

Standpoint Logic: Multi-Perspective Knowledge Representation

Lucía GÓMEZ ÁLVAREZ and Sebastian RUDOLPH
Computational Logic Group, TU Dresden, Germany

Abstract. Ontologies and knowledge bases encode, to a certain extent, the standpoints or perspectives of their creators. As differences and conflicts between standpoints should be expected in multi-agent scenarios, this will pose challenges for shared creation and usage of knowledge sources.

Our work pursues the idea that, in some cases, a framework that can handle diverse and possibly conflicting standpoints is more useful and versatile than forcing their unification, and avoids common compromises required for their merge. Moreover, in analogy to the notion of *family resemblance concepts*, we propose that a collection of standpoints can provide a simpler yet more faithful and nuanced representation of some domains.

To this end, we present *standpoint logic*, a multi-modal framework that is suitable for expressing information with *semantically heterogeneous* vocabularies, where a *standpoint* is a partial and acceptable interpretation of the domain. Standpoints can be organised hierarchically and combined, and complex correspondences can be established between them. We provide a formal syntax and semantics, outline the complexity for the propositional case, and explore the representational capacities of the framework in relation to standard techniques in ontology integration, with some examples in the *Bio-Ontology* domain.

Keywords. standpoint, perspectives, modal logic, ontology integration

1. Introduction

Natural language terms do not have precise, universally agreed definitions that fix their meanings [1]. Instead, their applicability is unclear in some instances, and it may vary depending on the context and pragmatics of use. In borderline cases, the speaker must make a *semantic commitment*: she must decide whether a term is applicable or not.

When an agent formalises a domain, the resulting conceptual model will also be shaped by this kind of semantic commitments, which are in turn influenced by her own worldview and by pragmatic factors, such as the intended granularity and scope.

Even simple domains like colours lend themselves to this issue: while different speakers may agree on most clear instances of *red*, they may disagree on the existence of certain colours (e.g. *vermilion*) and on systems of classification. For instance, consider two (partial) formalisations coming from different perspectives: an ontology of *colour theory* (*CT*), a discipline within the fine arts tradition, and a ‘common-sense’ representation of ink colours, used by an online *house painting* business (*HP*).




Figure 1. The conceptual models of s_{CT} and s_{HP} , each in a box. Each concept is represented with its initial. In dotted lines one can see the alignments, with their type and similarity in parentheses.

Example 1. It is universally accepted (thus also by CT and HP) that Yellow, Green, and Lime are Colour. According to CT , Colour is either a warm colour (WC) or a cold colour (CC) but not both¹, Yellow is a WC , and Green and Lime are CC . According to HP , Lime is Green and Yellow.

While ontologies are aimed at providing a common vocabulary for representing shared knowledge [2], the described semantic heterogeneity may hinder the interoperation of independently developed systems. Ontology integration is well-studied yet still challenging [3,4,5], and ontology merges often involve certain knowledge loss or weakening in order to avoid incoherence and inconsistency [6,7].

This is the case in Example 1, for which Figure 1 shows the two conceptual models in distinct boxes, as well as the alignments (in dotted lines) that could be found between them using some standard matching algorithm. While simple and intuitively closely aligned, the two conceptualisations cannot be trivially merged without the undesired consequence of Lime becoming unsatisfiable: to prevent this, we must either (1) give up on the JEPD relations of CT , or (2) let go one of the subsumptions of Lime in HP or (3) relinquish one of the alignments, such as Colour, and instead duplicate the concept: Colour. $_{CT}$ and Colour. $_{HP}$. Yet, it is preferable for the integration to preserve all entailments of the source ontologies and their alignments [8]. Moreover, in areas of growing interest, such as the research on complex alignments and *holistic* (many-source) ontology integration [9,10,11], even more inconsistencies have to be expected.

In this paper, we advocate a multi-perspective approach that can represent and reason with many – possibly conflicting – standpoints, instead of focusing on combining and merging different sources into a single conceptual model.

Beyond the challenges in knowledge integration, we believe that this is also useful in the process of formalising a domain [12]. Typically, heavy axiomatisations enable the derivation of interesting facts, but limit the interoperability of the system; the converse happens with shallow modelling that relies on little more than taxonomic relationships. With our *multi-standpoint* framework, we aim 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 associated with different users or interpretations of the domain.

We propose *standpoint logic* (Section 3): a simple formalism rooted in modal logic that supports the coexistence of multiple *standpoints* and the establishment of alignments between them. We highlight that our proposal retains good computational properties: in its

¹These are jointly exhaustive and pairwise disjoint (JEPD) categories of colour.

propositional version, reasoning in the logic is NP-complete (Section 4) in pleasant contrast to the PSPACE-completeness normally exhibited by multi-modal epistemic logics. Then, we proceed to show how the framework allows for the establishment of structures of standpoints and the expression of ‘complex alignments’ (Section 5), and we illustrate its use in an application in the bio-ontology domain (Section 6). We discuss the background and related work (Section 7) before concluding and providing an outlook on future research (Section 8).

2. Background and Framework Overview

Standpoint Logic is a lightweight multi-modal framework where labelled modalities \Box_s express information relative to a standpoint, and the set of formulae under a standpoint (sentences $\Box_s \phi$ or $\Diamond_s \phi$) represent the agent’s world-view or semantic commitments.² Standpoint logic draws from the philosophical theory of supervaluationism, according to which the phenomenon of vagueness (and more generally the semantic variability) can be explained by the fact that natural language can be interpreted in many different yet equally acceptable ways [13], commonly referred to as *precisifications*. Early proposals of this intuition were made by Mehlberg [14], and Fine [15] applied this model to the analysis of vagueness. Supervaluationism is a popular theory of vagueness, adopted by philosophers, logicians and linguists, yet scarce in KR. An exception is the earlier work of Bennett [16], from whom we borrow the notion of standpoint and together we propose a different treatment, already given in [17].

When using a modal infrastructure in a supervaluationistic framework, one replaces the usual structure of *possible worlds* by one of *precisifications*. To see the difference, consider a situation modelled with doxastic logic: Bob believes that there is a red apple at home ($B_{Bob}[\text{red_apple}]$) and Tim believes that there is a yellow apple instead, ($B_{Tim}[\text{yellow_apple}]$). Here, the possible worlds model the different possible states of affairs, and there is an *actual world* that dictates what is contingently true. In contrast, consider Tim and Bob facing a red/yellowish apple. Tim calls it a red apple ($\Box_{Tim}[\text{red_apple}]$), yet Bob calls it a yellow apple, ($\Box_{Bob}[\text{yellow_apple}]$). The latter is what standpoint logic models, where agents describe the state of the matter using different and equally acceptable interpretations of the vocabulary, so there is no ‘actual precisification’.

Different modal frameworks have been proposed in the supervaluationist literature [18, 19,20]. These focus on proposing modalities that can capture the linguistic behaviour of philosophical vagueness and the sorites paradox, and on the analysis of different forms of validity and logical consequence. In contrast, we use the supervaluationist model of natural language (in terms of a collection of admissible classical interpretations) but we focus on scenarios where different agents or different contexts are linked to different usages of such a semantically variable language.

Standpoints are modelled as non-empty sets of *precisifications*, which corresponds to the intuition that a standpoint is (typically) a *partial semantic commitment* that can be made fully precise in different ways. Practical uses of *standpoint logic* include, for instance, the representation of multiple (and possibly conflicting) symbolic conceptualisations of

²For clarity and better readability, we will sometimes use square brackets [...] to explicitly indicate the scope of the modal operators.

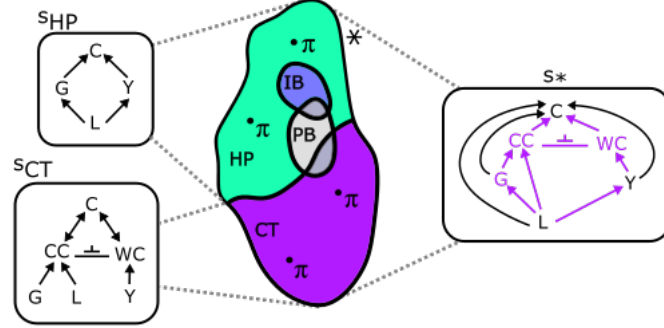


Figure 2. In the center, the precisification space, labelled with $*$, and the standpoints s_{CT} , s_{HP} , s_{PB} and s_{IB} (subsets of $*$). On the sides, the conceptual models of s_{CT} , s_{HP} , and s_{*} .

a domain, knowing what can be inferred from the consensual semantics of sets of agents and knowing which standpoints are compliant with a partial truth.

In Fig. 2 we can see a representation of the Example 1 in a standpoint framework. The shape in the centre, labelled with $*$, represents the set of all admissible precisifications, where the points (labelled π) are individual precisifications and the subsets (subareas of $*$) are the standpoints s_{CT} , s_{HP} , s_{PB} and s_{IB} . In boxes, on the sides, one can see the corresponding conceptual models that hold for s_{CT} , s_{HP} and s_{*} , linked with dotted lines to the relevant subset of precisifications that satisfies them. For instance, from the box s_{HP} we can see that $\Box_{HP}[L \rightarrow (G \wedge Y)]$, which means that $L \rightarrow (G \wedge Y)$ will hold in all the precisifications that belong to the area HP . In the box of s_{*} , diamond statements are represented in a lighter colour, and generally arrows stand for implications.

In terms of syntax, the standpoint framework exhibits a reasonable versatility through a very economic extension of the (in this case: propositional) base language by constants s for denoting standpoints – with the special constant $*$ denoting the global standpoint – and corresponding pairs of dual modal operators \Box_s and \Diamond_s .

$\Box_s \phi$ (or $\Box_* \phi$) reads as “It is unequivocal, [from the point of view s ,] that ϕ ”. This is the paradigmatic form of a semantic commitment (in the case of $*$, the statement is global, i.e., universally agreed upon).

$\Diamond_s \phi$ (or $\Diamond_* \phi$) reads as “[From the point of view s ,] in some sense ϕ ”. In practice, this means that the standpoint s (or $*$) has no argument to rule out ϕ , and hence, it is acceptable to interpret s (or $*$) in such a way that ϕ holds.

In scenarios with multiple perspectives like Example 1, we can use standpoint logic to represent and infer different kinds of facts, namely (a) global or standpoint-relative statements, (b) unequivocal, and ‘in some sense’ statements, (c) hierarchies and globality of standpoints, such as $(IB \preceq HP)$, which models that the standpoint IB is ‘subsumed’ by the standpoint HP , and (d) ‘complex alignments’ between standpoints, such as $\Box_{HP}[\text{Green}] \rightarrow (\Box_{CT}[\text{Green}] \vee \Box_{HP}[\text{Lime}])$. Intuitively the latter would mean that if Green holds according to HP , then either Green also holds according to CT or the colour is specifically Lime for HP (or both).

In the rest of this section, we demonstrate how to specify Example 1 using simple statements of the kinds (a) and (b), and we will present some of the inferences that we can obtain. Later, in Section 5, we will address more complex statements of the kinds (c) and

(d) by extending and enriching the same example, after presenting the formal specification of the language in Section 3. We now proceed with a simple standpoint formalisation of our example:

- (1). $\Box_*[(\text{Yellow} \vee \text{Green} \vee \text{Lime}) \rightarrow \text{Colour}]$
- (2). $\Box_{CT}[\text{Colour} \leftrightarrow (\text{WC} \vee \text{CC})]$
- (3). $\Box_{CT}[(\text{CC} \leftrightarrow \neg \text{WC})]$
- (4). $\Box_{CT}[((\text{Green} \vee \text{Lime}) \rightarrow \text{CC}) \wedge (\text{Yellow} \rightarrow \text{WC})]$
- (5). $\Box_{HP}[\text{Lime} \rightarrow (\text{Yellow} \wedge \text{Green})]$

From this, we can derive the expected facts for each theory/standpoint specifically, such as ‘Unequivocally, according to *CT*, if a colour is lime then it is not yellow’ ($\Box_{CT}[(\text{Lime} \rightarrow \neg \text{Yellow})]$), and general findings, such as ‘Under some interpretations, colours can be classified into cold and warm’ ($\Diamond_*[\text{Colour} \rightarrow (\text{CC} \vee \text{WC})]$).

The crucial aspect of this framework in contrast to the merging strategy is that the set of sentences (1)-(5) is not inconsistent, so all axioms and known alignments can be jointly represented. Standpoint logic escapes global inconsistency, without removing alignments or relations and avoiding the duplication of entities, because the model theory based on Kripke-structures (cf. Section 3) forces consistency only within standpoints and precisifications, but allows for the specification of sets of standpoints that are inconsistent among them, which are modelled as disjoint sets. With this, we overcome the *merge* tradeoff discussed in Section 1.

3. Syntax and Semantics of Standpoint Logic

We now formally specify propositional *Standpoint Logic*, denoted with \mathbb{S} . Essentially, we define it as a multi-modal normal logic satisfying **KD45** (and **S5** in the case of the universal standpoint),³ and additionally the stronger axioms **4'**, **5'** and **P**, describing the interaction of different standpoints. Hence, the formalism is an extension of well-known systems used for epistemic logics. We first define the syntax of \mathbb{S} .

Definition 1 (Syntax of Propositional Standpoint Logic). *A vocabulary is a tuple of the form $\mathcal{V} = \langle \mathcal{P}, \mathcal{S} \rangle$, where \mathcal{P} is a non-empty set of propositional variables and \mathcal{S} is a non-empty set of standpoint symbols, containing the distinguished symbol $*$. We denote elements of \mathcal{P} by p and elements of \mathcal{S} by s , potentially with extra decorations. The language $\mathcal{L}_{\mathbb{S}}$ of \mathbb{S} Propositions (denoted by ϕ, ϕ_1, ϕ_2) are defined by*

$$\phi ::= s' \preceq s \mid p \mid \neg \phi \mid (\phi_1 \wedge \phi_2) \mid \Box_s \phi.$$

where $s', s \in \mathcal{S}$, $p \in \mathcal{P}$, and \Box_s is called the standpoint operator for s . We call (sub)formulae of the form $s' \preceq s$ sharpening statements while those of the form p are referred to as atomic propositions. We also allow the connectives \vee , \rightarrow , \leftrightarrow , and \Diamond_s as shorthands with their usual definitions.

³These are systems of modal logic characterised by the axioms in Table 1. The system **KD45** is characterised by the axioms in its name and **S5** by K, D, T, 4 and 5.

Axiom	Schema	Property	Axiom	Schema	Property
K	$\Box_s(\phi \rightarrow \psi) \rightarrow (\Box_s \phi \rightarrow \Box_s \psi)$	All	4'	$\Box_s \phi \rightarrow \Box_{s'} \Box_s \phi$	Trans-transitive
D	$\Box_s \phi \rightarrow \Diamond_s \phi$	Serial	5'	$\Diamond_s \phi \rightarrow \Box_{s'} \Diamond_s \phi$	Trans-Euclidean
4	$\Box_s \phi \rightarrow \Box_s \Box_s \phi$	Transitive	T*	$\Box_* \phi \rightarrow \phi$	Reflexive
5	$\Diamond_s \phi \rightarrow \Box_s \Diamond_s \phi$	Euclidean	P	$(s' \preceq s) \rightarrow (\Box_s \phi \rightarrow \Box_{s'} \phi)$	

Table 1. Correspondence between the axioms in \mathbb{S} and the properties of the relations of the models $\mathfrak{M}_{\mathbb{S}}$.

In addition, we can define other useful modal operators in $\mathcal{L}_{\mathbb{S}}$:

$$\mathcal{I}_s \phi := (\Diamond_s \phi \wedge \Diamond_s \neg \phi) \quad \mathcal{D}_s \phi := (\Box_s \phi \vee \Box_s \neg \phi) = \neg \mathcal{I}_s \phi$$

$\mathcal{I}_s \phi$ expresses that the truth of ϕ is indeterminate (i.e. borderline) according to standpoint s , while $\mathcal{D}_s \phi$ expresses that it is determinate.

Semantics of modal logic are typically provided proof-theoretically, and in this spirit we now give a Hilbert-style axiomatic proof system for \mathbb{S} , before complementing it with a fitting model-theoretic semantics. We write $\vdash_{\mathbb{S}} \phi$ to mean that ϕ is a derivable theorem of \mathbb{S} (i.e. ϕ is derivable not requiring any premises). As standpoint logic is built upon an underlying classical logic, all classically valid propositional formulas are theorems. In addition, the proof system of \mathbb{S} provides the axiom schemas **K**, **D**, **T***, **4'**, **5'**, and **P** as displayed in Table 1. The inference rules of \mathbb{S} are the standard ones for modal logic: classical theorems are provable (**RC**), all instances of the axioms are provable (**RA**), the classical *modus ponens* (**MP**: if $\vdash_{\mathbb{S}} \phi$ and $\vdash_{\mathbb{S}} \phi \rightarrow \psi$, then $\vdash_{\mathbb{S}} \psi$), and the rule of necessitation (**RN**: if $\vdash_{\mathbb{S}} \phi$, then $\vdash_{\mathbb{S}} \Box_s \phi$, for all standpoints $s \in S$).

We proceed with a corresponding model-theoretic semantics of \mathbb{S} . We first propose a definition that captures the notion of *standpoint* well but deviates from the usual definition via Kripke models. Thereafter, we will justify that choice by arguing that there is a one-to-one correspondence between models according to our semantics and a particular class of Kripke models, allowing us to relate our work to other frameworks and results.

Definition 2 (Semantics of Standpoint Logic). *Given a vocabulary $\mathcal{V} = \langle \mathcal{P}, \mathcal{S} \rangle$, a model \mathcal{M} (over \mathcal{V}) is a triple $\langle \Pi, \sigma, \delta \rangle$, where Π is a non-empty set of precisifications, while $\sigma : \mathcal{S} \rightarrow 2^{\Pi}$ and $\delta : \mathcal{P} \rightarrow 2^{\Pi}$ are functions that assign sets of precisifications to standpoint symbols and propositional variables, respectively, satisfying $\sigma(s) \neq \emptyset$ for all $s \in \mathcal{S}$ as well as $\sigma(*) = \Pi$. The set of all such models is denoted by $\mathfrak{M}_{\mathbb{S}}$. For a model \mathcal{M} with a distinguished precisification $\pi \in \Pi$ we define satisfaction of formulae as follows:*

- $(\mathcal{M}, \pi) \models \phi_1 \wedge \phi_2$ iff $(\mathcal{M}, \pi) \models \phi_1$ and $(\mathcal{M}, \pi) \models \phi_2$,
- $(\mathcal{M}, \pi) \models p$ iff $\pi \in \delta(p)$,
- $(\mathcal{M}, \pi) \models \Box_s \phi$ iff $(\mathcal{M}, \pi') \models \phi$ for all $\pi' \in \sigma(s)$,
- $(\mathcal{M}, \pi) \models \neg \phi$ iff $(\mathcal{M}, \pi) \not\models \phi$,
- $(\mathcal{M}, \pi) \models s' \preceq s$ iff $\sigma(s') \subseteq \sigma(s)$.

We read $(\mathcal{M}, \pi) \models \phi$ as: ϕ is true at the precisification π in model \mathcal{M} . We write $\mathcal{M} \models \phi$ (read: \mathcal{M} is a model of ϕ) if $(\mathcal{M}, \pi) \models \phi$ holds for all $\pi \in \Pi$. We call ϕ satisfiable if it has a model and valid if every element of $\mathfrak{M}_{\mathbb{S}}$ is a model of it.

As immediate consequence from this definition, we obtain that $\mathcal{M} \models \phi$ if and only if $\mathcal{M} \models \Box_* \phi$. Moreover, ϕ is valid if and only if $\Diamond_* \neg \phi$ is unsatisfiable.

$\mathfrak{M}_{\mathbb{S}}$ is not defined in the usual Kripke-style way using accessibility relations. Rather, the latter are replaced by the function σ following the intuition of *standpoint*, where standpoint symbols are associated with non-empty sets of precisifications.

However, **Definition 2** can be easily recast in terms of a set of accessibility relations $\mathcal{R} = \{R_s \mid s \in \mathcal{S}\}$ over the set of precisifications Π (as in standard Kripke semantics), by letting

$$R_s := \Pi \times \sigma(s) = \{\langle \pi, \pi' \rangle \mid \pi \in \Pi, \pi' \in \sigma(s)\}$$

That is, for each standpoint s , the relations thus obtained connect *all* precisifications with all those in $\sigma(s)$. Then, we obtain the desired correspondence.

Lemma 1. $(\mathcal{M}, \pi) \models \Box_s \phi$ if and only if $(\mathcal{M}, \pi') \models \phi$ for all π' with $(\pi, \pi') \in R_s$.

Soundness and completeness of the presented proof system with respect to the model-theoretic semantics can be shown by standard arguments. This includes relating the properties of our models to the axiomatisation of \mathbb{S} using the well understood correspondence theory [21] (cf. Table 1). In particular, by construction, the standpoint accessibility relations are serial, transitive and euclidean, and the standpoint $*$ is a universal relation. This indeed corresponds to the **KD45** axiomatisation for the standpoints and the **S5** for $*$.

Discussion. As we have already mentioned, the main formal particularities of *standpoint logic* with regards to other modal frameworks are the axioms **P**, **4'** and **5'**. Axiom **P** captures the meaning of \preceq , by ensuring that any proposition considered definite in a given standpoint is also considered definite in any sharper standpoint. This simple mechanism is the basis for the construction of hierarchies of standpoints as well as for combinations. The axioms **4'** and **5'** are the interaction axioms, and are stronger than the well known modal axioms **4** and **5** (which are immediately derivable). This means that assertions such as $(\Box_a \Diamond_b \Box_c \dots \Box_s) \phi$ can be simplified into $\Box_s \phi$. While this may seem unrealistic at first glance, and only motivated by the reduction of complexity it brings about, it is in fact a desirable feature. The key is that standpoints do not model the epistemic state of agents (in which case an assertion like ‘Agent α knows that agent β knows ϕ ’ makes sense), they model the set of semantic commitments associated to a particular perspective. Moreover, we note that standpoints are not allowed to be empty (by axiom **D**), for they pick up compatible precisifications; if there are none, then the standpoint is deemed incoherent.

4. Translation into One-Variable First-Order Logic and Complexity

We next turn to the question regarding the difficulty of reasoning in \mathbb{S} . To this end, we will provide a polytime translation from \mathbb{S} into one-variable first-order logic, which not only settles the above question but is also interesting in its own right.

Definition 3. The function $\text{trans} : \mathcal{L}_{\mathbb{S}} \rightarrow \mathcal{L}_{FO1}$, mapping \mathbb{S} formulae to formulae in one-variable first-order predicate logic, is recursively defined as follows (with symbols from \mathcal{P} and \mathcal{S} repurposed as unary predicates, $s', s \in \mathcal{S}$ and $p \in \mathcal{P}$):

$$\begin{array}{ll} \text{trans}(p) &= p(x) & \text{trans}(\phi_1 \wedge \phi_2) &= \text{trans}(\phi_1) \wedge \text{trans}(\phi_2) \\ \text{trans}(\neg\phi) &= \neg\text{trans}(\phi) & \text{trans}(\Box_s \phi) &= \forall x.(s(x) \rightarrow \text{trans}(\phi)) \\ & & \text{trans}(s' \preceq s) &= \forall x.(s'(x) \rightarrow s(x)) \end{array}$$

Finally, for an \mathbb{S} formula ϕ with occurrences of standpoint constants s_1, \dots, s_k , let

$$\text{Trans}(\phi) := \forall x.(\text{trans}(\phi)) \wedge \forall x.(*(x)) \wedge \exists x.(s_1(x)) \wedge \dots \wedge \exists x.(s_k(x)).$$

We note that Trans produces formulae of linear size w.r.t. $|\phi|$. We now show that the translation is indeed satisfiability preserving, as intended.

Theorem 2. *An \mathbb{S} formula ϕ is satisfiable if and only if $\text{Trans}(\phi)$ is FO-satisfiable.*

Proof. (\Rightarrow) Assuming satisfiability ϕ consider some model $\mathcal{M} = \langle \Pi, \sigma, \delta \rangle$ of it. Let now $\text{FO}(\mathcal{M})$ denote the first-order interpretation with domain Π , satisfying $s^{\text{FO}(\mathcal{M})} = \sigma(s)$ for all $s \in \mathcal{S}$ and $p^{\text{FO}(\mathcal{M})} = \delta(p)$. Then it can be shown by a straightforward structural induction over ϕ that, for every $\pi \in \Pi$, $(\mathcal{M}, \pi) \models \phi$ if and only if $\text{FO}(\mathcal{M}), \{x \mapsto \pi\} \models \text{trans}(\phi)$. It is then a direct consequence, that $\text{FO}(\mathcal{M})$ is a model of the first conjunct of $\text{Trans}(\phi)$. Satisfaction of the other conjuncts follows from **Def. 2** via the definition of $\text{FO}(\mathcal{M})$. Hence $\text{FO}(\mathcal{M})$ is a model of $\text{Trans}(\phi)$, witnessing its satisfiability.

(\Leftarrow) Consider a first-order model \mathcal{M}' of $\text{Trans}(\phi)$. We now define the \mathbb{S} model $\mathbb{S}(\mathcal{M}') = \langle \Pi, \sigma, \delta \rangle$ by letting Π be the domain of \mathcal{M}' and stipulating $\delta(p) = p^{\mathcal{M}'}$ for all $p \in \mathcal{P}$ as well as $\sigma(s) = s^{\mathcal{M}'}$ whenever $s \in \{s_1, \dots, s_k\}$ and $\sigma(s) = \Pi$ otherwise. It is easily checked that, due to the second to last conjunct of Trans , the structure thus defined is indeed an \mathbb{S} model. In order to show that $\mathbb{S}(\mathcal{M}')$ is a model of ϕ , we can proceed as before and show by a straightforward structural induction over ϕ that, for every $\pi \in \Pi$, $(\mathbb{S}(\mathcal{M}'), \pi) \models \phi$ if and only if $\mathcal{M}', \{x \mapsto \pi\} \models \text{trans}(\phi)$. Thus, the established modelhood of $\mathbb{S}(\mathcal{M}')$ ensures satisfiability of ϕ . \square

It is folklore that the satisfiability problem of one-variable first-order logic is decidable and, in fact, in NP, membership even having been established for much more expressive logics [22]. As Trans realizes a polynomial reduction from satisfiability in \mathbb{S} to satisfiability to one-variable first-order logic, this membership carries over. On the other hand, \mathbb{S} subsumes propositional logic, which is known to have an NP-hard satisfiability problem. Hence we can conclude the following.

Corollary 3. *The satisfiability problem of \mathbb{S} is NP-complete.*

5. Integrating Different Perspectives with SL

So far, we have shown how we can formalise the conflicting perspectives of our Example 1 to overcome some limitations of traditional merging approaches. We now extend this example to briefly illustrate other capacities of the language, covering the representation of alignments, standpoint nesting and standpoint combination. When representing alignments, we will assume that the alignments themselves have been obtained via traditional ontology matching techniques or that they may be known to the domain experts.

‘Alignment’ representation. Standpoint logic allows for a reasonable nuance in the representation of correspondences between the entities of different standpoints compared to other formalisms. On the one hand, we usually interpret simple scenarios where we have

an equivalence relation like $(\equiv, 0.78)$ ⁴ as both entities being different perspectives of a single concept. This is the case that we have already widely covered, and it only involves, when necessary, the uniform renaming of entities.

The case of subsumed relations is more nuanced. What are called subsumptions in the context of knowledge integration can be either ‘genuine subsumptions’, such as $\langle O1:Lime, O2:Colour, \subseteq, 0.84 \rangle$ or ‘sharpenings’, where the concept is intuitively equivalent but one ontology has stricter criteria of application than the other, such as $\langle O1:Green, O2:Green, \subseteq, 0.73 \rangle$. Following the reported evidence in [11], we assume that these cases can be often recognised because an additional equivalence alignment is also found, $\langle O1:Green, O2:Green, \equiv, 0.56 \rangle$, as illustrated in Figure 1. In this case, we suggest that it is more appropriate to understand the relation as a subsumption between the standpoints on that concept, rather than a subsumption of different concepts:

$$(6). \quad \Box_{s_1}[\text{Green}] \rightarrow \Box_{s_2}[\text{Green}]$$

Beyond the more faithful representation of the correspondence, this approach limits the multiplication of entities in the merged ontology, which can otherwise hinder its usability. Finally, the framework allows for the representation of correspondences as complex as allowed by the base logic language. E.g., $\Box_{HP}[\text{Green}] \leftrightarrow (\Box_{CT}[\text{Green}] \vee \Box_{HP}[\text{Lime}])$ specifies that if Green holds according to *HP*, then either Green also holds according to *CT* or the colour is specifically Lime for *HP* (or both). Note that the behaviour of these sentences is in fact more similar to *bridge rules*, because rather than unifying the theories, they only establish a correspondence between what holds for different standpoints.

Standpoint Hierarchies and Combinations. Statements like (6) are radically different from those of the kind $(s_2 \preceq s_1)$, as the latter establishes a relation between full standpoints. Let us consider a use case scenario of \preceq ,

Example 2. An Ink brand with the standpoint s_{IB} reuses the categorisation of *HP*, and in addition it specifies that Ochre and Gold are (types of) Yellow.

$$(7). \quad \Box_{IB}[(\text{Gold} \vee \text{Ochre}) \rightarrow \text{Yellow}]$$

$$(8). \quad IB \preceq HP$$

Statement (8) specifies that *IB* is a sharpening of *HP*, that is, *IB* satisfies all the constraints of *HP* as well as its own commitments (7). Semantically, this is modelled as a subset relationship: the precisifications belonging to *IB* are a subset of those belonging to *HP*. This is relatable to the process of *importing* an ontology. Yet, sharpening a standpoint is slightly different from importing it: From the formulae (7) and (8), we can infer not only propositions that hold according to *IB* but also about *HP*. For instance, the two statements imply that, under *HP*’s perspective, it is admissible to interpret ochre as a kind of yellow: $\Diamond_{HP}[\text{Ochre} \rightarrow \text{Yellow}]$. As a consequence, if there was another standpoint *IB2* such that $\Box_{IB2}[\text{Ochre} \rightarrow \neg \text{Yellow}] \wedge IB2 \preceq HP$, and we knew $\Box_{HP}[\text{Ochre}]$, then one could infer that $\mathcal{I}_{HP}[\text{Yellow}]$, that is, that yellow is inherently borderline or indeterminate for s_{HP} since we have evidence that it can be sharpened in opposing ways.

⁴We use here the standard notation in ontology alignment, where matchings between two entities have a relation type, in this case \equiv , and a similarity or confidence measure, in this case 0.78.




Figure 3. Ambiguous equivalence correspondences leading to redundancies and cycles. Reprinted from [11].

With regards to the combination of standpoints, let us consider an example of reasoning in scenarios where two independently developed models are relevant, namely our previous *IB* and another standpoint *PB*.

Example 3. Assume there is a paint brand with standpoint *PB*, coming with some more definitions and axioms. If the standpoints *IB* and *PB* overlap, we can introduce a new, joint standpoint *IBPB* and define it as sharpening of both *IB* and *PB* to merge the theories.

$$(9). IBPB \preceq IB \wedge IBPB \preceq PB$$

Note that, so far, our framework does not allow to precisely refer to the intersection of the two standpoints (only at any possible subset of their intersection), but we will discuss an extension to this effect in the conclusion.

6. Application in the Biological Domain

In this section we consider two possible applications of the standpoint framework in the context of biological ontologies. First, it might serve as an alternative approach to reported integration challenges that arose in the *LargeBio* track of the Ontology Alignment Evaluation Initiative 2020 [23]. Second, it may help to address the semantic heterogeneity challenges around the concept *forest* in the EnvO ontology.

First, let us consider an alignment scenario extracted from an experiment in [11], where the goal is the holistic integration of three ontologies (using pairwise alignments) from the *LargeBio* track: FMA (Foundational Model of Anatomy), SNOMED-CT (Clinical Terms), and NCI (National Cancer Institute Thesaurus). In what follows, we will refer to the entities of these sources by their initials, for the sake of brevity.

One case of this scenario (extracted from [11]) is illustrated in Figure 3, which contains a snippet from the alignments found between FMA and SNOMED-CT on the left, and their naive integration on the right. The integration displays subsumption redundancies and a




Figure 4. Ambiguous equivalence correspondences leading to inconsistency. Reprinted from [11].

subsumption cycle formed because of the addition of two equivalence correspondences having the same source class ‘001#Abdominal_lymph_node’ (ALN), resulting in three equivalent classes merged together to constitute a single class. In [11], it is proposed to fix this merge by giving up one of the alignments, hence concluding either $ALN \equiv ALNG$ or $ALN \equiv ALNS$.

The Standpoint Logic Approach. Let us consider that the output of the merge process is represented under a standpoint (#001/2), that contains assertions relevant to all the statements that are not flagged as problematic. Then, we can merely state (10) and optionally (11), to strengthen the axiomatisation of the combination.

- (10). $(\#001) \preceq (\#001/2) \wedge (\#002) \preceq (\#001/2)$
 (11). $\Box_{(\#001/2)}[(ALN \leftrightarrow ALNS) \vee (ALNG \leftrightarrow ALN)]$

The same approach can be taken to address more problematic scenarios that provoke inconsistency instead of ‘malformation’, such as the one displayed in Figure 4. In this case, the traditional approaches are (1) to give up on one of the alignments and (2) to remove the disjointness restriction. With standpoint logic, we can model this again as a structure of three standpoints or, alternatively, we can add complex alignments directly between the initial ontologies, such as $\Box_{(\#002)}[CF] \leftrightarrow \Box_{(\#001)}[F \vee FG]$.

Let us now consider the second application scenario: that of a standpoint ontology by design. In the process of designing a general-purpose ontology, addressing the semantic heterogeneity of some terms is often challenging. For instance, the need for adding several forest characterisations in the Environment Ontology (*EnvO*) was reported [24]. In the absence of frameworks supporting ‘characterisations’, big ontologies seeking generality rely on (i) weakening the precision (e.g. *EnvO* mostly uses *part_of* and *is_a* roles, and avoids disjointness), and (ii) formalising different (or very similar) overlapping entities, that correspond to different standpoints on a concept and lead to convoluted taxonomies.

The scenario in *EnvO* is as follows: as of October 2019, forests in *EnvO* are represented via two main classes, namely *forested area* and *forest ecosystem*, both of which have the same textual description. *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. *Forest ecosystem* and *forest biome* are intended to characterise the forest as an ecosystem and as a biome respectively. Additional classes subsumed by *vegetated area* also refer to forests (e.g. *area of evergreen forest*), yet they are not related to any of the main forest concepts, possibly in order to avoid conflicts. Figure 5 is an overview of the most important ‘forest’ entities and some of the main superclasses.

Our proposal is that designing ontologies such as *EnvO* in a modular way by means of standpoints may help in overcoming the tradeoff between generality, convolution and




Figure 5. Fragment of the *EnvO* Ontology (2019).

precision. While radically redesigning EnvO forests based on thematic standpoints goes beyond the scope of this paper, we discuss one of many small details:

In Figure 5 we see that Forest_Biome is related to Ecosystem via parthood through Forest_Ecosystem and via subsumption through Biome. Yet, it does not seem intended for a biome to be a forest ecosystem and have a part that is also a forest ecosystem. Thus we assume that this could be better represented with two alternative standpoints, where one models that Forest_Biome has a part Forest_Ecosystem and Biome is not an Environment and the other models that Forest_Biome is a Forest_Ecosystem, and Biome is an Environment. With this, we'd avoid misuses and we capture the intuition that, depending on how we interpret the polysemous word Biome, we can think of a Forest_Biome being a Forest_Ecosystem or the latter just being part of it.

7. Related Work and Discussion

The importance of handling the interpretation of information in relation to its standpoint or context, and of understanding relationships between those, has been recognised by many researchers in AI [25]. This has led to the proposal of a variety of systems of representation, in rather overlapping areas of research and with diverse nomenclatures.

Contextual Ontologies. Semantic variation is often associated with differences of *context*. Our approach has some similarities to context logics in the style proposed by McCarthy [25], such as the modal framework in [26]. However, this tradition focuses on modelling contexts, and treats them as full-fledged formal objects over which one can express first-order properties. In contrast, Standpoint Logic is more suitable for the many cases in which a detailed formalisation of the contexts involved in a domain is either unnecessary or unfeasible, or where the interest resides in the perspectives or standpoints themselves rather than the context in which they occur.

A contextual framework where contexts are mere labels is [27]. However, in Benslimane's framework, context-labels can only be used in rather restricted scenarios and contexts are formally closer to the *ontology viewpoints* than to the standpoints of our logic, leaving no room for standpoint relations and compositions, which are crucial in many application scenarios as seen in the examples (cf. Section 6).

Ontology Viewpoints. The notion of ontology views is inherited from the well-known view mechanism in database theory. Most research in this area follows that tradition, and focuses on the presentation of partial (and consistent) views (partitions) on the content of a single ontology, which may be interesting to different agents [28,29].

However, some works consider potentially conflicting viewpoints, such as [30] and the similar and more developed [31]. Both share our motivation, yet they approach it differently and tailor their framework to description logics, in a style similar to Benslimane's work. Hemam and Boufaïda [31] define seven types of elements (viewpoints, global classes, local classes, global properties, local properties, bridge rules and individuals) in a nested structure, leading to a rather intricate logic, for which no complexity bounds are provided.

Instead of implementing the ad-hoc intuition of 'viewpoints', our work just extends the base language (in this paper: propositional logic) with modal operators and gives the resulting logic a Kripke semantics. This leads to a simpler, more recognisable and more

expressive framework, that allows for the unrestricted use of the modal operators in any kind of sentence and supports, for instance, hierarchies and combinations of standpoints, inferences of partial truths, the preservation of consistency with penumbral connections and inferences about the standpoints themselves. As for consistency, [31] develops a mechanism that works by bringing global facts down to the viewpoints, while the modal framework ensures consistency also in the inverse direction, allowing e.g. for disambiguation strategies. For instance, in our example, one can infer that $\diamond_*[\text{Lime} \rightarrow \text{Yellow}]$ and that $\diamond_*[\text{Yellow} \rightarrow \text{WC}]$ but not $\diamond_*[\text{Lime} \rightarrow \text{WC}]$ because *CT* and *HP* are not guaranteed to be compatible standpoints (cf. Fig. 2).

DDL Bridge Rules and ε -connections. In the area of ontology modularity, different formalisms such as DDL bridge rules [32] and ε -connections [33] have been proposed to specify the interaction between independent knowledge sources. These can be related to the present framework in that they provide mechanisms to establish links between conceptual models, similar to the role of assertions involving several standpoints such as $\square_{HP}[\text{Green}] \leftrightarrow (\square_{CT}[\text{Green}] \vee \square_{HP}[\text{Lime}])$, yet the motivation is inherently different: while the standpoint framework focuses on integrating possibly overlapping knowledge into a global source, DDL Bridge rules and ε -connections have been proposed to connect standalone modules. Moreover, both have been proposed for DL languages.

In contrast to standpoint statements, DDL bridge rules [32] are directional (INTO or ONTO) relations between concepts of different modules, such that an INTO rule between a concept *A* in module M_1 and a concept *B* in module M_2 , $i : A \xrightarrow{\text{INTO}} j : B$, does not entail the converse ONTO rule $j : B \xrightarrow{\text{ONTO}} i : A$.

ε -connections are a combination method that takes the union of the combined modules M_1 and M_2 , enriched with operators capable of talking about the link relations that the ε -connection establishes between them. This behaviour can be mimicked in a standpoint style by encapsulating the ε -connection into a standpoint s_ε that encodes those links and is subsumed by the standpoints s_{M_1} and s_{M_2} . ε -connections however require the vocabularies of the connected modules to be disjoint, which is reasonable in the context of ontology modularity but an important limitation for our main subject of interest: scenarios in which there are competing perspectives on the semantics of a shared vocabulary.

8. Conclusions and Future Work

The semantic heterogeneity of natural language together with the diversity of human world views are at the root of many knowledge interoperability scenarios. As an alternative to the mainstream unification strategy, this paper introduces a logic formalism based on the notion of standpoint that is suitable for knowledge representation and reasoning with sets of possibly conflicting perspectives or characterisations of a domain. We explore how different agents can establish their standpoints, which typically involves specifying constraints and relations (amounting to making ‘ontological commitments’) but not necessarily subscribing to a single sharp interpretation. Natural reasoning tasks over such multi-standpoint specifications include gathering unequivocal or undisputed knowledge, determining knowledge that is relative to a standpoint or a set of them, and contrasting the knowledge that can be inferred from different standpoints. The fact that the proposed formalism preserves the complexity of the propositional base language, having

an NP-complete algorithm for satisfiability in contrast to the PSPACE-completeness of multi-modal epistemic logics in general, indicates that the framework can have interesting, technically feasible applications in areas such as ontology alignment and concept negotiation, and knowledge aggregation.

Moreover, in contrast to other proposals of multi-perspective frameworks (such as *ontology viewpoints*), our framework is rooted in a well established philosophical theory of language, supervaluationism, and thus can be linked to a theoretic body of work. In addition, the use of modalities makes the language easily recognisable for a broad community of researchers and practitioners, and allows for the expression of assertions that are only guaranteed to hold in some sense or that are borderline, which becomes useful in scenarios involving collections of interpretations.

There are several directions of future research. First, the set-theoretic structure of standpoints not only facilitates the establishment of hierarchies of interpretations, but it also makes it possible to define an algebraic calculus allowing to define complex standpoints out of atomic ones by means of union ($s_1 \cup s_2$, representing the integration of knowledge coming from two different sources), intersection ($s_1 \cap s_2$, collecting the “agreed-on knowledge” shared between two standpoints) or difference ($s_1 \setminus s_2$, representing a sharpening of standpoint s_1 through the exclusion of all precisifications pertaining to s_2). It is easy to see that an extension of our formalism by such standpoint–algebraic expressions comes at no additional cost in the NP complexity of reasoning, since the translation to one-variable first order logic presented in Section 4 can be easily adapted.

A shortcoming of our current work is that, as a starting point for our investigations, we have chosen to define the language for a propositional base, while today’s knowledge representation formalisms use more expressive logics. Yet, because modal frameworks are well understood also for more advanced logics such as description logics and (fragments of) first-order logic, the corresponding adaptation of syntax and semantics should not be problematic and is object of current work. Of course, such extensions will also necessitate to determine the underlying complexities, to specify the corresponding proof-theoretic calculi and to develop strategies to employ off-the-shelf reasoners for standpoint-enhanced reasoning in practical scenarios. Given that the reason why the first order logic translation is well-behaved in terms of complexity is the small model property of the logic, we expect to maintain this good behaviour for more expressive base languages.

References

- [1] Wittgenstein L. *Philosophical Investigations*. tr. G.E.M. Anscombe ed. Blackwell; 1958.
- [2] Guarino N. *Formal Ontology and Information Systems*. In: *Procs. of the 1st Int. Conf. on Formal Ontology in Information Systems*. vol. 46. IOS Press; 1998. p. 3–15.
- [3] Kalfoglou Y, Schorlemmer M. *Ontology mapping: the state of the art*. *The knowledge engineering review*. 2003;18(1):1–31.
- [4] Euzenat J, Mocan A, Scharffe F. *Ontology Alignments*. In: Hepp M, Leenheer PD, Moor AD, Sure Y, editors. *Ontology Management. Computing for Human Experience*. vol. 7. Springer; 2008. p. 177–206.
- [5] Otero-Cerdeira L, Rodríguez-Martínez FJ, Gómez-Rodríguez A. *Ontology matching: A literature review*. *Expert Systems with Applications*. 2015;42(2).
- [6] Pesquita C, Faria D, Santos E, Couto FM. *To repair or not to repair: Reconciling correctness and coherence in ontology reference alignments*. In: *Procs. of the 8th Int. Conf. on Ontology Matching*. vol. 1111. CEUR; 2013. p. 13–24.

- [7] Solimando A, Jiménez-Ruiz E, Guerrini G. Minimizing conservativity violations in ontology alignments. *Knowledge and Information Systems*. 2017;51(3).
- [8] Jiménez-Ruiz E, Cuenca Grau B, Horrocks I, Berlanga R. Ontology integration using mappings: Towards getting the right logical consequences. In: *Procs. of the 6th European Semantic Web Conf.*. vol. 5554. Springer; 2009. p. 173–187.
- [9] Megdiche I, Teste O, Trojahn C. An extensible linear approach for holistic ontology matching. In: *Procs. of the 15th Int. Semantic Web Conf.*. vol. 9981. Springer; 2016. .
- [10] Thiéblin E, Haemmerlé O, Hernandez N, Trojahn C. Survey on complex ontology matching. *Semantic Web*. 2020;11(4).
- [11] Osman I, Ben Yahia S, Diallo G. Ontology Integration: Approaches and Challenging Issues. *Information Fusion*. 2021;71:38–63.
- [12] Bouquet P, Ghidini C, Giunchiglia F, Blanzieri E. Theories and uses of context in knowledge representation and reasoning. *Journal of Pragmatics*. 2003 3;35(3):455–484.
- [13] Keefe R, Smith P. Theories of vagueness. In: Keefe R, Smith P, editors. *Vagueness: A reader*. The MIT press; 1997. p. 1–57.
- [14] Mehlberg H. *The Reach of Science*. University of Toronto Press; 1958.
- [15] Fine K. Vagueness, truth and logic. *Synthese*. 1975;30(3-4):265–300.
- [16] Bennett B. Standpoint Semantics: a framework for formalising the variable meaning of vague terms. In: Cintula P, Fermüller C, Godo L, Hájek P, editors. *Understanding Vagueness: Logical, Philosophical and Linguistic Perspectives*. Studies in Logic. College Publications; 2011. p. 261–278.
- [17] Gómez Álvarez L, Bennett B. Dealing with conceptual indeterminacy: A framework based on supervaluation semantics. In: *Joint Procs. of MedRACER and WOMoCoE*. CEUR; 2018. .
- [18] Varzi AC. Supervaluationism and Its Logics. *Mind*. 2007;116(463):633–676.
- [19] Williams JRG. Supervaluationism and Logical Revisionism. *The Journal of philosophy*. 2008;105(4):192–212.
- [20] Cobreros P. Supervaluationism and classical logic. In: *Int. Workshop on Vagueness in Communication*. VIC 2009. Springer; 2011. p. 51–63.
- [21] Lemmon EJ, Scott D. *An introduction to modal logic*. Oxford: Blackwell; 1979.
- [22] Pratt-Hartmann I. On the Computational Complexity of the Numerically Definite Syllogistic and Related Logics. *Bulletin of Symbolic Logic*. 2008;14(1):1–28.
- [23] Pour N, Algergawy A, Amini R, Faria D, Fundulaki I, Harrow I, et al. Results of the ontology alignment evaluation initiative 2020. In: *Procs. of the 15th Int. Workshop on Ontology Matching*. vol. 2788. CEUR; 2020. p. 92–138.
- [24] Buttigieg PL, Pafilis E, Lewis SE, Schildhauer MP, Walls RL. The environment ontology in 2016: bridging domains with increased scope, semantic density, and interoperability. *Journal of Biomedical Semantics*. 2016 12;7(1):57.
- [25] McCarthy J, Buvac S. Formalizing context (expanded notes). 1997.
- [26] Klarman S, Gutiérrez-Basulto V. Description logics of context. *Journal of Logic and Computation*. 2013;26(3):817–854.
- [27] Benslimane D, Arara A, Falquet G, Maamar Z, Thiran P, Gargouri F. Contextual ontologies motivations, challenges, and solutions. In: *Procs. of the 4th Int. Conf. on Advances in Information Systems*. Springer; 2006. p. 168–176.
- [28] Rajugan R, Chang E, Dillon TS. Ontology views: A theoretical perspective. In: *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops*. vol. 4278 LNCS - II. Springer; 2006. p. 1814–1824.
- [29] Bolchini C, Curino C, Quintarelli E, Schreiber F, Tanca L. Context information for knowledge reshaping. *Int Journal of Web Engineering and Technology*. 2009;5(1):88–103.
- [30] Ribière M, Dieng-Kuntz R. A viewpoint model for cooperative building of an Ontology. In: *Procs. of the 10th Int. Conf. on Conceptual Structures*. Springer; 2002. p. 220–234.
- [31] Hemam M, Boufaïda Z. MVP-OWL: A multi-viewpoints ontology language for the Semantic Web. *Int Journal of Reasoning-based Intelligent Systems*. 2011;3(3-4):147–155.
- [32] Borgida A, Serafini L. Distributed Description Logics: Assimilating information from peer sources. *Journal on Data Semantics*. 2003;2800:153–184.
- [33] Kutz O, Lutz C, Wolter F, Zakharyashev M. E-connections of Description Logics. *Description Logics Workshop*. 2003;81.