mathcal{T}_m$ includes the set $\mathcal{T}_{\mathbb{D}}$ of decimal numerals, as well as terms formed by applying function symbols to nominal terms, which give the value of some scalar property of an entity (e.g. height) or tuple of entities (e.g. the distance between two entities).

4.5.2.3 Atomic Propositions.

The language has the following forms of atomic proposition:

- $P(\tau_1, \ldots \tau_n)$, where $\tau_1, \ldots \tau_n \in \mathcal{T}_n$,
- $\tau_1 = \tau_2$, where $\tau_1, \ldots \tau_2 \in (\mathcal{T}_n \cup \mathcal{T}_m)$,
- $\tau_1 \leq \tau_2$, where $\tau_1, \ldots \tau_2 \in \mathcal{T}_m$

$P(\tau_1, \ldots \tau_n)$ asserts that predicate P holds of the nominal terms $\tau_1, \ldots \tau_n$ (which are named entities and/or quantified variables).

$\tau_1 = \tau_2$ is the usual equality relation, that can hold either between named entities and/or variables.

$\mathcal{L}_{\mathbb{S}_1}$ contains all atomic propositions that can be formed using the vocabulary $\mathbb{V}$.

4.5.2.4 Complex Propositions.

For any $\phi, \psi \in \mathcal{L}_S$, and $x \in \mathcal{X}$ the following complex propositions are also in $\mathcal{L}_S$:

- $\neg\phi$, $(\phi \wedge \psi)$, $(\phi \vee \psi)$, $(\phi \rightarrow \psi)$, $(\phi \leftrightarrow \psi)$ — the standard boolean propositional operators,
- $\forall x[\phi]$, $\exists x[\phi]$ — the standard first-order quantifiers,
- $\Box_s \, \phi$ meaning ϕ is true in standpoint s, i.e. in all precisifications compatible with standpoint s.

$\mathcal{L}_{\mathbb{S}_1}$ is the smallest set containing all atomic propositions and all complex propositions formed by these constructions.

4.5.2.5 Semantics

A Kripke model for our language $\mathcal{L}_{\mathbb{S}_1}$ is a tuple $\mathcal{M} = \langle \Pi, D, (R, \preceq), \rho, \delta \rangle$ where:

- Π is the set of precisifications,
- D is a non-empty set, the domain of individuals,
- $(R, \preceq)$ is a partial order of $R_{s_1}, ..., R_{s_n}, R_* \in R$ under the subset relation,
  - $R_{s_i} : (2^{\Pi}/\,\emptyset)$ is an accessibility relation mapping to a non-empty set of precisifications,
- $\rho = \rho_n \cup \rho_m$, where:
  - $\rho_n : \mathcal{T}_n \to D$ maps each nominal term to an element of the domain of individuals,
  - $\rho_m : \mathcal{T}_m \to \mathbb{Q}$ maps each magnitude term to a rational number.
- $\delta : P \times D \times \Pi \to \{\mathbf{t}, \mathbf{f}\}$. So δ maps, each n-ary predicate, n-tuple of individuals and precisification to a truth value.

4.5.2.6 Semantic Interpretation Function

With respect to an interpretation structure $\mathcal{M} = \langle \Pi, D, (R, \preceq), \rho, \delta \rangle$, formulas of $\mathcal{L}_{\mathbb{S}_1}$ are interpreted as follows:

- $[\![P(\tau_1, \ldots \tau_n)]\!]^{\pi}_{\mathbb{S}_1} = \ \delta(P, \langle \rho(\tau_1), \ldots \rho(\tau_n \rangle, \pi)),$
- $[\![(\tau_1 = \tau_2)]\!]^{\pi}_{\mathbb{S}_1} \quad = \ \mathbf{t}$ if $\rho(\tau_1) = \rho(\tau_2)$, else $= \mathbf{f}$,
- $[\![(\tau_1 \leq \tau_2)]\!]^{\pi}_{\mathbb{S}_1} \quad = \ \mathbf{t}$ if $\rho(\tau_1) \leq \rho(\tau_2)$, else $= \mathbf{f}$,
- $[\![\neg\phi]\!]^{\pi}_{\mathbb{S}_1} \quad = \ \mathbf{t}$ if $[\![\phi]\!]^{\pi}_{\mathbb{S}_1} = \mathbf{f}$, else $= \mathbf{f}$,
- $[\![\phi \wedge \psi]\!]^{\pi}_{\mathbb{S}_1} \ = \ \mathbf{t}$ if $[\![\phi]\!]^{\pi}_{\mathbb{S}_1} = \mathbf{t}$ and $[\![\psi]\!]^{\pi}_{\mathbb{S}_1} = \mathbf{t}$, else $= \mathbf{f}$,
- $[\![\forall x[\phi]]\!]^{\pi}_{\mathbb{S}_1} \ = \ \mathbf{t}$ if ($[\![\phi]\!]^{\pi}_{\mathbb{S}'} = \mathbf{t}$, for every interpretation structure $\mathcal{M} = \langle \Pi, D, (R, \preceq), \rho, \delta \rangle$, such that ρ' is identical to ρ, except that $\rho'(x)$ *may* have a different value from $\rho(x)$), else $= \mathbf{f}$,
- $[\![\Box_s\, \phi]\!]^{\pi}_{\mathbb{S}_1} \ = \ \mathbf{t}$ if $[\![\phi]\!]^{\pi'}_{\mathbb{S}_1} = \mathbf{t}$ for all $\pi' \in R_s$, else $= \mathbf{f}$,

4.5.2.7 Quantifier Axioms.

The universal quantifier satisfies its classical axioms, which are covered by axiom **C** above. Moreover, since our semantics is based on a single domain of individuals, its models will satisfy the Barcan formula:

AB $\forall x[\Box_s \phi(x)] \to \Box_s \forall x[\phi(x)]$

The single domain of quantification can be disputed. One can argue that if we change the precisification according to which the world is classified, the set of entities in the domain is likely to change. For example, under one precisification π of a vocabulary about forestry, a particular tree-covered area might form a single forest ($forest(o_1)$), whereas under another precisification π' the same area may consist of two forests separated by a band of heathland ($forest(o_2) \wedge forest(o_3)$). Hence, there are two objects (o_2 and o_3) in π' where there is a single one (o_1) in π.

However, we take the contrary approach, in which the set of entities (or candidate objects) can be regarded as the same even though their classification changes. This is consistent with a *de dicto* view of vagueness, in which it is linguistic descriptions that are vague, not the reality that they describe. In the previous example, the domain of quantification in both π and π' is the same, and consists of the three objects o_1, o_2 and o_3. Further, in π it is the case that $forest(o_1) \wedge \neg forest(o_2) \wedge \neg forest(o_3)$ and in π' it is the case that $\neg forest(o_1) \wedge forest(o_2) \wedge forest(o_3)$. Under the later view, **AB** is appropriate for the $\Box_s$ and $\Box_*$ operators.

4.5.3 Two-dimensional standpoint logics, $\mathbb{S}_{1A}$

4.5.3.1 Incorporation of necessity and contingency

The logics discussed so far in this chapter are modal in nature and consider the evaluation of truth with respect to points in a Kripke model, which are interpreted as precisifications. Hence, these points do not represent possible worlds as in other modal logics. Instead, they represent possible interpretations of the language with respect to a single state of affairs.

Beyond the possible limitations for reasoning about scenarios that require the representation of multiple states of affairs, such logics do not support the distinction between the intensional and extensional semantics of interpretations. In order to account for intensional meaning we are required to consider possible worlds and differentiate necessity from contingency. Hence, in this subsection we briefly present a two-dimensional modal logic that incorporates possible worlds and the usual operators for necessity and possibility (which we name $\Box_\circ$ and $\Diamond_\circ$ to avoid confusion with the standpoint operators). In a similar but even more expressive line [LYV16] proposes a three-dimensional semantics.

4.5.3.2 Syntax

$\mathcal{L}_{\mathbb{S}_{1A}}$ is an extension of $\mathcal{L}_{\mathbb{S}_1}$ including the complex propositions formed with the necessity operator, here named $\Box_\circ$ for a clearer distinction with the standpoint operators $\Box_s$.

- $\Box_\circ \phi$ — means ϕ is true in all possible worlds.

4.5.3.3 Semantics

A Kripke model for our language $\mathcal{L}_{\mathbb{S}_{1A}}$ is a tuple $\mathcal{M} = \langle \Pi, W, D, (R, \preceq), \rho, \delta \rangle$ where:

- Π is the set of precisifications,
- W is the set of possible worlds,
- D is a non-empty set, the domain of individuals,
- $(R, \preceq)$ is a partial order of $R_{s_1}, ..., R_{s_n}, R_* \in R$ under the subset relation,
  - $R_{s_i} : (2^\Pi / \emptyset)$ is an accessibility relation mapping to to a non-empty set of precisifications,
- $\rho = \rho_n \cup \rho_m$, where:
  - $\rho_n : \mathcal{T}_n \to D$ maps each nominal term to an element of the domain of individuals,
  - $\rho_m : \mathcal{T}_m \to \mathbb{Q}$ maps each magnitude term to a rational number.
- $\delta : P \times D \times \Pi \times W \to \{\mathbf{t}, \mathbf{f}\}$. So δ maps, each n-ary predicate, n-tuple of individuals, precisification and world to a truth value.

4.5.3.4 Semantic Interpretation Function

With respect to an interpretation structure $\mathcal{M} = \langle \Pi, W, D, (R, \preceq), \rho, \delta \rangle$, formulas of $\mathcal{L}_{\mathbb{S}_{1A}}$ are interpreted as follows:

- $[\![P(\tau_1, \ldots \tau_n)]\!]^{w\pi}_{\mathbb{S}_{1A}} = \delta(P, \langle \rho(\tau_1), \ldots \rho(\tau_n \rangle, \pi, w)),$
- $[\![(\tau_1 = \tau_2)]\!]^{w\pi}_{\mathbb{S}_{1A}} = \mathbf{t}$ if $\rho(\tau_1) = \rho(\tau_2)$, else $= \mathbf{f}$,
- $[\![(\tau_1 \leq \tau_2)]\!]^{w\pi}_{\mathbb{S}_{1A}} = \mathbf{t}$ if $\rho(\tau_1) \leq \rho(\tau_2)$, else $= \mathbf{f}$,
- $[\![\neg\phi]\!]^{w\pi}_{\mathbb{S}_{1A}} = \mathbf{t}$ if $[\![\phi]\!]^{w\pi}_{\mathbb{S}_{1A}} = \mathbf{f}$, else $= \mathbf{f}$,
- $[\![\phi \wedge \psi]\!]^{w\pi}_{\mathbb{S}_{1A}} = \mathbf{t}$ if $[\![\phi]\!]^{w\pi}_{\mathbb{S}_{1A}} = \mathbf{t}$ and $[\![\psi]\!]^{w\pi}_{\mathbb{S}_{1A}} = \mathbf{t}$, else $= \mathbf{f}$,
- $[\![\forall x[\phi]]\!]^{w\pi}_{\mathbb{S}_{1A}} = \mathbf{t}$ if ($[\![\phi]\!]^{w\pi}_{\mathbb{S}'} = \mathbf{t}$, for every interpretation structure $\mathcal{M} = \langle \Pi, W, D, (R, \preceq), \rho, \delta \rangle$, such that ρ' is identical to ρ, except that $\rho'(x)$ *may* have a different value from $\rho(x)$), else $= \mathbf{f}$,
- $[\![\Box_\circ \phi]\!]^{w\pi}_{\mathbb{S}_{1A}} = \mathbf{t}$ if $[\![\phi]\!]^{w'\pi}_{\mathbb{S}_{1A}} = \mathbf{t}$ for all $w' \in W$, else $= \mathbf{f}$,
- $[\![\Box_s \phi]\!]^{w\pi}_{\mathbb{S}_{1A}} = \mathbf{t}$ if $[\![\phi]\!]^{w\pi'}_{\mathbb{S}_{1A}} = \mathbf{t}$ for all $\pi' \in R_s$, else $= \mathbf{f}$,

4.5.3.5 Proof theory

Besides the axioms holding for the predicate standpoint operators and the common axiomatisation of the aletic operators, $\Box_\circ$ and $\Diamond_\circ$, which satisfy **S5**, it is questionable whether there are interaction axioms holding between them. In particular, [LYV16] suggests *commutativity*, such that 'Necessarily definitely ϕ iff definitely necessarily ϕ'. In our terms this would mean that $\Box_\circ \Box_s \phi \leftrightarrow \Box_s \Box_\circ \phi$. This should not extend, however, to other interactions, so the interaction between operators is not expected to be trivial. For instance the following should not hold $\Box_\circ \Diamond_s \phi \leftrightarrow \Diamond_s \Box_\circ \phi$.

In order to see this, let us consider two propositions, $\Box_\circ \Diamond_* \phi$ and $\Diamond_* \Box_\circ \phi$. Let us then consider in 4.2 some of the models that satisfy them. As we can see, the models (A) and (B) are the same for both propositions. Yet, the column (C) shows models that do not hold interchangeably. In a model where ϕ holds in $\langle w_1, \pi_1 \rangle$ and $\langle w_2, \pi_2 \rangle$ (but not in $\langle w_1, \pi_2 \rangle$ nor $\langle w_2, \pi_1 \rangle$), the proposition 'In all worlds, in some sense ϕ' $\Box_\circ \Diamond_* \phi$ holds but the proposition 'It is arguable that in all worlds ϕ' does not, for it requires that,

Proposition	Models (A)			Models (B)			Models (C)		
$\Box_\circ \Diamond_* \phi$	π/w	w_1	w_2	π/w	w_1	w_2	π/w	w_1	w_2
	π_1	ϕ	ϕ	π_1	ϕ	ϕ	π_1	ϕ	$\neg\phi$
	π_2	ϕ	ϕ	π_2	$\neg\phi$	$\neg\phi$	π_2	$\neg\phi$	ϕ
$\Diamond_* \Box_\circ \phi$	π/w	w_1	w_2	π/w	w_1	w_2	π/w	w_1	w_2
	π_1	ϕ	ϕ	π_1	ϕ	ϕ	π_1	ϕ	ϕ
	π_2	ϕ	ϕ	π_2	$\neg\phi$	$\neg\phi$	π_2	ϕ	$\neg\phi$

Table 4.2: Iteration of the standpont and necessity operators.

at least according to one precisification (π_1 in the last model of the table), ϕ holds in all worlds.

Moreover, whether the Barcan formula (**AB**) is suitable for the necessity operator $\Box_\circ$ is another issue. It is arguably unreasonable to expect the set of entities to be the same at every possible world. However, taking account of this would require a more elaborate semantics than we have given. In this thesis we do not commit to a full specification of the interaction axioms between the two dimensions of modal operators. We do believe, however, that this issue can be further investigated and that further rules of interaction may arise in closer inspection.

Chapter 5

Expressive Capabilities of Standpoint Logic

5.1 Introduction

So far, we have introduced the problem in chapter 1 and reviewed the literature, focusing on logic frameworks to model vagueness in chapter 2. We have then presented *standpoint logic*, a multi-modal logic for reasoning with multiple perspectives or standpoints on the semantics of vague terms, and specified both the theoretical interpretation of the elements of the framework (in chapter 3) and the formalisation of its syntax and semantics with different underlying logics in chapter 4, together with the proofs for soundness, completeness and complexity for the case of predicate logic.

In this chapter, we proceed to show how such logics, based on the supervaluationistic view of vagueness, can be used to reason about the different standpoints that an agent may have on the meanings of vague or semantically variable terms.

5.2 A propositional example

Let us begin by considering a set of propositions that display the perspectives of different agents (*Lewis Carroll, Brandon, Lucia* and *Judi*) with regards to the semantics of vague predicates, in particular the pair *sane/lunatic* and the notion of *being able to do logic*.

Example 5.2.1. Consider the following syllogism:

(1). [It is unequivocal that] If Alice is a lunatic then she is not fit to serve on a jury.

(2). According to Lewis Carroll, if Alice is sane then she can do Logic.

(3). Lucia and Brandon agree with Lewis Carroll's standpoint.

(4). According to Lucia, if Alice did not get distinction[1] in Knowledge Representation, then she can't do logic.

(5). According to Brandon, if Alice got distinction in Knowledge Representation then she can do logic.

[1]Distinction and merit are qualifications used in some universities in the United Kingdom. Distinction is equivalent to a numerical mark greater than or equal to 70, and merit is a numerical mark greater than or equal to 60 and smaller than 70.

(6). According to Brandon, it is indeterminate (borderline) whether Alice can do logic if she got merit.

(7). According to Judi, if Alice screams in the corridors then she is a lunatic.

(8). [It is unequivocal that] Alice screams in the corridors.

(9). [It is unequivocal that] Alice got merit.

(10). [It is unequivocal that] If Alice got merit, then she did not get distinction. *(penumbral connection)*

The former is a variation of a problem proposed by Lewis Carroll in [Car00] that includes additional standpoints. In the propositions that do not display semantic variability we add the prefix '[It is unequivocal that]' for the sake of clarity. Then, we add Brandon and Lucia's standpoints, that agree with Lewis Carroll's original judgement that being sane is a sufficient condition to be able to do logic, and Judi's, for whom it is screaming in the corridors that is sufficient for being a lunatic (not sane).

Moreover, Brandon's and Lucia's standpoints on whether Alice can do logic depend on the grades that she got in the Knowledge Representation exam. For Lucia, it is certain that if she got less than 70 (distinction) then she cannot do logic, while Brandon, more generously, judges that it is definite that with more than 70 (distinction) she can, and that it is indeterminate if she got more than 60 (merit). It must be remarked that, in this example, the applicability of the predicate *Alice-can-do-logic* (L) depends on the understanding that Lucia and Brandon have of 'being able to do logic', which may involve her having a different set of necessary skills and to a different level of expertise, and that they have the capacity to assess those. Consequently, the problem that we intend to represent relates to the semantic variation of the predicates rather than the epistemic lack of knowledge about the state of affairs, in this case, Alice's actual skills.

5.2.1 Representation

The following is a formalisation of the syllogism (**Eg 5.2.1**) in standpoint propositional logic, $\mathbb{S}_0$:

Example 5.2.2. The formalisation goes as follows,

(1). $\Box_*[\neg\mathsf{S} \to \neg\mathsf{J}]$.

(2). $\Box_{Carroll}[\mathsf{S} \to \mathsf{L}]$

(3). $s_{Lucia} \preceq s_{Carroll}$ and $s_{Bran} \preceq s_{Carroll}$

(4). $\Box_{Lucia}[\neg\mathsf{D} \to \neg\mathsf{L}]$

(5). $\Box_{Bran}[\mathsf{D} \to \mathsf{L}]$

(6). $\Box_{Bran}[\mathsf{M} \to \mathfrak{I}_{Bran}\,\mathsf{L}]$

(7). $\Box_{Judi}[\mathsf{C} \to \neg\mathsf{S}]$

(8). $\Box_{Judi}\,\mathsf{C}$

(9). $\Box_*\,\mathsf{M}$

(10). $\Box_*[\mathsf{M} \to \neg\mathsf{D}]$

We now consider the representation of the standpoints, the relations between standpoints and the notion of universal-truth.

5.2.1.1 Standpoints

Statements relative to a *standpoint* are generally represented by means of the different modal operators indexed with the corresponding identification label. Determinate propositions relative to a standpoint, such as those expressed in propositions **EG5.2.1.2** or **EG5.2.1.4**, express semantic commitments and hence they are formalised with the box operator $\Box_s$ (see **EG5.2.2.2** and **EG5.2.2.4**). For example, **EG5.2.1.2** expresses the semantic commitment that forces all precisifications contained in Carrol's standpoint to satisfy the proposition 'if Alice is sane then she can do Logic', thereby ensuring consistency within the standpoint.

One can also make non-definite assertions relative to a standpoint, such as expressing its *possible truth*, formalised with a diamond operator $\Diamond_s$, or *indefiniteness*, formalised with the $\mathcal{I}_s$ operator, as we can see in **EG5.2.2.6**. Moreover, in the formalisation of proposition (6), **EG5.2.2.6**, we see an example of nested operators, which can be further simplified with the interaction axioms **AS4** and **AS5**.

5.2.1.2 Relations between standpoints

The relations between standpoints expressed by proposition **EG5.2.1.3** are presented in our formalised example as $s_{Lucia} \preceq s_{Carroll}$ and $s_{Bran} \preceq s_{Carroll}$ (**EG5.2.2.3**), indicating that Lucia's and Brandon's standpoints are subsumed by Carrol's standpoint.

However, the reader may note that neither the lone standpoints s_{Carrol}, s_{Bran} and s_{Lucia} nor the *sharper* symbol $\preceq$ are part of the object language. Indeed, while we use such notation in **EG5.2.2.3** for the sake of convenience and expressivity, in practice, we syntactically formalise the standpoint relations by instantiating the axiom **AP** for the corresponding standpoints. Hence, the relations $s_{Lucia} \preceq s_{Carroll}$ and $s_{Bran} \preceq s_{Carroll}$ imply the following schemas:

EG5.2.2.(11). $\Box_{Carroll}\,\phi \rightarrow \Box_{Lucia}\,\phi$

EG5.2.2.(12). $\Box_{Carroll}\,\phi \rightarrow \Box_{Bran}\,\phi$

5.2.1.3 Universal standpoint

In our example, propositions **EG5.2.1.3**, **EG5.2.1.3** and **EG5.2.1.3** are preceded by 'It is universally agreed that', denoting that, rather than referring to a particular standpoint, they are propositions that hold according to all standpoints and precisifications of the language. They can hence be considered non-vague propositions given that all agents using the language agree on their truth conditions. Consequently, we use the universal standpoint $\Box_*$ to formalise these formulas, as we can see in the examples **EG5.2.2.1**, **EG5.2.2.9** and **EG5.2.2.10**.

5.2.1.4 Naked propositions

In standpoint logic, we call *naked propositions* to those that are not preceded by any standpoint operator in a formula, and hence, according to the semantics of standpoint logic, they are modelled by the current precisification. Unlike in other modal logics where one may want to have a privileged point for capturing notions such as 'the actual state of affairs' (in epistemic logic) or the 'the current world' (in some temporal logics), the standpoint logic model of vagueness does not have a notion of 'the actual precisification'. Instead, all precisifications are admissible in the same right and, consequently, the precisification that 'happens to be the current one' is arbitrary.

With that in consideration, the natural way of expressing facts in our framework involves avoiding using naked propositions when reasoning about standpoints. Instead, we use the universal standpoint modalities to state facts that are unequivocal and distinct standpoint modalities when representing judgements relevant to particular interpretations of the language, as illustrated in our propositional example.

It is possible, however, to formalise the same scenario making use of naked propositions and statements with the *determinate* operator, $\mathcal{D}_*$. We could do this by stating that a proposition (e.g. *Alice got merits*, M) is determinate (not vague) and then stating whether it holds or not in the actual world:

Example 5.2.3.

(1). $\mathcal{D}_* \mathsf{M}$

(2). M (or $\neg\mathsf{M}$)

It is easy to see that, by axiom **AT***, we can infer $\Diamond_*(\neg)\mathsf{M}$, and with the definiteness definition $\Diamond$**Df** we can get either $\Box_* \mathsf{M}$ or $\Box_* \neg\mathsf{M}$. Consequently, if it were deemed more intuitive, one could formalise unequivocal statements by asserting the determinacy of the non-vague propositions using $\mathcal{D}_* \phi$, and, subsequently, stating whether they hold or not using naked propositions (ϕ or $\neg\phi$). We are thereby relying on the fact that if they are the case in the current arbitrary precisification, then they must be the case in all the precisifications of the model.

Finally, one could easily restrict the language $\mathcal{L}_{\mathbb{S}_0}$ so that naked propositions are not allowed, in a similar vein as in [BD14] for the language of the modal logic MEL and in [PS07] for consensus logic.

5.2.1.5 Definiteness

The definiteness operators are useful for reasoning in a variety of scenarios, such as the one described in the section above. Recall that $\mathcal{D}_s \phi$ means that ϕ is either true in every precisification in standpoint s or false in every precisification of standpoint s.

Thus, definiteness allows us to infer from possibility to necessity ($\mathcal{D}_s\,\phi,\ \Diamond_s\,\phi \vdash_{\mathbb{S}_0} \Box_s\,\phi$) and from negated necessity to necessary negation ($\mathcal{D}_s\,\phi,\ \neg\Box_s\,\phi \vdash_{\mathbb{S}_0} \Box_s\,\neg\phi$).

Proof — $\mathcal{D}_s\,\phi,\ \Diamond_s\,\phi \vdash_{\mathbb{S}_0} \Box_s\,\phi$.

(1) $\Diamond_s\,\phi$ (assumption)

(2) $\mathcal{D}_s\,\phi$ (assumption)

(3) $\mathcal{D}_s\,\phi \leftrightarrow (\Box_s\,\phi \vee \Box_s\,\neg\phi)$ ($\mathcal{D}$**Df**)

(4) $\Box_s\,\phi \vee \Box_s\,\neg\phi$ (2,3)

(5) $\neg\Box_s\,\neg\phi$ (1, $\Diamond$**Df**)

(6) $\Box_s\,\phi$ (4,5)

$\Box$

Proof — $\mathcal{D}_s\,\phi,\ \neg\Box_s\,\phi \vdash_{\mathbb{S}_0} \Box_s\,\neg\phi$.

(1) $\neg\Box_s\,\phi$ (assumption)

(2) $\mathcal{D}_s\,\phi$ (assumption)

(3) $\mathcal{D}_s\,\phi \leftrightarrow (\Box_s\,\phi \vee \Box_s\,\neg\phi)$ ($\mathcal{D}$**Df**)

(4) $\Box_s\,\phi \vee \Box_s\,\neg\phi$ (2,3)

(5) $\Box_s\,\neg\phi$ (1,4)

$\Box$

The universal definiteness operator, $\mathcal{D}_*$, states that a proposition is 'not vague'. In other words, it either true in every precisification or false in every precisification. Hence, it is also either definitely true in all standpoints or defintely false in all standpoints of the model. In some scenarios, one may want to make assertions of this kind to separate those parts of the language that do not display semantic variability (SVoNLT).

In a similar vein to their universal counterpart, standpoint definiteness operators $\mathcal{D}_s$ assert that a predicate or proposition is not vague or semantically variable within one standpoint (i.e. across all the precisifications in the standpoint). This is useful to represent standpoints that have a sharp interpretation of generally variable terms. For instance, there may be a standpoint $s_{sharpTall}$ such that every person is judged to be either tall or not tall (with certainty with regards to that standpoint). This can occur where some legislation or institutional framework makes precise stipulations regarding the interpretation of terminology that is vague in ordinary natural language.

5.2.2 Reasoning

Let us examine now, using the example provided in section 5.2, some inferences that we can make with propositional standpoint logic. In particular, we may ask ourselves:

'*Is Alice fit to serve in a jury?*'

We can infer, from the previous sentences, that according to Lucia's and Judi's standpoints Alice is definitely not fit to serve in a jury, while according to Brandon (and hence generally) in some sense she is not fit to serve: $\square_{Lucia} \neg \mathsf{J}$, $\lozenge_{Bran} \neg \mathsf{J}$, $\square_{Judi} \neg \mathsf{J}$ and hence $\lozenge_{*} \neg \mathsf{J}$.

Proof — $\square_{Lucia} \neg \mathsf{J}$.

(1) $\square_{*} \mathsf{M} \rightarrow \square_{*} \neg \mathsf{D}$ (**EG5.2.2.10**, **AK**)

(2) $\square_{*} \neg \mathsf{D}$ (1, **EG5.2.2.9**)

(3) $\square_{Lucia} \neg \mathsf{D}$ (2, **AP**)

(4) $\square_{Lucia} \neg \mathsf{D} \rightarrow \square_{Lucia} \neg \mathsf{L}$ (**EG5.2.2.4**, **AK**)

(5) $\square_{Lucia} \neg \mathsf{L}$ (3, 4)

(6) $\square_{Lucia}[\mathsf{S} \rightarrow \mathsf{L}]$ (**EG5.2.2.2**, **AP**)

(7) $\square_{Lucia}[\neg \mathsf{L} \rightarrow \neg \mathsf{S}]$ (6, rw)

(8) $\square_{Lucia} \neg \mathsf{L} \rightarrow \square_{Lucia} \neg \mathsf{S}$ (7, **AK**)

(9) $\square_{Lucia} \neg \mathsf{S}$ (5, 8)

(10) $\square_{Lucia}[\neg \mathsf{S} \rightarrow \neg \mathsf{J}]$ (**EG5.2.2.1**,**AP**)

(11) $\Box_{Lucia} \neg\mathsf{S} \rightarrow \Box_{Lucia} \neg\mathsf{J}$ (10, **AK**)

(12) $\Box_{Lucia} \neg\mathsf{J}$ (9, 11)

$\Box$

(In the following proofs we skip some obvious steps for the sake of brevity)

Proof — $\Diamond_{Bran} \neg\mathsf{J}$.

(1) $\Box_{Bran} \mathsf{M} \rightarrow \Box_{Bran} \mathfrak{I}_{Bran} \mathsf{L}$ (**EG5.2.2.6**, **AK**)

(2) $\Box_{Bran} \mathfrak{I}_{Bran} \mathsf{L}$ (1, **EG5.2.2.9**)

(3) $\Box_{Bran}[\Diamond_{Bran} \mathsf{L} \wedge \Diamond_{Bran} \neg\mathsf{L}]$ (2, $\mathfrak{I}$**Df**)

(4) $\Box_{Bran} \Diamond_{Bran} \mathsf{L} \wedge \Box_{Bran} \Diamond_{Bran} \neg\mathsf{L}$ (3, **AC**)

(5) $\Diamond_{Bran} \mathsf{L} \wedge \Diamond_{Bran} \neg\mathsf{L}$ (4, **AS5**)

(6) $\Diamond_{Bran} \neg\mathsf{L}$ (5, PL)

(7) $\Box_{Bran}[\neg\mathsf{L} \rightarrow \neg\mathsf{S}]$ (**EG5.2.2.2**, **AP**)

(8) $\Diamond_{Bran} \neg\mathsf{S}$ (6, 7, **AK**$\Diamond$)

(9) $\Box_{Bran}[\neg\mathsf{S} \rightarrow \neg\mathsf{J}]$ (**EG5.2.2.1**, **AP**)

(10) $\Diamond_{Bran} \neg\mathsf{J}$ (8, 9, **AK**$\Diamond$)

$\Box$

Proof — $\Box_{Judi} \neg\mathsf{J}$.

(1) $\Box_{Judi} \neg\mathsf{S}$ (**EG5.2.2.8**, **EG5.2.2.7**)

(2) $\Box_{Judi}[\neg\mathsf{S} \rightarrow \neg\mathsf{J}]$ (**EG5.2.2.1**,**AP**)

(3) $\Box_{Judi} \neg\mathsf{J}$ (1, 2)

$\Box$

In Chapter 4 it was proved that the satisfiability problem of propositional standpoint logic $\mathbb{S}_0$ is NP-complete, which makes it a reasonably efficient logic in the context of modal frameworks[1]. However, its lack of expressivity imposes limitations into the kind of statements that can be represented with it, for instance propositions explicitly tackling aspects of the semantic variability of certain predicates.

In the rest of this chapter, we examine the different varieties of vagueness that can be represented with more expressive underlying logics, namely description and first-order standpoint logics.

5.3 Sorites vagueness and judgements regarding graded predicates

The problem of formalising and reasoning with sorites susceptible predicates is the subject of a vast amount of debate in the philosophical literature. In the sorites variety of vagueness, as introduced in the literature review (section 2.3.2), the applicability of a vague predicate depends on the value of one or more graded properties for which there is no clear-cut threshold that determines truth or falsity, giving raise to borderline cases. Some of the examples commonly given to illustrate this are predicates such as *tall*, whose applicability depends on the property *height*, or *heap*, whose applicability depends on the number of grains (assuming an appropriate spatial arrangement). A common formulation of the semantics of graded adjectives is in terms of an uncertain threshold value defined with respect to a particular measurement scale [Cre76].

Let us consider the example of tallness to illustrate how can we formalise the semantic variation of sorites susceptible predicates with our framework *standpoint logic*, in this case with the first-order variant $\mathbb{S}_1$.

Example 5.3.1.

(1). $\square_*[\boldsymbol{height}(\mathsf{Tara}) = 186cm]$

(2). $\square_*[\boldsymbol{height}(\mathsf{Nena}) = 160cm]$

(3). $\square_{s_1}[\forall x[\mathsf{Tall}(x) \leftrightarrow \boldsymbol{height}(x) > \mathsf{t\_tall}]]$

[1] For instance, the well known framework of epistemic logics is $PSPACE$ for multiple modalities, and the underlying propositional calculus is NP itself.

(4). $\square_{s_1}[\mathsf{t\_tall} < 185cm]$

(5). $\square_{s_2}[\neg\mathsf{tall}(\mathsf{Nena})]$

(6). $s_2 \preceq s_1$

Formulae **EG5.3.1.1** and **EG5.3.1.2** express objective, non-vague facts that are taken to be true for all precisifications: the height of Tara and Nena. Following, **EG5.3.1.3** fixes the conditions of standpoint s_1 for tallness, in particular a height greater than an unspecified threshold, namely $\mathsf{t\_tall}$, whose possible values are restricted in **EG5.3.1.4**. Finally, standpoint s_2 is a sharpening of s_1 (**EG5.3.1.6**) according to which Nena is unequivocally not tall according to standpoint s_2 in **EG5.3.1.5**.

Given **EG5.3.1.3**, **EG5.3.1.1** and **EG5.3.1.4** we can see that $\square_{s_1}[\mathsf{Tall}(\mathsf{tara})]$. **EG5.3.1.6** tells us that s_2 is sharper than s_1, so by using **AP** we can infer that $\square_{s_2}[\mathsf{Tall}(\mathsf{tara})]$.

Proof. — $(\square_{s_1}[\mathsf{Tall}(\mathsf{tara})] \wedge \square_{s_2}[\mathsf{Tall}(\mathsf{tara})])$

(1) $\square_{s_1}[\forall x[\mathsf{Tall}(x) \leftrightarrow \boldsymbol{height}(x) \leq 185]$ (**EG5.3.1.3**, **EG5.3.1.4**)

(2) $\square_{s_1}[\mathsf{Tall}(\mathsf{tara})]$ (1, **EG5.3.1.1**)

(3) $\square_{s_2}[\mathsf{Tall}(\mathsf{tara})]$ (2, **EG5.3.1.6**-**AP**)

$\square$

One of the features that make *single-dimensional sorites* cases the paradigmatic examples of vagueness is that there is a total order governing the applicability of the vague predicate (on the property in which it depends), creating what is known as the 'slippery slope' effect. Consequently, statements of the kind 'person a is "more tall" than person b' are always definite, and so, if person b is tall, then we can infer that a is tall as well.

This is illustrated in the example 5.3.1, which uses the total order governing the property *height*: The fact that *Nena* is not tall (**EG5.3.1.5**) entails the narrowing of the variability of the predicate tall in the standpoint $s2$, so that its threshold can not be lower than Nena's height.

Proof. — ($\Box_{s_2}[\mathsf{t\_tall} \geq 160cm]$)

(1) $\Box_{s_2}[\neg\mathsf{Tall}(\mathsf{Nena})]$ (**EG5.5.1.1**, **EG5.3.1.5**)

(2) $\Box_{s_2}[\forall x[\neg\mathsf{Tall}(x) \leftrightarrow \neg(\boldsymbol{height}(x) > \mathsf{t\_tall})]$ (**EG5.3.1.3**, Neg)

(3) $\Box_{s_2}[\neg(\boldsymbol{height}(\mathsf{Nena}) > \mathsf{t\_tall})]$ (1, 2)

(4) $\Box_{s_2}[\neg(160cm > \mathsf{t\_tall})]$ (3, **EG5.3.1.2**)

(5) $\Box_{s_2}[\mathsf{t\_tall} \geq 160cm]$ (4, rw)

□

Finally, we must remark that, while predicate standpoint logic $\mathbb{S}_1$ offers expressivity for the representation of sorites predicates with variable thresholds of applicability, its complexity makes it an inappropriate formalism for many computational applications. In contrast, description standpoint logics $\mathbb{S}_{ALC}$ with numerical domains (see [BHS08]) are expressive enough to represent such scenarios, only giving up on modelling the threshold of applicability as a variable on its own right. Along these lines, a similar example is as follows:

Example 5.3.2.

(1). $\Box_* \, \mathsf{hasHeight}(Tina, 186)$

(2). $\Box_* \, \mathsf{hasHeight}(Nena, 160)$

(3). $\Box_{s_1}[\exists \mathsf{hasHeight}.(\geq, 185) \sqsubseteq \mathsf{Tall}]$

(4). $\Box_{s_2}[\neg\mathsf{Tall}(Nena)]$

(5). $s_2 \preceq s_1$

From the previous we can deduce that

(6). $\Box_{s_1} \, \mathsf{Tall}(Tina)$

(7). $\Box_{s_2} \, \mathsf{Tall}(Tina)$

Moreover, if we want to be able to infer that any subject shorter than Nena is also not tall from $\Box_{s_2}(Nena : \neg\mathsf{Tall})$, then we need to add further constraints. A formalisation for this goes as follows:

- $\Box_*(\mathsf{shorterThan} \equiv (\mathsf{hasHeight} \circ < \circ\, \mathsf{hasHeight}^-)$
- $\Box_*(\exists \mathsf{shorterThan}.\neg \mathsf{Tall} \sqsubseteq \neg \mathsf{Tall})$

It must be noted that the latter formula requires our description logic to allow for cyclic definitions. This can be done with a description logic with fixpoint semantics (see [CDG03, BHS08]).

5.4 Conceptual or non-numerical variation

As discussed in section 2.3.3, the semantic variability of most natural language terms gives rise to borderline cases that cannot be expressed exclusively in terms of the variation in degree of a reduced set of properties or dimensions. Instead, the lack of clarity on which attributes or conditions are essential to the meaning of a given term brings about qualitative differences in the characterisations coming from different perspectives.

For example, the scenario where we have two conceptualisations of a domain, and we have partial knowledge about how they are related is common in computational domains. Let us see a simplified example in forestry concerning the standpoints of two fictional forestry organisations, $FO1$ and $FO2$:

Example 5.4.1.

1. According to $FO1$,
   (a) Land can be classified into either Forestland or Shrubland (but not both).
   (b) Palms are not considered to be trees.
2. According to $FO2$,
   (a) Land can be classified into either Forestland, Savanah or Shrubland.
   (b) Borderline cases may occur between Forestland and Savanah and between Savanah and Shrubland, but not between Forestland and Shrubland.
   (c) Palms are classed as trees.
3. Forestland is predominantly covered by trees.
4. Land that is definitely Forestland according to $FO1$ is at least possibly Forestland according to $FO2$.

A formalisation is as follows:

Example 5.4.2 (Propositional formalisation of example 5.4.1)**.**

(1). $\Box_{FO1}[\mathsf{L} \to (\mathsf{F} \vee \mathsf{S})]$

(2). $\Box_{FO1}[\mathsf{F} \leftrightarrow \neg\mathsf{S}]$

(3). $\Box_{FO1}[\mathsf{P} \to \neg\mathsf{T}]$

(4). $\Box_{FO2}[\mathsf{L} \to (\mathsf{F} \vee \mathsf{S} \vee \mathsf{Sa})]$

(5). $\Diamond_{FO2}\,\mathsf{F} \leftrightarrow \neg\Diamond_{FO2}\,\mathsf{S}$

(6). $\Box_{FO2}[\mathsf{P} \to \mathsf{T}]$

(7). $\Box_{*}[\mathsf{F} \to \mathsf{T}]$

(8). $\Box_{FO1}\,\mathsf{F} \to \Diamond_{FO2}\,\mathsf{F}$

From the previous statements, if we also know that an area of land is (unequivocally) predominantly covered by palms ($\Box_{*}\,\mathsf{P}$), we can infer that $\Box_{FO1}\,\mathsf{S}$ and $\Box_{FO2}[\mathsf{S} \vee \mathsf{Sa}]$.

Proof — ($\Box_{FO1}\,\mathsf{S}$).

(A) $\Box_{*}\,\mathsf{P}$ (assumption)

(1) $\Box_{FO1}\,\mathsf{P}$ (A, **AP**)

(2) $\Box_{FO1}[\neg\mathsf{T} \to \neg\mathsf{F}]$ (**EG5.4.2.7**, **AP** (and transposed))

(3) $\Box_{FO1}[\mathsf{P} \to \neg\mathsf{F}]$ (follows easily from **EG5.4.2.3**, 2)

(4) $\Box_{FO1}\,\mathsf{S}$ (1,3 then **EG5.4.2.2**)

□

Proof — ($\Box_{FO2}[\mathsf{S} \vee \mathsf{Sa}]$).

(5) $\Diamond_{FO2}[\mathsf{S}]$ (3, **EG5.4.2.8**)

(6) $\neg\Diamond_{FO2}[\mathsf{F}]$ (**EG5.4.2.5**, 4)

(7) $\Box_{FO2}[\mathsf{S} \vee \mathsf{Sa}]$ (**EG5.4.2.4**, 5)

□

In cases like this, where multiple and overlapping classification systems coexist within a domain, propositional logic can be a sufficiently expressive formalism for representation and reasoning, as long as we can reason about each piece of land independently. In more complex scenarios, where either quantification or relations are necessary (for example to quantify over trees or to relate adjacent pieces of land) or where sorites parameters play a role in some or all the interpretations, more expressive logics may be necessary.

For example, we may re-examine the propositional example provided in section 5.2, where the characterisations of notions like Sane or $\mathsf{LogicCapacity}$ are qualitatively different across standpoints. The following is a slight adaptation of our previous example to illustrate this.

Example 5.4.3.

(1). It is unequivocal that no lunatics are fit to serve on a jury.

(2). According to Lewis Carroll, everyone who is sane can do Logic.

(3). Lucia and Brandon agree with Lewis Carroll's standpoint.

(4). According to Lucia, no student with less than distinction in Knowledge Representation can do logic.

(5). According to Brandon, no student with less than merits in Knowledge Representation can do logic.

(6). According to Judi, those who scream in the corridors are lunatics.

(7). According to Judi, those who have lunatic friends are lunatics.

If a more expressive logic such as $\mathbb{S}_{ALC}$ is used, then we can further reason about individual students that may have got different marks as well as about relations. Moreover, the formalisation expresses more explicitly the semantic variation of the vague terms predicates involved. A formalisation is as follows.

Example 5.4.4 (Description logic formalisation of example 5.4.3)**.**

(1). $\square_{*}[\neg\mathsf{Sane} \sqsubseteq \neg\mathsf{Jury}]$

(2). $\square_{Carroll}[\mathsf{Sane} \equiv \mathsf{LogicCapacity}]$

(3). $s_{Lucia} \preceq s_{Carroll}$ and $s_{Brandon} \preceq s_{Carroll}$

(4). $\Box_{Lucia}[\exists \mathsf{hasMark}.(<, 70) \sqsubseteq \neg \mathsf{LogicCapacity}]$

(5). $\Box_{Brandon}[\exists \mathsf{hasMark}.(>, 70) \sqsubseteq \mathsf{LogicCapacity}]$

(6). $\Box_{Brandon}[\exists \mathsf{hasMark}.(>, 60) \sqsubseteq \mathfrak{I}_{Bran}\, \mathsf{LogicCapacity}]$

(7). $\Box_{Judi}[\mathsf{ScreamsCorridors} \sqsubseteq \neg \mathsf{Sane}]$

(8). $\Box_{Judi}[\exists \mathsf{hasFriend}.(\neg \mathsf{Sane}) \sqsubseteq \neg \mathsf{Sane}]$

(9). $\Box_{*}(\mathsf{ScreamsCorridors}(Tina))$

(10). $\Box_{*}\, \mathsf{hasMark}(Tina, 78)$

(11). $\Box_{*}\, \mathsf{hasFriend}(Tina, Nena)$

(12). $\Box_{*}\, \mathsf{hasMark}(Nena, 65)$

5.5 Penumbral connections

Standpoint semantics regards precisifications as applying to the whole language. This is a strategy to prevent the meaning of related vague concepts from varying independently, as a means to ensure consistency. For instance, we do not allow for precisifications in which related concepts are given conflicting interpretations (e.g. someone can be both 'tall' and 'short') — in fact, retaining this kind of *penumbral connection* is recognised as one of the main advantages of supervaluationistic frameworks [KS96a].

Along these lines, *standpoint logic* enables us to impose penumbral connections between concepts, both applying to all possible interpretations, using the $\Box_{*}$ operator, and also from the point of view of a particular standpoint s, by means of the $\Box_{s}$ operator. An example of this can be found in our previous example 5.2.2, where the formula **EG5.2.2.10** expresses the condition that the two predicates M and D, standing for merits and distinction, are mutually exclusive over all precisifications. Similarly, the following is an extension of example 5.3.1 that includes the predicate Short and its connection to Tall.

Example 5.5.1.

(1). $\square_* \forall x[(\mathsf{Tall}(x) \leftrightarrow \neg\mathsf{Short}(x))]$ (a penumbral axiom).

(2). $\square_*[\boldsymbol{height}(\mathsf{Tara}) = 186cm]$

(3). $\square_*[\boldsymbol{height}(\mathsf{Nena}) = 160cm]$

(4). $\square_{s_1}[\forall x[\mathsf{Tall}(x) \leftrightarrow \boldsymbol{height}(x) > \mathsf{t\_tall}]$

(5). $\square_{s_1}[\mathsf{t\_tall} < 185cm]$

(6). $\square_{s_1}[\mathsf{Short}(\mathsf{Nena})]$

(7). $s_2 \preceq s_1$

(8). $\square_{s_2}[\mathsf{t\_tall} = 180cm]$

All precisifications must now satisfy the penumbral connection axiom **EG5.5.1.1**: someone cannot be both Tall and Short. Thus, **EG5.5.1.1** will be true in all precisifications and will force the threshold for tallness to be higher than $160cm$ (the proof is identical to that of example 5.3.1 after inferring $\square_{s_1}[\neg\mathsf{Tall}(\mathsf{Nena})]$ from the penumbral axiom **EG5.5.1.1** and **EG5.5.1.6**). Finally, for the example 5.3.2, in description standpoint logic, we can formalise the same penumbral connection with:

(1). $\square_*(\mathsf{Tall} \sqsubseteq \neg\mathsf{Short})$

(2). $\square_*(\mathsf{Short} \sqsubseteq \neg\mathsf{Tall})$

5.6 Generality, context and ambiguity

In chapter 2, we introduced the concept of the SVoNLT and the notions that can be subsumed in it. Moreover, in section 2.3.5, we discussed phenomena that is sometimes confused with the semantic variability or vagueness of natural language terms, but which we differentiate from it. In this section, we go back to these phenomena and discuss the possibilities and limitations of the *standpoint framework* to represent scenarios in which they play a role.

5.6.1 The role of context

In this work we have taken the frequent stance that vagueness (or more generally the SVoNLT) is an independent phenomena from *context-dependence*, even though they often co-occur. One of the common reasons to justify this position is the prevalence of the sorites paradox even within a well delimited context:

> Fix on a context which can be made as definite as you like (in particular, choose a specific comparison class, e.g. current professional American basketball players): 'tall' will remain vague, with borderline cases and fuzzy boundaries, and the sorites paradox will retain its force. [KS96a]

Yet, contextualism has been proposed as a philosophical approach to vagueness, using indexicality to provide an account of the sorites paradox. A general introduction to the contextualist approach to vagueness can be found in [Sor18]. An interesting example is that of Shapiro, who incorporates the concept of *open-texture*[1] to his view in contextuality [Sha06]. According to him, borderline cases are predications that are dependent on the judgement of the speaker: 'they come out true in virtue of the speaker judging them to be true'. Moreover, competent speakers of the language are aware that such borderline cases could be judged differently by other agents. In this work, Shaphiro proposed the use of Kleene's three-valued logic [Kle52, Kle38] in order to model vague predicates.

While the *standpoint logic* is not developed to implement contextualist theories of vagueness, it is easy to see that it can accommodate, to some extent, the representation of scenarios where contextuality is understood along these lines. We could do this by linking each standpoint to a pair of agent and context, whose contents are the *open-textured* (not fully or precisely defined) interpretations of terms, including any judgements on borderline cases. With *standpoint logic* we could then reason with different arrangements of standpoints, belonging to pairs or sets of agents and contexts, for instance with the set of contexts attached to one agent or with the set of agents attached to one context.

In order to illustrate this, let us recover the propositional example of the Carrol sylogism (example 5.2.2). We know that, *according to Judi, those who scream in the*

[1] The term, due to Waismann, denotes 'the inability of certain concepts to be fully or precisely defined or of regulations to be exhaustive and leave no room for interpretation.'[Bla08]

corridors are lunatics. Yet, we may want to specify the context, hence having: *according to Judi when she works in the School of Computing, those who scream in the corridors are lunatics*. We could model this with the intersection of the two relevant standpoints, $\square_{Judy \cap WSoC})$[1]. For an example of more elaborate modelling with agents and contexts, we may want to include another judgement: *according to Judi when she goes to a rugby match, those who scream in the corridors are possibly not lunatics*, which can be formalised with $\square_{Judy \cap rugby}$. Further, if we intersect the standpoint of *working in the School of computing* (s_{WSoc}) to the standpoints of Lucia and Brandon, we can then also reason about the set of judgements from the general s_{WSoc} standpoint in opposition to others. The combination operations are explored in more detail in section 5.7.

Finally, further work could be done in extensions explicitly formalising contexts and or agents associated with standpoints, rather than only allowing for custom standpoint labels. However, we are not pursuing further such a line of research.

5.6.2 Generality

As we have seen in section 2.3.5.1, we do not regard generality as contributing to the SVoNLT. Instead, in this thesis *generality* is understood as a relation between the semantics of terms or interpretations rather than a property of a term itself (as in 'p is a general predicate'). On the one hand, we can express that a term is more general than another: i.e. 'this month' is more general than 'today', as we would normally do in classical logic. On the other hand, we can also express that a standpoint or perspective is more general than another with the partial order of standpoints.

It must be noted that, whenever we say that a standpoint is more general than another ($s_1 \preceq s_2$), we mean that s_1 has the same or more borderline cases than s_2, and those apply to the whole language rather than to a term in particular (generality holds between standpoints here, not between terms). In contrast, when we say that a term is more general than another, it is normally meant that the more precise one is subsumed by the more general one, so that whenever the general one holds, so does the more precise one. This can be represented in the standard way with any of the underlying logic used in a *standpoint framework*, such as propositional, e.g.: $\square_* \, \mathsf{Today} \rightarrow \mathsf{ThisMonth}$, or first-order logic, e.g.: $\square_* \, \forall x \, [\mathsf{OakForest}(x) \rightarrow \mathsf{Forest}(x)]$.

[1] Here we have used the intersection operation, which will be described subsequently in section 5.7. We could, alternatively, provide a custom tag $s_{Judy@WSoC}$

5.6.3 Ambiguity

Last but not least, we consider the formalisation issues regarding ambiguous predicates. We recall, from section 2.3.5.3, that ambiguity is understood here as homonymy.

A traditional test to detect ambiguity, *conjunction reduction*, can be used to illustrate the reasons why the phenomenon can not be directly represented in standpoint semantics. For this example, let us consider two sentences, 'The colours are light' and 'The feathers are light'. It would be plausible for an agent in a particular context to make both judgements without it being regarded as contradictory or nonsensical; hence, both of them could be stated from a particular standpoint on the SVoNLT, say $s1$. However, if both sentences are formalised using the same predicate, say $\mathsf{Light}(\mathsf{x})$, then, from '[According to $s1$] The colours are light' and '[According to $s1$] The feathers are light', the following can be inferred: '[According to $s1$] The colours and the feathers are light'. This is what is called *zeugma* in the linguistic literature, which refers to a certain absurdity of the merged meaning. Hence, it is not desirable for a logic theory to make this kind of inference.

One intuition would be to presume that the agent may be taking two different standpoints in the same context, one for the tone of colours and one for the weight. While this would prevent the framework from deducing a zeugmatic fact, it would imply that each precisification can only formalise either colour lightness or weight lightness, but not both. This is not ideal. In contrast, we suggest that the most sensible solution for many scenarios of this kind is to represent homonym terms as different predicates, say $\mathsf{LightColour}(\mathsf{x})$ and $\mathsf{LightWeight}(\mathsf{x})$.

Finally, we acknowledge that, in some cases, the line between polysemy and ambiguity can be blurry. In such cases, the strategy for formalisation (as SVoNLT or as ambiguity) must be chosen in relation to the context and intended use. When the terms are homonyms with distinct meanings, the standpoint framework does not offer more tools for reasoning than those provided by the underlying logics that it uses. The strength of *standpoint logic* is its capacity to represent different standpoints regarding the meanings of terms, while respecting penumbral connections and the self-consistency of the interpretations. This is irrelevant in the case of ambiguity, where the different senses are not at all semantically connected.

5.7 Standpoint relations and combinations

When working with systems with a variety of standpoints, it is often useful to reason not only about what can be inferred from individual standpoints, but also about the relations that hold between them and the consequences of these relationships. For that purpose, we will examine some significant relations and combinations of standpoints.

5.7.1 Relations

Our framework supports the relatively straightforward analysis of the relations that hold between standpoints. In particular, we identify four main logic relations between the standpoints, roughly analogous to set-theoretic relations, to be inferred by the system with respect to a formula ψ (typically consisting of our knowledge base). These are:

1. *Equivalence*: $s_1 \equiv s_2$. The set of precisifications of s_1 is the same as the set of s_2 at all models. Hence, for all ϕ, then $\psi \models_{\mathbb{S}_0} \Box_{s_1} \phi$ if and only if $\psi \models_{\mathbb{S}_0} \Box_{s_2} \phi$.
2. *Subsumption*: $s_2 \subseteq s_1$. The set of precisifications of s_2 is a subset of the set of s_1 at all models. Hence, for all ϕ, if $\psi \models_{\mathbb{S}_0} \Box_{s_2} \phi$ then $\psi \models_{\mathbb{S}_0} \Box_{s_1} \phi$.
3. *Disjointness*: $s_1 \preceq *\backslash s_2$. The set of precisifications of s_1 is disconnected from the set of s_2 at all models. Hence, their theories are contradictory. This can be verified with the addition of the schema $\Box_{s_1} \phi \leftrightarrow \Box_{s_2} \phi$ to ψ, if then ψ becomes $\mathbb{S}_0$-inconsistent.
4. *Overlap*: *None of the previous (NP)* relations hold between s_1 and s_2. Thus, at least some models will display an overlap between the sets of s_1 and s_2.

Relations between standpoints show, to some degree, the connections among them. While overlap is the most common scenario, other relations such as subsumption and disjointness provide valuable information to the agent, to the point of inter-operation becoming trivial (such as when information is linked to a standpoint that is subsumed by the agent's standpoint) or not feasible because of explicitly conflicting commitments (in the case of disjoint standpoints). Moreover, the previous relations, in particular subsumption, may hold for a formula ψ even if the standpoints have not been said to satisfy the partial order explicitly. In this sense, the relation is discoverable.

Combinations		Logic definitions				
		$s_1 \equiv s_2$	$s_2 \subseteq s_1$	$s_1 \subseteq s_2$	$s_1 \subseteq *\backslash s_2$	NP
Union	$s_1 \cup s_2$	s_1, s_2	s_1	s_2	$s_1 \cup s_2$*	$s_1 \cup s_2$
Intersection	$s_1 \cap s_2$	s_1, s_2	s_2	s_1	$\emptyset$	$s_1 \cap s_2$
Complement	$s_2 \setminus s_1$	$\emptyset$	$\emptyset$	$s_2 \backslash s_1$*	s_2	$s_2 \backslash s_1$

Table 5.1: The combinations (columns) between s_1 and s_2 and the result with respect to the relations holding between them.

It must be noted that the former relations are logical, i.e. they constrain the possible arrangements of standpoints in the models but they do not fix them completely. Consequently, the relations can not be inferred from the set-theoretical counterpart relations that may hold between the standpoints in a single model of ψ. For instance, if s_1 and s_2 overlap, it may be that in a model $\mathcal{M}_1$ the sets R_{s_1} and R_{s_2} are disconnected while in another model $\mathcal{M}_2$ they may overlap. Similarly, if s_1 is subsumed under s_2, in a model $\mathcal{M}_3$ the sets R_{s_1} and R_{s_2} may be equal while in another model $\mathcal{M}_4$ R_{s_1} may be smaller than R_{s_2}. This occurs particularly in reduced models such as those constructed in the complexity section 4.4.2. This is because standpoints correspond to sets of precisifications in the models $\mathcal{M}_i$, but there are different possible models/sets, even for the same set of precisifications Π.

5.7.2 Combinations

In order to support the additional specification of custom standpoints from existing ones, we define some basic combinations between standpoints. Let us consider the models $\mathcal{M}$ of a formula ϕ in the standpoint class of models $\mathfrak{M}_{\mathbb{S}_0}$.

1. *Union:* $s_1 \cup s_2$. For all models $\mathcal{M}$, for all the precisifications $\pi \in (s_1 \cup s_2)$ then either $\pi \in s_1$ or $\pi \in s_2$ (or both). Hence, $s_1 \preceq s_{1\cup2}$ and $s_2 \preceq s_{1\cup2}$, and for all formula ψ, then $\phi \models_{\mathbb{S}_0} \Diamond_{s_{1\cup2}} \psi$ if and only if $\phi \models_{\mathbb{S}_0} \Diamond_{s_1} \psi$ or $\phi \models_{\mathbb{S}_0} \Diamond_{s_2} \psi$.
2. *Intersection:* $s_1 \cap s_2$. For all models $\mathcal{M}$, for all the precisifications $\pi \in (s_1 \cap s_2)$ then $\pi \in s_1$ and $\pi \in s_2$. Hence, $s_1 \succeq s_{1\cup2}$ and $s_2 \succeq s_{1\cup2}$, and for all formula ψ, then $\phi \models_{\mathbb{S}_0} \Box_{s_{1\cap2}} \psi$ if and only if $\phi \models_{\mathbb{S}_0} \Box_{s_1} \psi$ and $\phi \models_{\mathbb{S}_0} \Box_{s_2} \psi$.

3. *Difference:* $s_2 \setminus s_1$. For all models $\mathcal{M}$, for all the precisifications $\pi \in (s_2 \backslash s_1)$ then $\pi \in s_2$ and $\pi \notin s_1$. Hence, $s_2 \succeq s_{2\backslash 1}$, and for all formula ψ, if $\phi \models_{\mathbb{S}_0} \Box_{s_2} \psi$ then $\phi \models_{\mathbb{S}_0} \Box_{s_{2\backslash 1}} \psi$ (**AP**), otherwise, if $\phi \not\models_{\mathbb{S}_0} \Box_{s_2} \psi$ and $\phi \models_{\mathbb{S}_0} \Box_{s_1} \psi$ then $\phi \models_{\mathbb{S}_0} \Box_{s_{2\backslash 1}} \neg\psi$.

Table 5.1 shows the analysis of these combinations with respect to the relations holding between s_1 and s_2, which allows for the simplification of the result in some scenarios and shows the combination to be trivial or the empty set ($\emptyset$) in others. In the table, two cases are marked with a star (*), the union of two disjoint standpoints and the complement of a standpoint s_1 subsumed in s_2. Albeit legal, both combinations are potentially problematic and the agent should be aware of the explicit inconsistencies between s_1 and s_2 in the former, and the subsumption of s_1 in s_2 in the latter.

Moreover, we must recall that whenever the operation results in the empty set, the calculated standpoint is not legal in the system.

5.8 Standpoint normal form

While modal logics are central formalisms for multi-agent systems, the intractability of major reasoning tasks limits their applications. Normal forms are sub-languages of logical languages consisting of formulae with a limited range of structures. Typically they are chosen so that every formula is equivalent to some formula in the normal form. Normal forms can be useful for simplifying proof procedures.

We finish this chapter by providing a normal form in which all formulas in $\mathcal{L}_{\mathbb{S}_0}$ can be expressed. Several normal forms have been proposed for modal logics in general, and for multi-agent modal logics in particular. For instance, a disjunctive normal form is defined in [BFM10] for a (single-agent) modal logic satisfying **S5**. Moreover, the prime implicate normal forms (PINFs) [Bie08] have been proposed for the description logic *ALC*, which is a syntactic variant of a multi-agent modal logic **K**, and the cover disjunctive normal forms (CDNFs) [CCMV06] support single-modal logics **KD45** and **S5** as shown in [HFD11] and have been extended for a multi-agent modal logic **KD45** in [HFD12].

In what follows, we extend the single-agent CDNFs in [HFD11] to the standpoint modalities, and we name the extension Cover logic Multi-modal Prenex Normal Form (CMPNF). For this purpose we will first introduce a generalisation of the prenex normal form to the multiple modalities of standpoint logic.

Definition 18 (Multi-modal prenex normal form)**.** A formula in multi-modal prenex normal form is specified by the following abstract syntax:

$$\alpha = \delta \quad \text{or} \quad \alpha = (\alpha' \vee \alpha'')$$

That is, a formula α is either a disjunction of formulas $\alpha' \vee \alpha''$ or a term δ such that:

$$\delta = r \quad \text{or} \quad \delta = \Box_s r \quad \text{or} \quad \delta = \Diamond_s r \quad \text{or} \quad \delta = (\delta' \wedge \delta'')$$

where r is a propositional formula.

Lemma 8. *Every formula in $\mathbb{S}_0$ is equivalent to a formula in multi-modal prenex normal form.*

Lemma 9. *We have the following equivalences in $\mathbb{S}_0$:*

- $\Box_s(\phi \vee (\psi \wedge \Box_{s'} \beta)) \leftrightarrow (\Box_s(\phi \vee \psi) \wedge \Box_{s'} \beta) \vee (\Box_s \phi \wedge \neg \Box_{s'} \beta))$
- $\Box_s(\phi \vee (\psi \wedge \Diamond_{s'} \beta)) \leftrightarrow (\Box_s(\phi \vee \psi) \wedge \Diamond_{s'} \beta) \vee (\Box_s \phi \wedge \neg \Diamond_{s'} \beta))$

Proof. This is proven in [MH04] for every formula in a logic **S5** with a single modality (hence s and s' are the same) and extended to **KD45** in [HFD12]. It is easy to see that with the stronger standpoint interaction axioms **AS4** and **AS5** instead of **4** and **5**, the exact same proof holds for $\mathbb{S}_0$. □

Proof of Lemma 8. A proof that every formula in **S5** is equivalent to a formula in prenex normal form is provided in [MH04]. Such a proof applies also to the case of a **KD45** single modal logic, as shown in [HFD11], and also to *standpoint logic* following **Lemma 9** and the multi-modal theorems $\vdash_{\mathbb{S}_0} \Box_{s'}\Box_s \phi \rightarrow \Box_s \phi$ and $\vdash_{\mathbb{S}_0} \Box_{s'}\neg\Box_s \phi \rightarrow \neg\Box_s \phi$ instead of $\vdash \Box\Box\phi \rightarrow \Box\phi$ and $\vdash \Box\neg\Box\phi \rightarrow \neg\Box\phi$. □

Next, we introduce the Cover logic Multi-modal Prenex Normal Form.

Definition 19 (Cover logic multi-modal prenex normal form)**.** A formula in Cover logic Multi-modal Prenex Normal Form (CMPNF) is specified by the following abstract syntax:

$$\alpha = \delta \wedge \bigwedge_{s \in S} \nabla_s \Gamma \quad \text{or} \quad \alpha = (\alpha' \vee \alpha'')$$

where δ is a propositional formula, Γ is a set of propositional formulae and the cover operator ∇ is defined as:

$$\nabla_s\Gamma \;\equiv_{def}\; \Box_s \bigvee_{\gamma\in\Gamma} \gamma \wedge \bigwedge_{\gamma\in\Gamma} \Diamond_s \gamma$$

Lemma 10. *Every formula in $\mathbb{S}_0$ is equivalent to a formula in Cover logic Multi-modal Prenex Normal Form.*

This has been proven for the prenex normal form in [HFD11] and applies in exactly the same way to the multimodal case as shown in [HFD12]. In the latter paper, additional work is done in order to account for nested modalities. However, this is not necessary in the case of *standpoint logic* as we have previously shown.

Chapter 6

Applications in the Forestry Domain

In the thesis so far we have reviewed the problem of the semantic variability of natural language, introduced the *standpoint logic framework*, presented the syntax and semantics of the logic and generally illustrated its expressivity to represent and reason about different aspects of the SVoNLT, such as with graded predicates and in the presence of penumbral connections.

In this chapter, we engage in a case study in the domain of forestry, with the aim of demonstrating some applications of *standpoint logic* in a real world scenario. In particular, we hope to provide both direct insight into how to apply our framework to address some of the challenges that the SVoNLT presents for the forestry domain, but also more generally we hope to illustrate its possibilities in the broader field of Geographical Information Systems (GIS) and other domains where it becomes relevant to reason within different points of view.

It is well known that forest definitions have a fundamental impact on the measurement and reporting of forest dynamics and processes, such as global changes in forest extension and occurrences of deforestation and degradation. Moreover, with ecological and climate concerns being at the top of the current international agenda, interdisciplinary collaboration and data dissemination and interoperation are key to gain better understanding of the state and dynamics of our forests.

Moreover, it must be noted that the particular case of the term 'forest' is interesting in several aspects. On the one hand, there is an actual debate that has motivated researchers to discuss appropriate definitions and the SVoNLT is a recognised challenge for practitioners that make use of forestry data. The political, ecological, environmental and economic implications of *what a forest is* highlights the importance of handling the variability. On the other hand, forests are geographical objects and, as such, they display particular features that make them challenging to model and particularly affected by vagueness in many aspects, as reflected in the wide literature in the domain of Geographic Ontology.

Our case study is focused on the formal representation of the standpoints of various institutions on the semantics of 'forest' (and related concepts such as 'tree') and the reasoning that can be done with them. We first provide a short literature review on the issues and challenges derived from the SVoNLT. There is a relatively wide body of research on this topic, ranging from the general discussion and assessment of the impact of the vagueness of 'forest' in the multidisciplinary scale, to specific case studies of the benefits and disadvantages of using different definitions in a particular area and for a particular purpose. We provide a broad overview and point at relevant literature in section 6.1.1, and subsequently, we select two prominent public resources to derive knowledge from: the forest data repository Global Forest Watch (GFW) and the environmental ontology EnvO, both of which are discussed in section 6.1.2. We then proceed to present our case study, which focuses on the reported need for analysing the interaction between a subset of forest definitions and their associated data at a global and local level in section 6.1.3.

In order to address the scenario and use cases considered in section 6.1.3, we consider representational aspects of geographical objects in 6.2. We differentiate between the functions of *classification*, *individuation* and *demarcation*, that comprise our interpretation of predicative terms, and we apply them to the interpretations of the term 'forest' relative to the standpoints under consideration. We then present in section 6.3 the full formalisation of all the standpoints relevant to our case studies, and we proceed, in section 6.4, to dissect some application scenarios for the framework where reasoning about the aforementioned standpoints serves to address several use cases in the domain. We finish the chapter with an analysis of the opportunities of the framework in relation to formal ontologies in section 6.5 and a conclusion.

6.1 Preliminaries

6.1.1 Forest Definitions

There are many definitions of forest in the literature [Lun07, CBL16] (more than 600 were reported in [Lun07]), which have been specified for different purposes, thereby leading to very different estimates [Gra08].

Traditionally, two main categories have been discussed: *land cover* and *land use* definitions [Lun02]. While the former characterise forest in terms of the ecological layer and the physical characteristics of the land, the latter do it concerning the purpose to which the land is put to use by humans [MT94]. Definitions favouring one or another (or both) approaches, together with other relevant features, are linked to different perspectives and management objectives, the most relevant ones being timber management, conservation of ecosystems, increasing carbon stocks and landscape restoration [CBL16]. For example, definitions used for the analysis of carbon stocks generally focus on land cover features but ignore aspects like connectedness or distinctions between natural or planted forests, because they are not relevant to describe the carbon potential. The opposite happens when defining forest for landscape restoration purposes, where together with land use information, they become crucial aspects for understanding the effects on ecosystem services and forest-based livelihoods. These differences are linked to scale and disciplinary compartmentalisation and, while justified by the specific needs of the purposes for which the definitions were created [CBL16], they pose limitations on the construction of global knowledge [Gra10] and data interoperability.

Moreover, beyond the semantic variability of the concept itself, reflected in a wide variety of specifications, most definitions of forest (and other geographic features), both in the academic literature and in administrative regulations, are not 'precise'[1] [HTL14, TSG09], which questions consistency even within a particular research community or monitoring project. This limits the understanding of data and may impair management decisions or distort research findings.

Awareness of these issues exists and there is an extensive literature reporting it, both in academia [BRP57, BW88, Lun07] and in policy [HP03, MJP15]. However, a global agreement on the meaning of such words does not seem reachable, as the variation is relevant for different contexts of application. Part of the community has focused on precisely defining 'forest' and other natural resource terms for different purposes. Other research has focused on examining the reasons for definitional problems. Many papers have advocated the need to accommodate a variety of definitions [CBL16, Gra10] and have also pointed to consequent challenges, particularly for data integration and multidisciplinary work [Gra10]. This chapter follows these lines and proposes the use of the *standpoint framework* to enable the coexistence, analysis and comparison of formal knowledge and information referring to different partial characterisations of 'forest'.

6.1.2 Resources

Efforts on measuring the location, extent and evolution of forests and other geographical features have contributed to both the increase of available data and to the development

[1]Logically one would rather say that they are partial, in that they do not commit, for instance, to the semantics of secondary concepts used for the characterisation, nor to aspects that do not play an important role in the intended applications of the definition.

and formalisation of semantic infrastructures to represent the domain. In order to illustrate the applications of *standpoint logic* in this context, we will construct our examples using well established definitions in the literature in Forestry and two prominent infrastructures, Global Forest Watch and the Environment Ontology [BMS+13], so that the representations developed in this chapter refer to data structures and ontology concepts that are publicly available and widely used in research.

6.1.2.1 Global Forest Watch

Global Forest Watch (GFW) is an online platform that distributes and visualises data produced by different institutions about the world's forests [Wor02]. It publishes annual maps of tree cover, tree cover loss and tree cover gain derived from Landsat satellite observations [HPM+13], a collection of global and local land cover and land use maps (e.g. [RAW+13] for Indonesia) and instant deforestation alert maps among others. In this chapter, we will look at the representations and inferences that we can make using the data structures and definitions from [HPM+13] and [RAW+13]. These studies will be described in the scenarios **S1** (section 6.1.3.1) and **S2** (section 6.1.3.2) respectively.

GFW can be seen as a tool that facilitates the use and access to forest data beyond the scientific community to include decision-makers from governments, companies, and civil society organisations [CSH+18]. However, the fact that each dataset responds to a different conceptualisation of 'forest' hinders the capacity of its users to interpret and integrate the information as a whole.

6.1.2.2 Ontologies and semantic data: EnvO

The Environment Ontology, EnvO, is an ontology conceived to specify a wide range of environments relevant to multiple life science disciplines and to accommodate the terminological requirements of all those needing to annotate data using ontology classes [BMS+13]. Especially since its expansion in the representation of habitats and environmental processes in [BPL+16], EnvO has become one of the main ontologies formalising the domain of forestry at a reasonable level of detail. We shall note that EnvO is expressed in OWL, based on a description logic language.

However, as one might expect, in the process of defining a general-purpose ontology, addressing the semantic heterogeneity of some terms may be challenging, as is the case with 'forest'. Indeed, the specific need for adding several forest characterisations in

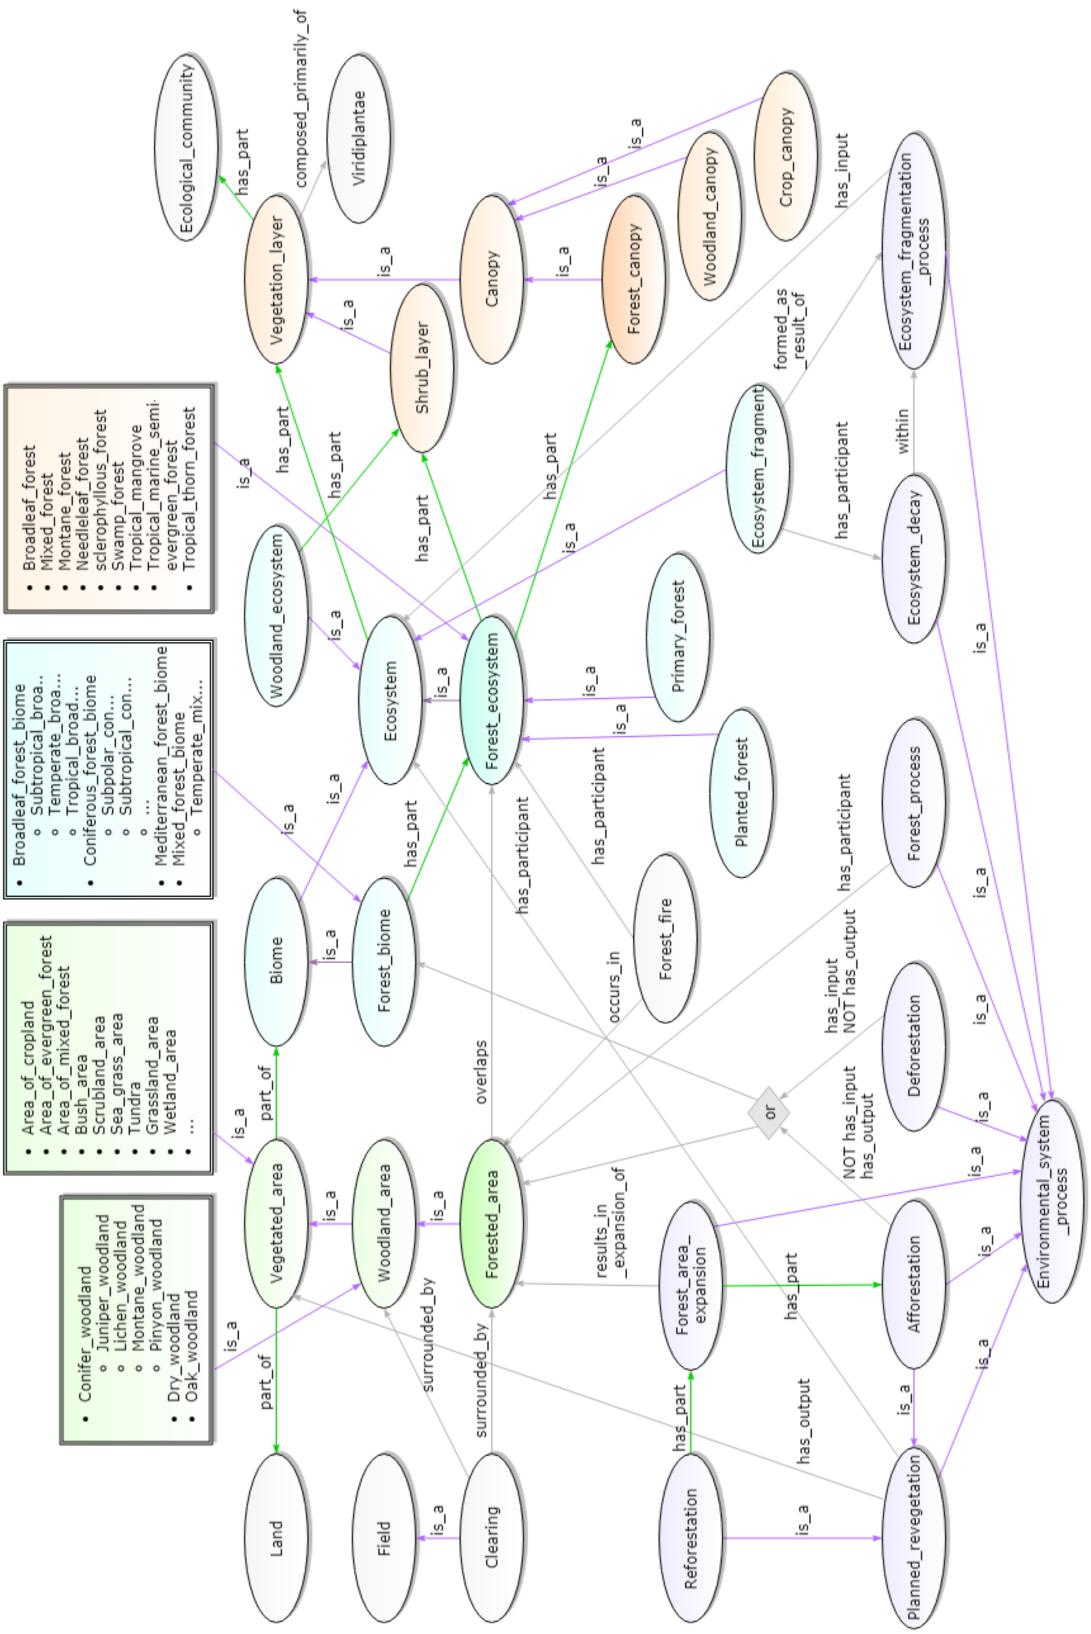

Figure 6.1: General overview of the main concepts and relations in EnvO on the domain of forestry.

EnvO was noted in [BPL+16] (in a similar vein to the more elaborate analysis in [CGV11][1]), but it does not seem to have been addressed to date. In that respect, our use of *standpoint logic* with the concepts of EnvO in our examples can be seen as a strategy to address this issue.

As for October 2019, forests in EnvO are represented via two main classes, namely *forested area* and *forest ecosystem*. *Forested area* has 'forest' as a related synonym and links to the forest entry of Wikipedia among other database cross-references. It is hence the 'de facto' concept for forest. Moreover, a *Forested area* is said to overlap with a *forest ecosystem*, for which the same textual definition is provided. The latter concept is reportedly intended at characterising the communities of plants that constitute the ecosystem, rather than its spatial extent. Finally, additional classes subsumed in *vegetated area* seem to refer to forests (e.g. *area of evergreen forest*), yet they are not explicitly related to the class *forested area*. A general overview of the main concepts and relations in EnvO on the domain of forestry can be seen in figure 6.1.

Across this chapter we will model forest standpoints using EnvO concepts and, like EnvO itself, relations from the Relation Ontology (RO). The reader is encouraged to refer to the figure 6.1 for a better understanding and context to the representations provided here.

6.1.3 The Scenarios

We finish this section by narrowing our attention to two scenarios which have been reported in the literature, and around which our representation examples and use cases revolve. They will be referred as **(S1:Global scenario)** and **(S1:Indonesia scenario)** in the rest of the chapter.

6.1.3.1 S1: Global scenario

In the first place, we consider the quantification of global forest extent, degradation and forest gain, and how the divergences in the estimates produced by both academic

[1]Carrara [CGV11] holds that, for family resemblance concepts (conceptual or non-numerical vagueness), including the representations of a variety of meanings or interpretations of the same term may be useful for the users of an ontology, and favours a 'descriptive strategy' that consists on producing a series of formalisations of the primary meanings of a term to deal with its semantic variability, as if there were a series of homonymic terms in the ontology.

and policy institutions challenge the acquisition of knowledge about the current state of affairs.

It has been recently claimed in different pieces of research both that 'global forest cover' has increased and that 'forests' have decreased: Song et al. [SHS+18] show that 'contrary to the prevailing view that forest area has declined globally, tree cover has increased by 2.24 million km2 (+7.1% relative to the 1982 level)'. On the other hand, the NYDF 2019 Report on the New York Declaration on Forests [SSR19] states that 'on average, an area of tree cover the size of the United Kingdom was lost every year between 2014 and 2018'. While these are in appearance contradictory statements, several factors explain the extreme divergences in the highlighted estimates. One that plays a major role is the different semantics attributed to the term forest and its variants (such as forest cover) and the consequent differences between the actual phenomena that are examined in those pieces of research.

The work in [SHS+18] uses Vegetation Continuous Fields (VCF), consisting of percentages of tree canopy (TC) cover, short vegetation (SV) cover and bare ground (BG) cover, to represent the land surface. Trees are defined as all vegetation taller than 5 meters in height. TC refers to the proportion of the ground covered by the vertical projection of tree crowns. In related work ([HPM+13]), the category of *tree cover* is generally attributed to land with a minimum of 30% of canopy cover, although estimates are provided for several thresholds. While both [SHS+18] and [HPM+13] clarify that TC is not equivalent to 'forest', they use the terms 'forest cover' and 'tree cover' interchangeably across the research.

Conversely, in the NYDF 2019 report [SSR19] it is noted that, although the definitions used vary by government, organisation, and intended use, they must satisfy that a *forest* is 'an area of land of minimum 0.5-1 hectares with a tree cover density of 10-30%, where trees have potential to reach a minimum height of 2-5 meters at maturity in place'.

6.1.3.2 S2: The case of REDD+ in Indonesia

For our second scenario, we narrow our attention to Indonesia. We examine a piece of research that analyses the impact of the use of several forest conceptualisations on the estimation of forest emission levels and the distribution of the drivers of deforestation in the context of REDD+ monitoring[1] [RAW+13]. Three forest definitions are considered, the FAO definition, the current national definition of the Ministry of Environment and Forestry of Indonesia (MoFI) and the 'natural forest definition', which is often employed as the preferred forest definition by conservation agencies.

[RAW+13] provides the mappings to classify the areas associated to different land cover types published by the MoFI (and available in GFW) into different classes, most notably *forested area* and *not forested area*, according to each of the considered definitions. By formally representing this, we explore how *standpoint logic* can be used to infer new knowledge from their analysis.

6.1.3.3 Set of Standpoints and definitions

We now establish the following set of standpoints with the semantic commitments based on the definitions and constraints about 'forest' and related terms that can be found in the aforementioned publications:

D1.1. s_{Song} - Tree canopy: proportion of the ground covered by the vertical projection of tree crowns.

D1.2. s_{Song} - Tree: all vegetation taller than 5 meters in height.

D1.3. s_{Hansen} - Tree/forest cover: Area with tree canopy greater than a percentage, by default 30%.

D1.4. s_{NYDF} - Forest: area of land of minimum 0.5-1 hectares with a tree canopy of 10-30 percent.

D1.5. s_{NYDF} - Tree: trees have potential to reach a minimum height of 2-5 meters at maturity in place.

D2.1. s_{FAO} - Forest: area of land of minimum 0.5 hectares with a canopy cover of more than 10 percent. [2].

D2.2. s_{FAO} - Tree: trees have potential to reach a minimum height of 5 meters at maturity in place.

[1] REDD+ is a program developed by the UN that aims at diminishing, halting and reversing forest cover loss and carbon emissions in developing countries

[2] The paper uses the FAO definition from 2000 [RAW+13], which has been updated in [MJP15]

D2.3. s_{FAO} - Land-use categories: forest land, cropland, grassland, wetlands, settlements and other land. Land categories are mutually exclusive.

D2.4. s_{MoFI} - Forest: vegetation cover dominated by intertwined tree crowns with canopy cover of more than 60%.

D2.5. s_{MoFI} - Bush: Vegetation coverage dominated by trees, with 25-60% canopy cover

D2.6. s_{MoFI} - Shrubland as shrubs with height of more than 0.5 m and more than 25% coverage.

D2.7. s_{NFD} - Forest: Excludes tree plantations.

6.2 Representation of forests

In this section, we tackle the problem of representing the semantic variability of the term 'forest', first in a broad way, taking into account both the challenges reported in the forestry literature and building on the research of the ontological issues of geography, and then proceeding to provide representations for the scenarios considered in section 6.1.3, namely **(S1:Global scenario)** and **(S2:Indonesia scenario)**.

In our study of the particular characteristics of the geographical domain, we consider three main aspects that may be present in the representations of the different interpretations of 'forest' encoded by standpoints, namely the classification, individuation and demarcation of geographical objects. We analyse them and provide formal examples implementing some of the standpoints considered in section 6.1.3.3. Finally, we present a formal representation of the set of standpoints from section 6.1.3.3.

6.2.1 Representation of the geographic domain

Given the complexity of the geographic space, its special characteristics [Ege93] and the variety of 'things' it can include [SM98, Var01], much discussion has been raised when trying to answer concisely the question of what is a geographic concept [TK04]. Some of the most interesting characteristics pointed by ontologists and geographic information scientists include location, topology, boundaries and mereology. It is surprising thus that, in 'forest' definitions, little attention has been paid to those aspects.

Among the most challenging issues affecting forest representations are the following:

Aspects of forest concept definitions	
1. Classification	
1.1 Qualitative characteristics (of the whole object)	A typical example would be the land use
1.2 Presence (or absence) of features	E.g. roads, trees of more than 5m, shrubs, ...
1.3 Density, uniformity and scale of features	Canopy cover should be measured not only in terms of the density of trees, but also in terms of the uniformity and or scale, given that the predicate can be applied to regions with different characteristics.
1.4 Location restrictions	Some definitions are contextualised in one area, like tropical forest
2. Individuation	
2.1 Morphological restrictions.	Such as shape or minimal area
2.2 Metrical restrictions	We may want to evaluate the proximity of constituents
2.3 Topological restrictions	Is the forest necessarily self-connected? Does the forest have holes?
2.4 Mereological restrictions	Is the forest the same forest (whole) if it loses one part?
2.5 Rough location	Part of the identity of the object is linked to its geographical rough position
3. Demarcation	
3.1 Fine grained threshold	Determines the precise boundary of the forest
3.2 Fuzzy threshold?	We may allow for fuzzy boundaries

Table 6.1: Compilation of features of forest definitions from different sources.

- *The dichotomy of the object-field representations.* In order to represent geographic phenomena, ontologies have to encapsulate not only the meanings linked to specific concepts but also the way these meanings are handled and represented [Aga05]. Thus, a precisification of a concept such as forest must embed information about its mode of specification, typically either in terms of an object model (as in the case of the FAO definition) or a field (as in the *Hansen* definition and data).

- *Individuation criteria.* How are entities such as mountains, rivers and forests individuated within a landscape? Although the possession of a boundary is one mark of individuality, in the geographical domain boundaries give rise to a number of ontological conundrums and may themselves be difficult to individuate [CSV98].

- *Topology, mereology and location.* In the geographic space, topology is considered to be first-class information, whereas metric properties, such as distances and shapes, are used as refinements that are frequently less exactly captured [EM95]. A general theory of spatial location is necessary to relate an entity with the spatial region that it occupies and, finally, topology is crucial as mereology alone cannot account for some very basic spatial relations, such as the relationship of continuity between two adjacent objects or the relation of one thing being entirely inside or surrounding some other thing [CSV98]. For instance, a clearance is distinguished from other bare land by being entirely inside a forest or entirely surrounded by it, depending on whether the standpoint on forest considers the clearance to be part of the forest or not.

- *Scale and granularity.* A conceptualisation of the geographic space may have several levels of granularity, each of which will be appropriate for problem solving at different levels of detail [EM95]. The inclusion of such information in the formalisation of the relevant standpoint may, in some cases, allow for relevant inferences, particularly in field representations.

6.2.1.1 Characterisation of geographical objects

With the previous in mind, we have compiled a set of relevant features for forest definitions in Table 6.1, which we have grouped in three main categories, namely classification, individuation and demarcation. These categories refer to three main *purposes of* *definition* identified as particularly relevant for characterising geographical objects in [GÁB17]. It must be noted that our understanding of *classification* and *individuation* matches with the notions suggested in [GW09] for the general domain of ontology, and the additional category of *demarcation* is specific to the geographical domain.

To illustrate the three purposes, we may consider three simple questions using the term 'forest' and the different aspects of its meaning to which they relate. These questions are used as guidelines to link cognitive conceptualisations of the geographical space with relevant notions in the domain of philosophy and ontology, as well as with actual research questions around the topic of global forest monitoring.

(a) Is this land a forest? – Classification
(b) What forests are there in this region? – Individuation
(c) What area is occupied by *this* forest? – Demarcation

Question (a) should be interpreted as *Is this land of the type 'forest'?*, where the mass noun 'forestland' is typically interpreted in terms of a field conceptualisation of the geographical space. This is the implicit approach in both the papers [HPM+13] and [SHS+18] (and which is encoded in the standpoints s_{Hansen} and s_{Song}), where the global land-area is organised as a grid of land-pieces which are then systematically classified in terms of the canopy cover. It must be stressed that the entities to be classified are land parcels. Hence, a relevant standpoint for 'forest' (or rather 'forestland') for this study is exclusively concerned with fixing the characteristics that a piece of land needs to display in order to satisfy a certain classification.

Question (b), however, requires the *individuation* of forests in order to be able to count them, thus taking an object-model approach. Characterising individuation criteria of objects is hard and, as seen in Table 6.1, relies strongly on spatial (and also temporal) factors. As Chazdon exposes in her analysis of forest definitions [CBL16], identity criteria are necessary in order to characterise forest for many management objectives. For instance, for landscape restoration purposes and conservation of natural ecosystems, it is important to understand the dynamics of individual forests, and to track whether they merge, split, appear or disappear, even if the global amount of forestland remains constant. In our scenarios, identity criteria is necessary, for example, to infer what forests are there according to s_{FAO} from the data published in [HPM+13].

Question (c) asks for the demarcation (or extension) of *a forest*. In order to give a precise answer, appropriate thresholds and footprint algorithms need to be selected. Demarcation criteria are particularly relevant for assessing forest loss and gain and for generating visualisations, among others. It is often presumed that forests need to be demarcated in order to answer the question (b). However, this need not necessarily be the case: in some cases, we may be able to differentiate (and hence count) forests, without committing to their exact boundaries.

In the following sections, we analyse these functions in-depth and illustrate how can *standpoint logic* be used to provide accurate interpretations of the term 'forest' for them. As we will see, the generality of the *standpoint framework* enables us to adapt to the particularities of the geographical domain and to fix fine-grained standpoints that respond to different needs and use cases. We hope this, in turn, supports the value of the *standpoint logic* for reasoning in scenarios where rigorous reasoning is fundamental, yet different interpretations of the SVoNLT must be managed.

6.2.2 Classification

Almost all predicates, among them 'forest', incorporate some classification features. In this work, we consider that an object x is classified under a predicate ϕ if it satisfies the necessary and sufficient conditions that govern ϕ's applicability. Formally we can express the classification of the objects in a particular domain of individuals as:

$$\forall x[\phi(x) \rightarrow \Phi(x)] \text{ and } \forall x[\Psi(x) \rightarrow \phi(x)] \text{ (where } \phi \text{ does not occur in either } \Phi \text{ or } \Psi).$$

Common examples of classification tasks in the geographic domain include both the assignation of a category to an already individuated geographical object, such as classifying a particular forest f into a forest type $\mathsf{tropicalForest}(f)$ or a tree t into a species $\mathsf{oak}(t)$, and the assignation of a category to a portion of a mass term, typically a region *landpixel*, for example into $\mathsf{forestland}(\mathit{landpixel})$. The latter, which focuses on the properties that characterise whether the concept *foresthood* applies to a given land parcel, is at the basis of the analysis in [SHS+18] and [HPM+13].

This kind of characterisation does not incorporate any specification of individuality, which is not required to answer questions of the sort (a). It assumes that an appropriate division of land into parcels has already been made (e.g. as raster cells) and characterises 'forest' or 'forestland' as a mass term. Thus, the predicate is not concerned with forest *objects*. Moreover, although ϕ can be used to determine the total area of forestland over the entire domain under consideration, it does not determine the amount or the extension of individual forests, and, as it happens in section 6.1.3.1 **(S1:Global scenario)**, the total area of forestland will be different from the total area of the forests contained, for example when some parcels of forestland are isolated from any significant forest.

6.2.2.1 Representation of classification problems with *standpoint logic*

The representation of classification aspects in *standpoint logic* is considerably natural and consists of the encoding of its necessary and sufficient conditions, which may or may not be the same. For example, using an underlying description logic, we can formalise the standpoint s_{NYDF} on what constitutes a Forested_area (in the EnvO ontology) in the following way:

- $\square_*(\exists \mathsf{hasCanopy}.\ envo{:}\mathsf{Canopy} \equiv\, = 1\ ro{:}\mathsf{has\_part}.\ envo{:}\mathsf{Canopy})$
- $\square_*(\mathsf{hasCanopyRatio} \equiv \mathsf{hasCanopy} \circ \mathsf{ratio})$
- $(\square_{NYDF}\ envo{:}\mathsf{Forested\_area}) \equiv\ envo{:}\mathsf{Vegetated\_area} \sqcap \mathsf{size\_ha}.(\geq, 0.5) \sqcap \mathsf{hasCanopyRatio}.(\geq, 30)$
- $(\lozenge_{NYDF}\ envo{:}\mathsf{Forested\_area}) \equiv\ envo{:}\mathsf{Vegetated\_area} \sqcap \mathsf{size\_ha}.(\geq, 0.5) \sqcap \mathsf{hasCanopyRatio}.(\geq, 10)$

Moreover, we may want to differentiate necessary from sufficient conditions, for example:

- $\square_{NYDF}(\ envo{:}\mathsf{Forested\_area} \sqsubseteq\ envo{:}\mathsf{Woodland\_area} \sqcap \mathsf{size\_ha}.(\geq, 0.5))$
- $\square_{NYDF}(\ envo{:}\mathsf{Forested\_area} \sqsubseteq\ \exists \mathsf{hasCanopyRatio}.(\geq, 30))$
- $\lozenge_{NYDF}(\ envo{:}\mathsf{Forested\_area} \sqsubseteq\ \exists \mathsf{hasCanopyRatio}.(\geq, 10))$
- $\square_{NYDF}(envo{:}\mathsf{Vegetated\_area} \sqcap \mathsf{overlaps}.envo{:}\mathsf{Primary\_forest} \sqsubseteq envo{:}\mathsf{Forested\_area})$

In some cases, more fine-grained specifications for the classification criteria of *forestland* could be required. In the framework proposed in Table 6.1 we consider certain aspects of the classification that tend to be overlooked, even in attempts to provide concise definitions, such as [MJP15]. Following the framework, a precisification of forestland may combine, first, (1.1) qualitative attributes, such as the legal land use of the area. Then, (1.2) the presence of certain features, such as the trees and the absence of other elements such as roads or buildings (these classificatory features may, of course, make reference to other kinds of object or land cover defined in the ontology, which in turn may also be subject to issues of vagueness and of finding an appropriate individuation). Next, (1.3) the density, uniformity and scale of features. Finally, (1.4) some location restrictions (e.g. the area must be within the tropics for *tropical-forestland*(x)) may be added in order to improve contextual adaptation.

6.2.3 Individuation

The notion of individuality is fundamental in the study of ontology and essential when adopting an object model of representation. However, formally characterising the full criteria for the individuation of particular types of object tends to be extremely hard [GW09], and is not addressed in the majority of actual ontologies. In contrast, different standpoints can be adequate to adapt the individuation criteria of objects to the intended use of the conceptualisation that they refer to. Studies in Cognitive Science show that humans identify and individuate objects using at least three sources of information: spatio-temporal information, property (featural) information, and sortal information [Xu07]. Moreover, among them, spatial features such as shape are typically more salient than other properties [Xu07].

Within the philosophical literature, it is considered that individuation requires both identity and unity. The former involves distinguishing a specific instance from others (by means of a characteristic property unique to that object), and the latter requires discriminating the parts or constituents of an instance from the rest of the world.

Existence conditions differ from classification in that the latter express the necessary and sufficient conditions for an object to be an instance of a class while the former explicitly specify the necessary and sufficient conditions to infer the 'existence' of an object. Below is a constructive existential axiom that specifies that, whenever a set of conditions $\Phi(x_1, \dots, x_n)$ are satisfied for some objects of the domain $x_1, \dots, x_n$, then an object of kind K must exist in that domain, and a relation holds between the original group and the existent object $\Psi(x_1, ..., x_n, y)$.

$$\forall x_1 ... \forall x_n [\Phi(x_1, \dots, x_n) \rightarrow \exists y [K(y) \wedge \Psi(x_1, \dots, x_n, y)]]$$

$$\forall y[K(y) \rightarrow \exists x_1 ... \exists x_n[\Phi(x_1, \ldots, x_n) \wedge \Psi(x_1, \ldots, x_n, y)]]$$

The identity criteria I of a concept determine the conditions under which it can be established that two references refer to the same object, that is, the characteristics that are unique to a single specific instance [GW00].

$$\forall x \forall y[(K(x) \wedge K(y)) \rightarrow (I_k(x, y) \leftrightarrow (x = y))]$$

Finally, the notion of unity refers to the problem of describing the parts of objects and the specific conditions (UC) under which the object constitutes a whole. A general axiomatic characterisation of this, in terms of a unifying relation among the parts of a whole is given in [GW09]. In modelling the standpoints of a particular domain or type of object, it is likely that more specific unity criteria will be required.

Example 6.2.1. A forest may be regarded, by a hypothetical standpoint s_u, as a spatially connected region of forested land, which is of maximal extent (i.e. is not part of a larger spatially connected forested region).

Assuming a predicate $\mathsf{Forested}$ has been defined and applied to all parcels of forest-land, then the following assertion expresses the content of the standopint (where P is the parthood relation and SCON is the property of being spatially self connected), capturing the unity condition for a possible standpoint of forest:

$$\Box_u \forall x[\mathsf{Forest}(x) \rightarrow \mathsf{Forested}(x) \wedge \mathsf{SCON}(x) \wedge \neg\exists y[\mathsf{P}(x, y)$$

$$\wedge \neg(x = y) \wedge \mathsf{Forested}(y) \wedge \mathsf{SCON}(y)]$$

A variety of situations can hinder the specification of a unified criteria for identity and unity. Some of them are drastic evolutions of objects through time, situations in which objects merge or split and objects whose boundaries are ill defined or affected by sorites vagueness, thus creating confusion about self-connectedness and parthood. It is in these scenarios that the specification of tailored commitments associated to different standpoints becomes relevant.

Underlying axiomatisations of the space and mereotopology are key to provide appropriate notions of parthood. In some cases, a set theoretical view where two sets are the same if and only if they have exactly the same elements is appropriate to model the space. However, for most objects, a looser identity criteria that allows one to accept the continued existence of an object even after the loss of certain parts is necessary.

6.2.3.1 Individuation of geographical objects

The consideration of the individuation of forests entangles in the extensive bibliography about the ontology of geographical features, their characterisation and their boundaries. Difficulties tend to arise both regarding the unification of geographical features (e.g. deciding whether something is part or not of a forest) and their identity (e.g. deciding whether a forest now is the same forest as one that existed 100 years ago) particularly if there have been substantial changes in vegetation or location [Ben02]. Moreover, while most of the objects in the physical world have a *bona fide* boundary that acts as one of the main marks of their individuality, geographical boundaries are often fuzzy or otherwise indeterminate [CSV98], which makes the individuation even more challenging and the demarcation of most geographical objects non trivial.

A simple example of individuation relevant for our use cases is as follows:

1. $\Box_*(\exists\mathsf{hasCanopy}.\,envo\!:\!\mathsf{Canopy} \equiv = 1\, ro\!:\!\mathsf{has\_part}.\,envo\!:\!\mathsf{Canopy})$
2. $\Box_*(\mathsf{hasCanopyRatio} \equiv \mathsf{hasCanopy} \circ \mathsf{ratio})$
3. $\Box_{NYDF}[\,envo\!:\!\mathsf{Site} \sqcap \exists\mathsf{hasCanopyRatio}.(\geq, 30) \sqsubseteq \mathsf{Tree\_covered\_region}]$
4. $\Diamond_{NYDF}[\,envo\!:\!\mathsf{Site} \sqcap \exists\mathsf{hasCanopyRatio}.(\geq, 10) \sqsubseteq \mathsf{Tree\_covered\_region}]$
5. $\Box_*(\mathsf{part\_of\_tree\_covered\_area} \sqsubseteq\ ro\!:\!\mathsf{part\_of})$
6. $\Box_*(\mathsf{connected\_tree\_cover} \sqsubseteq\ ro\!:\!\mathsf{connected\_to})$
7. $\Box_*(\mathsf{Tree\_covered\_region} \sqsubseteq = 1\mathsf{part\_of\_tree\_covered\_area}.\,envo\!:\!\mathsf{Woodland\_area})$
8. $\Box_*(\exists\mathsf{connected\_tree\_cover}.\top \sqsubseteq \mathsf{Tree\_covered\_region} \sqcap$
$\forall\mathsf{connected\_tree\_cover}.\mathsf{Tree\_covered\_region})$
9. $\Box_*(\mathsf{connected\_tree\_cover} \circ \mathsf{part\_of\_tree\_covered\_area} \sqsubseteq \mathsf{part\_of\_tree\_covered\_area})$
10. $\Diamond_{NYDF}(\,envo\!:\!\mathsf{Woodland\_area} \sqcap \mathsf{size\_ha}.(\geq, 0.5) \sqsubseteq\ envo\!:\!\mathsf{Forested\_area})$
11. $\Box_{NYDF}(\,envo\!:\!\mathsf{Woodland\_area} \sqcap \mathsf{size\_ha}.(\geq, 1) \sqsubseteq\ envo\!:\!\mathsf{Forested\_area})$

This formalisation encodes the individuation for the s_{NYDF} definition, according to which a forest is a self-connected area with a minimum canopy cover and size (see section 6.1.1). It fixes the unity conditions in (6) by establishing that any area previously classified as Tree_covered_region must be part of a single *envo*: Woodland_area[1]. Subsequently, identity criteria are represented in lines (7) and (8), stating that whenever a Tree_covered_region is connected to another, then, by (8), they are parts of the same *envo*:Woodland_area (this is achieved via the composition operator $\circ$, by which $Rxy \circ Ryz \rightarrow Rxz$). Finally, the *envo*:Woodland_area is classified as a *envo*:Forested_area if its size is greater than the given threshold.

6.2.4 Demarcation

Finally, by *demarcation* we mean the act of determining the spatial extension of an object, or equivalently, of establishing its boundary. Once this extension/boundary is established, it may be referred to as 'the demarcation' of the object; or, in cases where the boundary is unclear or debatable, it may be regarded as one of many possible demarcations of the object, relative to one precisification.

Establishing an object's demarcation may be straightforward or extremely problematic. Although this may be because of the characteristics of the particular object under consideration, it is usually strongly related to the ontological category of the object. For example, physical artefacts (e.g. cups, tables) are typically easily demarcated because they consist of solid matter, forming an integral whole that is not physically connected to any other matter. On the contrary, for aggregate objects such as 'forest' demarcation is often problematic, both because it may not be clear which entities should be counted as constituents and because there is no unique way to determine the spatial extension of something that is made up of many disconnected constituents (e.g., trees and other elements of the ecosystem, or pieces of land in abstracted representations). Distinguished regions within field-like objects, which are again prevalent in geography (e.g. soil type regions) also give rise to significant demarcation problems [MC89].

Although the study of suitable algorithms for the demarcation goes beyond the scope of this chapter, it must be noted that many standpoints will be associated to different demarcation strategies for 'forests', yet the explicit semantic commitments are often limited to the establishment of of sets of thresholds (e.g. [MJP15]). A more careful analysis of some strategies for demarcation that can be used to provide more precise standpoints on the demarcation of forests can be found in [DG09], and useful

[1] In EnvO, *envo*:Forested_area (i.e. forest) is a *envo*:Woodland_area that has a size greater than a certain threshold.

ontological analysis considering the sort of data that is available for this purpose can be found in [CFW04, CFW05].

However, given that computing the demarcation of spatial objects tends to be a computationally expensive task, it is expected that non-logical systems will implement such strategies, and that the calculated areas will be subsequently associated with the relevant objects. We will, therefore, not provide formalised examples in this area. However, we consider that the additional encoding of axioms and rules regarding the demarcation of the objects that a standpoint refers to can, potentially, enhance the inferences that we can make from our data and standpoint knowledge base.

6.2.4.1 Interpreting the extension of standpoints

In the case of geographical information the spatial projection of the instances can be particularly relevant to complement the analysis of the relations that hold between standpoints. One might expect that the possible relations holding between standpoints (equivalence, subsumption, inverse subsumption, overlap and disjunction) would map to the analogous RCC5 relations (equal, proper part, inverse proper part, partial overlap and discrete) between the total spatial area covered by sets of instances satisfying each standpoint. However, as with the models of the standpoint relations, that does not necessarily need to be the case. Instead, relations between standpoints restrict the space of possibilities while leaving some variation open, and this is in fact what provides insight to the user.

In Table 6.2 we show possible relations holding in the different scenarios. While the analogous to RCC5 is the 'expected' relation to hold, other possibilities are allowed, which may be symptomatic of different circumstances. Take, for instance, the case where the spatial projection of $\square_{s_1}(\mathsf{Forest})$ is *equivalent* to that of $\square_{s_2}(\mathsf{Forest})$. It may be the case that s_1 and s_2 are equivalent. But it could also be the case that s_1 is broader (less *sharp* than) than s_2, by allowing some borderline cases which however never manifest in the available data about the state of the world (this can help, for instance, assessing what semantic commitments have a considerable impact in the interpretation of the data and which do not). In fact, it could even be the case that s_1 and s_2 describe forest in a logically different way (e.g. one uses tree proximity and the other canopy cover) that is highly correlated, and therefore both map to the same objects while only holding an *overlap* relation. In other cases this may be due to mere exemplar

Spatial ext. of $\Box_{s_i}\,\phi$	**Logic relations**			
	$s1 \equiv s2$	$s2 \preceq s1$	$s1 \preceq s2$	NP
	Expected	The sharpness of s_1 with respect to s_2 does not manifest in any instance.	The sharpness of s_2 with respect to s_1 does not manifest in any instance.	The only unequivocal instances of s_1 and s_2 lay in their intersection.
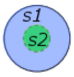	Impossible	Expected	Impossible	The only unequivocal instances of s_2 lay in its intersection with s_1.
	Impossible	Impossible	Expected	The only unequivocal instances of s_1 lay in its intersection with s_2.
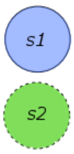	Impossible	Impossible	Impossible	Despite the standpoins logically overlapping, there are no known unequivocal instances on the intersection between them.
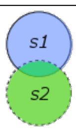	Impossible	Impossible	Impossible	Expected

Table 6.2: Relations between the logical relations between two standpoints, s_1 and s_2 (columns), and the RCC5 spatial relations holding between the projections of the instances satisfying a formula $\Box_{s_i}\,\phi$ for $i = 1$ represented in blue and $i = 2$ represented in green.

clustering[1]. Even when the expected spatial relation holds between the projections, it is expected to be useful for the agent to know how populated are the intersections or borderline scenarios of the definitions under consideration.

However, we must add a word of caution. The former type of analysis is indeed useful but can only be carried out whenever the demarcation of the spatial objects is constant across the precisification space. For instance, in our application scenario this applies to LandPixels. Conversely, whenever the demarcation (which in turn is likely to depend in the unity conditions among others) is variable, the logic relations become less straightforward and can not be interpreted by these simple observations.

6.3 Formal representation of the use cases

We now proceed to present the full representation of the set of standpoints introduced in section 6.1.3.3. We intend to maximise the use of the EnvO ontology in our formalisations (for reference, figure 6.1 reveals the main classes and relations that relate to forestry). In what follows, relations from the Relation Ontology will be preceded by the prefix 'ro:', and concepts from EnvO by the prefix '$envo$:'. Moreover, knowledge from EnvO is considered supertrue and therefore, any statement will be assumed to be preceded by $\square_*$. While this is appropriate for our use cases, it may not be adequate if we were modelling a system in which we used more than one ontology. In the latter situation, it would be more appropriate to consider that knowledge from EnvO relates to a standpoint s_{EnvO}, which some other standpoints may subsume if they are known to comply with the ontology.

Let us begin with the standpoints relevant to the scenario **(S1:Global scenario)** (section 6.1.3.1). In the first place, we have the representation used in [SHS+18], denoted by standpoint s_{Song} and described in **D1.1**. According to it, *forest cover* is measured as the ratio of canopy cover present in the terrestrial surface. Their data represents the vegetation composition of every 'land pixel' by associating it with its ratio of tree cover, short vegetation cover and bare ground.

[1] In ordinary situations, objects that exhibit one property, will very often also exhibit another property and vice versa, even though there is no necessary connection between the properties. The cause may be because of patterns and regularities that are essentially contingent [Ben05].

The relevant standpoint commitments are as follows:

R1.1. $\square_{Song}[\mathsf{LandPixel} \sqsubseteq \mathit{envo}{:}\mathsf{Site}]$

R1.2. $\square_{Song}[\mathsf{LandPixel} \sqsubseteq = 1\ \mathit{ro}{:}\mathsf{has\_part}.\ \mathit{envo}{:}\mathsf{Canopy}.\mathsf{ratio}.(\geq, 0)]$

R1.3. $\square_{Song}[\mathsf{LandPixel} \sqsubseteq = 1\ \mathit{ro}{:}\mathsf{has\_part}.\ \mathit{envo}{:}$
$(\mathsf{Vegetation\_layer} \sqcap \neg\, \mathit{envo}{:}\mathsf{Canopy}).\mathsf{ratio}.(\geq, 0)]$

R1.4. $\square_{Song}[\mathsf{LandPixel} \sqsubseteq = 1\ \mathit{ro}{:}\mathsf{has\_part}.\neg(\mathit{envo}{:}\mathsf{Vegetation\_layer}).\mathsf{ratio}.(\geq, 0)]$

Note that we use the class $\mathit{envo}{:}\mathsf{Canopy}$ rather than $\mathit{envo}{:}\mathsf{Woodland\_canopy}$ or $\mathit{envo}{:}\mathsf{Forest\_canopy}$ because those entail parthood in a Forest or Woodland ecosystem, which may or may not be the case.

Moreover, in [SHS+18], it is specified that the requirement for the vegetation cover to be considered tree cover is a height greater than or equal to $5m$.

R1.5. $s_{Song\_t} \preceq (s_{Song} \cap s_{Tree\_sh})$

R1.6. $\square_{Tree\_sh}[\mathit{envo}{:}\mathsf{Canopy} \sqsubseteq \mathit{envo}{:}\mathsf{Vegetation\_layer}.\mathsf{height\_m}.(\geq, 5)]$

Additional penumbral connections to enable further inferences could be established, for example, stating that two different LandPixels can not partially overlap or that they only have three parts, whose ratios sum up 100. This sort of knowledge is only relevant to the standpoints sharper than or equal to s_{Song} but can be used to deduce new facts from its data. We will nevertheless omit the formalisation of those rules as they are not useful for our use cases.

Departing from the same conceptual framework, Hansen et al. [HPM+13] further classify land pixels into the class *tree cover* or *forest cover* (both names are used interchangeably) whenever the canopy ratio is higher than a certain threshold, by default set to 30% and at least 10%.

R1.7. $\square_{*}(\exists \mathsf{hasCanopy}.\ \mathit{envo}{:}\mathsf{Canopy} \equiv = 1\ \mathit{ro}{:}\mathsf{has\_part}.\ \mathit{envo}{:}\mathsf{Canopy})$

R1.8. $\square_{*}(\mathsf{hasCanopyRatio} \equiv \mathsf{hasCanopy} \circ \mathsf{ratio})$

R1.9. $\square_{Hans}[\mathsf{Tree\_covered\_region} \sqsubseteq \mathsf{LandPixel} \sqcap \exists \mathsf{hasCanopyRatio}.(\geq, 10)]$

R1.10. $\lozenge_{Hans}[\mathsf{Tree\_covered\_region}] \equiv \mathsf{LandPixel} \sqcap \exists \mathsf{hasCanopyRatio}.(\geq, 10)$

We must note that the s_{Hans} definition, like most definitions, do not specify what sizes of land are admissible for their categories to be applicable. While we understand that a fair variation is admissible, if we establish the minimum standpoint semantic constraint, using the concept $envo:\mathsf{Site}$, then we could infer that vast expanses of land with zones of high canopy density are tree-covered areas as a whole. For the sake of example, the Amazon rainforest has enough canopy density such that the whole of Peru would be classed as a $\mathsf{Tree\_covered\_region}$ for a canopy ratio higher than 30%. Consequently, when implementing logic-based standpoints, these aspects, captured in the table on forest definitions, must be taken into account and added to the common-sense definitions conceived to be applied by humans. For this reason, in this case we use the specific predicate $\mathsf{LandPixel}$, for which appropriate restrictions in dimensions should be specified.

Note that, as mentioned in chapter 5, section 5.3, description logic is not expressive enough to allow for the use of variables, which is the most natural way to represent sorites thresholds. This limits what we can formally represent with regards to the s_{Hans} standpoint when using $\mathbb{S}_{ALC}$. A way to circumvent this is by creating subsumed standpoints with the desired threshold sharpenings, which could be established 'on demand' by an agent.

Following, we consider the classification criteria for the standpoint s_{NYDF}.

R1.11. $\Diamond_{NYDF}[\,envo:\mathsf{Woodland\_area} \sqcap \mathsf{size\_ha}. \geq 0.5 \sqsubseteq envo:\mathsf{Forested\_area}]$

R1.12. $\Box_{NYDF}[\,envo:\mathsf{Woodland\_area} \sqcap \mathsf{size\_ha}. \geq 1 \sqsubseteq envo:\mathsf{Forested\_area}]$

R1.13. $\Box_{NYDF}[\,envo:\mathsf{Woodland\_area} \sqsubseteq \exists\mathsf{hasCanopyRatio}.(\geq, 30)]$

R1.14. $\Diamond_{NYDF}[\,envo:\mathsf{Woodland\_area} \sqsubseteq \exists\mathsf{hasCanopyRatio}.(\geq, 10)]$

R1.15. $\Box_{NYDF}[\,envo:\mathsf{Canopy} \equiv envo:\mathsf{Vegetation\_layer.height\_m}.(\geq, 5)]$

R1.16. $\Diamond_{NYDF}[\,envo:\mathsf{Canopy} \equiv envo:\mathsf{Vegetation\_layer.height\_m}.(\geq, 2)]$

The former characterisation requires us to have a candidate $\mathsf{Woodland\_area}$ from which we can evaluate the size, to subsequently classify it as a $\mathsf{Forested\_area}$ (i.e. forest), and which can not be trivially done unless we establish individuation criteria. For this, we recover the formalisation provided and discussed in section 6.2.3:

R1.17. $\Box_{NYDF}[\,envo:\mathsf{Site} \sqcap \exists\mathsf{hasCanopyRatio}.(\geq, 30) \sqsubseteq \mathsf{Tree\_covered\_region}]$

R1.18. $\Diamond_{NYDF}[\,envo\!:\!\mathsf{Site} \sqcap \exists \mathsf{hasCanopyRatio}.(\geq, 10) \sqsubseteq \mathsf{Tree\_covered\_region}]$

R1.19. $\square_{*}[\mathsf{part\_of\_tree\_covered\_area} \sqsubseteq\ ro\!:\!\mathsf{part\_of}]$

R1.20. $\square_{*}[\mathsf{connected\_tree\_cover} \sqsubseteq\ ro\!:\!\mathsf{connected\_to}]$

R1.21. $\square_{*}[\mathsf{Tree\_covered\_region} \sqsubseteq = 1\mathsf{part\_of\_tree\_covered\_area}.\ envo\!:\!\mathsf{Woodland\_area}]$

R1.22. $\square_{*}[\exists \mathsf{connected\_tree\_cover}.\top \sqsubseteq \mathsf{Tree\_covered\_region} \sqcap$
$\forall \mathsf{connected\_tree\_cover}.\mathsf{Tree\_covered\_region}]$

R1.23. $\square_{*}[\mathsf{connected\_tree\_cover} \circ \mathsf{part\_of\_tree\_covered\_area} \sqsubseteq \mathsf{part\_of\_tree\_covered\_area}]$

Finally, for the second scenario, **(S2:Indonesia scenario)** in section 6.1.3.2, we focus our attention on the case of Indonesia. In [RAW+13], Romijn et al. analyse the different interpretations of forest that the FAO (s_{FAO}), the Ministry of Environment and Forestry of Indonesia (s_{MoFI}) and the natural forest definition (s_{NDF}) have, and their impact on the development of the REDD+ reference emission levels. In their study, they use a map of the land cover data released by the government (whose land cover base classification will be subsequently referred as *lci* :) and they explore what categories fit under the forest definition according to the different interpretations. As reported in the paper, we can formalise the following mapping between land cover classes and forest standpoints:

R1.24. $s_{FIN} = s_{(FAO\ \cup\ MoFI\ \cup\ NFD)}$

R1.25. $\square_{FIN}[(\mathsf{lci}\!:\!\mathsf{Primary\_Upland\_Forest}\ \sqcup\ \mathsf{lci}\!:\!\mathsf{Primary\_Mangrove\_Forest} \sqcup$
$\mathsf{lci}\!:\!\mathsf{Primary\_Swamp\_Forest}) \sqsubseteq$
$(\,envo\!:\!\mathsf{Forested\_area}\ \sqcap\ \exists\, ro\!:\!\mathsf{overlaps}.\ envo\!:\!\mathsf{Primary\_forest})]$

R1.26. $\square_{FIN}[(\mathsf{lci}\!:\!\mathsf{Secondary\_Upland\_Forest}\ \sqcup\ \mathsf{lci}\!:\!\mathsf{Secondary\_Mangrove\_Forest} \sqcup$
$\mathsf{lci}\!:\!\mathsf{Secondary\_Swamp\_Forest}) \sqsubseteq\ envo\!:\!\mathsf{Forested\_area}]$

R1.27. $\square_{FAO}[(\mathsf{lci}\!:\!\mathsf{Forest\_Plantation}\ \sqcup\ \mathsf{lci}\!:\!\mathsf{Shrubland}\ \sqcup\ \mathsf{lci}\!:\!\mathsf{Bushy\_swamp} \sqcup$
$\mathsf{lci}\!:\!\mathsf{Palm\_oil\_plantation}) \sqsubseteq\ envo\!:\!\mathsf{Forested\_area}]$

R1.28. $\square_{MoFI}[(\mathsf{lci}\!:\!\mathsf{Forest\_Plantation} \sqsubseteq\ envo\!:\!\mathsf{Forested\_area}]$

R1.29. $\Box_{MoFI}[(\mathsf{lci{:}Shrubland} \sqcup \mathsf{lci{:}Bushy\_swamp} \sqcup \mathsf{lci{:}Palm\_oil\_plantation}) \sqsubseteq \neg\, envo{:}\mathsf{Forested\_area}]$

R1.30. $\Box_{NFD}[(\mathsf{lci{:}Forest\_Plantation} \sqcup \mathsf{lci{:}Shrubland} \sqcup \mathsf{lci{:}Bushy\_swamp} \sqcup \mathsf{lci{:}Palm\_oil\_plantation}) \sqsubseteq \neg\, envo{:}\mathsf{Forested\_area}]$

R1.31. $\Box_{FIN}[(\mathsf{lci{:}Savanna} \sqcup \mathsf{lci{:}Upland\_agriculture} \sqcup \mathsf{lci{:}Upland\_agriculture\_mixed\_bush} \sqcup \mathsf{lci{:}Rice\_field} \sqcup \mathsf{lci{:}Fishpond} \sqcup \mathsf{lci{:}Settlement} \sqcup \mathsf{lci{:}Transportation} \sqcup \mathsf{lci{:}Open\_land} \sqcup \mathsf{lci{:}Mining} \sqcup \mathsf{lci{:}Water\_body} \sqcup \mathsf{lci{:}Swamp} \sqcup \mathsf{lci{:}Airport}) \sqsubseteq \neg\, envo{:}\mathsf{Forested\_area}]$

R1.32. $\Box_{NFD \,\cup\, MoFI}[(\mathsf{lci{:}Shrubland} \sqcup \mathsf{lci{:}Bushy\_swamp}) \sqsubseteq envo{:}\mathsf{Scrubland\_area}]$

We define the standpoint s_{FIN} as the union of the three standpoints s_{FAO}, s_{MoFI} and s_{NFD} to represent the common ground between them. Moreover, all the land cover classes of *lci* are, by definition of the Ministry of Environment and Forestry of Indonesia, definite ($\mathcal{D}_*$), mutually exclusive and collectively exhaustive (**MECE**), but we omit the formalisation of these features here for the sake of brevity (the formalisation is easy but verbose).

The advantage of departing from an already individuated dataset is that the formalisation of the standpoint commitments is substantially simplified. On the other hand, some features that determine the classification of the instances into one or another standpoint are not explicit (i.e. the actual underlying characteristics of the land), and hence the inferential power is reduced.

6.4 Use cases

In this section, we consider a set of use cases that are relevant to the scenarios **(S1:Global scenario)** and **(S2:Indonesia scenario)**, as introduced in section 6.1.3. They illustrate how we can reason and query the formal representations provided in the previous section 6.3, together with the underlying use of the resources introduced in 6.1.2.

For this purpose, let us assume that we have a *standpoint logic* endpoint from which a user can modify and query a *standpoint knowledge base* consisting of the representations in 6.3 and the EnvO ontology, and which is populated with the publicly available data in GFW. Moreover, let us also assume that we have a GIS tool that visualises the spatial projection of the elements that are retrieved from each query.

6.4.1 Use Case 1: Analysing the differences between standpoints

Given that in the global scenario **(S1:Global scenario)** described in section 6.1.3.1 both pieces of research use different data sources and methods that lead to different estimates, it may be difficult to assess to which extent the semantic differences play a role in the final quantification. Consequently, different agents working in either research or policy may want to analyse what are the semantic commitments that play a significant role on the estimates in forest cover that can be visualised in platforms such as GFW.

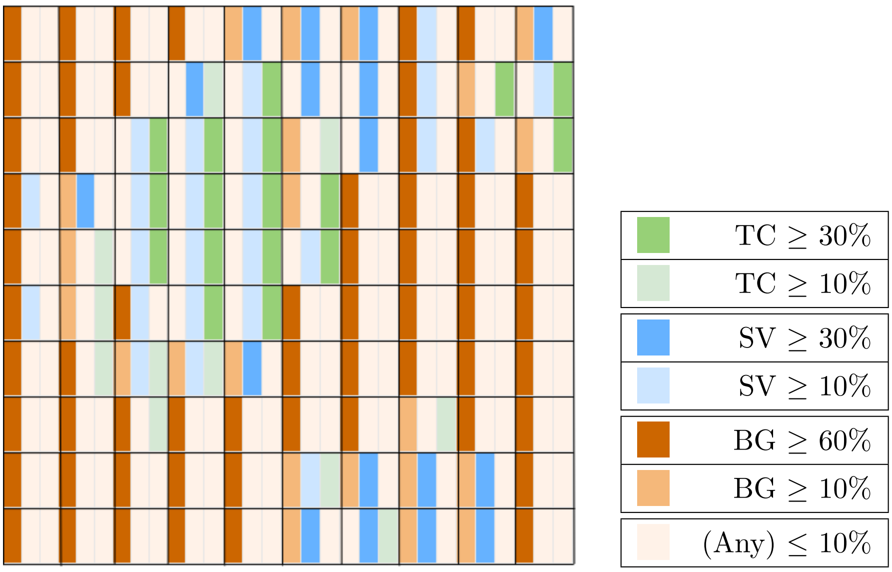

Figure 6.2: A sample area, A_α, divided in 10×10 landPixels. The cover of each landPixel is represented via three vertical stripes, indicating its composition in BG, SV and TC following the legend.

In particular the agent may want to compare the maps provided in [SHS+18], and exemplified for a hypothetical area A_α in figure 6.2, to the following three queries/visualisations:

UC1.1:- The visualisation of *Hansen*'s account of tree/forest cover (**D1.3**) at different thresholds.

UC1.2:- The visualisation of the previous if we relax the tree height condition down to $2m$, in line with **D1.5**.

UC1.3:- The visualisation if we add the minimum size requirement for an area to be a forest (i.e, if we use definition **D1.4**).

For **UC1.1**, the user may want to establish, in addition to the standpoints in 6.3, the following:

UC1.S1:- $s_{Hans\_10-30} \preceq s_{Hans}$

UC1.S2:- $\square_{Hans\_10-30}[\mathsf{Tree\_covered\_region}] \equiv$
$$\square_*[\,envo\!:\!\mathsf{LandPixel} \sqcap \exists \mathsf{hasCanopyRatio}.(\geq, 30)]$$

UC1.S3:- $s_{Hans\_20-30} \preceq s_{Hans\_10-30}$

UC1.S4:- $\Diamond_{Hans\_20-30}[\mathsf{Tree\_covered\_region}] \equiv$
$$\square_*[\,envo\!:\!\mathsf{LandPixel} \sqcap \exists \mathsf{hasCanopyRatio}.(\geq, 20)]$$

The former are the bridging commitments that the agent may want to assess. Subsequently, they simply may query:

UC1.Q1:- $[\mathsf{MO}]_{Hans\_10-30\,\cap\,Song\_t}[\mathsf{Tree\_covered\_region}]?$

UC1.Q2:- $[\mathsf{MO}]_{Hans\_10-30\,\cap\,Song\_t}[\neg\mathsf{Tree\_covered\_region}]?$

UC1.Q3:- $[\mathsf{MO}]_{Hans\_20-30\,\cap\,Song\_t}[\mathsf{Tree\_covered\_region}]?$

UC1.Q4:- $[\mathsf{MO}]_{Hans\_20-30\,\cap\,Song\_t}[\neg\mathsf{Tree\_covered\_region}]?$

Where $[MO]_s$ is a shortcut that stands for all modal operators ($\square_s$, $\Diamond_s$, $\mathcal{I}_s$ and $\mathcal{D}_s$). Hence, each query of the form $[MO]_s[\phi]$ is in fact a set of four queries, one with each of the operators. One must note that both $\square_s\,\phi Impossible \Diamond_s\,\phi$ and $\mathcal{I}_s\,\phi Impossible \Diamond_s\,\phi$.

Hence, we will only display $\Diamond_s$ in the visualisation whenever the knowledge base entails $\Diamond_s\, \phi$ but not $\Box_s\, \phi$ nor $\mathfrak{I}_s\, \phi$.

The visualisation of the queries **UC1.Q1** and **UC1.Q2** on the area A_α (whose land pixels were represented in figure 6.2), can be seen in figure 6.3. In particular we show, as specified in the lengend, which areas are unequivocally tree covered, which areas are borderline and which areas are unequivocally not tree covered according to the s_{Hans} standpoint with a tree cover threshold between the 10% and the 30%.

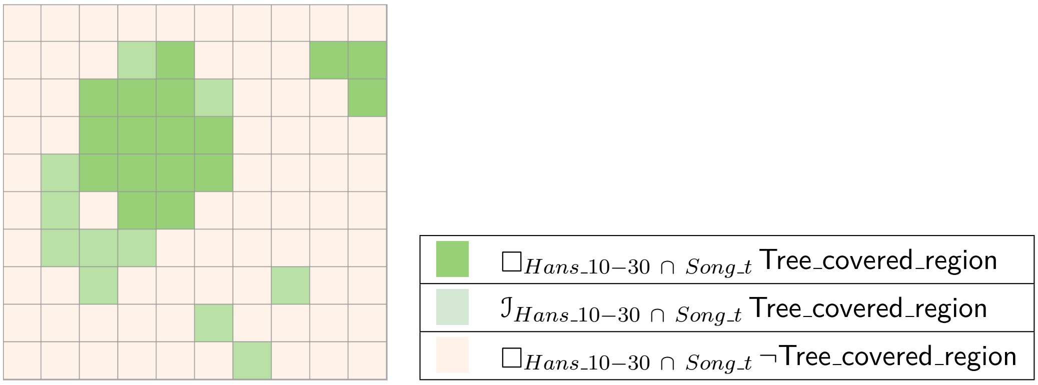

Figure 6.3: Visualisation of the areas that are unequivocally tree covered, indeterminately tree covered and unequivocally not tree covered according to s_{Hans} in A_α. The colour indicates that the corresponding fact in the legend can be derived from the østandpoint knowledge base.

For the second visualisation, **UC1.2**, the following standpoints may be established:

UC1.S5:- $s_{Song\_lt} \preceq s_{Song}$

UC1.S6:- $\Box_{Song\_lt}[\,envo\!:\!\mathsf{Canopy}] \equiv [\,envo\!:\!\mathsf{Vegetation\_layer.height\_m.}(\geq, 5)]$

UC1.S7:- $\Diamond_{Song\_lt}[\,envo\!:\!\mathsf{Canopy}] \equiv [\,envo\!:\!\mathsf{Vegetation\_layer.height\_m.}(\geq, 2)]$

Here s_{Song_lt} encodes a standpoint that is sharper than s_{Song} for which the criteria for a plant to be a tree is vague, in contrast to the previously used s_{Song_t}. In this case, according to s_{Song_lt} something is unequivocally a tree if it is higher than 5m as for s_{Song_t}, and it is in some sense a tree if it is higher than 2m.

Next, the agent can perform the following queries, with the previous pattern:

UC1.Q5:- $[\mathsf{MO}]_{Hans\_10-30\ \cap\ Song\_lt}[\mathsf{Tree\_covered\_region}]$?

UC1.Q6:- $[\mathsf{MO}]_{Hans\_10-30\ \cap\ Song\_lt}[\neg\mathsf{Tree\_covered\_region}]$?

UC1.Q7:- $[\mathsf{MO}]_{Hans\_20-30\ \cap\ Song\_lt}[\mathsf{Tree\_covered\_region}]$?

UC1.Q8:- $[\mathsf{MO}]_{Hans\_20-30\ \cap\ Song\_lt}[\neg\mathsf{Tree\_covered\_region}]$?

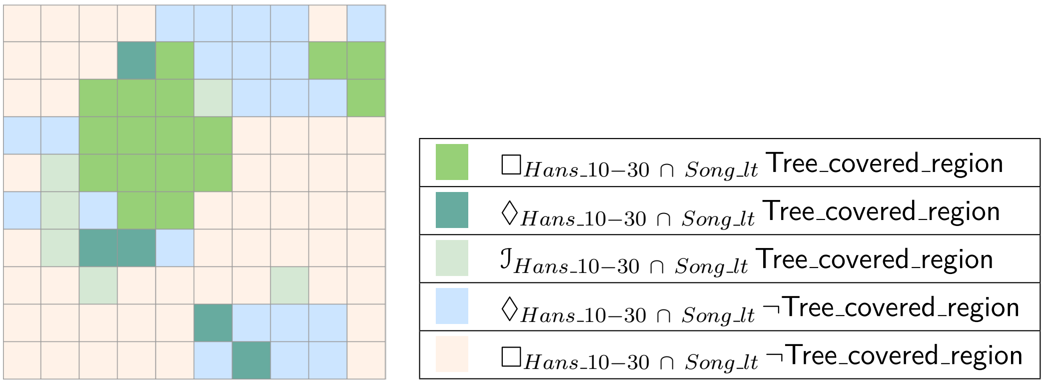

Figure 6.4: Visualisation of the areas that are tree covered according to the standpoint $s_{Hans\_10-30\ \cap\ Song\_lt}$ in the area A_α. The colour indicates that the corresponding fact in the legend can be derived from the østandpoint knowledge base.

As we can see in figure 6.4, despite the fact that figure 6.2 does not contain information about the tree cover of vegetation from $2m$ to $5m$, we can still infer a variety of facts. In particular, we use the knowledge that BG is not an $envo$: Vegetation_layer and hence it can determine unequivocally not Tree_covered_region, and those areas that have certain $envo$:Canopy for vegetation of more than $5m$ certainly retain at least that ratio for the lower threshold. In the figure 6.4 we represent the most specific modality that can be inferred. In particular it must be reminded that both necessity ($\Box_s$) and indeterminacy ($\mathfrak{I}_s$) entail possibility ($\Diamond_s$).

Finally, for the visualisation **UC1.3** we may use the standpoints s_{NYDF}, that go through the individuation of the forests in function of the forest pixels as we described in the representation section 6.2. We could, of course, consider further standpoints with different minimum sizes, canopy covers and tree sizes. For this example, we create the

s_{NYDF_lc} and s_{NYDF_mc} standpoints, with determinate thresholds on canopy cover at the 10% and 20% respectively, and the full s_{NYDF} admissible range of minimum sizes:

UC1.S8:- $s_{NYDF\_lc} \preceq s_{NYDF}$

UC1.S9:- $\Box_{NYDF}[\, envo\!:\!\mathsf{Site} \sqcap \exists \mathsf{hasCanopyRatio}.(\geq, 10) \sqsubseteq \mathsf{Tree\_covered\_region}]$

UC1.S10:- $s_{NYDF\_mc} \preceq s_{NYDF}$

UC1.S11:- $\Box_{NYDF}[\, envo\!:\!\mathsf{Site} \sqcap \exists \mathsf{hasCanopyRatio}.(\geq, 20) \sqsubseteq \mathsf{Tree\_covered\_region}]$

Together with the standard standpoint for trees, s_{Song_t}, we perform the following queries:

UC1.Q9:- $[\mathsf{MO}]_{NYDF\_lc \,\cap\, Song\_t}[\, envo\!:\!\mathsf{Forested\_area}]?$

UC1.Q10:- $[\mathsf{MO}]_{NYDF\_mc \,\cap\, Song\_t}[\, envo\!:\!\mathsf{Forested\_area}]?$

As we can see in figure 6.5, in which the area of each land pixel is assumed to be $0.25ha$, the sparse tree covered areas are unequivocally not forests after the individuation process, regardless of their individual extent of canopy cover.

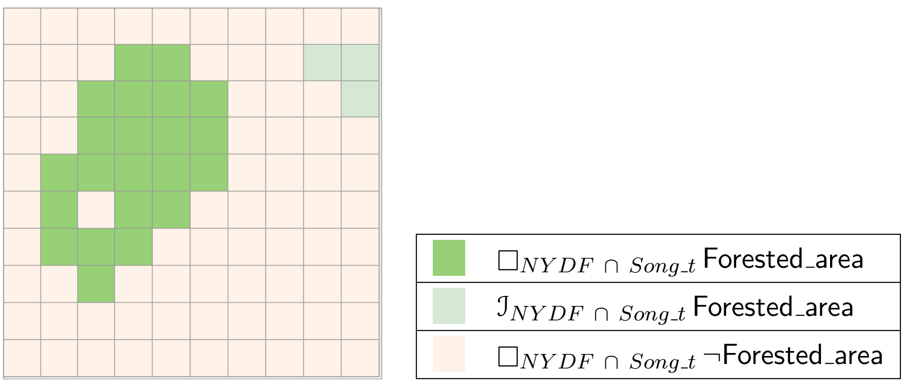

Figure 6.5: Representation of **UC1.Q9** for the query, showing its parts. The colour indicates, in this case, that the land pixels are part of, or not, of a detected forested area (as shown in the legend) and can be derived from the østandpoint knowledge base.

It must be noted that more in-depth comparative studies on the impact of using terms with different semantics can be (and are) undertaken by experts in the domain, using a variety of GIS tools and techniques. [RAW+13] is an example of that, which we have used in 6.3. However, integrating knowledge representation formalisms like *standpoint logic*, which allow for automated reasoning on semantic heterogeneity, may not only facilitate this sort of analysis but also enable a broader audience to perform them and tailor them to the specific needs of their research.

For example, the in-depth study in [RAW+13] for **(S2:Indonesia scenario)** analyses the general implications of using one or another conceptual framework on a certain scale, in this national, and concludes that the most appropriate standpoint for the **REDD** program is the adoption of the natural forest definition (s_{NFD}). However, another hypothetical team researching the livelihoods of communities in a particular region of the country may find that the results of the national study (the insight that the natural forest definition is best suited) may not be relevant to the specific geographic features of their area of interest. However, if using a standpoint framework, the latter can instead reuse the formalised standpoints, hence facilitating subsequent analysis over the same interpretations.

6.4.2 Use Case 2: Sharing knowledge derived from a standpoint

Data repositories such as GFW contain datasets produced by different institutions and provide tools for their geographic visualisation. In GFW, each dataset encodes information represented in either a field-model or in an object-model (see section 6.2.1), and can be visualised as an independent layer on a GIS map.

Moreover, if the data is linked to a semantic framework such as the EnvO Ontology, one can query and retrieve complex information derived from it. However, inferences that rely on the spatial co-occurrence of phenomena (for example the presence of forest in a region x at time t_1 and the lack of it at time t_2) could be deemed sufficient to infer that an $envo:\mathsf{deforestation\_process}$ takes place at the region x. However, this would not be adequate adequate whenever the conceptual frameworks behind the facts from x at t_1 and x at t_2 may not be compatible[1]. In this case, such inference should only be made if $\Box_{t\_1}\,\mathsf{Forest} \subseteq \Box_{t\_2}\,\mathsf{Forest}$, i.e. both the case where the conceptual framework is the same and when the latter is less strict than the former for the forest category.

[1] For instance, the data produced by the FAO is not temporally consistent, because their conceptual framework and requirements for admissibility of definitions of forest from different countries have changed within the last ten years.

With this, we want to illustrate both that it is generally desirable to share knowledge that is either represented (e.g. an area of forest) or has been inferred (e.g. a deforestation process) from one or more datasets by linking it to an ontology such as EnvO, and also that by establishing their relevant standpoints and appropriate penumbral connections, one can enhance the inferencing power of the knowledge base without compromising the generality of the concepts represented in an ontology.

A particular use case for the subsequent retrieval of the generated knowledge from establishing the standpoints s_{FAO}, s_{MoFI} and s_{NFD} in relation to the land types of (ici :) is as follows:

Let us consider a research project interested in the study of the impact of forest change on the ecological stability of sites of the Alliance for Zero Extinction (AZE)[1]. They want to target three areas to perform their case studies in Indonesia: one with an unperturbed ecosystem, one with confirmed habitat degradation and one with potential habitat degradation, which is not recognised by the Ministry of Environment and Forestry of Indonesia. In order to narrow the areas in which to look for a location for each of their case studies, they want to:

1. Retrieve the intersection of the *AZE sites* and the areas that are unequivocally[2] *primary forest* and are not experiencing *deforestation* or *degradation* processes.
2. Retrieve the intersection of the *AZE sites* and the areas that are unequivocally forests experiencing *deforestation* or *degradation* processes.
3. Retrieve the intersection of the *AZE sites* and the areas that are possibly forests experiencing *deforestation* or *degradation* processes but not according to the MoFI.

For this purpose, they may perform the following queries using the standpoint s_{FIN} and its subsumptions, as represented in section 6.3:

UC2.Q1:- $\square_{FIN}(\mathsf{AZE\_site} \sqcap \mathit{envo}{:}\mathsf{Forest\_area} \sqcap \exists \mathit{ro}{:}\mathsf{overlaps}.\mathit{envo}{:}\mathsf{Primary\_forest} \sqcap \neg\exists\, \mathit{ro}{:}\mathsf{overlaps}.(\mathit{envo}{:}\mathsf{deforestation} \sqcup \mathit{envo}{:}\mathsf{ecosystem\_decay}))$?

[1] The protection of AZE sites is a recognised indicator for the Convention on Biological Diversity's Aichi Targets [FCLF17].

[2] In this case we refer to unequivocally for all the specified standpoints.

UC2.Q2:- $\square_{FIN}(\mathsf{AZE\_site} \sqcap \mathit{envo}\!:\!\mathsf{Forest\_area} \sqcap \exists \mathit{ro}\!:\!\mathsf{overlaps}.(\mathit{envo}\!:\!\mathsf{deforestation} \sqcup \mathit{envo}\!:\!\mathsf{ecosystem\_decay}))?$

UC2.Q3:- $\square_{FIN \setminus MoFI}(\mathsf{AZE\_site} \sqcap \mathit{envo}\!:\!\mathsf{Forest\_ecosystem} \sqcap \exists \mathit{ro}\!:\!\mathsf{overlaps}.(\mathit{envo}\!:\!\mathsf{deforestation} \sqcup \mathit{envo}\!:\!\mathsf{ecosystem\_decay}))?$

Three sample areas that would satisfy these queries can be visualised in figure 6.4.2.

6.4.3 Use Case 3: Agents using sets of standpoints

In some scenarios, different conceptualisations may entail the same facts. While it is sometimes useful to dissect the semantic differences between different representation choices, in many cases an agent may want to reason about what is unequivocally the case, without needing to consider the different ways in which this is materialised.

For instance, imagine an agent that must pick any box that is either red or orange (or both). For this task, the threshold of applicability between red and orange is not relevant, because any specific threshold will validate that a borderline case between red and orange is either red or orange. Similarly, agents handling complex scenarios may benefit from the same feature. A particular use case is as follows:

An automated agent produces deforestation alerts for an agency in Indonesia, using a set of information sources. In the first place, it uses the GLAD deforestation alerts, that detect loss of tree cover [PFVB19]. However, as they declare, GLAD alerts do not differentiate actual deforestation from other phenomena such as regular harvesting in plantations.

Consequently, the agent uses additional data about forest area according to three standpoints: The FAO definition, the MoFI definition and the Natural Forest Definition. In particular, the agency is interested in detecting the loss of tree cover in natural areas, which can be forests or scrublands, but not in other land types. The automated agent must fire a *definite* alarm whenever the detected event unequivocally occurs in a forest or scrubland, and an *arguable* alarm whenever the loss of tree cover occurs in an area that is only in some sense a forest or scrubland. In the latter case, the agency may use an expert to assess the environmental relevance of the alert and manually determine whether to confirm the alarm or not.

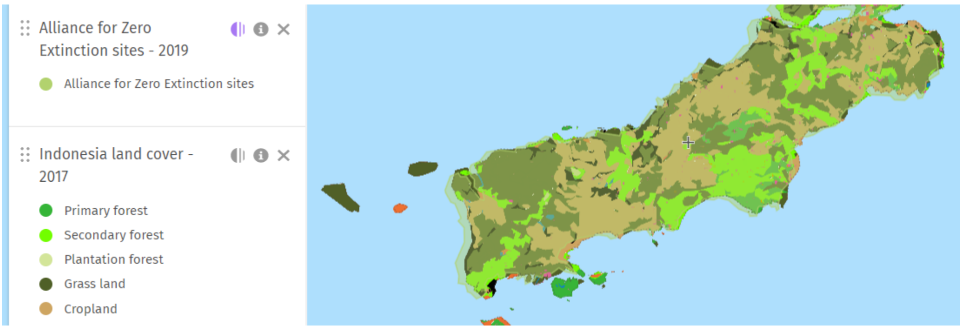

(a)

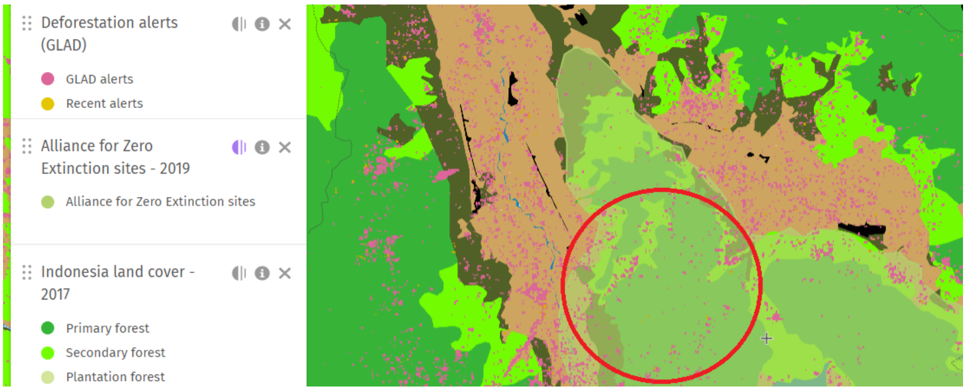

(b)

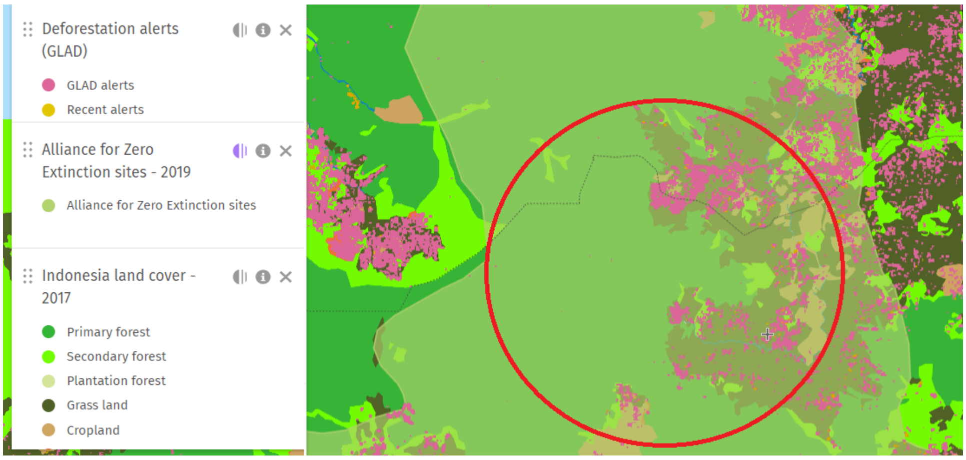

(c)

Figure 6.6: Representation of sample areas for queries **UC2.Q1-UC2.Q3**. AZE areas are the light green shade overlapping large areas of the map. (a) shows an area satisfying query **UC2.Q1**, (b) shows an unequivocal forest area (primary forest) with deforestation alerts, satisfying query **UC2.Q2**, and (c) shows an area where deforestation is restricted to areas that are in some sense a forest but not according to the MoFI, namely grasslands (lci:Shrubland), satisfying query **UC2.Q3**.

The automated agent, in order to detect the events of deforestation, regularly performs the following queries:

UC3.Q1:- $\square_{FIN}(\,envo\!:\!\mathsf{Forested\_area} \sqcup envo\!:\!\mathsf{Scrubland\_area})$

$$\sqcap\ \square_{*}\,\exists\, ro\!:\!\mathsf{overlaps.GLAD\_alert}$$

UC3.Q2:- $\lozenge_{FIN}(\,envo\!:\!\mathsf{Forested\_area} \sqcup envo\!:\!\mathsf{Scrubland\_area})$

$$\sqcap\ \square_{*}\,\exists\, ro\!:\!\mathsf{overlaps.GLAD\_alert}$$

Note that an instance of a GLAD alert in an area that is a forest according to some interpretations and a scrubland according to others, such as the instances of the original *ici* : *shrublands*, will prompt a definite alert, while instances from plantations or crops will be only in some sense true, as they are not $(\,envo\!:\!\mathsf{Forested\_area} \sqcup envo\!:\!\mathsf{Scrubland\_area})$ in all precisifications.

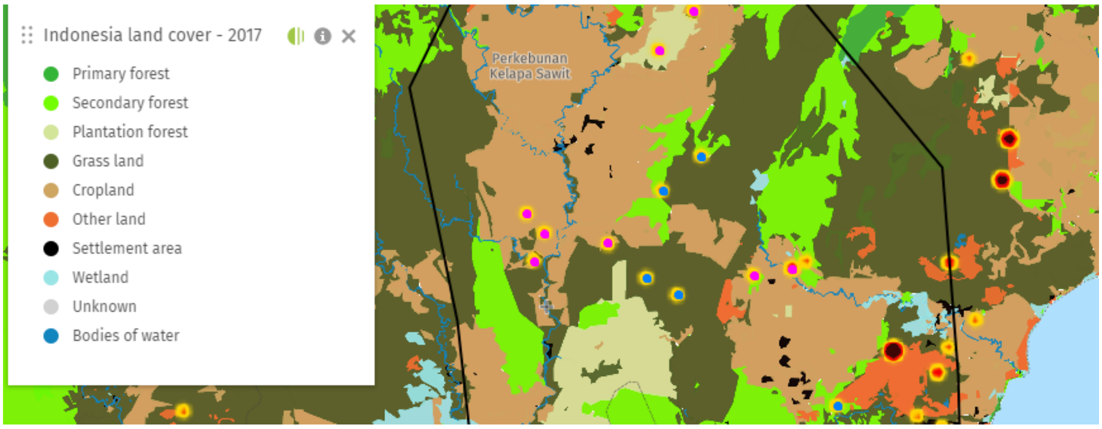

Figure 6.7: GLAD alerts over different land cover type areas in a region of Indonesia.

In the figure 6.7 we can see some definite fires (blue dots) that are located in unequivocal forests and forest-scrublands (note that the legend refers to the icl: classification), and some arguable alerts in areas that are in some sense a forest, in the legend plantation forests and croplands.

6.5 ‘Intra’ and ‘inter’ ontology applications

Following the discussion in this chapter, we suggest that, beyond purpose built multiagent systems with standpoint knowledge bases, there are two promising applications for standpoint logic as a tool to enrich existing formal ontologies . We name them the *intra-standpoint ontology* approach and the *inter-standpoint ontology* approach.

While the value of ontologies is now well established [Gua98], their support of vagueness and semantic heterogeneity[1] remains challenging. One of the main advantages of ontologies is that they improve the interoperability of systems, acting to enforce a consensus view reached by a community regarding a certain domain by formalising the semantics of the terminology in some logic formalism [GÁB17].

Typically this process involves the cooperation of domain experts and results in a unified decision on the formalisation of the semantics of the terminology. However, as a result of the semantic heterogeneity and vagueness of the concepts to define, strong semantic commitments favour specific interpretations of language and involve a loss of generality, thus restricting the opportunities for interoperability. On the other hand, approaches with shallower semantics rely fundamentally on taxonomic relationships, such as subsumption and mutual exclusivity, thus leaving uncertainty on the specific semantics of instances of these terms, and potentially compromising the sound reuse of information.

The *intra-standpoint ontology* approach consists on 'supervaluating' the concepts of an ontology, to preserve, on the one hand, the advantages of providing a common high-level conceptual structure for a domain while, on the other hand, making explicit the more fine-grained semantic commitments linked to different interpretations that are relevant to different uses of the ontology.

We thus see the potential for *supervaluated ontologies* by design, that are not the fruit of the integration of different models but a direct formalisation of a domain in which the semantic variablility of its terms is sufficiently meaningful and/or relevant to be represented in the ontology. Our claim is that *stanpoint logic* can provide a natural framework to accommodate such discrepancies in meaning.

On the other hand, following an *inter-standpoint ontology* approach, we suggest that *stanpoint logic* can be an adequate representation language to bridge between the conceptualisations linked to different ontologies, each of which can be associated to a particular standpoint.

Moreover, in the context of knowledge sharing, operations between standpoints are

[1] Occurs when ontologies, schemas or datasets of the same domain present differences in meaning and interpretation of categories and/or data values, thus challenging interoperability.

expected to be useful for two main purposes, namely concept negotiation and specification of combined standpoints. In practice, concept negotiation not only involves the analysis of the objective meaning and ontological commitments formalised on the different standpoints, but also of the models of such commitments. With an expressive enough framework, an agent can gain insight from both the logical relation between the interpretations themselves and from what can be inferred from the available data, for example, how many instances actually fall within the borderline areas and which instances are consistent with all the relevant standpoints.

With regards to the domain of individuals, we must note that often we will be dealing with ontologies that have been applied to formalising the meaning of certain datasets. In some cases, as we have considered in this chapter, we will know that two standpoints have been applied to the same objects. In practice, this may not be the case, and identifying objects between different ontologies may be non-trivial and involve complex issues. These issues are related to the problems of establishing correspondences between entities at different possible worlds and/or precisifications that were noted in the section 6.2.3 (and briefly commented at the end of chapter 4, section 4.5). For present purposes we simply assume that certain entities can be identified between ontologies, be that trough formal constraints that determine the identity[1] or trough external strategies used in the domain of ontology matching.

Summary

In this chapter we have explored the potential of the proposed logic to enable communication between agents holding diverse standpoints, in this case by relying on a shared ontology and supporting the establishment of standpoints on its vague terms. In particular, we have shown how the standpoint logic can be an adequate formalism to represent multiple aspects of the semantic variability of geographical concepts, which can manifest in the specification of the classification, individuation and demarcation criteria.

Moreover, in this application scenario we have considered uses of the standpoint framework both to enable agents to link data to the ontology according to their precise interpretation and also to query and reason within the ontology according to a spe-

[1]For instance by the use of the same proper name and rough location or because certain parts of the objects in all precisifications refer to the same entities in some database

cific standpoint, which can be modified during the interaction. In a context of sharing information, the proposed framework can offer tools that enable the analysis of the relations that hold between standpoints or interpretations and the execution of modifying operations such as calculating the intersection or union of a pair of standpoints. This, together with information on the instances satisfying these standpoints, serves as a support to the agent for the specification of the semantics with which vague terms are used, guaranteeing integrity and enabling interoperation with the ontology.

Chapter 7

Conclusions and Further Work

7.1 Thesis overview and contributions

In this thesis we have developed a formalised theory of the notion of *standpoint* and explored its capabilities for representing the semantics of terms within a supervaluationistic theory of vagueness. According to this view, vagueness can be explained in terms of the existence of a range of possible precise interpretations of a language. Each of these interpretations corresponds to some reasonable and mutually compatible interpretation of its terms, a precisification. In such a setting, *Standpoints* play the role of narrowing the semantic variability of the language in order for it to suit a specific context or an agent's viewpoint; and the capability of referring to several different standpoints allows these viewpoints to be compared and contrasted.

After setting the background and providing a literature review (chapter 2) we presented a formal language of *Standpoint Logic* (chapter 3), which we believe is well-suited for describing the variability of vague concepts and also for expressing relationships between different points of view regarding concept meaning, in a way that is suitable for use in KR applications where such issues are important.

We have developed our formalism as a modal logic. Such logics have been used for many purposes in knowledge representation, especially when dealing with information that cannot be simply expressed in terms of truth and falsity. Indeed many approaches to vagueness have employed modal logic, and some have been given different forms of supervaluation semantics. This is an obvious combination because of the similarity between the idea of a precisification and the notion of a possible world, which forms the basis of the standard Kripke semantics for modal logics. Nevertheless, the idea of formalising the notion of standpoint as a modal operator is a distinctive feature of our logic. Although other modal approaches do explain some of the kinds of reasoning that standpoint logic supports, they do not represent the notion explicitly. In considering the idea of a standpoint, we immediately recognise that there can be multiple standpoints. Hence standpoint logic is a multi-modal framework in which particular relationships between its modalities hold.

We have considered detailed examples of the capacities of the framework to reason about different forms of semantic variability, such as graded predicates like 'tall' and qualitative variations as in 'forest', and also how it can be used to reason about relationships between different standpoints. The logic can easily incorporate complex logical constraints that must hold between terms even though their interpretation is vague (i.e. penumbral connections). This can either be done generally (using $\Box_*$) or with respect to a particular standpoint (with $\Box_s$). We have also seen how the logic can represent and reason about particular combinations of standpoints and the differences between them.

In chapter 4 we gave both a semantics and a set of axioms for our formal language. In its current form, we do not envisage our axiom set being used directly as practical inference mechanism. But it can provide a framework within which one could define more limited sub-languages, suitable for particular data interpretation tasks; and it also provides a specification of a proof theory for which it may be possible to develop more practical proof algorithms.

The most characteristic features of *standpoint logic* are the interaction axioms **AS4** and **AS5**. These are stronger than the well known **A4** and **A5** axioms and correspond to what we call *trans-transitive* and *trans-euclidean* relations in the Kripke models. Moreover, a partial order on standpoints may be enforced using axiom **AP**, which encodes the *sharper* relation that can hold between standpoints (notion **N9**).

We have proved the soundness, completeness and decidability of propositional standpoint logic, and established the complexity of the satisfiability problem as **NP**-complete. This is a significant 'good' result. It may be regarded as a consequence of the stronger interaction axioms **AS4** and **AS5** and contrasts with the higher complexity of other multi-modal logics such as epistemic logic, that is **PSPACE**-complete. Finally, we also provide the syntax and semantics for standpoint logic with more expressive underlying logics, namely description logic and first-order logic.

In order to illustrate the applicability of *standpoint logic*, we have explored an application scenario in chapter 6 that tackles the difficulties that the semantic variability of the term *forest* poses in the context of the scientific multi-disciplinarity of the forestry domain. In addressing it, we proceed to analyse the literature on ontological issues of geographical objects and we determine three main aspects in which the semantics of *forest* vary, namely the classification, individuation and demarcation criteria. We then provide formal representations for standpoints corresponding to selected definitions and characterisations of *forest* from leading organisations that map into publicly available data. Finally, we illustrate how the framework can handle four use cases through different reasoning tasks on the created standpoint knowledge base.

7.2 Further work

In this section we highlight a number of areas that we believe are deserving of further research, some of which we intend to carry out in the near future. The proposed future work concentrates on two areas. On the one hand, on the development of methods for automated reasoning with the *propositional standpoint logic*. On the other, on the further analysis of the framework with richer underlying languages.

An automated theorem prover for propositional standpoint logic.

The existence of models that are small in terms of the number of precisifications, and also the possibility to greatly simplify formulae into a normal forms, indicate promising routes to the design of an automated theorem proving algorithm. The small models suggest the possibility of a tableau-style procedure to test for satisfiability, whereas the normal forms could provide the basis for a system based on a generalisation of resolution.

In particular, in section 4.4.2 we showed that not only the validity problem is NP-Complete, but also that there is an algorithm that, given a finite Kripke structure $\mathcal{M}$ with at most $|\phi|$ states, determines if a formula ϕ is satisfied by the model in time $O(|\phi|^2)$.

In further work we plan to tackle the guessing strategy of such small models, and we believe that there are efficient strategies to find good candidates. Moreover, for this purpose we will consider the guessing strategy of models both for general formulas $\phi \in \mathcal{L}_{\mathbb{S}_0}$ and for formulas in the Cover logic Multi-modal Prenex Normal Form (CMPNF) (see section 5.8). This should pave the way for the development of an automated theorem prover, which we expect to implement in Prolog.

Complexity of standpoint formulas with more expressive underlying logics.

It is well known that expressivity and complexity tend to pair together. In this thesis we have explored in depth the *propositional standpoint logic*, but we have also defined syntax and semantics for building standpoint logics based on more expressive underlying logics, in particular for a *description standpoint logic* and a *first-order standpoint logic*. Moreover, we highlighted in chapters 5 and 6 that many of the natural applications of a standpoint framework do in fact benefit from more expressive capabilities than those offered by an underlying propositional language. In this regard, an important line of further research is to investigate the decidability and complexity of more expressive forms of standpoint logic. In particular we want to explore the logic $\mathcal{ALC}$ as well as more expressive description logics, notably $\mathcal{SHOIN}$ [HPS04] and $\mathcal{SROIQ}$ [HKS06]. In addition, we want to establish the complexity of known decidable fragments of modal first order logic (see [WZ01]) for the standpoint framework.

Exploiting the notion of penumbral independence.

In supervaluationistic frameworks, accounting for penumbral connections is regarded as a core feature. Building on that, we intend to put forward the additional notion of *penumbral independence*. This notion attempts to capture the intuition that certain concepts are semantically independent within a standpoint (or globally). This would allow for the modularisation of a vocabulary into sets of terms that do not impinge on each other (e.g. concepts relating to forests can be modelled without worrying about books or buildings), which is expected to both facilitate the design and integration of systems and to improve their computational properties.

We say that two vague propositions ϕ and ψ are *penumbrally independent* when every sense in which ϕ can be interpreted is compatible with every sense which ψ can be interpreted. Building on the notion of logical separability, as investigated by [Lev98], we intend to provide a formal treatment of the notion that effectively contributes to the automated modularisation of vague vocabularies. Moreover, we plan to develop an adaptation of the normal form Cover logic Multi-modal Prenex Normal Form (CMPNF) to integrate such notion, along the lines of [FWWW19], so that terms in the disjunctive form are penumbrally independent.

Computation of standpoint combinations.

In this thesis we have confined ourselves to the establishment of a definition of the combination. In further work, we aim to establish effective methods for computing relationships and combinations of standpoints (for example different classifications of the same domain objects proposed by different organisations).

This idea is related to the work in [LD12], which introduces a framework for combining agents' beliefs and establishes four operators for combining different viewpoints, which are represented by means of bipolar valuation pairs. Using subsequent work comparing different representation formats for vague propositions [CDL14], which highlights the connections of this formalism with our modal logic, we intend to explore the parallels between standpoint logic and the operators suggested in [LD12].

Identity and the domain of individuals

One can argue that, when reasoning with logics with a domain of individuals, if we change the precisification according to which the world is classified, the set of entities is likely to change. For example, a particular set of tree-covered pixels might be considered to form a single forest under one precisification π, whereas under another, π', it might consist of two forests separated by a band of heathland. However, one could take the view that the set of entities can be regarded as the same, even though their classification changes. This is consistent with a *de dicto* view of vagueness, in which it is only linguistic descriptions that are vague, not the objects that they describe, and it is the view taken (implicitly) in this thesis. According to this view, the set of tree-covered pixels classed as a forest in π also 'exists''It is an element on he domain of individuals, yet it may not belong to any particular class.' in π', but it is not a forest (it may be an unclassified section of land).

While we have investigated individuation, classification and demarcation issues in chapter 6, our logic framework is agnostic to these issues. Consequently, we rely on the assumption that a reasonable criteria for identity, establishing the correspondence between entities in different precisifications, can be explicitly provided by the standpoints and be encoded in the underlying language.

In further work we intend to investigate these issues in more detail, and consider the incorporation of mechanisms that can explicitly account for the individuation of objects across precisifications and worlds. In the philosophical literature on vagueness, this issue is known as 'the problem of the many' [Ung80] and some frameworks such as [Aki00] have considered the representation of vague objects and trans-world identity. Moreover, work in the domain of applied ontology has considered representational issues of indentity and individuality in less expressive frameworks, namely description logics [GW00, GW09].

7.3 Conclusions

We shall conclude this thesis with some general considerations regarding the distinctive features and potential advantages of Standpoint Logic as a vehicle for representing and reasoning in scenarios that involve vague terminology.

7.3.1 Our subject of inquiry

In this thesis we have used the phrase *Semantic Variability of Natural Language Terms (SVoNLT)* to describe our subject. This may seem to be unnecessarily unconventional jargon, especially since most of the philosophical and logical literature, upon which we draw, describes its topic simply as *vagueness*; and we are studying much the same class of phenomena, including *sorites* susceptible graded predicates (as exemplified by *tall*) as well as multi-dimensional and non-numerical forms of vagueness.

We made this move to highlight the shift from the philosophical study of the sorites paradox towards a more practical aim of modelling scenarios involving semantically heterogeneous languages. This phrasing is, in our view, more aligned with contemporary research in applied domains, and avoids extended discussion of the nuances of the philosophical notion of vagueness itself, which we understand is contested. For instance, whereas philosophical analysis typically seeks to distinguish vagueness from polysemy, we consider both polysemy and vagueness to fall under the umbrella of the *Semantic Variability of Natural Language Terms (SVoNLT)*.

While the listing of senses is a common strategy to informally convey the semantics of a natural language term, drawing the attention of the reader to its interpretation in different contexts and hopefully capturing its most relevant features, research in lexicography [RL00] tends to take the view that there is rarely such a set of well defined senses; instead, the phenomenon is better explained through theories such as open-texture [Wai45] or family resemblance concepts [Wit09], which can then be formally represented, to a certain extent, through the notions of multidimensional and non-numerical vagueness. We consequently hope, in that respect, that we have been successful at drawing the attention to this understanding of *vagueness*, in its widest and applied sense.

7.3.2 Relation to other approaches

In the development of standpoint logic we have adopted an underlying supervaluationistic framework and we have focused on a particular problem: representation and reasoning in relation to contexts in which the semantic variability of terms is restricted. Such restricted contexts, or *sharpenings*, are a result of semantic commitments either explicitly expressed with the language or of principles that are implicit in a particular viewpoint. We have called these contexts of restricted variability *standpoints*. We consider that standpoints offer an interesting and useful tool for the modelling of vagueness, and are particularly useful for modelling non-numerical vagueness.

Although problematic for classical logics based on truth conditions, sorites vagueness is a relatively uniform phenomenon for which a variety of styles of representation have been proposed and explored in detail. However, the symbolic characterisation of conceptual or non-numerical vagueness seems to be less well developed.

Most frameworks dealing with the latter type of vagueness follow one of two approaches. One approach is to extend mechanisms that have been useful in representing homogeneous sorites vagueness to cover conceptual vagueness. In this direction, terms might be represented using graded membership functions which rely on similarity measures over a fixed set of parameters. Approaches based on *prototype theory*, where the distance from an instance to a prototype is computed within a space of parameters, are of this kind.

The other common approach does not involve modelling the semantic variability of the term itself, but instead focuses on providing a framework for handling it, by specifying mechanisms for merging and negotiating diverging meanings. These mechanisms may then be used to facilitate a multi-agent transaction of information or knowledge, yet they can not be considered a knowledge representation strategy as such.

The first approach faces different representational limitations from a symbolic point of view, as it cannot easily represent logical features of the semantics and the relationships between interpretations, except in terms of compound measures of similarity. These may be difficult to interpret unless based in a small number of dimensions, in which case the framework only supports the modelling of multi-dimensional vagueness but not of non-numerical or conceptual vagueness, that involves complex semantic interactions between terms and different feature arrangements (see section 2.3.3 for recap).

The standpoint framework can be seen as unifying some key aspects these two approaches. The supervaluationistic interpretation of vagueness provides a framework in which vagueness, whether it be sorites or conceptual, is modelled in terms of a set of fully precise interpretations, in what was described by Fine as 'ambiguity on a grand and systematic scale'. The flexibility of the approach relies on providing such a space of interpretations ranging over possible meanings but without specifying any particular organisation of this space. Yet, from the point of view of modelling the semantics of vague terms symbolically, this lack of organisation makes it difficult to form a strategy of how to use it for actual representation and reasoning involving vague terms.

By means of standpoints, we can impose order on the space of possible interpretations. Specifically, we can delimit and segregate the variability in meaning in the form of alternative sets of semantic commitments and constraints, which may be adapted to particular communities and contexts. In this way, the more (and more diverse) standpoints we establish, the better we will characterise the semantics of the term, which will be given by the whole set and the relations between them. And, from a practical point of view, we provide a framework in which a variable set of standpoints can be represented, which are relevant for the intended applications of the different users of the system and facilitate their interoperation. Moreover, small subsets of standpoints can be selected for reasoning in different scenarios, and operations between them can be easily established. Hence, we see that the standpoint approach is, in some respects, closely aligned to the transactional approaches to handling vague terminology.

In this way, by putting the focus on the representational issues we take a somewhat different and complementary approach to the main strands of work on supervaluationism, such as the seminal work of Fine [Fin75] or more recent frameworks like [LT12]. These lines of research focus more on the characterisation of the truth of a proposition and on the notion of validity in terms different logical principles, as well as on the exploration of the interactions between the variability of language and other phenomena such as uncertainty and beliefs.

7.3.3 Relation to other logics

In this project we have used a multi-modal normal logic as the tool to model the SVoNLT under a supervaluationistic interpretation of vagueness. This is not an unconventional move, and indeed there is a substantial amount of work examining mostly philosophical considerations on this, as discussed in section 2.5.2. It is, however, not the only approach. For instance, Lawry and Tang [LT12] argue that the simplest formulation of supervaluationism in a propositional logic framework is in terms of supervaluation pairs, which are sets of admissible classical valuations akin to epistemic states.

As discussed in section 2.5.3, several pieces of research explore the connections between the major representational frameworks that have been proposed to account for vagueness. With regards to the position of standpoint logic within that landscape, it is easy to see that one could express the models of individual standpoints via sets of supervaluation pairs. However, Lawry and Tang's logic develops the supervaluationistic semantics for a propositional language, therefore it is, alone, not sufficient for the representation of different standpoints. In that direction, one option would be to follow the approach of Multi-modal Epistemic Logic (MEL), which despite being a modal logic uses a semantics akin to Lawry's supervaluation pairs. In fact, it is easy to see the similarities between our semantics (if we consider the case of a single modality) and MEL's, which rests on the results that **K45** modal logics can be represented in terms of simplified Kripke frames, in which the relations are sets [Pie09]. We can then envisage the extension for multiple modalities and the partial order of relations so that an alternative semantics is provided in the form of sets of epistemic sets. This can then be related to the rest of formalisms as shown in [CDL14].

On the other hand, the simplified Kripke semantics for multiple modalities that we have provided in this dissertation has distinct advantages. In the first place, we have shown that, by maintaining the "de facto" standard of modal logic we provide a framework that would seem natural to the community working on multiagent systems, where the use of multi-modal epistemic logics is pervasive and hence well understood. Moreover, we are particularly interested in exploring the computational properties of standpoint logic with more expressive underlying logics that preserve good computational properties. A particularly interesting case is that of description logics, which has interesting applications in the area of applied ontology. Given that normal modal logics have been extensively studied, there is existing work in both modal description logics and modal horn clauses that can support these further developments, where existing automated reasoning systens generally use Kripke semantics. In addition, Kripke semantics for multi-dimensional modal logics are also well understood, paving the road for the formulation of different extensions of the system.

From a formal point of view, our logic is not ground-breaking. It has strong similarities to other modal logics, such as multimodal episemic logic [FHMV95] and MEL logic [BD14]. Partial orders are formalised as in [AH10] and we have provided the proofs of completeness, decidability and complexity following the lines of [HM92], adapting them to the particularities of our models $\mathfrak{M}_{\mathbb{S}_0}$. Further work of ours includes the study of the complexity of standpoint logic with different underlying description logics and fragments of FOL, which are in our opinion very interesting for this project because, for a logic that is committed with the representation of vagueness, the expressivity of the underlying language plays a major role, as has been highlighted in chapter 5.

7.3.4 Simplicity and expressivity

Our framework also contrasts with an important part of the existing literature in the simplicity of the model presented. Across this thesis, with the exception of the brief introduction of a two-dimensional standpoint logic in 4.5, we only consider an elementary scenario: a single state of the world that can be interpreted via multiple precisifications of the language. Moreover, we have not considered the representation of higher order vagueness, we have not accounted for cognitive aspects of concept acquisition and evolution, and we presume absolute lack of uncertainty both about the state of affairs and the language (we do not model the agent's mental state. In fact, standpoints can belong to communities rather than individual agents).

Instead, we have focused on the development of the most simple framework that can support the representational expressivity of standpoints within a supervaluationistic framework. We have attempted to provide a well developed characterisation which also shows good computational properties and has applications in a variety of domains.

7.3.5 Possible extensions

Our framework provides a basis for further work in a variety of directions, and that can be easily extended to incorporate expressivity in other aspects that we consider different from (but that may interact with) the semantic variability of the language.

For instance, if we wanted to support higher-order vagueness we could consider providing an alternative set of axioms and/or modalities to model this phenomena. The relevance of the problem of higher-order vagueness is itself contested (see, for instance, [Sor01] for Sorensen's defence of its importance and [Wri10] for Wright's argument that it is an illusion). Beyond the philosophical considerations, we take the view that, in most applications, higher-order vagueness is not a relevant (not even desirable) feature, and adds complexity to the formalisation. Yet, one could investigate standpoint extensions that supported the phenomenon, which has been generally considered for modal interpretations of supervaluationsim in [Var07].

Another possible extension would be to model epistemic lack of knowledge and its interaction with vagueness. One could establish a framework along the lines of the two-dimensional extension in 4.5, where instead of (or in addition to) the aletic modalities, we use epistemic operators. With this in mind, the most interesting part is the establishment of interaction axioms across dimensions. For instance, if we consider exclusively uncertainty regarding the state of the world (e.g. we model an agent's beliefs as a set of possible worlds compatible with her belief state), then epistemic operators would commute with standpoint operators. However if we allow the agents to have beliefs about language not only about the worlds, then there would be more complex interactions between the two kinds of modality. In such setting, we would accept that an agent α can believe that according to a standpoint s something is the case, in which the uncertainty may be both epistemic and linguistic, in a similar way to the framework presented in [LT12].

7.3.6 Applications to supporting reasoning about the forestry domain

In chapter 6 we used the forestry domain to illustrate applications of standpoint logic. It should be noted that, rather than choosing a contrived or well circumscribed example, we have chosen a very complex domain for representation and reasoning. There are different reasons for this. On the one hand, we already gave examples of a more tailored nature in chapter 5, that were intended to illustrate the expressive capabilities of the framework and exemplify how to model different phenomena.

In contrast, in chapter 6 we aim to explore the usefulness of the framework in a real-world scenario, where there are large open problems for representation and reasoning. Vagueness is indeed known to be pervasive in the spatial domain in general and the geographical domain in particular. On top of that, the domain of forestry widely reports the challenges of the SVoNLT even outside of the computational domain. Thus, it is clear that analysis of information relating to forestry could benefit from a framework that enables representation and reasoning in detail about the variable semantics of the vague terms. This gives us breath to explore the challenges and opportunities that such a framework would face. Hence, the domain of forestry has many attractions as an application scenario.

This part of our research has driven us to the exploration of additional problems of representation, which result in the distinction of three fundamental aspects in which the semantics of terms referring to geographic objects tend to be vague. These are the classification, identity and demarcation criteria, all of which seem to be essential to resolving issues that arise in the interpretation and computational manipulation of geographic data. Concerns relating to of variability in these criteria can be found scattered throughout the ontology literature, and in particular in the ontology of geography, and of course, more specifically still, in the analysis of specifications or definitions of forest (a wide collection is reported in [Lun02]). Nevertheless, prior research has not systematically considered the semantic variability of 'forest' in terms of these aspects. Rather, it has been mainly concerned with difficulties arising at the level of human communication (i.e. mostly between research teams, administrations and stakeholders): they discuss terminologies and definitions that are expressed in natural language and used by humans.

Through the development of the use cases, we have illustrated how issues of vagueness in individuation, classification and demarcation arise in the representation of standpoints associated with different datasets, research communities or organisational standards, and how we can deal with them using standpoint logic. Although many issues relating to terminology and definitions can be handled by our framework, others are difficult to capture. Indeed, the issues of individuation, classification and demarcation are particularly tricky and have been introduced but not explored in great detail in this thesis. Development in this direction would involve further explication of the nature of vaguely defined objects (e.g. is there *de re* vagueness of objects as well as *de dicto* vagueness in the descriptions?), and is likely to require certain enrichment of the semantics that we have given.

Beyond these philosophical considerations, in the use cases of the chapter 6 we have examined a scenario in which we have an integrated system that allows us to acquire information from a well established ontology, EnvO, and from a well established forest data repository, GFW, and we have taken account of existing pieces of research reporting the use of such data. Thus, the chapter has illustrated some possibilities for using the framework in a realistic and complex setting. We acknowledge, however, that further work with experts in the domain could help us greatly in directing the capabilities of standpoint logic to better support practical application scenarios.

References

[Ade13] F. Ademollo. *Ancient Philosophy of Language.* Oxford Bibliographies, 2013.

[ADP09] N. Asher, J. Dever, and C. Pappas. Supervaluations debugged. *Mind*, 118(472):901–933, 2009.

[AE07] E. Agirre and P. G. Edmonds. *Word sense disambiguation : Algorithms and applications.* Springer Science & Business Media, 2007.

[Aga05] P. Agarwal. Ontological Considerations in GIScience. *International Journal of Geographical Information Science*, 19(5):501–536, 2005.

[AH10] G. Allwein and W. L. Harrison. Partially-ordered Modalities. In *Proceedings of the 8th conference on Advances in Modal Logic (AiML)*, pages 1–21, 2010.

[AJ11] B. Adams and K. Janowicz. Constructing geo-ontologies by reification of observation data. In A. Agrawal, I. Cruz, C. A. Jensen, E. Ofek, and E. Tanin, editors, *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 309–318, Association for Computing Machinery, 2011.

[Aki00] K. Akiba. Vagueness as a Modality. *The Philosophical Quarterly*, 50(200):359–370, 2000.

[Ari55] Aristotle. *On sophistical refutations [and] On coming-to-be and passing away.* Loeb classical library, Greek authors, 1955.

[ASPPBL06] H. Anaya-Sánchez, A. Pons-Porrata, and R. Berlanga-Llavori. Word Sense Disambiguation Based on Word Sense Clustering. In J. S. Sichman, H. Coelho, and S. O. Rezende, editors, *Advances in Artificial Intelligence-IBERAMIA-SBIA*, pages 472–481, Springer, 2006.

[Ath93] C. Atherton. The Stoics on ambiguity. *Cambridge classical studies*, 1993.

[Auc10] G. Aucher. An Internal Version of Epistemic Logic. *Studia Logica*, 94(1):1–22, 2010.

[BA91] W. Bechtel and A. Abrahamsen. *Connectionism and the mind: An introduction to parallel processing in networks.* Basil Blackwell, 1991.

[Baa03] F. Baader. *The description logic handbook : Theory, implementation, and applications.* Cambridge University Press, 2003.

[Bal02] J. M. Baldwin. *Dictionary of Philosophy and Psychology.* The MacMillan Company, 1902.

[BC12] O. Bueno and M. Colyvan. Just what is vagueness? *Ratio*, 25(1):19–33, 2012.

[BD14] M. Banerjee and D. Dubois. A simple logic for reasoning about incomplete knowledge. *International Journal of Approximate Reasoning*, 55(2):639–653, 2014.

[BdRV01] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic.* Cambridge University Press, 2001.

[Ben98] B. Bennett. Modal Semantics for Knowledge Bases Dealing with Vague Concepts. In A. G. Cohn, L. Schubert, and S. C. Shapiro, editors, *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 234–244, Morgan Kaufmann Publishers, 1998.

[Ben01] B. Bennett. Application of supervaluation semantics to vaguely defined spatial concepts. In D. Montello, editor, *Proceedings of the 5th International Conference on Spatial Information Theory (COSIT'01)*, pages 108–123, Springer, 2001.

[Ben02] B. Bennett. Physical objects, identity and vagueness. In D. Fensel, D. McGuinness, and M. A. Williams, editors, *Proceedings of the 8th International Conference on Principles of Knowledge Representation and Reasoning (KR'02)*, pages 395–406, Morgan Kaufmann, 2002.

[Ben05] B. Bennett. Modes of Concept Definition and Varieties of Vagueness. *Applied Ontology*, 1(1):17–26, 2005.

[Ben06] B. Bennett. A theory of vague adjectives grounded in relevant observables. In P. Doherty, J. Mylopoulos, and C. A. Welty, editors, *Proceedings of the 10th International Conference in Principles of Knowledge Representation and Reasoning (KR'06)*, pages 36–45, AAAI Press, 2006.

[Ben11] B. Bennett. Standpoint semantics: a framework for formalising the variable meaning of vague terms. In P. Cintula, C. Fermüller, L. Godo, and P. Hájek, editors, *Understanding Vagueness - Logical, Philosophical and Linguistic Perspectives*, Studies in Logic, pages 261–278. College Publications, 2011.

[BFM10] M. Bienvenu, H. Fargier, and P. Marquis. Knowledge Compilation in the Modal Logic S5. 2010.

[BHHR18] H. Böttcher, L. Herrmann, M. Herold, and E. Romijn. Independent Monitoring: Building trust and consensus around GHG data for increased accountability of mitigation in the land use sector. Technical report, Publications Office of the European Union, 2018.

[BHS08] F. Baader, I. Horrocks, and U. Sattler. Description Logics. In F. V. Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, volume 3, chapter 3, pages 135–179. Elsevier, 2008.

[Bie08] M. Bienvenu. Prime Implicate Normal Form for ALC Concepts. In R. C. Holte and A. Howe, editors, *Proceedings of the 23th Conference on Artificial Intelligence (AAAI'08)*, AAAI Press, 2008.

[BKW10] F. Baader, R. Küsters, and F. Wolter. Extensions to Description Logics. In F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. Patel-Schneider, editors, *The Description Logic Handbook: Theory,* *Implementation, and Applications*, pages 219–261. Cambridge University Press, 2010.

[Bla37] M. Black. Vagueness. An Exercise in Logical Analysis. *Philosophy of Science*, 4(4):427–455, 1937.

[Bla63] M. Black. Reasoning with Loose Concepts. *Dialogue*, 2(01):1–12, 1963.

[Bla08] S. Blackburn. *The Oxford dictionary of philosophy.* Oxford University Press, 2008.

[BMS+13] P. L. Buttigieg, N. Morrison, B. Smith, C. J. Mungall, and S. E. Lewis. The environment ontology: contextualising biological and biomedical entities. *Journal of Biomedical Semantics*, 4(1):43, 2013.

[BMT08] B. Bennett, D. Mallenby, and A. Third. An ontology for grounding vague geographic terms. In C. Eschenbach and M. Gruninger, editors, *Proceedings of the 5th International Conference on Formal Ontology in Information Systems (FOIS-08)*, IOS Press, 2008.

[BPL+16] P. L. Buttigieg, E. Pafilis, S. E. Lewis, M. P. Schildhauer, and R. L. Walls. The environment ontology in 2016: bridging domains with increased scope, semantic density, and interoperation. *Journal of Biomedical Semantics*, 7(1):57, 2016.

[BRP57] D. C. Buckland, D. R. Redmond, and R. Pomerleau. Definitions of Terms in Forest and Shade Tree Diseases. *Canadian Journal of Botany*, 35(5):675–679, 1957.

[Bur95] L. C. Burns. Something to do with Vagueness. *Southern Journal of Philosophy*, 33(S1):23–47, 1995.

[BW88] D. Bruce and L. Wensel. Modelling forest growth: approaches, definitions, and problems. Technical report, USDA Forest Service, 1988.

[BW11] E. Barnes and J. R. G. Williams. A theory of metaphysical indeterminacy. In K. Bennett and D. W. Zimmerman, editors, *Oxford Studies in Metaphysics: Volume 6*, pages 509–522. Oxford University Press, 2011.

[Car97] L. Carroll. *Symbolic logic: Part 1, Elementary.* Macmillan, 1897.

[Car69] J. Cargile. The Sorites Paradox. *The British Journal for the Philosophy of Science*, 20(3):193–202, 1969.

[Car00] L. Carroll. *Symbolic logic and the game of logic.* Dover Publications and Berkeley Enterprises, 2000. Original text: [Car97].

[CBL16] R. Chazdon, P. Brancalion, and L. Laestadius. When is a forest a forest? Forest concepts and definitions in the era of forest and landscape restoration. *Ambio*, 45(5):538–550, 2016.

[CC18] V. Cross and S. Chen. Fuzzy Ontologies: State of the Art Revisited. In Barreto, G. A., and R. Coelho, editors, *37th Conference of the North American Fuzzy Information Processing Society*, pages 230–242, Springer, 2018.

[CCMV06] B. T. Cate, W. Conradie, M. Marx, and Y. Venema. Definitorially Complete Description Logics. In P. Doherty, J. P. Mylopoulos, J. Mylopoulos, and C. A. Welty, editors, *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning (KR'06)*, pages 79–89, AAAI Press, 2006.

[CDG03] D. Calvanese and G. De Giacomo. Expressive description logics. In F. Baader, D. Calvanese, D. L. McGuinness, D. Nardi, and P. Patel-Schneider, editors, *The Description Logic Handbook: Theory, Implementation, and Applications*, pages 178–218. Cambridge University Press, 2003.

[CDL14] D. Ciucci, D. Dubois, and J. Lawry. Borderline vs. unknown: Comparing three-valued representations of imperfect information. *International Journal of Approximate Reasoning*, 55(9):1866–1889, 2014.

[CFW04] A. Comber, P. Fisher, and R. Wadsworth. Integrating land-cover data with different ontologies: Identifying change from inconsistency. *International Journal of Geographical Information Science*, 18(7):691–708, 2004.

[CFW05] A. Comber, P. Fisher, and R. Wadsworth. What is land cover? *Environment and Planning B: Planning and Design*, 32(2):199–209, 2005.

[CGV11] M. Carrara, P. Garbacz, and P. E. Vermaas. If engineering function is a family resemblance concept: assessing three formalization strategies. *Applied Ontology*, 6(2):141–163, 2011.

[Che80] B. F. Chellas. *Modal Logic: an introduction.* Cambridge University Press, 1980.

[CHN11] P. Cintula, P. Hájek, and C. Noguera. Handbook of Mathematical Fuzzy Logic. In *Studies in Logic, Mathematical Logic and Foundations*, volume 37. College Publications, 2011.

[Cob08] P. Cobreros. Supervaluationism and logical consequence: A third way. *Studia Logica*, 30(3):291–312, 2008.

[Cob11] P. Cobreros. Supervaluationism and classical logic. In R. Nouwen, R. v. Rooij, U. Sauerland, and H. Schmitz, editors, *Vagueness in Communication: Lecture Notes in Computer Science*, pages 51–63, Springer, 2011.

[Cob16] P. Cobreros. Supervaluationism and Fara's Argument Concerning Higher-Order Vagueness. In P. Égré and N. Klinedinst, editors, *Vagueness and Language Use*, pages 207–221. Palgrave Macmillan, 2016.

[Cre76] M. Creswell. The semantics of degree. In B. H. Partee, editor, *Montague Grammar*, pages 261–292. Academic Press, 1976.

[Cri17] W. Crispin. Is Higher Order Vagueness Coherent? In T. Williamson and D. Graff, editors, *Vagueness*, chapter 15, pages 295–305. Routledge, 1st edition, 2017.

[CSH+18] P. G. Curtis, C. M. Slay, N. L. Harris, A. Tyukavina, and M. C. Hansen. Classifying drivers of global forest loss. *Science*, 361(6407):1108–1111, 2018.

[CSV98] R. Casati, B. Smith, and A. C. Varzi. Ontological Tools for Geographic Representation. In N. Guarino, editor, *Formal Ontology in Information Systems*, pages 77–85. IOS Press, 1998.

[DDDÉ13] I. Douven, L. Decock, R. Dietz, and P. Égré. Vagueness: A Conceptual Spaces Approach. *Journal of Philosophical Logic*, 42(1):137–160, 2013.

[Dea88] P. D. Deane. Polysemy and cognition. *Lingua*, 75(4):325–361, 1988.

[dFL+97] M. d'Inverno, M. Fisher, A. Lomuscio, M. Luck, M. De Rijke, M. Ryan, and M. Wooldridge. Formalisms for multi-agent systems. *The knowledge engineering review*, 12(3):315–321, 1997.

[DG09] M. Dupenois and A. Galton. Assigning Footprints to Dot Sets: An Analytical Survey. In K. Hornsby, C. Claramunt, M. Denis, and G. Ligozat, editors, *Proceedings of the 9th international Conference on Spatial Information Theory (COSIT'09)*, pages 227–244, Springer, 2009.

[DNR97] F. M. Donini, D. Nardi, and R. Rosati. Autoepistemic description logics. In M. E. Pollack, editor, *Proceedings of the 15th International Joint Conference on Artifical Intelligence (IJCAI'97)*, pages 136–141, Morgan Kaufmann Publishers, 1997.

[DP97] D. Dubois and H. Prade. The three semantics of fuzzy sets. *Fuzzy Sets and Systems*, 90(2):141–150, 1997.

[DP14] D. Dubois and H. Prade. Possibilistic Logic - An Overview. In D. M. Gabbay, J. H. Siekmann, and J. Woods, editors, *The Handbook of the History of Logic*, volume 9, pages 283–342. North Holland, 2014.

[Dum75] M. Dummett. Wang's paradox. *Synthese*, 30(3-4):301–324, 1975.

[EdGR13] P. Egré, V. de Gardelle, and D. Ripley. Vagueness and Order Effects in Color Categorization. *Journal of Logic, Language and Information*, 22(4):391–420, 2013.

[Ege93] M. J. Egenhofer. What's Special About Spatial?: Database Requirements for Vehicle Navigation in Geographic Space. In P. Buneman and S. Jajodia, editors, *Proceedings of the 1993 ACM SIGMOD international conference on Management of data (SIGMOD '93)*, pages 398–402, Association for Computing Machinery, 1993.

[EM95] M. J. Egenhofer and D. Mark. Naive geography. In A. U. Frank and W. Kuhn, editors, *Proceedings of the 2nd International Conference on Spatial Information Theory (COSIT'95)*, pages 1–16, Springer, 1995.

[Far10] D. G. Fara. Scope Confusions and Unsatisfiable Disjuncts: Two Problems for Supervaluationism. In R. Dietz and S. Moruzzi, editors, *Cuts and Clouds: Vagueness, its Nature, & its Logic*. Oxford University Press, 2010.

[FCLF17] S. M. Funk, D. Conde, J. Lamoreux, and J. E. Fa. Meeting the Aichi targets: Pushing for zero extinction conservation. *Ambio*, 46(4):443–455, 2017.

[FHMV95] R. Fagin, J. Y. Halpern, Y. Moses, and M. Vardi. *Reasoning about knowledge*. MIT Press, 1995.

[Fin75] K. Fine. Vagueness, truth and logic. *Synthese*, 30(3-4):265–300, 1975.

[FK06] C. Fermüller and R. Kosik. Combining supervaluation and degree based reasoning under vagueness. In M. Hermann and A. Voronkov, editors, *Proceedings of the International Conference on Logic for Programming Artificial Intelligence and Reasoning (LPAR'06)*, pages 212–226, Springer, 2006.

[FR09] C. Fermüller and C. Roschger. Bridges between Contextual Linguistic Models of Vagueness and T-Norm Based Fuzzy Logic. In T. Kroupa and J. Vejnarova, editors, *Proceedings of the 8th Workshop on Uncertainty Processing (WUPES'09)*, pages 71–79, Liblice, 2009.

[Fre48] F. L. G. Frege. Sense and reference. *The philosophical review*, 57(3):209–230, 1948.

[Ful06] D. Fuller. Tropical forest monitoring and remote sensing: A new era of transparency in forest governance? *Singapore Journal of Tropical Geography*, 27(1):15–29, 2006.

[FWWW19] L. Fang, K. Wang, Z. Wang, and X. Wen. Disjunctive Normal Form for Multi-Agent Modal Logics Based on Logical Separability. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI'19)*, AAAI Press, 2019.

[GÁB17] L. Gómez Álvarez and B. Bennett. Classification, Individuation and Demarcation of Forests: formalising the multi-faceted semantics of geographic terms. In E. Clementini, M. Donnelly, M. Yuan, C. Kray, P. Fobliarone, and A. Ballatore, editors, *Proceedings of the 13th International Conference on Spatial Information Theory (COSIT'17)*, pages 1–8, Leibniz International Proceedings in Informatics, 2017.

[Gär04] P. Gärdenfors. Conceptual Spaces as a Framework for Knowledge Representation. *Mind and Matter*, 2(2):9–27, 2004.

[GC10] M. García-Carpintero. Supervaluationism and the Report of Vague Contents. In R. Dietz and S. Moruzzi, editors, *Cuts and Clouds: Vagueness, its Nature, & its Logic*. Oxford University Press, 2010.

[GL05] V. Gallese and G. Lakoff. The Brain's concepts: the role of the Sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology*, 22(3-4):455–479, 2005.

[Göd32] K. Gödel. Zum intuitionistischen aussagenkalkül. *Anzeiger der Akademie der Wissenschaften in Wien*, 69:65–66, 1932.

[Gra08] A. Grainger. Difficulties in tracking the long-term global trend in tropical forest area. *Proceedings of the National Academy of Sciences*, 105(2):818–823, 2008.

[Gra10] A. Grainger. Uncertainty in the construction of global knowledge of tropical forests. *Progress in Physical Geography*, 34(6):811–844, 2010.

[Gri75] H. Grice. *Logic and Conversation*. Unpublished manuscript, 1975.

[Gri12] T. W. Grinsell. Avoiding predicate whiplash: social choice theory and linguistic vagueness. In A. Chereches, editor, *Proceedings of the 22nd Semantics and Linguistic Theory Conference (SALT'12)*, pages 424–440, Linguistic Society of America, 2012.

[GSD04] A. Gliozzo, C. Strapparava, and I. Dagan. Unsupervised and supervised exploitation of semantic domains in lexical disambiguation. *Computer Speech & Language*, 18(3):275–299, 2004.

[Gua98] N. Guarino. Formal Ontology and Information Systems. In N. Guarino, editor, *Proceedings of the 1st International Conference on Formal Ontology in Information Systems (FOIS-98)*, pages 81–97, IOS Press, 1998.

[GW00] N. Guarino and C. Welty. Identity, unity, and individuality: Towards a formal toolkit for ontological analysis. In W. Horn, editor, *Proceedings of the 14th European Conference on Artificial Intelligence (ECAI'00)*, pages 219–223, IOS Press, 2000.

[GW01] P. Gärdenfors and M. A. Williams. Reasoning about categories in conceptual spaces. In *Proceedings of the 17th International Joint Conference on Artifical Intelligence (IJCAI'01)*, pages 385–392, Morgan Kaufmann Publishers, 2001.

[GW09] N. Guarino and C. Welty. An Overview of OntoClean. In S. Staab and R. Studer, editors, *Handbook on Ontologies*, pages 201–220. Springer, 2009.

[GZR05] C. Giri, Z. Zhu, and B. Reed. A comparative analysis of the Global Land Cover 2000 and MODIS land cover data sets. *Remote Sensing of Environment*, 94(1):123–132, 2005.

[Haa79] S. Haack. Do we need fuzzy logic? *International Journal of Man-Machine Studies*, 11(4):437–445, 1979.

[HFD11] J. Hales, T. French, and R. Davies. Refinement Quantified Logics of Knowledge. *Electronic Notes in Theoretical Computer Science*, 278:85–98, 2011.

[HFD12] J. Hales, T. French, and R. Davies. Refinement Quantified Logics of Knowledge and Belief for Multiple Agents. *Advances in Modal Logic*, 9:317–338, 2012.

[Hin59] J. Hintikka. Aristotle and the ambiguity of ambiguity. *Inquiry*, 2(1-4):137–151, 1959.

[Hin62] J. Hintikka. *Knowledge and Belief: An Introduction to the Logic of the Two Notions.* King's College Publications, 1962.

[HKS06] I. Horrocks, O. Kutz, and U. Sattler. The even more irresistible SROIQ. In P. Doherty, J. P. Mylopoulos, J. Mylopoulos, and C. A. Welty, editors, *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning (KR'06)*, pages 57–67, AAAI Press, 2006.

[HM92] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54(3):319–379, 1992.

[HP98] A. Hunter and S. Parsons. *Applications of uncertainty formalisms.* Springer, 1998.

[HP03] P. Holmgren and R. Persson. Evolution and prospects of global forest assessments. Global forest assessments move towards the goal of addressing a full range of benefits from forests and tree resources. Technical report, UNASYLVA (FAO), 2003.

[HPM+13] M. C. Hansen, P. V. Potapov, R. Moore, M. Hancher, S. A. Turubanova, A. Tyukavina, D. Thau, S. Stehman, S. J. Goetz, T. R. Loveland, et al. High-resolution global maps of 21st-century forest cover change. *Science*, 342(6160):850–853, 2013.

[HPS04] I. Horrocks and P. Patel-Schneider. Reducing owl entailment to description logic satisfiability. *Journal of Web Semantics*, 1(4):345–357, 2004.

[HS15] V. Hendricks and J. Symons. Epistemic Logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy (Fall 2015 Edition)*, Stanford University, 2015.

[HTL14] Z. Hu, G. Tang, and G. Lu. A new geographical language: A perspective of GIS. *Journal of Geographical Sciences*, 24(3):560–576, 2014.

[IW07] N. Ide and Y. Wilks. Making Sense About Sense. In E. Agirre and P. Edmonds, editors, *Word Sense Disambiguation Algorithms and Applications*, pages 47–73. Springer, 2007.

[Jaś36] S. Jaśkowski. Recherches sur le système de la logique intuitioniste. *Actes du Congrès International de Philosophie Scientifique*, 6:58–61, 1936.

[Joh87] M. Johnson. *The body in the mind : The bodily basis of meaning, imagination, and reason.* University of Chicago Press, 1987.

[Kam75] H. Kamp. Two theories about adjectives. In E. Keenan, editor, *Formal Semantics of Natural Language.* Cambridge University Press, 1975.

[Kee00] R. Keefe. Supervaluationism and Validity. *Philosophical Topics*, 28(1):93–105, 2000.

[Kee10] R. Keefe. Supervaluationism, Indirect Speech Reports, and Demonstratives. In R. Dietz and S. Moruzzi, editors, *Cuts and Clouds: Vagueness, its Nature, & its Logic.* Oxford University Press, 2010.

[Kle38] S. C. Kleene. On a notation for ordinal numbers. *Journal of Symbolic Logic*, 3:150–155, 1938.

[Kle52] S. C. Kleene. *Introduction to Metamathematics.* North-Holland, 1952.

[KM02] D. E. Klein and G. L. Murphy. Paper has been my ruin: Conceptual relations of polysemous senses. *Journal of Memory and Language*, 47(4):548–570, 2002.

[KP95] H. Kamp and B. Partee. Prototype theory and compositionality. *Cognition*, 57(2):129–191, 1995.

[Kri16] H. U. Krieger. Integrating graded knowledge and temporal change in a modal fragment of owl. In J. Herik and J. Filipe, editors, *Proceedings of the 8th International Conference on Agents and Artificial Intelligence (ICAART'16)*, pages 75–95, Springer, 2016.

[KS96a] R. Keefe and P. Smith. Theories of vagueness. In R. Keefe and P. Smith, editors, *Vagueness: A reader*, pages 1–57. The MIT press, 1996.

[KS96b] R. Keefe and P. Smith. *Vagueness: A reader.* MIT Press, 1996.

[Kyb00] A. Kyburg. When Vague Sentences Inform: A Model Of Assertability. *Synthese*, 124(2):175–191, 2000.

[Lak70] G. Lakoff. A Note on Vagueness and Ambiguity. *Linguistic Inquiry*, 1(3):357–359, 1970.

[Lak87] G. Lakoff. *Women, fire, and dangerous things.* University of chicago press, 1987.

[Las11] D. Lassiter. Vagueness as probabilistic linguistic knowledge. In R. Nouwen, R. V. Rooij, U. Sauerland, and H. C. Schmitz, editors, *Proceedings of the International Workshop on Vagueness in Communication (ViC'09)*, pages 127–150, Springer-Verlag, 2011.

[Law06] J. Lawry. *Modelling and reasoning with vague concepts.* Springer, 2006.

[Law08] J. Lawry. Appropriateness measures: An uncertainty model for vague concepts. *Synthese*, 161(2):255–269, 2008.

[LC06] L. Liu and L. Chen. A modal supervaluation description logic for characterization of vague concepts: Its semantics and a tableau algorithm for it. *Logic Journal of the IGPL*, 14(6):873–888, 2006.

[LD12] J. Lawry and D. Dubois. A Bipolar Framework for Combining Beliefs about Vague Propositions. In G. Brewka, T. Eiter, and S. A. McIlraith, editors, *Proceedings of the 13th International Conference on the Principles of Knowledge Representation and Reasoning (KR'12)*, AAAI Press, 2012.

[Lev98] H. Levesque. A Completeness Result for Reasoning with Incomplete First-Order Knowledge Bases. In A. G. Cohn, L. K. Schubert, and S. S. C, editors, *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 14–23, Morgan Kaufmann Publishers, 1998.

[Lew70] D. K. Lewis. General semantics. *Synthèse*, 22:18–67, 1970.

[LGR11] J. Lawry and I. González-Rodríguez. A bipolar model of assertability and belief. *International Journal of Approximate Reasoning*, 52(1):76–91, 2011.

[LJ17] J. Lawry and O. James. Vagueness and Aggregation in Multiple Sender Channels. *Erkenntnis*, 82(5):1123–1160, 2017.

[LT93] C. Leacock and G. Towell. Corpus-based statistical sense resolution. In *Proceedings of the Workshop on Human Language Technology (HLT'93)*, pages 260–265, Association for Computational Linguistics, 1993.

[LT12] J. Lawry and Y. Tang. On truth-gaps, bipolar belief and the assertability of vague propositions. *Artificial Intelligence*, 191-192:20–41, 2012.

[Łuk20] J. Łukasiewicz. O logice trojwartosciowej. *Ruch Filozoficny*, 5(23):170–171, 1920. English translation in [Łuk70].

[Łuk70] J. Łukasiewicz. Selected works. In L. Borkowski, editor, *Studies in Logic and the Foundations of Mathematics*. North-Holland, 1970.

[Lun02] H. G. Lund. When Is a Forest not a forest? *Journal of Forestry*, 100(8):21–27, 2002.

[Lun07] H. G. Lund. Definitions of Forest, Deforestation, Afforestation and Reforestation. Technical report, VA: Forest Information Services, 2007.

[LWZ08] C. Lutz, F. Wolter, and M. Zakharyaschev. Temporal Description Logics: A Survey. In *Proceedings of the 15th International Symposium on Temporal Representation and Reasoning (TIME'08)*, pages 3–14, IEEE, 2008.

[LYV16] J. E. Litland and J. Yli-Vakkuri. Vagueness and Modality. *Philosophical Perspectives*, 30(1):1–39, 2016.

[MC89] D. Mark and F. Csillag. The nature of boundaries on 'area-class' maps. *Cartographica*, 26:65–78, 1989.

[Meh58] H. Mehlberg. *The Reach of Science*. University of Toronto Press, 1958. Extract on Truth and Vagueness, pp. 427-55, reprinted in [KS96b].

[MH04] J. Meyer and W. Hoek. *Epistemic logic for AI and computer science*. Cambridge University Press, 2004.

[Mil95] G. A. Miller. WordNet: A Lexical Database for English. *Communications of the ACM*, 38(11):39–41, 1995.

[MJP15] M. MacDicken, O. Jonsson, and L. Pina. Global forest resources assessment. Technical report, Food and Agriculture Organization of the United Nations, 2015.

[MM13] M. Minock and J. Mollevik. Context-dependent 'near' and 'far' in spatial databases via supervaluation. *Data & Knowledge Engineering*, 86:295–305, 2013.

[MT94] W. Meyer and B. Turner. *Changes in land use and land cover : a global perspective*. Cambridge University Press, 1994.

[Nav09] R. Navigli. Word sense disambiguation. *ACM Computing Surveys*, 2009.

[PFVB19] P. Perbet, M. Fortin, A. Ville, and M. Béland. Near real-time deforestation detection in Malaysia and Indonesia using change vector analysis with three sensors. *International Journal of Remote Sensing*, 40(19):7439–7458, 2019.

[Pie05] C. Pierce. Issues of pragmatism. *The monist*, 15(4):481–99, 1905.

[Pie09] A. Pietruszczak. Simplified Kripke style semantics for modal logics K45, KB4 and KD45. *Bulletin of the Section of Logic*, 38(3/4):163–171, 2009.

[PKP19] A. Pietruszczak, M. Klonowski, and Y. Petrukhin. Simplified Kripke-Style Semantics for Some Normal Modal Logics. *Studia Logica*, 2019.

[PS07] A. Pauly and M. Schneider. Vague spatial data types. In S. Shekhar and H. Xiong, editors, *Encyclopedia of GIS*, pages 1213–1216. Springer, 2007.

[PTG12] S. T. Piantadosi, H. Tily, and E. Gibson. The communicative function of ambiguity in language. *Cognition*, 122(3):280–291, 2012.

[Qiz06] M. Qizilbash. Philosophical accounts of vagueness, fuzzy poverty measures and multidimensionality. In J. Silber, editor, *Economic Studies in Inequality, Social Exclusion and Well-Being*. Springer, 2006.

[Raf05] D. Raffman. Borderline Cases and Bivalence. *The Philosophical Review*, 114(1):1–31, 2005.

[Raf13] D. Raffman. *Unruly words : a study of vague language*. Oxford University Press, 2013.

[RAW+13] E. Romijn, J. J. H. Ainembabazi, A. Wijaya, M. Herold, A. Angelsen, L. Verchot, and D. Murdiyarso. Exploring different forest definitions and their impact on developing REDD+ reference emission levels: A case study for Indonesia. *Environmental Science and Policy*, 33:246–259, 2013.

[RGMW02] J. Rodd, G. Gaskell, and W. Marslen-Wilson. Making Sense of Semantic Ambiguity: Semantic Competition in Lexical Access. *Journal of Memory and Language*, 46(2):245–266, 2002.

[RL00] Y. Ravin and C. Leacock. Polysemy: an overview. In Y. Ravin and C. Leacock, editors, *Polysemy: Theoretical and computational approaches*, pages 1–29. Oxford University Press, 2000.

[Ron11] G. Ronzitti. *Vagueness : a guide*. Springer, 2011.

[Ros88] E. Rosch. Principles of Categorization. In A. Collins and E. E. Smith, editors, *Readings in Cognitive Science, a Perspective From Psychology and Artificial Intelligence*, pages 312–22. Morgan Kaufmann Publishers, 1988.

[Ros04] T. J. Ross. *Fuzzy Logic with Engineering Applications*. John Wiley & Sons, 2004.

[Rus23] B. Russell. Vagueness. *Australasian Journal of Psychology and Philosophy*, 1(2):84–92, 1923.

[Sau11] U. Sauerland. Vagueness in Language: The Case Against Fuzzy Logic Revisited. In P. Cintula, C. Fermüler, L. Godo, and P. Hájek, editors, *Understanding Vagueness: Logical, Philosophical, and Linguistic Perspectives*, pages 185–198. College Publications, 2011.

[SBS05] P. Santos, B. Bennett, and G. Sakellariou. Supervaluation semantics for an inland water feature ontology. In M. Zaanen, T. Oates, G. Paliouras, and C. Higuera, editors, *Proceedings of the 19th International Joint*

Conference on Artificial Intelligence (IJCAI'05), pages 564–569, Morgan Kaufmann Publishers, 2005.

[Sha06] S. Shapiro. *Vagueness in context.* Oxford University Press, 2006.

[SHS+18] X. Song, M. C. Hansen, S. V. Stehman, P. V. Potapov, A. Tyukavina, E. F. Vermote, and J. R. Townshend. Global land change from 1982 to 2016. *Nature*, 560(7720):639–643, 2018.

[SM98] B. Smith and D. M. Mark. Ontology and geographic kinds. In T. K. Poiker and N. Chrisman, editors, *Proceedings of the 8th International Symposium on Spatial Data Handling (SDH'98)*, pages 308–320, International Geographic Union, 1998.

[Sor88] R. A. Sorensen. *Blindspots.* Oxford: Claredon Press, 1988.

[Sor01] R. A. Sorensen. *Vagueness and Contradiction.* Oxford University Press, 2001.

[Sor18] R. Sorensen. Vagueness. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy (Summer 2018 Edition)*, Stanford University, 2018.

[SSR19] I. Schulte, C. Streck, and S. Roe. Protecting and Restoring Forests: A Story of Large Commitments yet Limited Progress. New York Declaration on Forests Five-Year Assessment Report. Technical report, NYDF Assessment Partners, 2019.

[Str01] U. Straccia. Reasoning within fuzzy description logics. *Journal of Artificial Intelligence Research*, 14(1):137–166, 2001.

[Swa14] E. Swanson. Ordering Supervaluationism, Counterpart Theory, and Ersatz Fundamentality. *Journal of Philosophy*, 111(6):289–310, 2014.

[Tar80] A. Tarski. *Logic, semantics, metamathematics: papers from 1923 to 1938.* Hackett Publishing, 1980.

[TK04] E. Tomai and M. Kavouras. From Onto-GeoNoesis to Onto-Genesis: The Design of Geographic Ontologies. *GeoInformatica*, 8(3):285–302, 2004.

[TSG09] D. Tagliapietra, M. Sigovini, and V. Ghirardini. A review of terms and definitions to categorise estuaries, lagoons and associated environments. *Marine and Freshwater Research*, 60(6):497–509, 2009.

[Tug93] D. Tuggy. Ambiguity, polysemy, and vagueness. *Cognitive Linguistics*, 4(3):273–290, 1993.

[Tye94] M. Tye. Sorites Paradoxes and the Semantics of Vagueness. *Philosophical Perspectives*, 8:189, 1994.

[Ung80] P. Unger. The Problem of the Many. *Midwest Studies In Philosophy*, 5(1):411–468, 1980.

[Var01] A. C. Varzi. Vagueness in geography. *Philosophy & Geography*, 4(1):49–65, 2001.

[Var04] A. C. Varzi. Higher-Order Vagueness and the Vagueness of 'Vague'. *Mind*, 112(446):295–299, 2004.

[Var07] Supervaluationism and Its Logics. *Mind*, 116(463):633–676, 2007.

[vF68] B. C. van Frassen. Presupposition, implication and self-reference. *Journal of Philosophy*, 65(5):136–52, 1968.

[VP06] C. Van Petten. Lexical Ambiguity Resolution. In *Encyclopedia of Cognitive Science*. John Wiley & Sons, 2006.

[Wai45] F. Waismann. Verifiability. In H. R., editor, *How I See Philosophy*, pages 39–66. Palgrave Macmillan, 1945.

[Wil92] T. Williamson. Vagueness and ignorance. *Proceedings of the Aristotelian Society*, Suplementary Volume 66:145–162, 1992. Reprinted in [KS96b].

[Wil94] T. Williamson. *Vagueness*. Routledge, 1994.

[Wil99] T. Williamson. On the structure of higher-order vagueness. *Mind*, 108(429):127–143, 1999.

[Wil08] J. R. G. Williams. Supervaluationism and Logical Revisionism. *The Journal of philosophy*, 105(4):192–212, 2008.

[Wil18] T. Williamson. Supervaluationism and good reasoning. *Theoria*, 33(3):521–537, 2018.

[Wit53] L. Wittgenstein. *Philosophical Investigations*. G.E.M. Anscombe (trans.) and R. Rhees, editors, Blackwell, 1953.

[Wit09] L. Wittgenstein. *Philosophical Investigations*. P.M.S. Hacker and J. Schulte, editors, Wiley-Blackwell, 2009. Reprint from: [Wit53].

[WJ94] M. Wooldridge and N. R. Jennings. Agent theories, architectures, and languages: A survey. In *International Workshop on Agent Theories, Architectures, and Languages*, pages 1–39, Springer, 1994.

[WL16] C. Wurm and T. Lichte. The proper treatment of linguistic ambiguity in ordinary algebra. In A. Foret, G. Morrill, R. Muskens, R. Osswald, and S. Pogodalla, editors, *Proceedings of the International Conference on Formal Grammar (FG'16)*, pages 306–322, Springer, 2016.

[Wol10] M. P. Wolf. Philosophy of Language. Encyclopedia of Philosophy Pub, 2010.

[Wol13] M. Wolski. Science and Semantics: A Note on Rough Sets and Vagueness. In A. Skowron and Z. Suraj, editors, *Rough Sets and Intelligent Systems*, chapter 24, pages 623–643. Springer, 2013.

[Wöl15] S. Wölfl. *Introduction to Modal Logics*. Unpublished, 2015.

[Wor02] Global forest watch. Technical report, World Resources Institute, 2002.

[WPB05] T. Wasow, A. Perfors, and D. Beaver. The puzzle of ambiguity. In C. O. Orgun and P. Sells, editors, *Morphology and the web of grammar: Essays in memory of Steven G. Lapointe*, pages 265–282. CSLI Publications, 2005.

[Wri87] C. Wright. Further Reflections on the Sorites Paradox. 15:227–290, 1987.

[Wri10] C. Wright. The Illusion of Higher-Order Vagueness. In R. Dietz and S. Moruzzi, editors, *Cuts and Clouds: Vagueness, its Nature, & its Logic*, chapter 31. Oxford University Press, 2010.

[WZ01] F. Wolter and M. Zakharyaschev. Decidable fragments of first-order modal logics. *Journal of Symbolic Logic*, 66(3):1415–1438, 2001.

[Xu07] F. Xu. Sortal concepts, object individuation, and language. *Trends in Cognitive Sciences*, 11(9):400–406, 2007.

[YCF01] D. Yarowsky, S. Cucerzan, and R. Florian. The Johns Hopkins senseval2 system descriptions. In *The Proceedings of the 2nd International Workshop on Evaluating Word Sense Disambiguation Systems*, pages 163–166, Association for Computational Linguistics, 2001.

[Zad65] L. Zadeh. Fuzzy sets. *Information and Control*, 8(3):338–353, 1965.

[Zad75] L. Zadeh. Fuzzy Logic and Approximate Reasoning. *Synthese*, 30:407–428, 1975.

[Zha98] Q. Zhang. Fuzziness - vagueness - generality - ambiguity. *Journal of Pragmatics*, 29(1):13–31, 1998.

[Zip49] G. K. Zipf. Human Behavior and the Principle of Least Effort. *Addisson-Wesley Press*, 1949.