

How to Agree to Disagree

Managing Ontological Perspectives using Standpoint Logic

Lucía Gómez Álvarez^[0000–0002–2525–8839], Sebastian Rudolph^[0000–0002–1609–2080],
and Hannes Strass^[0000–0001–6180–6452]

Computational Logic Group, Faculty of Computer Science, TU Dresden, Germany
{lucia.gomez_alvarez, sebastian.rudolph, hannes.strass}@tu-dresden.de

Abstract The importance of taking individual, potentially conflicting perspectives into account when dealing with knowledge has been widely recognised. Many existing ontology management approaches fully merge knowledge perspectives, which may require weakening in order to maintain consistency; others represent the distinct views in an entirely detached way.

As an alternative, we propose *Standpoint Logic*, a simple, yet versatile multi-modal logic “add-on” for existing KR languages intended for the integrated representation of domain knowledge relative to diverse, possibly conflicting *standpoints*, which can be hierarchically organised, combined, and put in relation with each other. Starting from the generic framework of *First-Order Standpoint Logic* (FOSL), we subsequently focus our attention on the fragment of *sentential* formulas, for which we provide a polytime translation into the standpoint-free version. This result yields decidability and favourable complexities for a variety of highly expressive decidable fragments of first-order logic. Using some elaborate encoding tricks, we then establish a similar translation for the very expressive description logic *SROIQb_s* underlying the OWL 2 DL ontology language. By virtue of this result, existing highly optimised OWL reasoners can be used to provide practical reasoning support for ontology languages extended by standpoint modelling.

Keywords: knowledge integration · ontology alignment · conflict management

1 Introduction

Artefacts of contemporary knowledge representation (ontologies, knowledge bases, or knowledge graphs) serve as means to conceptualise specific domains, with varying degrees of expressivity ranging from simple classifications and database schemas to fully axiomatised theories. Inevitably, such specifications reflect the individual points of view of their creators (be it on a personal or an institutional level) along with other contextual aspects, and they may also differ in modelling design decisions, such as the choice of conceptual granularity or specific ways of axiomatising information. This semantic heterogeneity is bound to pose significant challenges whenever the interoperability of independently developed knowledge specifications is required.

This paper proposes a way to address the interoperability challenge while at the same time preserving the varying perspectives of the original sources. This is particularly important in scenarios that require the simultaneous consideration of multiple, potentially contradictory, viewpoints.

Example 1. A broad range of conceptualisations and definitions for the notion of *forest* have been specified for different purposes, giving rise to diverging or even contradictory statements regarding forest distributions. Consider a knowledge integration scenario involving two sources adopting a *land cover* (LC) and a *land use* (LU) perspective on forestry. LC characterises a *forest* as a “forest ecosystem” with a minimum area (F1) where a *forest ecosystem* is specified as an ecosystem with a certain ratio of tree canopy cover (F2). LU defines a forest with regard to the purpose for which an area of land is put to use by humans, i.e. a forest is a maximally connected area with “forest use” (F3).¹

Both sources LC and LU agree that forests subsume broadleaf, needleleaf and tropical forests (F4), and they both adhere to the Basic Formal Ontology (BFO, [1]), an upper-level ontology that formalises general terms, stipulating for instance that *land* and *ecosystem* are disjoint categories (F5). Using standard description logic notation and providing “perspective annotations” by means of correspondingly labelled box operators borrowed from multi-modal logic, the above setting might be formalised as follows:

- (F1) $\Box_{LC}[\text{Forest} \equiv \text{ForestEcosystem} \sqcap \exists \text{hasLand.Area}_{\geq 0.5\text{ha}}]$
- (F2) $\Box_{LC}[\text{ForestEcosystem} \equiv \text{Ecosystem} \sqcap \text{TreeCanopy}_{\geq 20\%}]$
- (F3) $\Box_{LU}[\text{Forest} \equiv \text{ForestlandUse} \sqcap \text{MCON}] \wedge \Box_*[\text{ForestlandUse} \sqsubseteq \text{Land}]$
- (F4) $\Box_{LC \cup LU}[(\text{BroadleafForest} \sqcup \text{NeedleleafForest} \sqcup \text{TropicalForest}) \sqsubseteq \text{Forest}]$
- (F5) $(LC \preceq \text{BFO}) \wedge (LU \preceq \text{BFO}) \wedge \Box_{\text{BFO}}[\text{Land} \sqcap \text{Ecosystem} \sqsubseteq \perp]$

In the case of Example 1, *ecosystem* and *land* are disjoint categories according to the overarching BFO (F5), yet forests are defined as ecosystems according to LC (F1) and as lands according to LU (F3). These kinds of disagreements result in well-reported challenges in the area of Ontology Integration [6,21] and make ontology merging a non-trivial task, often involving a certain knowledge loss or weakening in order to avoid incoherence and inconsistency [22,28]. In Example 1, to merge LU and LC, there are two typical options to resolve the issue: (Opt-Weak) one may give up on the disjointness axiom (F5), or (Opt-Dup) one could duplicate all the conflicting predicates [20], in this case not only *Forest* (into *Forest_LC* and *Forest_LU*), but also the forest subclasses in (F4): *BroadleafForest*, *NeedleleafForest* and *TropicalForest*. In contrast, 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 conflict-free conceptual model.

Standpoint logic [8] is a formalism inspired by the theory of supervaluationism [7] and rooted in modal logic that supports the coexistence of multiple standpoints and the establishment of alignments between them, by extending the base language with labelled modal operators. Propositions $\Box_{LC} \phi$ and $\Diamond_{LC} \phi$ express information relative to the *standpoint* LC and read, respectively: “according to LC, it is *unequivocal/conceivable* that ϕ ”. In the semantics, standpoints are represented by sets of *precisifications*,² such that $\Box_{LC} \phi$ and $\Diamond_{LC} \phi$ hold if ϕ is true in all/some of the precisifications in LC.

The logical statements (F1)–(F5), which formalise Example 1 by means of a standpoint-enhanced description logic, are not inconsistent, so all axioms can be jointly

¹ “Forest use” areas may qualify for logging and mining concessions as well as be further classified into, e.g. agricultural or recreational land use.

² Precisifications are analogous to the *worlds* of frameworks with possible-worlds semantics.

represented. Let us now illustrate the use of standpoint logic for reasoning with the individual perspectives. First, assume the following (globally agreed) facts about three instances, an ecosystem e , a parcel of land l , and a city c :

- (F6) $\text{ForestEcosystem}(e) \quad \text{hasLand}(e, l) \quad \text{ForestlandUse}(l)$
(F7) $\text{Area}_{\geq 0.5\text{ha}}(l) \quad \text{MCON}(l) \quad \text{in}(l, c) \quad \text{City}(c)$

It is clear from (F1) that according to LC, e is a forest, written as $\Box_{\text{LC}}[\text{Forest}(e)]$, since it is a forest ecosystem (F6) with an area larger than 0.5ha (F7). On the other hand, it is clear from (F3) that according to LU, l is a forest, $\Box_{\text{LU}}[\text{Forest}(l)]$, since it has a forest land use (F6) and it is a maximally connected area (F7). More interestingly, we can also obtain *joint inferences*: assuming the (generally accepted) background knowledge expressed by $\text{hasLand} \circ \text{in} \sqsubseteq \text{in}$, we can infer

$$\Box_{\text{LC} \cup \text{LU}}[(\text{City} \sqcap \exists \text{in}^- . \text{Forest})(c)],$$

which means that “according to both LC and LU there is some forest in City c .” This holds for LU since l is a forest and is in c (F7); and it holds for LC because e is a forest in the land l , which is in turn in c .

In contrast to the options of the ontology merging approach, using standpoint logic prevents the multiplication (and corresponding “semantic detachment”) of predicates from (Opt-Dup). It also avoids unintended consequences arising when knowledge sources are weakened just enough to maintain satisfiability: In the corresponding (Opt-Weak) scenario, after merging the knowledge sources of LU and LC and removing (F5), we can consistently infer $\text{Forest}(e)$ from the standpoint-free versions of (F1), (F6) and (F7) and $\text{Forest}(l)$ from (F3), (F6) and (F7), similar to the standpoint framework. But on top of that, reapplying (F1) and (F3) also yields “ e is a forest, and its land l is also a forest and an ecosystem, and has some other associated land, bigger than 0.5ha” through the following derivable assertions:

$$\text{Forest}(e) \quad \text{hasLand}(e, l) \quad \text{Forest}(l) \quad \text{ForestEcosystem}(l) \quad \exists \text{hasLand.Area}_{\geq 0.5\text{ha}}(l)$$

This illustrates how, beyond the problem of inconsistency, naively merging different models of a domain may lead to erroneous reasoning. In fact, other non-clashing differences between the forest definitions (F1) and (F3) respond to relevant nuances that relate to each standpoint and should also not be naively merged. For instance, from the land cover perspective (F1), there is no spatial connectedness requirement, since there are “mosaic forest ecosystems” where the landscape displays forest patches that are sufficiently close to constitute a single ecosystem. On the other hand, for LU, there is no minimum tree canopy (F3), since a temporarily cleared area still has a “forest use”.

Standpoint logic preserves the independence of the perspectives and escapes global inconsistency – without weakening the sources or duplicating entities – because its model theory (cf. Section 2.1) requires consistency only within standpoints and precisifications. Notwithstanding, it allows for the specification of structures of standpoints and alignments between them. 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.

Let us get an idea of the expressivity of the proposed logic. In spite of its simple syntax, the language is remarkably versatile; it allows for specifying knowledge relative (a) to a standpoint, e.g. (F1), (b) to the global standpoint, denoted by $*$, e.g. (F3), and (c) to set-theoretic combinations of standpoints, e.g. (F4). Additional language features can be defined in terms of the former: $\mathcal{I}_{LC}\phi$, which means that, “according to LC, it is inherently *indeterminate* whether ϕ ” can be defined by $\mathcal{I}_{LC}\phi := \diamond_{LC}\phi \wedge \diamond_{LC}\neg\phi$. The *sharper* operator \preceq is used to establish hierarchies of standpoints and constraints on the structure of precisifications, e.g. (F5), and can be defined via $s_1 \preceq s_2 := \square_{s_1 \setminus s_2}[\top \sqsubseteq \perp]$. Intuitively, $s_1 \preceq s_2$ expresses that standpoint s_1 inherits the propositions of s_2 , by virtue of “ $s_1 \subseteq s_2$ ” holding for the corresponding sets of precisifications. This type of statement comes handy to “import” background knowledge from some ontology, such as the foundational ontology BFO in our example. In combination, these modelling features allow for expressing further constraints useful for knowledge integration scenarios, e.g.,

$$(F8) \quad * \preceq (LC \cup LU) \wedge \diamond_{LC}[\top \sqsubseteq \top] \wedge \diamond_{LU}[\top \sqsubseteq \top],$$

where the first conjunct allows us to specify that no interpretations beyond the standpoints of interest are under consideration, by stating that the universal standpoint is a subset of the union of LC and LU. The other two conjuncts enforce the non-emptiness of the standpoints of interest, LC and LU, ensuring that each standpoint by itself is coherent. To illustrate a use case, consider the statement $\diamond_*[\text{Forest}(f) \wedge \neg\text{MCON}(f)]$, expressing that it is conceivable that f is a non-spatially-connected forest. From this, we can infer together with (F8) and the unfulfilled requirement of connectedness of LU (F3), that f must be conceivable for LC instead, and thus f must be a forest ecosystem (F1):

$$\diamond_{LC}[\text{ForestEcosystem}(f) \wedge (\exists \text{hasLand.Area}_{\geq 0.5\text{ha}})(f)]$$

Gómez Álvarez and Rudolph [8] have introduced the standpoint framework over a propositional base logic. While they showed favourable complexity results (standard reasoning tasks are NP-complete just like for plain propositional logic), the framework is not expressive enough for knowledge integration scenarios employing contemporary ontology languages. In this paper, we widen the scope by (1) introducing the very general framework of *first-order* standpoint logic (FOSL) and (2) allowing for more modelling flexibility on the side of standpoint descriptions by introducing support for set-theoretical combinations of standpoints (Section 2). We provide the syntax and semantics of this generic framework, before focusing on the identification of FOSL fragments with beneficial computational properties. To this end, we define the *sentential* fragment, which imposes restrictions on the use of standpoint operators and guarantees a small model property (Section 2.2). Tailored to this case, we introduce a polynomial satisfiability-preserving translation (Section 2.3) that does not affect membership in diverse decidable fragments of FO. This allows us to immediately obtain decidability and tight complexity bounds for the standpoint versions of diverse FO fragments (e.g. the 2-variable counting fragment, the guarded negation fragment and the triguarded fragment) (Section 3). In addition, it provides a way to leverage off-the-shelf reasoners for practical reasoning in standpoint versions of popular ontology languages. We demonstrate this by extending our results to a standpoint logic based on the description logic $SR\mathcal{OIQb}_s$, a semantic fragment of FO closely related to the OWL 2 DL ontology language (Section 4). Finally, we revisit our example to discuss and illustrate properties of our proposal (Section 5). An extended version of this paper including proofs is available as a technical report [30].

2 First-Order Standpoint Logic

In this section we introduce the general framework of first order standpoint logic (FOSL) as well as its sentential fragment and establish model theoretic and computational properties. In addition to establishing various worthwhile decidability and complexity results, this approach also provides us with a clearer view on the underlying principles of our arguments, while avoiding distractions brought about by some peculiarities of expressive ontology languages, which we will address separately later on.

2.1 FOSL Syntax and Semantics

Definition 1. *The syntax of first-order standpoint logic (\mathbb{S}_{FO}) is based on a signature $\langle \mathcal{P}, \mathcal{C}, \mathcal{S} \rangle$, consisting of predicate symbols \mathcal{P} (each associated with an arity $n \in \mathbb{N}$), constant symbols \mathcal{C} and standpoint symbols \mathcal{S} , usually denoted with s, s' , as well as a set \mathcal{V} of variables, typically denoted with x, y, \dots (possibly annotated). These four sets are assumed to be countably infinite and pairwise disjoint. The set \mathcal{T} of terms contains all constants and variables, that is, $\mathcal{T} = \mathcal{C} \cup \mathcal{V}$.*

The set $\mathcal{E}_{\mathcal{S}}$ of standpoint expressions is defined as follows:

$$e_1, e_2 ::= * \mid s \mid e_1 \cup e_2 \mid e_1 \cap e_2 \mid e_1 \setminus e_2$$

The set \mathbb{S}_{FO} of FOSL formulas is then given by

$$\phi, \psi ::= P(t_1, \dots, t_k) \mid \neg\phi \mid \phi \wedge \psi \mid \forall x\phi \mid \Box_e \phi,$$

where $P \in \mathcal{P}$ is an k -ary predicate symbol, $t_1, \dots, t_k \in \mathcal{T}$ are terms, $x \in \mathcal{V}$, and $e \in \mathcal{E}_{\mathcal{S}}$.

For a formula ϕ , we denote the set of all of its subformulas by $\text{Sub}(\phi)$. The size of a formula is $|\phi| := |\text{Sub}(\phi)|$. The connectives and operators \mathbf{t} , \mathbf{f} , $\phi \vee \psi$, $\phi \rightarrow \psi$, $\exists x\phi$, and $\Diamond_e \phi$ are introduced as syntactic macros as usual. As further useful syntactic sugar, we introduce *sharpening* statements $e_1 \preceq e_2$ to denote $\Box_{e_1 \setminus e_2} \mathbf{f}$, the *indeterminacy operator* via $\mathcal{I}_e \phi := \Diamond_e \phi \wedge \Diamond_e \neg\phi$, and the *determinacy operator* via $\mathcal{D}_e \phi := \neg \mathcal{I}_e \phi$.

Definition 2. *Given a signature $\langle \mathcal{P}, \mathcal{C}, \mathcal{S} \rangle$, a first-order standpoint structure \mathfrak{M} is a tuple $\langle \Delta, \Pi, \sigma, \gamma \rangle$ where:*

- Δ is a non-empty set, the domain of \mathfrak{M} ;
- Π is the non-empty set of precisifications;
- σ is a function mapping each standpoint symbol from \mathcal{S} to a set of precisifications (i.e., a subset of Π);
- γ is a function mapping each precisification from Π to an ordinary first-order structure \mathcal{I} over the domain Δ , whose interpretation function $\cdot^{\mathcal{I}}$ maps:
 - each predicate symbol $P \in \mathcal{P}$ of arity k to an k -ary relation $P^{\mathcal{I}} \subseteq \Delta^k$,
 - each constant symbol $a \in \mathcal{C}$ to a domain element $a^{\mathcal{I}} \in \Delta$.

Moreover, for any two $\pi_1, \pi_2 \in \Pi$ and every $a \in \mathcal{C}$ we require $a^{\gamma(\pi_1)} = a^{\gamma(\pi_2)}$.

Note that all first-order structures in all precisifications implicitly share the same interpretation domain Δ given by the overarching first-order standpoint structure \mathfrak{M} , that is, we adopt the *constant domain assumption*.³ Moreover, the last condition of Definition 2

³ This is not a substantial restriction, as other variants – expanding domains, varying domains – can be emulated using constant domains [29, Theorem 6].

also enforces *rigid constants*, that is, constants denote the same objects in different standpoints (while clearly their properties could differ).

Definition 3. Let $\mathfrak{M} = \langle \Delta, \Pi, \sigma, \gamma \rangle$ be a first-order standpoint structure for the signature $\langle \mathcal{P}, \mathcal{C}, \mathcal{S} \rangle$ and \mathcal{V} be a set of variables. A variable assignment is a function $v : \mathcal{V} \rightarrow \Delta$ mapping variables to domain elements. Given a variable assignment v , we denote by $v_{\{x \mapsto \delta\}}$ the function mapping x to $\delta \in \Delta$ and any other variable y to $v(y)$.

An interpretation function $\cdot^{\mathcal{I}}$ and a variable assignment specify how to interpret terms by domain elements: We let $t^{\mathcal{I},v} = v(x)$ if $t = x \in \mathcal{V}$, and $t^{\mathcal{I},v} = a^{\mathcal{I}}$ if $t = a \in \mathcal{C}$.

To interpret standpoint expressions, we lift σ from \mathcal{S} to all of $\mathcal{E}_{\mathcal{S}}$ by letting $\sigma(*) = \Pi$ and $\sigma(e_1 \boxtimes e_2) = \sigma(e_1) \boxtimes \sigma(e_2)$ for $\boxtimes \in \{\cup, \cap, \setminus\}$.

The satisfaction relation for formulas is defined in the usual way via structural induction. In what follows, let $\pi \in \Pi$ and let $v : \mathcal{V} \rightarrow \Delta$ be a variable assignment; we now establish the definition of the satisfaction relation \models for first-order standpoint logic using pointed first-order standpoint structures:

$$\begin{array}{ll}
\mathfrak{M}, \pi, v \models \mathsf{P}(t_1, \dots, t_k) & \text{iff } (t_1^{\gamma(\pi),v}, \dots, t_k^{\gamma(\pi),v}) \in \mathsf{P}^{\gamma(\pi)} \\
\mathfrak{M}, \pi, v \models \neg \phi & \text{iff } \mathfrak{M}, \pi, v \not\models \phi \\
\mathfrak{M}, \pi, v \models \phi \wedge \psi & \text{iff } \mathfrak{M}, \pi, v \models \phi \text{ and } \mathfrak{M}, \pi, v \models \psi \\
\mathfrak{M}, \pi, v \models \forall x \phi & \text{iff } \mathfrak{M}, \pi, v_{\{x \mapsto \delta\}} \models \phi \text{ for all } \delta \in \Delta \\
\mathfrak{M}, \pi, v \models \Box_e \phi & \text{iff } \mathfrak{M}, \pi', v \models \phi \text{ for all } \pi' \in \sigma(e) \\
\mathfrak{M}, \pi \models \phi & \text{iff } \mathfrak{M}, \pi, v \models \phi \text{ for all } v : \mathcal{V} \rightarrow \Delta \\
\mathfrak{M} \models \phi & \text{iff } \mathfrak{M}, \pi \models \phi \text{ for all } \pi \in \Pi
\end{array}$$

As usual, \mathfrak{M} is a *model* for a formula ϕ iff $\mathfrak{M} \models \phi$. As an aside, note that the modal-logic nature of FOSL may become more evident upon realizing that an alternative definition of its semantics via Kripke structures can be given (with \Box_e interpreted in the standard way) by assigning every $e \in \mathcal{E}_{\mathcal{S}}$ the accessibility relation $\{(\pi, \pi') \mid \pi, \pi' \in \Pi, \pi' \in \sigma(e)\}$.

Later in this paper, we will consider cases where the number of precisifications is fixed. Thus, we conclude this section by a corresponding definition.

Definition 4. For a natural number $n \in \mathbb{N}$, a FOSL formula ϕ is *n-satisfiable* iff it has a model $\langle \Delta, \Pi, \sigma, \gamma \rangle$ with $|\Pi| = n$.

2.2 Small Model Property of Sentential Formulas

One interesting aspect of standpoint logic is that its simplified Kripke semantics brings about convenient model-theoretic properties that do not hold for arbitrary (multi-)modal logics. For propositional standpoint logic, it is known that standard reasoning tasks (such as checking satisfiability) are NP-complete [8], in contrast to PSPACE-completeness in related systems such as $\mathsf{K45}_n$, $\mathsf{KD45}_n$ and $\mathsf{S5}_n$. This result is in fact linked to a *small model property*, according to which every satisfiable formula has a model with a “small” number of precisifications. This beneficial property only holds in the single-modal $\mathsf{K45}$, $\mathsf{KD45}$ and $\mathsf{S5}$ [23] but applies to the multi-modal propositional standpoint logic because of its stronger modal interaction. Fortunately, it can also be shown to carry over to some fragments of FOSL and to the use of standpoint expressions. In particular, in this section,

we will show that if we restrict the language to those formulas with no free variables in subformulas of the form $\Box_e \phi$, then we can indeed guarantee that every satisfiable FOSL formula has a model whose number of precisifications is linear in the size of the formula.

Definition 5. Let ϕ be a formula of FOSL. We say that ϕ is *sentential* iff for all subformulas of ϕ of the form $\Box_e \psi$, all variables occurring in ψ are bound by a quantifier.

Theorem 1. A sentential FOSL formula ϕ is satisfiable iff it has a model with at most $|\phi|$ precisifications. That is, for sentential FOSL, satisfiability and $|\phi|$ -satisfiability coincide.

In the following, it will be convenient to assume that formulas are in *standpoint standard normal form* (SSNF), where no modal operator \Box_e occurs inside the scope of another. Any FOSL formula ϕ can be transformed into SSNF in polynomial time.

2.3 Translation to Plain First-Order Logic

In this section, we present a translation Trans_n mapping any FOSL formula ϕ to a plain FO formula $\text{Trans}_n(\phi)$ such that n -satisfiability of ϕ coincides with satisfiability of $\text{Trans}_n(\phi)$. The translation will make explicit use of a fixed, finite set I_n of precisifications with $|I_n| = n$.

Our translation will map any ϕ into a formula of (standpoint-free) first-order logic. The basic idea is to “emulate” standpoint structures $\langle \Delta, I_n, \sigma, \gamma \rangle$ in plain first-order structures over Δ by means of a “superposition” of all $\gamma(\pi)$, which requires to introduce n “copies” of the original set of predicates. To this end, we define our first-order vocabulary by $\mathbb{V}_{\text{FO}}(\mathcal{P}, \mathcal{C}, \mathcal{S}, I_n) = \langle \mathcal{P}', \mathcal{C} \rangle$ where \mathcal{P}' contains

- for each predicate $P \in \mathcal{P}$ and precisification $\pi \in I_n$, a predicate of the form P_π of the same arity as P , intuitively expressing that P_π should capture $P^{\gamma(\pi)}$;
- for each standpoint constant $s \in \mathcal{S}$ and every precisification $\pi \in I_n$, a nullary predicate of the form s_π , intuitively expressing that $\pi \in \sigma(s)$.

The top-level translation is then defined to set:

$$\text{Trans}_n(\phi) = \bigwedge_{\pi \in I_n} \text{trans}_n(\pi, \phi) \wedge \bigwedge_{\pi \in I_n} *_\pi,$$

where trans_n is inductively defined by

$$\begin{aligned} \text{trans}_n(\pi, P(t_1, \dots, t_k)) &= P_\pi(t_1, \dots, t_k) \\ \text{trans}_n(\pi, \neg\psi) &= \neg\text{trans}_n(\pi, \psi) \\ \text{trans}_n(\pi, \psi_1 \wedge \psi_2) &= \text{trans}_n(\pi, \psi_1) \wedge \text{trans}_n(\pi, \psi_2) \\ \text{trans}_n(\pi, \forall x\psi) &= \forall x(\text{trans}_n(\pi, \psi)) \\ \text{trans}_n(\pi', \Box_e \psi) &= \bigwedge_{\pi \in I_n} (\text{trans}_\mathcal{E}(\pi, e) \rightarrow \text{trans}_n(\pi, \psi)) \end{aligned}$$

Therein, $\text{trans}_\mathcal{E}$ implements the semantics of standpoint expressions, providing for each expression $e \in \mathcal{E}_\mathcal{S}$ a propositional formula $\text{trans}_\mathcal{E}(\pi, e)$ over $\{s_\pi \mid \pi \in I_n\}$ as follows:

$$\begin{aligned} \text{trans}_\mathcal{E}(\pi, s) &= s_\pi \\ \text{trans}_\mathcal{E}(\pi, e_1 \cup e_2) &= \text{trans}_\mathcal{E}(\pi, e_1) \vee \text{trans}_\mathcal{E}(\pi, e_2) \\ \text{trans}_\mathcal{E}(\pi, e_1 \cap e_2) &= \text{trans}_\mathcal{E}(\pi, e_1) \wedge \text{trans}_\mathcal{E}(\pi, e_2) \\ \text{trans}_\mathcal{E}(\pi, e_1 \setminus e_2) &= \text{trans}_\mathcal{E}(\pi, e_1) \wedge \neg\text{trans}_\mathcal{E}(\pi, e_2) \end{aligned}$$

A routine inspection of the translation ensures that it can be done in polynomial time and its output is of polynomial size, provided it is applied to formulas in SSNF.

Theorem 2. *A formula ϕ is n -satisfiable in FOSL if and only if the formula $\text{Trans}_n(\phi)$ is satisfiable in first-order logic.*

In fact, Theorem 2 provides us with a recipe for a satisfiability-preserving translation for any formula that comes with a “small model guarantee”, whenever a bound on the number of precisifications can be computed upfront. In particular, leveraging Theorem 1, we obtain the following corollary.

Corollary 1. *A formula ϕ is satisfiable in sentential first-order standpoint logic if and only if the formula $\text{Trans}_{|\phi|}(\phi)$ is satisfiable in first-order logic.*

3 Expressive Decidable FOSL Fragments

We note that even for the sentential version, first-order standpoint logic is still a generalization of plain first-order logic, whence reasoning in it is undecidable. Therefore, we will next look into some popular decidable FO fragments and establish decidability and complexity results for reasoning in their sentential standpoint versions.

Definition 6. *Let \mathcal{F} denote some FO fragment. Then the logic sentential Standpoint- \mathcal{F} , denoted $\mathbb{S}_{[\mathcal{F}]}$, contains the sentential FOSL formulas ϕ where:*

- all variables inside ϕ are bound by some quantifier,
 - for every subformula $\psi \in \text{Sub}(\text{SSNF}(\phi))$ preceded by a quantifier, $\psi \in \mathcal{F}$ holds.
- Fragment \mathcal{F} is standpoint-friendly iff every $\phi \in \mathbb{S}_{[\mathcal{F}]}$ satisfies $\text{Trans}_{|\phi|}(\text{SSNF}(\phi)) \in \mathcal{F}$.*

Lemma 1. *Let \mathcal{F} be a standpoint-friendly fragment of FOL. Then the following hold:*

1. *Satisfiability for $\mathbb{S}_{[\mathcal{F}]}$ is decidable if and only if it is for \mathcal{F} .*
2. *If the satisfiability problem in \mathcal{F} is at least NP-hard, then the satisfiability problem in $\mathbb{S}_{[\mathcal{F}]}$ is of the same complexity as in \mathcal{F} .*

It turns out that many popular formalisms are standpoint-friendly. For propositional logic (PL), this is straightforward: quantifiers and variables are absent altogether, which is also the reason why sentential Standpoint-PL and proper Standpoint-PL coincide. Thus, Lemma 1 yields an alternative argument for the NP-completeness of the latter, which was established previously [8].

On the expressive end of the logical spectrum, it is worthwhile to inspect fragments of FO that are still decidable (as a minimal requirement for the feasibility of automated reasoning). In fact, standpoint-friendliness can be established by structural induction for many of those. Notable examples are:

- the *counting 2-variable fragment* C^2 [10,24], which subsumes many *description logics* and serves as a mathematical backbone for related complexity results,
- the *guarded negation fragment* GNFO [3,2], which encompasses both the popular *guarded fragment* as well as the ubiquitous class of (*unions of*) *conjunctive queries* also known as *existential positive FO*, and

- the triguarded fragment TGF [27,14] (a more recent formalism subsuming both the two-variable and the guarded fragment without equality).

Intuitively, standpoint-friendliness for all these (and presumably many more) fragments follows from the fact that they are closed under Boolean combinations of sentences and that the transformation does not affect the structure of quantified formulas. We therefore immediately obtain that these four popular decidable fragments of FOL allow for accommodating standpoints without any increase in complexity.

Corollary 2. *The sentential FOSL fragments $\mathbb{S}_{[PL]}$ ($= \mathbb{S}_{PL}$), $\mathbb{S}_{[C^2]}$, $\mathbb{S}_{[GNFO]}$, and $\mathbb{S}_{[TGF]}$ are all decidable and the complexity of their satisfiability problem is complete for NP, NEXPTIME, 2EXPTIME, and N2EXPTIME, respectively.*

As an aside, we note that all these results remain valid when considering *finite satisfiability* (i.e., restricting to models with finite Δ), because for all considered fragments, companion results for the finite-model case exist and the equisatisfiability argument for our translation preserves (finiteness of) Δ .

4 Sentential Standpoint- $SRQIb_s$

We next present the highly expressive yet decidable logic (*Sentential*) *Standpoint- $SRQIb_s$* , which adds the feature of standpoint-aware modelling to *SRQIb_s*, a description logic (DL) obtained from the well-known DL *SRQIQ* [13] by a gentle extension of its expressivity, allowing safe Boolean role expressions over simple roles [26].⁴ The *SRQIQ* family serves as the logical foundation of popular ontology languages like OWL 2 DL. In view of the fact that *SRQIb_s* is a semantic fragment of FO, we can leverage the previously established results and present a satisfiability-preserving polynomial translation from Standpoint-*SRQIb_s* into plain *SRQIb_s* knowledge bases. On the theoretical side, this will directly provide us with favourable and tight complexity results for reasoning in Standpoint-*SRQIb_s*. On the practical side, this paves the way towards practical reasoning in “Standpoint-OWL”, since it allows us to use highly optimised OWL 2 DL reasoners off the shelf.

4.1 *SRQIb_s*: Syntax and Semantics

Let \mathcal{C} , \mathcal{P}_1 , and \mathcal{P}_2 be finite, mutually disjoint sets called *individual names*, *concept names* and *role names*, respectively. \mathcal{P}_2 is subdivided into *simple role names* \mathcal{P}_2^s and *non-simple role names* \mathcal{P}_2^{ns} , the latter containing the *universal role* u and being strictly ordered by some strict order \prec . In the original definition of *SRQIb_s*, simplicity of roles and \prec are not given a priori, but meant to be implicitly determined by the set of axioms. Our choice to fix them explicitly upfront simplifies the presentation without restricting expressivity. Then, the set \mathcal{R}^s of *simple role expressions* is defined by

$$r_1, r_2 ::= s \mid s^- \mid r_1 \cup r_2 \mid r_1 \cap r_2 \mid r_1 \setminus r_2,$$

with $s \in \mathcal{P}_2^s$, while the set of (arbitrary) *role expressions* is $\mathcal{R} = \mathcal{R}^s \cup \mathcal{P}_2^{ns}$. The order \prec

⁴ Focusing on the mildly stronger *SRQIb_s* instead of the more mainstream *SRQIQ* allows for a more coherent and economic presentation, without giving up the good computational properties and the availability of optimised algorithms and tools.

Name	Syntax	Semantics	Name	Syntax	Semantics
inverse role	s^-	$\{(\delta, \delta') \in \Delta \times \Delta \mid (\delta', \delta) \in s^{\mathcal{I}}\}$	concept assertion	$C(a)$	$a^{\mathcal{I}} \in C^{\mathcal{I}}$
role union	$r_1 \cup r_2$	$r_1^{\mathcal{I}} \cup r_2^{\mathcal{I}}$	role assertion	$r(a, b)$	$(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$
role intersection	$r_1 \cap r_2$	$r_1^{\mathcal{I}} \cap r_2^{\mathcal{I}}$	equality	$a \doteq b$	$a^{\mathcal{I}} = b^{\mathcal{I}}$
role difference	$r_1 \setminus r_2$	$r_1^{\mathcal{I}} \setminus r_2^{\mathcal{I}}$	inequality	$a \not\doteq b$	$a^{\mathcal{I}} \neq b^{\mathcal{I}}$
universal role	u	$\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$	general concept inclusion (GCI)	$C \sqsubseteq D$	$C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$
nominal	$\{a\}$	$\{a^{\mathcal{I}}\}$	role inclusion axioms (RIAs)	$r_1 \circ \dots \circ r_n \sqsubseteq r$	$r_1^{\mathcal{I}} \circ \dots \circ r_n^{\mathcal{I}} \subseteq r^{\mathcal{I}}$
top	\top	$\Delta^{\mathcal{I}}$		$r_1 \circ \dots \circ r_n \circ r \sqsubseteq r$	$r_1^{\mathcal{I}} \circ \dots \circ r_n^{\mathcal{I}} \circ r^{\mathcal{I}} \subseteq r^{\mathcal{I}}$
bottom	\perp	\emptyset		$r \circ r_1 \circ \dots \circ r_n \sqsubseteq r$	$r^{\mathcal{I}} \circ r_1^{\mathcal{I}} \circ \dots \circ r_n^{\mathcal{I}} \subseteq r^{\mathcal{I}}$
negation	$\neg C$	$\Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$		$r \circ r \sqsubseteq r$	$r^{\mathcal{I}} \circ r^{\mathcal{I}} \subseteq r^{\mathcal{I}}$
conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$	In RIAs, $r \in \mathcal{P}_2^{\text{ss}}$, while $r_i \in \mathcal{R}$ and $r_i \prec r$ for all $i \in \{1, \dots, n\}$.		
disjunction	$C \sqcup D$	$C^{\mathcal{I}} \cup D^{\mathcal{I}}$			
univ. restriction	$\forall r.C$	$\{\delta \mid \forall y. (\delta, \delta') \in r^{\mathcal{I}} \rightarrow \delta' \in C^{\mathcal{I}}\}$			
exist. restriction	$\exists r.C$	$\{\delta \mid \exists y. (\delta, \delta') \in r^{\mathcal{I}} \wedge \delta' \in C^{\mathcal{I}}\}$			
<i>Self</i> concept	$\exists r.Self$	$\{\delta \mid (\delta, \delta) \in r^{\mathcal{I}}\}$			
qualified number restrictions	$\leq n r.C$	$\{\delta \mid \#\{\delta' \in C^{\mathcal{I}} \mid (\delta, \delta') \in r^{\mathcal{I}}\} \leq n\}$			
	$\geq n r.C$	$\{\delta \mid \#\{\delta' \in C^{\mathcal{I}} \mid (\delta, \delta') \in r^{\mathcal{I}}\} \geq n\}$			

Table 1: $SR\mathcal{OIQ}b_s$ role, concept expressions and axioms. $C \equiv D$ abbreviates $C \sqsubseteq D$, $D \sqsubseteq C$.

is then extended to \mathcal{R} by making all elements of \mathcal{R}^s \prec -minimal. The syntax of *concept expressions* is given by

$$C, D ::= \mathbf{A} \mid \{a\} \mid \top \mid \perp \mid \neg C \mid C \sqcap D \mid C \sqcup D \mid \forall r.C \mid \exists r.C \mid \exists r'.Self \mid \leq n r'.C \mid \geq n r'.C,$$

with $\mathbf{A} \in \mathcal{P}_1$, $a \in \mathcal{C}$, $r \in \mathcal{R}$, $r' \in \mathcal{R}^s$, and $n \in \mathbb{N}$. We note that any concept expression can be put in negation normal form, where negation only occurs in front of concept names, nominals, or *Self* concepts. The different types of $SR\mathcal{OIQ}b_s$ sentences (called *axioms*) are given in Table 1.⁵

Similar to FOL, the semantics of $SR\mathcal{OIQ}b_s$ is defined via interpretations $\mathcal{I} = (\Delta, \cdot^{\mathcal{I}})$ composed of a non-empty set Δ called the *domain of \mathcal{I}* and a function $\cdot^{\mathcal{I}}$ mapping individual names to elements of Δ , concept names to subsets of Δ , and role names to subsets of $\Delta \times \Delta$. This mapping is extended to role and concept expressions and finally used to define satisfaction of axioms (see Table 1).

4.2 Standpoint- $SR\mathcal{OIQ}b_s$

The set $\mathbb{S}_{[SR\mathcal{OIQ}b_s]}$ of *sentential Standpoint- $SR\mathcal{OIQ}b_s$ sentences* is now defined inductively as follows:

- if \mathbf{Ax} is a $SR\mathcal{OIQ}b_s$ axiom then $\mathbf{Ax} \in \mathbb{S}_{[SR\mathcal{OIQ}b_s]}$,
- if $\phi, \psi \in \mathbb{S}_{[SR\mathcal{OIQ}b_s]}$ then $\neg\phi$, as well as $\phi \wedge \psi$ and $\phi \vee \psi$ are in $\mathbb{S}_{[SR\mathcal{OIQ}b_s]}$,
- if $\phi \in \mathbb{S}_{[SR\mathcal{OIQ}b_s]}$ and $e \in \mathcal{E}_S$ then $\Box_e \phi \in \mathbb{S}_{[SR\mathcal{OIQ}b_s]}$ and $\Diamond_e \phi \in \mathbb{S}_{[SR\mathcal{OIQ}b_s]}$.

The semantics of sentential Standpoint- $SR\mathcal{OIQ}b_s$ is defined in the obvious way, by “plugging” the semantics of $SR\mathcal{OIQ}b_s$ axioms into the semantics of $\mathbb{S}_{[FO]}$. We say a $\mathbb{S}_{[SR\mathcal{OIQ}b_s]}$ sentence ϕ is in *negation normal form* (NNF), if negation occurs only inside or directly in front of $SR\mathcal{OIQ}b_s$ axioms; obviously every Standpoint- $SR\mathcal{OIQ}b_s$ sentence can be efficiently transformed into an equivalent one in NNF.

⁵ The original definition of $SR\mathcal{OIQ}$ contained more axioms (role transitivity, (a)symmetry, (ir)reflexivity and disjointness), but these are syntactic sugar in our setting.

4.3 Coping with Peculiarities of $SR\mathcal{OIQ}b_s$

In the following, we will provide a polynomial translation, mapping any $\mathbb{S}_{[SR\mathcal{OIQ}b_s]}$ sentence ϕ to an equisatisfiable set of $SR\mathcal{OIQ}b_s$ axioms. This translation is very much in the spirit of the one presented for sentential FOSL, however, $SR\mathcal{OIQ}b_s$ comes with diverse syntactic impediments that we need to circumvent. Thus, before presenting the translation, we will briefly discuss these issues and how to solve them.

First, $SR\mathcal{OIQ}b_s$ does not provide nullary predicates (i.e., propositional symbols). As a surrogate, we use concept expressions of the form $\forall u.A$ which have the pleasant property of holding either for all domain individuals or for none. Second, $SR\mathcal{OIQ}b_s$ does not directly allow for arbitrary Boolean combinations of axioms. For all non-RIA axioms, a more or less straightforward equivalent encoding is possible using nominals and the universal role; for instance the expression $\neg[r(a, b)] \vee [A \sqsubseteq B]$ can be converted into $\top \sqsubseteq \neg\exists u.(\{a\} \sqcap \exists r.\{b\}) \sqcup \forall u.(\neg A \sqcup B)$.

Dealing with RIAs requires auxiliary vocabulary; for negated RIAs, we introduce a fresh nominal, say $\{x\}$, to mark the end of a “violating” role chain, so $\neg[s \circ s \sqsubseteq r]$ essentially becomes $\top \sqsubseteq \exists u.((\exists s.\exists s_1.\{x\}) \sqcap (\neg\exists r.\{x\}))$.

Unnegated RIAs are even trickier. There is no way of converting them into GCIs, so we have to keep them, but we attach an additional “guard”, which allows us to disable them whenever necessary. This guard can then be triggered from within a GCI. For an example, consider the expression $[t \circ t' \sqsubseteq r] \vee [t' \circ t \sqsubseteq r]$. Then, introducing fresh “guard roles” s_1 and s_2 , we assert the three axioms $\top \sqsubseteq (\forall u.\exists s_1.Self) \sqcup (\forall u.\exists s_2.Self)$ as well as $s_1 \circ t \circ t' \sqsubseteq r$ and $s_2 \circ t' \circ t \sqsubseteq r$. With this arrangement, the first axiom will ensure that all domain elements carry an s_1 -loop or all domain elements carry an s_2 -loop. Depending on that choice, the corresponding RIA in the second line will behave like its original, unguarded version, while the other one may be entirely disabled.

The introduced strategy for handling positive RIAs has a downside: due to the restricted shapes of RIAs (governed by \prec), axioms of the shape $r \circ r \sqsubseteq r$ (expressing transitivity) cannot be endowed with guards. In order to overcome this nuisance, every nonsimple role r has to be accompanied by a subrole \underline{r} , which acts as a “lower approximation” of r and – whenever r is defined transitive – “feeds into” r via tail recursion. This way of reformulating $r \circ r \sqsubseteq r$ allows to attach the wanted guard, but requires adjustments in some axioms that mention r .

4.4 Translation into Plain $SR\mathcal{OIQ}b_s$

We now assume a given $\mathbb{S}_{[SR\mathcal{OIQ}b_s]}$ sentence ϕ , w.l.o.g. in NNF, and provide the formal definition of the translation. As before, we fix $\Pi_{|\phi|}$ and let our translation’s vocabulary $\mathbb{V}_{[SR\mathcal{OIQ}b_s]}(\phi)$ consist of all individual names inside ϕ , plus, for each $\pi \in \Pi_{|\phi|}$, the following symbols: (a) a concept name A^π for each $A \in \mathcal{P}_1$; (b) a simple role name s^π for each $s \in \mathcal{P}_2^s$; (c) non-simple role names r^π and \underline{r}^π for each $r \in \mathcal{P}_2^{ns} \setminus \{u\}$; (d) a simple role name s_ρ^π for each unnegated RIA ρ inside ϕ ; (e) a fresh constant name a_ρ^π for each negated RIA ρ inside ϕ ; (f) a concept name M_π^s for each $s \in \mathcal{S}$. Thereby, the non-simple role names inherit their ordering \prec from \mathcal{P}_2^{ns} and we also let $\underline{r}^\pi \prec r^\pi$ for each $r \in \mathcal{P}_2^{ns} \setminus \{u\}$.

The translation $\text{Trans}(\phi)$ of ϕ is then a set of \mathcal{SROIQ} axioms defined as follows: First, $\text{Trans}(\phi)$ contains the RIA $\underline{\mathbf{r}}^\pi \sqsubseteq \mathbf{r}^\pi$ for every $\mathbf{r} \in \mathcal{P}_2^{\text{ns}} \setminus \{\mathbf{u}\}$ and each $\pi \in \Pi_{|\phi|}$. Second, for every unnegated RIA ρ inside ϕ and each $\pi \in \Pi_{|\phi|}$, $\text{Trans}(\phi)$ contains the RIA $BG_\pi(\rho)$, with BG_π defined by

$$\begin{aligned} r_1 \circ \dots \circ r_n \sqsubseteq \mathbf{r} &\mapsto \mathbf{s}_\rho^\pi \circ r_1^\pi \circ \dots \circ r_n^\pi \sqsubseteq \underline{\mathbf{r}}^\pi & r_1 \circ \dots \circ r_n \circ \mathbf{r} \sqsubseteq \mathbf{r} &\mapsto \mathbf{s}_\rho^\pi \circ r_1^\pi \circ \dots \circ r_n^\pi \circ \underline{\mathbf{r}}^\pi \sqsubseteq \underline{\mathbf{r}}^\pi \\ \mathbf{r} \circ r_1 \circ \dots \circ r_n \sqsubseteq \mathbf{r} &\mapsto \underline{\mathbf{r}}^\pi \circ r_1^\pi \circ \dots \circ r_n^\pi \circ \mathbf{s}_\rho^\pi \sqsubseteq \underline{\mathbf{r}}^\pi & \mathbf{r} \circ \mathbf{r} \sqsubseteq \mathbf{r} &\mapsto \mathbf{s}_\rho^\pi \circ \underline{\mathbf{r}}^\pi \circ \mathbf{r}^\pi \sqsubseteq \mathbf{r}^\pi, \end{aligned}$$

whereby the role expression r^π is obtained from r by substituting every role name \mathbf{s} with \mathbf{s}^π (except \mathbf{u} , which remains unaltered). Third and last, $\text{Trans}(\phi)$ contains the GCI

$$\top \sqsubseteq \prod_{\pi \in \Pi_{|\phi|}} \text{trans}(\pi, \phi) \sqcap \prod_{\pi \in \Pi_{|\phi|}} \forall \mathbf{u}. \mathbf{M}_\pi^*$$

where, by inductive definition,

$$\begin{aligned} \text{trans}(\pi, \mathbf{Ax}) &= \text{trans}^+(\pi, \mathbf{Ax}) \\ \text{trans}(\pi, \neg \mathbf{Ax}) &= \text{trans}^-(\pi, \mathbf{Ax}) \\ \text{trans}(\pi, \psi_1 \wedge \psi_2) &= \text{trans}(\pi, \psi_1) \sqcap \text{trans}(\pi, \psi_2) \\ \text{trans}(\pi, \psi_1 \vee \psi_2) &= \text{trans}(\pi, \psi_1) \sqcup \text{trans}(\pi, \psi_2) \\ \text{trans}(\pi', \square_e \psi) &= \prod_{\pi \in \Pi_{|\phi|}} (\neg \text{trans}_\mathcal{E}(\pi, \mathbf{e}) \sqcup \text{trans}(\pi, \psi)) \\ \text{trans}(\pi', \diamond_e \psi) &= \bigsqcup_{\pi \in \Pi_{|\phi|}} (\text{trans}_\mathcal{E}(\pi, \mathbf{e}) \sqcap \text{trans}(\pi, \psi)) \end{aligned}$$

We next present the translation of unnegated and negated \mathcal{SROIQ} axioms (ρ stands for an RIA $r_1 \circ \dots \circ r_m \sqsubseteq \mathbf{r}$):

$$\begin{aligned} \text{trans}^+(\pi, \rho) &= \forall \mathbf{u}. \exists \mathbf{s}_\rho^\pi. \text{Self} & \text{trans}^-(\pi, \rho) &= \exists \mathbf{u}. ((\forall \mathbf{r}^\pi. \neg \{a_\rho^\pi\}) \sqcap (\exists \underline{\mathbf{r}}_1^\pi \dots \exists \underline{\mathbf{r}}_m^\pi. \{a_\rho^\pi\})) \\ \text{trans}^+(\pi, C \sqsubseteq D) &= \forall \mathbf{u}. (\neg C \sqcup D)^\pi & \text{trans}^-(\pi, C \sqsubseteq D) &= \exists \mathbf{u}. (C \sqcap \neg D)^\pi \\ \text{trans}^+(\pi, C(a)) &= \exists \mathbf{u}. (\{a\} \sqcap C^\pi) & \text{trans}^-(\pi, C(a)) &= \exists \mathbf{u}. (\{a\} \sqcap (\neg C)^\pi) \\ \text{trans}^+(\pi, r(a, b)) &= \exists \mathbf{u}. (\{a\} \sqcap \exists \underline{\mathbf{r}}^\pi. \{b\}) & \text{trans}^-(\pi, r(a, b)) &= \exists \mathbf{u}. (\{a\} \sqcap \forall \mathbf{r}^\pi. \neg \{b\}) \\ \text{trans}^+(\pi, a \doteq b) &= \exists \mathbf{u}. (\{a\} \sqcap \{b\}) & \text{trans}^-(\pi, a \doteq b) &= \exists \mathbf{u}. (\{a\} \sqcap \neg \{b\}) \end{aligned}$$

Therein, for any role expression r , we let \underline{r} denote $\underline{\mathbf{r}}$ if $r = \mathbf{r}$ is a non-simple role name, and otherwise $\underline{r} = r$. Moreover, C^π denotes the concept expression that is obtained from C by transforming it into negation normal form, replacing concept names \mathbf{A} with \mathbf{A}^π and role expressions \mathbf{r} by \mathbf{r}^π , and replacing every $\exists \mathbf{r}$ for non-simple \mathbf{r} with $\exists \underline{\mathbf{r}}$.

As before, $\text{trans}_\mathcal{E}$ implements the semantics of standpoint expressions, but now adjusted to the new framework: each expression $\mathbf{e} \in \mathcal{E}_\mathcal{S}$ is transformed into a concept expression $\text{trans}_\mathcal{E}(\pi, \mathbf{e})$ over the vocabulary $\{\mathbf{M}_\pi^s \mid s \in \mathcal{S}, \pi \in \Pi_{|\phi|}\}$ as follows:

$$\begin{aligned} \text{trans}_\mathcal{E}(\pi, \mathbf{s}) &= \forall \mathbf{u}. \mathbf{M}_\pi^s \\ \text{trans}_\mathcal{E}(\pi, \mathbf{e}_1 \cup \mathbf{e}_2) &= \text{trans}_\mathcal{E}(\pi, \mathbf{e}_1) \sqcup \text{trans}_\mathcal{E}(\pi, \mathbf{e}_2) \\ \text{trans}_\mathcal{E}(\pi, \mathbf{e}_1 \cap \mathbf{e}_2) &= \text{trans}_\mathcal{E}(\pi, \mathbf{e}_1) \sqcap \text{trans}_\mathcal{E}(\pi, \mathbf{e}_2) \\ \text{trans}_\mathcal{E}(\pi, \mathbf{e}_1 \setminus \mathbf{e}_2) &= \text{trans}_\mathcal{E}(\pi, \mathbf{e}_1) \sqcap \neg \text{trans}_\mathcal{E}(\pi, \mathbf{e}_2) \end{aligned}$$

With all definitions in place, we obtain the desired result.

Theorem 3. *Given $\phi \in \mathbb{S}_{[\mathcal{SROIQb}_s]}$, the set $\text{Trans}(\phi)$ (i) is a valid \mathcal{SROIQb}_s knowledge base, (ii) is equisatisfiable with ϕ , (iii) is of polynomial size wrt. ϕ , and (iv) can be computed in polynomial time.*

5 Example in the Forestry Domain

We consider an extension of Example 1 in Sentential Standpoint- \mathcal{SROIQ}_s to illustrate the main reasoning tasks in more detail. The following additional axiom specifies that forest land use and urban land use are disjoint subclasses of land (F9).

$$(F9) \quad \Box_*[\text{ForestLandUse} \sqcup \text{UrbanLandUse} \sqsubseteq \text{Land} \wedge \text{ForestLandUse} \sqcap \text{UrbanLandUse} \sqsubseteq \perp]$$

Now, let us see how, through inferences in $\mathbb{S}_{[\mathcal{SROIQ}_s]}$, we can gather unequivocal knowledge (Uneq), obtain knowledge that is relative to a standpoint (Rel), and contrast the knowledge that can be inferred from different standpoints (Cont). For unequivocal knowledge (Uneq), we can infer unambiguously that forests are no urban-use lands:

$$\Box_*[\text{Forest} \sqsubseteq \neg \text{UrbanLandUse}]$$

This holds because each precisification must comply with LC or LU (F8), and we have $\Box_{LC}[\text{Forest} \sqsubseteq \neg \text{UrbanLandUse}]$ from (F1) and (F5), and $\Box_{LU}[\text{Forest} \sqsubseteq \neg \text{UrbanLandUse}]$ from (F3) and (F9). Regarding relative (Rel) and contrasting (Cont) knowledge, if we now wanted to query our knowledge base for instances of forest, we would obtain

$$\Box_{LC}[\text{Forest}(e)] \wedge \Box_{LC}[\neg \text{Forest}(l)] \quad \mathcal{I}_*[\text{Forest}(e)] \wedge \mathcal{I}_*[\text{Forest}(l)]$$

The first deduced formula contains knowledge relative to LC, showing its stance on whether the instances constitute a forest, which happens to be conclusive in both cases. The second formula states the global indeterminacy of both l 's and e 's membership to the concept Forest. This stems from the disagreement between the interpretations LC and LU, whose overall incompatibility ($\Box_{LC \cap LU}[\top \sqsubseteq \perp]$) can also be inferred.

Finally, it is worth looking at the limitations of the sentential fragment of Standpoint \mathcal{SROIQ}_s . In a non-sentential setting, where modalities can be used at the concept level, “complex alignments” or bridges can be established between concepts according to possibly many standpoints. For instance, one can write

$$\Box_{LU}[\text{Forest}] \sqsubseteq \Box_{LC}[\exists \text{hasLand}^- . \text{Forest}] \sqcup \Box_*[\text{Cleared}]$$

to express that the areas classified as forest according to LU belong to a forest according to LC or have been cleared (in which case LC does not recognise them as forest). It is an objective of future work of ours to study decidable fragments for which the restrictions on the use of modalities are relaxed to express such kinds of axioms.

6 Related Work

A variety of formal representation systems have been proposed to model perspectives in rather diverse areas of research and with heterogeneous nomenclatures. Standpoint logic bears some similarities to *context logic* in the style proposed by McCarthy and Buvac [19], which has also been applied in a description logic setting [15]. This tradition treats contexts as “first-class citizens” of the logic, i.e., full-fledged formal objects over which one can express first-order properties. In contrast, standpoint logic is suitable when a formalisation of the contexts involved is unfeasible, or when the interest resides in the content of the standpoints rather than the context in which they occur.

Another related notion is that of *ontology views*, where some works consider potentially conflicting viewpoints [25, 12, 11]. Ribière and Dieng [25] and Heman et al. [12, 11]

implement the intuition of “viewpoints” via ad-hoc extensions of the syntax and semantics of description logics, in a style similar to the work on contextuality by Benslimane et al. [4]. Gorshkov et al. [9] implement them using named graphs. Instead, the standpoint approach extends the base language with modalities and provides a Kripke-style semantics for it. This leads to a simpler, more recognisable and more expressive framework that supports, for instance, hierarchies and combinations of standpoints, inferences of partial truths, the preservation of consistency with the established alignments and inferences about the standpoints themselves. On a technical level, first-order standpoint logic can be seen as a many-dimensional (multi-)modal logic [16], whence results from that area apply to our setting. In particular, the search for non-trivial fragments of first-order modal logics that are still decidable and even practically relevant is an important endeavour, for which we believe that standpoint logic can play useful role.

Finally, in the area of ontology modularity, different formalisms such as *DDL bridge rules* [5] and ε -connections [17,18] 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 that do not need to be entirely coherent with each other. Yet the motivation is inherently different: while the standpoint framework focuses on integrating possibly overlapping knowledge into a global source (while preserving “standpoint-provenance” and thus enabling a peaceful coexistence of conflicting information), DDL bridge rules and ε -connections have been devised to establish a certain synchronisation between modules that are and will remain separate. DDL bridge rules could, however, be simulated within a standpoint framework.

7 Conclusions and Future Work

The diversity of human world views along with the semantic heterogeneity of natural language are at the heart of well-recognised knowledge interoperability challenges. As an alternative to the common strategy of merging, we proposed the use of a logical formalism based on the notion of *standpoint* that is suitable for knowledge representation and reasoning with sets of possibly conflicting characterisations of a domain.

Using first-order logic as an expressive underlying language, we proposed a multi-modal framework by means of which different agents can establish their individual standpoints (which typically involves specifying constraints and relations), but which also allows for combining standpoints and establishing alignments between them. Reasoning tasks over such multi-standpoint specifications include gathering unequivocal 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.

Remarkably, the simplified Kripke semantics allows us to establish a small model property for the *sentential* fragment of FOSL. This result gave rise to a polynomial, satisfiability-preserving translation into the base logic, which also maintains membership in diverse decidable fragments, immediately implying that for a range of logics, reasoning in their standpoint-enhanced versions does not increase their computational complexity. This indicates that the framework can be applied to ontology alignment, concept negotiation, and knowledge aggregation with inference systems built on top of existing, highly optimised off-the-shelf reasoners.

Future work includes the study of the complexity of FOSL fragments allowing the presence of free variables within the scope of modalities. Note however that in the general case of FOSL, this leads to the loss of the small model property.

Example 2. Consider the following (non-sentential) FOSL sentence, axiomatising "better" (Btt) to be interpreted as a non-well-founded strict linear order and requiring for every domain element x (of infinitely many) the existence of some precisification where x is the (one and only) "best":

$$\begin{aligned} \forall xyz((\text{Btt}(x, y) \wedge \text{Btt}(y, z)) \rightarrow \text{Btt}(x, z)) \quad \wedge \quad \forall xy\neg(\text{Btt}(x, y) \wedge \text{Btt}(y, x)) \wedge \\ \forall xy(x \neq y \rightarrow (\text{Btt}(x, y) \vee \text{Btt}(y, x))) \quad \wedge \quad \forall x\exists y\text{Btt}(x, y) \wedge \forall x\Diamond_*\neg\exists y\text{Btt}(y, x) \end{aligned}$$

Obviously, this sentence is satisfiable, but only in a model with infinitely many precisifications; that is, the small model property is violated in the worst possible way.

On the other hand, it is desirable from a modelling perspective to allow for some interplay between FO quantifiers and standpoint modalities. E.g., the non-sentential FOSL sentence $\forall x_1 \cdots x_k (\text{P}(x_1, \dots, x_k) \rightarrow \Box_* \text{P}(x_1, \dots, x_k))$ expresses the rigidity of a predicate P , thereby "synchronising" it over all precisifications.

Consequently, we will study how by imposing syntactic restrictions, we can guarantee the existence of small (or at least reasonably-sized) models for non-sentential standpoint formulas. Results in the field of many-dimensional modal logics [16] show that reasoning is decidable for diverse fragments of first-order modal logic such as the *monodic* fragment (where modalities occur only in front of formulas with at most one free variable). However, Example 2 already shows (by virtue of being within the monodic fragment) that we cannot hope for a small model property even for this slight extension of the sentential fragment. A detailed analysis of these issues as they apply to the simplified semantics of standpoint logic is the object of current work.

Additionally, we intend to implement the proposed translations and perform experiments to test the performance of the standpoint framework in scenarios of Knowledge Integration. While sentential standpoints can be added at no extra cost in complexity for the discussed fragments in this paper, we intend to run experiments to assess the runtime impact on large knowledge bases with off-the-shelf reasoners.

As another important topic toward the deployment of our framework, we will look into conceptual modelling aspects. Reviewing documented recurrent scenarios and patterns in the area of knowledge integration, we intend to establish guiding principles for conveniently encoding those by using novel strategies possible with structures of standpoints. Examples for such scenarios include the disambiguation of knowledge sources by using combinations of standpoints, and the establishment of bridge-like rules for alignment. For the latter we will investigate their relationship to similar constructs from other frameworks such as ε -connections, distributed description logics, and others.

Supplemental Material Statement: Proofs can be found in the extended version [30].

Acknowledgments: Lucía Gómez Álvarez was supported by the *Bundesministerium für Bildung und Forschung* (BMBF) in the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI). Sebastian Rudolph has received funding from the European Research Council (Grant Agreement no. 771779, DeciGUT).

References

1. Arp, R., Smith, B., Spear, A.D.: Building ontologies with basic formal ontology. MIT Press (2015)
2. Bárány, V., Benedikt, M., ten Cate, B.: Some model theory of guarded negation. *Journal of Symbolic Logic* **83**(4), 1307–1344 (2018)
3. Bárány, V., ten Cate, B., Segoufin, L.: Guarded negation. *Journal of the ACM* **62**(3) (2015)
4. 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*. pp. 168–176. Springer (2006)
5. Borgida, A., Serafini, L.: Distributed Description Logics: Assimilating information from peer sources. *Journal on Data Semantics* **2800**, 153–184 (2003)
6. Euzenat, J., Mocan, A., Scharffe, F.: Ontology alignments. In: Hepp, M., Leenheer, P.D., Moor, A.D., Sure, Y. (eds.) *Ontology Management. Computing for Human Experience*, vol. 7, pp. 177–206. Springer (2008)
7. Fine, K.: Vagueness, truth and logic. *Synthese* **30**(3-4), 265–300 (1975)
8. Gómez Álvarez, L., Rudolph, S.: Standpoint logic: Multi-perspective knowledge representation. In: Neuhaus, F., Brodaric, B. (eds.) *Procs. of the 12th Int. Conf. on Formal Ontology in Information Systems. FAIA*, vol. 344, pp. 3–17. IOS Press (2021)
9. Gorshkov, S., Kralin, S., Miroshnichenko, M.: Multi-viewpoint ontologies for decision-making support. In: Ngonga Ngomo, A.C., Křemen, P. (eds.) *Knowledge Engineering and Semantic Web*. pp. 3–17. Springer (2016)
10. Grädel, E., Otto, M., Rosen, E.: Two-variable logic with counting is decidable. In: *Procs. of the 12th Annual IEEE Symposium on Logic in Computer Science*. pp. 306–317. IEEE (1997)
11. Hemam, M.: An extension of the ontology web language with multi-viewpoints and probabilistic reasoning. *Int. Journal of Advanced Intelligence Paradigms* **10**(3), 247–265 (2018)
12. Hemam, M., Boufaïda, Z.: MVP-OWL: A multi-viewpoints ontology language for the Semantic Web. *Int. Journal of Reasoning-based Intelligent Systems* **3**(3-4), 147–155 (2011)
13. Horrocks, I., Kutz, O., Sattler, U.: The even more irresistible SROIQ. In: Doherty, P., Mylopoulos, J., Welty, C.A. (eds.) *Procs. 10th Int. Conf. on Principles of Knowledge Representation and Reasoning*. pp. 57–67. AAAI Press (2006)
14. Kieronski, E., Rudolph, S.: Finite model theory of the triguarded fragment and related logics. In: *Procs. of the 36th Annual ACM/IEEE Symposium on Logic in Computer Science*. pp. 1–13. IEEE (2021)
15. Klarman, S., Gutiérrez-Basulto, V.: Description logics of context. *Journal of Logic and Computation* **26**(3), 817–854 (2013)
16. Kurucz, A., Wolter, F., Zakharyashev, M., Gabbay, D.M.: *Many-Dimensional Modal Logics: Theory and Applications*. *Studies in Logics and the Foundations of Mathematics*, Elsevier, 1st edn. (2003)
17. Kutz, O., Lutz, C., Wolter, F., Zakharyashev, M.: E-connections of description logics. *Description Logics Workshop* **81** (2003)
18. Kutz, O., Lutz, C., Wolter, F., Zakharyashev, M.: E-connections of abstract description systems. *Artificial intelligence* **156**(1), 1–73 (2004)
19. McCarthy, J., Buvac, S.: Formalizing context (expanded notes). *CSLI Lecture Notes* **81**, 13–50 (1998)
20. Osman, I., Ben Yahia, S., Diallo, G.: Ontology Integration: Approaches and Challenging Issues. *Information Fusion* **71**, 38–63 (2021)
21. Otero-Cerdeira, L., Rodríguez-Martínez, F.J., Gómez-Rodríguez, A.: Ontology matching: A literature review. *Expert Systems with Applications* **42**(2) (2015)

22. Pesquita, C., Faria, D., Santos, E., Couto, F.M.: 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, pp. 13–24. CEUR (2013)
23. Pietruszczak, A.: Simplified Kripke style semantics for modal logics K45, KB4 and KD45. *Bulletin of the Section of Logic* **38**(3/4), 163–171 (2009)
24. Pratt-Hartmann, I.: Complexity of the two-variable fragment with counting quantifiers. *Journal of Logic, Language, and Information* **14**(3), 369–395 (2005)
25. Ribière, M., Dieng-Kuntz, R.: A viewpoint model for cooperative building of an Ontology. In: *Procs. of the 10th Int. Conf. on Conceptual Structures*. pp. 220–234. Springer (2002)
26. Rudolph, S., Krötzsch, M., Hitzler, P.: Cheap Boolean role constructors for description logics. In: Hölldobler, S., Lutz, C., Wansing, H. (eds.) *Procs. of the 11th European Conf. on Logics in Artificial Intelligence*. vol. 5293, pp. 362–374. Springer (2008)
27. Rudolph, S., Simkus, M.: The triguarded fragment of first-order logic. In: Barthe, G., Sutcliffe, G., Veanes, M. (eds.) *Procs. of the 22nd Int. Conf. on Logic for Programming, Artificial Intelligence and Reasoning*. EPiC Series in Computing, vol. 57, pp. 604–619 (2018)
28. Solimando, A., Jiménez-Ruiz, E., Guerrini, G.: Minimizing conservativity violations in ontology alignments. *Knowledge and Information Systems* **51**(3) (2017)
29. Wolter, F., Zakharyashev, M.: On the decidability of description logics with modal operators. In: Cohn, A.G., Schubert, L.K., Shapiro, S.C. (eds.) *Procs. of the Sixth Int. Conf. on Principles of Knowledge Representation and Reasoning*. pp. 512–523. Morgan Kaufmann (1998)
30. Álvarez, L.G., Rudolph, S., Strass, H.: How to Agree to Disagree: Managing Ontological Perspectives using Standpoint Logic. Tech. rep., arXiv.org (2022)