

Automating Reasoning with Standpoint Logic via Nested Sequents

Tim S. Lyon , Lucía Gómez Álvarez

Computational Logic Group, Faculty of Computer Science, TU Dresden

{timothy_stephen.lyon, lucia.gomez_alvarez}@tu-dresden.de

Abstract

Standpoint logic is a recently proposed formalism in the context of knowledge integration, which advocates a multi-perspective approach permitting reasoning with a selection of diverse and possibly conflicting standpoints rather than forcing their unification. In this paper, we introduce nested sequent calculi for propositional standpoint logics—proof systems that manipulate trees whose nodes are multisets of formulae—and show how to automate standpoint reasoning by means of non-deterministic proof-search algorithms. To obtain worst-case complexity-optimal proof-search, we introduce a novel technique in the context of nested sequents, referred to as *coloring*, which consists of taking a standpoint formula as input, guessing a certain coloring of its subformulae, and then running proof-search in a nested sequent calculus on the colored input. Our technique lets us decide the validity of standpoint formulae in CoNP since proof-search only produces a *partial* proof relative to each permitted coloring of the input. We show how all partial proofs can be fused together to construct a complete proof when the input is valid, and how certain partial proofs can be transformed into a counter-model when the input is invalid. These “certificates” (i.e. proofs and counter-models) serve as explanations of the (in)validity of the input.

1 Introduction

Standpoint Logic. The fact that knowledge bases (KBs) encode the standpoints of their creators (e.g. in the form of viewpoints, contextual factors or semantic commitments) is the source of well-known challenges in the area of knowledge integration. Since semantic heterogeneity between the sources is to be expected, inconsistencies may arise if we attempt to combine them into a single conflict-free conceptual model. To illustrate this, consider three KBs: C, a ‘common-sense’ representation of colours; H, a KB by a *house painting* business that reuses and extends C; and R, a KB that formalises the *RYB color model*, from the fine arts tradition.

Example 1. *According to C, basic colours such as Blue and Green are disjoint. H complies with C and further specifies that Teal is Green. In contrast, according to R it is unequivocal that Teal is both Green and Blue. Generally, it is conceivable that Blue holds.*

These sources cannot be merged without the undesired effect of inconsistency, and circumventing it requires either




Figure 1: Diagrams of C, H and R. \preceq indicates that H extends C.

knowledge weakening or duplication (Pesquita et al. 2013). Instead, one may wish to jointly reason with the KBs, treating them as alternative standpoints on a domain.

Standpoint logic (Gómez Álvarez and Rudolph 2021) is a simple multi-modal logic intended for the representation of knowledge relative to different, possibly conflicting, perspectives. The framework introduces the labeled modalities \Box_s and \Diamond_s for each standpoint s , where $\Box_s \varphi$ is read as “according to s , it is *unequivocal* that φ ” and $\Diamond_s \varphi$ as “according to s , it is *conceivable* that φ ”. In addition, $s \preceq s'$ indicates that the standpoint s is *sharper* than s' , that is, s complies with s' and further specifies it.

$$(F1) \Box_C \neg(\text{Blue} \wedge \text{Green}) \wedge \Box_R(\text{Teal} \rightarrow (\text{Blue} \wedge \text{Green}))$$

$$(F2) (H \preceq C) \quad (F3) \Box_H(\text{Teal} \rightarrow \text{Green}) \quad (F4) \Diamond_*(\text{Blue})$$

The formulae (F1)-(F4) formalize Example 1, (illustrated in Figure 1) in propositional standpoint logic. (F1) encodes that Blue and Green are unequivocally disjoint according to standpoint C, while according to standpoint R it is unequivocal that Teal implies both Blue and Green. (F2) encodes that H includes the knowledge of C, and (F3) that Teal is Green according to H. Last, (F4) encodes that Blue holds under some interpretations by using the universal standpoint $*$, which sits atop any hierarchy of standpoints and is used to reference knowledge that is unequivocally true or conceivable among all perspectives.

In addition to representing unequivocal and conceivable facts (e.g. (F3) and (F4)), which may be relative to standpoints (e.g. (F1)), hold universally (e.g. (F4)), or establish a hierarchy of standpoints (e.g. (F2)), one may also express (in)determinate knowledge by means of the (definable) dual operators \mathcal{I}_s and \mathcal{D}_s . The indeterminacy operator

$\mathcal{I}_s\varphi := \diamond_s\varphi \wedge \diamond_s\neg\varphi$ makes explicit that both φ and $\neg\varphi$ are conceivable in the context of s , thus making φ inherently indeterminate. Finally, the framework can be used to establish correspondences or bridges between the standpoints themselves. For instance, (F5) encodes that if something is *Teal* according to R, then it is *Green* for C and R.

$$(F5) \quad \Box_R(\text{Teal}) \rightarrow (\Box_C \text{Green} \wedge \Box_R \text{Green})$$

Natural reasoning tasks over 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. To illustrate, let us assume $\Box_*\text{Teal}$ and examine some inferences that we can draw from this in the setting of Example 1. On the one hand, from (F5), (F3) and $\Box_*\text{Teal}$ we obtain that green is unequivocal for the three standpoints: $\Box_C \text{Green}$, $\Box_H \text{Green}$ and $\Box_R \text{Green}$. On the other hand, we can infer the global indeterminacy of blue $\mathcal{I}_*\text{Blue}$, because (i) *Teal* holds universally, (ii) it is unequivocal for R that *Teal* implies *Blue* (F1), hence $\diamond_*\text{Blue}$, and (iii) we know $\Box_C \text{Green}$, which together with (F1) implies $\Box_C \neg\text{Blue}$ and thus $\diamond_* \neg\text{Blue}$.

Conveniently, the satisfiability problem in propositional standpoint logic is known to be NP-complete (Gómez Álvarez and Rudolph 2021), in pleasant contrast to the PSPACE-completeness normally exhibited by multi-modal epistemic logics. This result, obtained via a translation to one-variable first-order logic, makes the framework attractive in applied scenarios, and prompts our work to provide a suitable proof-theory for standpoint logic. Not only can our proof systems be leveraged to provide a proof-search procedure deciding the validity of standpoint formulae, but our proof-theoretic approach yields *witnesses*, that is, proofs of valid formulae and counter-models of invalid formulae. Such “certificates” (i.e. proofs and counter-models) possess explanatory value, and may be used, for instance, to trace the standpoints involved in a certain inference; e.g. when a global indeterminacy such as $\mathcal{I}_*\text{Blue}$ is inferred from a large collection of standpoints, we may want to gather the standpoints that hold contrasting views (in this case R and $\{H, C\}$, which can be easily extracted from a proof). Thus, our reliance on proof theory provides essential information that may be used to answer “why” a certain piece of information holds while still allowing “low” complexity reasoning.

Nested Sequents and Proof Theory. Since their inception, sequent systems—which consist of inference rules that syntactically manipulate pairs of multisets of formulae—have proven themselves fruitful in writing decision algorithms for logics (Dyckhoff 1992; Gentzen 1935a; Gentzen 1935b; Slaney 1997). A crucial feature of such systems, and their use in decidability, is the so-called *subformula property*, which a sequent system has *iff* the premise(s) of each inference rule only contain subformulae of the conclusion of the rule. (NB. Systems with the subformula property are also referred to as *analytic*.) With the goal of securing this property for proof systems for theories *beyond classical propositional logic* (e.g. the modal logics *Kt* and *S5*), more sophisticated sequent systems extending Gentzen’s original formalism were eventually proposed;

e.g., see (Belnap 1982; Simpson 1994; Wansing 2002). In this paper, we employ one such extended formalism, viz. the *nested sequent formalism* (Brünnler 2009; Bull 1992; Kashima 1994; Poggiolesi 2009), which utilizes trees of multisets of formulae in deriving theorems. Such systems have proven well-suited for automated reasoning with modal and related logics, being used (for instance) in the writing of decision/proof-search algorithms (Brünnler 2009; Tiu, Ianovski, and Goré 2012) and the extraction of interpolants (Fitting and Kuznets 2015; Lyon et al. 2020).

Drawing on ideas from the *structural refinement* methodology, detailed in (Lyon 2021a) and used to provide nested sequent systems for diverse classes of modal and constructive logics (see (Lyon and van Berkel 2019; Lyon 2020; Lyon 2021b)), our first contribution in this paper is the introduction of analytic nested sequent systems (each dubbed $\text{NS}(\mathcal{V})$ with \mathcal{V} a certain parameter) for propositional standpoint logics (Gómez Álvarez and Rudolph 2021). For our second contribution, we exploit our nested systems to write concrete, *worst-case complexity-optimal* proof-search algorithms (deciding the validity of propositional standpoint formulae in CoNP), which apply inference rules from $\text{NS}(\mathcal{V})$ *in reverse* on an input formula with the goal of building a proof thereof. Whereas typical proof-search algorithms operate deterministically and attempt to build a *complete* proof of the input, we introduce a novel technique (our third contribution) referred to as *coloring*, which performs proof-search *non-deterministically* and which only constructs a *partial* proof of the input relative to each non-deterministic choice. The technique of coloring involves first guessing a particular labeling of the subformulae of an input formula with *active* \circ and *inactive* \bullet labels, with the proof-search algorithm subsequently only processing data deemed active. An interesting consequence of this technique is the attainment of a CoNP proof-search algorithm as the partial proofs constructed during proof-search are at most polynomially larger than the input and only require polynomial time to compute. Moreover, in the instance where the input formula is invalid, we show how to construct a counter-model from failed proof-search, and in the instance where our input formula is valid, we provide a procedure that generates a *complete* proof witnessing the validity of the input formula by patching together all *partial* proofs (our fourth contribution).

Organization of Paper. Our paper is organized as follows: Sect. 2 presents the syntax and semantics of propositional standpoint logic. In Sect. 3, we introduce our nested sequent systems for propositional standpoint logics, proving such systems sound and concluding their completeness. In the penultimate section (Sect. 4), we introduce the method of coloring and show how to automate reasoning with standpoint logics, that is, we provide a (worst-case complexity-optimal) proof-search algorithm deciding the validity of propositional standpoint formulae in CoNP. The final section (Sect. 5) concludes the paper and discusses future work.

2 Standpoint Logic

Let us now specify the syntax of propositional standpoint logic (SL), denoted by \mathbb{S} .

Definition 1 (Syntax of Standpoint Logic). Let $\mathcal{V} = \langle \mathcal{P}, \mathcal{S} \rangle$ be a vocabulary where \mathcal{P} is a non-empty set of propositional variables and \mathcal{S} is a set of standpoint symbols containing the distinguished symbol $*$, i.e. the universal standpoint. We define the language $\mathcal{L}_{\preceq} := \{s \preceq s' \mid s, s' \in \mathcal{S} \setminus \{*\}\}$, and the language $\mathcal{L}_{\mathcal{V}}$ via the following grammar in BNF:

$$\varphi ::= p \mid \neg p \mid (\varphi \vee \psi) \mid (\varphi \wedge \psi) \mid \Box_s \varphi \mid \Diamond_s \varphi$$

where $p \in \mathcal{P}$ and $s \in \mathcal{S}$. We also use \top and \perp as shorthands with the usual definitions.

Last, for $\Gamma \subseteq \mathcal{L}_{\preceq}$ and $\varphi \in \mathcal{L}_{\mathcal{V}}$, we define a standpoint implication to be a formula of the form $\bigwedge \Gamma \rightarrow \varphi$, where $\bigwedge \Gamma$ is a conjunction of all elements of Γ , which equals \top when Γ is empty.

We make use of formulae in negation normal form as this will simplify the structures present in our nested systems and enhance the readability of our proof theory. To further simplify, we also assume w.l.o.g. that sets of sharpening statements are (1) free of cycles $s_1 \preceq s_2, \dots, s_n \preceq s_1$ and (2) omit occurrences of $*$. Assumption (1) is permitted since any standpoint implication containing a cycle $s_1 \preceq s_2, \dots, s_n \preceq s_1$ of standpoints is equivalent to one where the cycle is deleted and all occurrences of s_1, \dots, s_{n-1} are replaced by s_n in the formula. Regarding assumption (2), any sharpening statement with $*$ is either of the form $s \preceq *$, and is thus trivial (see Def. 4 below), or is of the form $* \preceq s$, in which case s can be systematically replaced by $*$ in a standpoint implication to obtain an equivalent one.

Definition 2 (Subformula and Size). We define the set of subformulae of φ , denoted $\text{sufo}(\varphi)$, recursively as follows:

- $\text{sufo}(p) := \{p\}$ and $\text{sufo}(\neg p) := \{\neg p\}$;
- $\text{sufo}(\heartsuit \psi) := \{\heartsuit \psi\} \cup \text{sufo}(\psi)$;
- $\text{sufo}(\psi \otimes \chi) := \{\psi \otimes \chi\} \cup \text{sufo}(\psi) \cup \text{sufo}(\chi)$.

with $\heartsuit \in \{\Diamond_s, \Box_s \mid s \in \mathcal{S}\}$ and $\otimes \in \{\vee, \wedge\}$. We say that ψ is a subformula of φ iff $\psi \in \text{sufo}(\varphi)$, and define the size of a formula φ in $\mathcal{L}_{\mathcal{V}}$, denoted $|\varphi|$, to be equal to $|\text{sufo}(\varphi)|$, i.e. to the number of its subformulae.

In what follows, we introduce the semantics of SL, defined over a structure of precisifications instead of the more usual structure of possible worlds. A *precisification* is a complete and consistent way in which the current state of affairs can be described with a given vocabulary, and standpoints are modelled as sets of precisifications that they consider admissible. This strategy of modelling the variability of natural language as hyper-ambiguity is based on the supervaluationist theory of natural language (Fine 1975; Keefe and Smith 1997), which standpoint logic draws from (Gómez Álvarez and Bennett 2018).

Definition 3 (Standpoint Model). Given a vocabulary \mathcal{V} , a model \mathcal{M} (over \mathcal{V}) is a triple $\langle \Pi, \sigma, \delta \rangle$, where Π is a non-empty set of precisifications, $\sigma : \mathcal{S} \rightarrow 2^\Pi$, and $\delta : \mathcal{P} \rightarrow 2^\Pi$ with $\sigma(s) \neq \emptyset$ for all $s \in \mathcal{S}$ and $\sigma(*) = \Pi$. The set of all such models is denoted by $\mathfrak{M}_{\mathcal{S}}$.

Definition 4 (Semantic Clauses). Let $\Gamma \subseteq \mathcal{L}_{\preceq}$ and $\varphi, \psi \in \mathcal{L}_{\mathcal{V}}$. Moreover, let $\mathcal{M} = \langle \Pi, \sigma, \delta \rangle$ be a standpoint model with $\pi \in \Pi$. We recursively define the satisfaction of a formula on \mathcal{M} at π accordingly:

- $\mathcal{M}, \pi \models p$ iff $\pi \in \delta(p)$;
- $\mathcal{M}, \pi \models \neg p$ iff $\pi \notin \delta(p)$;
- $\mathcal{M}, \pi \models \varphi \wedge \psi$ iff $\mathcal{M}, \pi \models \varphi$ and $\mathcal{M}, \pi \models \psi$;
- $\mathcal{M}, \pi \models \varphi \vee \psi$ iff $\mathcal{M}, \pi \models \varphi$ or $\mathcal{M}, \pi \models \psi$;
- $\mathcal{M}, \pi \models \Diamond_s \varphi$ iff for some $\pi' \in \sigma(s)$, $\mathcal{M}, \pi' \models \varphi$;
- $\mathcal{M}, \pi \models \Box_s \varphi$ iff for all $\pi' \in \sigma(s)$, $\mathcal{M}, \pi' \models \varphi$;
- $\mathcal{M}, \pi \models s \preceq s'$ iff $\sigma(s) \subseteq \sigma(s')$.
- $\mathcal{M}, \pi \models \bigwedge \Gamma$ iff $\mathcal{M}, \pi \models s \preceq s'$ for all $s \preceq s' \in \Gamma$;
- $\mathcal{M}, \pi \models \bigwedge \Gamma \rightarrow \varphi$ iff $\mathcal{M}, \pi \models \bigwedge \Gamma$ implies $\mathcal{M}, \pi \models \varphi$;
- $\mathcal{M} \models \bigwedge \Gamma \rightarrow \varphi$ iff $\mathcal{M}, \pi \models \bigwedge \Gamma \rightarrow \varphi$ for all $\pi \in \Pi$.

A standpoint implication $\bigwedge \Gamma \rightarrow \varphi$ is defined to be valid (relative to a vocabulary \mathcal{V}) iff it is true on each model $\mathcal{M} \in \mathfrak{M}_{\mathcal{S}}$; it is defined to be invalid (relative to \mathcal{V}) otherwise.

For a vocabulary \mathcal{V} , the standpoint logic $\mathbb{S}(\mathcal{V})$ is the set of all valid standpoint implications $\bigwedge \Gamma \rightarrow \varphi$ over $\mathfrak{M}_{\mathcal{S}}$.

It is worth remarking that the specification of sharpening statements in a separate language (viz. \mathcal{L}_{\preceq}) and the above definition of satisfiability and validity contrast with the original presentation in (Gómez Álvarez and Rudolph 2021). However, this specification simplifies our treatment of sharpening statements, which before served as atomic propositions in the language $\mathcal{L}_{\mathcal{V}}$. In fact, these statements are obsolete in extensions of the language allowing set theoretical combinations of standpoints in modalities (which is the object of current research). Moreover, in these extensions, the natural requirement of inner consistency (i.e. the non-emptiness of $\sigma(s)$, for each $s \in \mathcal{S}$) of standpoints is relaxed, which can be easily reflected in our nested sequent systems by dropping the (n_s) rule (see Fig. 2 in Section 3).


3 Nested Sequent Systems

We define a *nested sequent* (which we will also refer to as a *sequent*) to be a formula of the form $\Gamma \vdash \Delta$ with Γ and Δ defined via the following grammars in BNF:

$$\Gamma ::= s \preceq s' \mid \emptyset \mid \Gamma, \Gamma \quad \Delta ::= \Sigma \mid \Delta, (s)[\Sigma]_{\pi}$$

$$\Sigma ::= \varphi \mid \emptyset \mid \Sigma, \Sigma$$

where $s, s' \in \mathcal{S} \setminus \{*\}$, $\varphi \in \mathcal{L}_{\mathcal{V}}$, and π is among a countably infinite set of labels $\{\pi_i \mid i \in \mathbb{N} \setminus \{0\}\}$. We use Φ and Ψ (occasionally annotated) to denote nested sequents and note that we employ the use of labels as this proves useful in extracting a counter-model from failed proof-search (see Thm. 3). Moreover, each nested sequent $\Gamma \vdash \Delta$ with $\Delta = \Sigma_0, (s_1)[\Sigma_1]_{\pi_1}, \dots, (s_n)[\Sigma_n]_{\pi_n}$ possesses a special structure; namely, the *antecedent* Γ is a set of sharpening statements of the form $s \preceq s'$, and the *consequent* Δ is a multiset encoding a tree of depth 1 whose nodes are multisets of formulae from $\mathcal{L}_{\mathcal{V}}$. Δ can be expressed graphically as follows:



We refer to a multiset Σ_i occurring in the consequent of a nested sequent as a *component*, and note that components (along with the antecedent and consequent) are permitted to be empty \emptyset . It is also worthwhile to define the relation \preceq_Γ^* on standpoints as this will be used as a side condition dictating applications of certain inference rules:

Definition 5. For a nested sequent $\Gamma \vdash \Delta$, let $\preceq_\Gamma^* \subseteq \mathcal{S} \times \mathcal{S}$ be the minimal reflexive and transitive relation such that

- $s \preceq_\Gamma^* s'$ for every $s \preceq s' \in \Gamma$, and
- $s \preceq_\Gamma^* *$ for every $s \in \mathcal{S}$.

A nice feature of nested sequents is that such objects typically permit a formula translation, e.g. (Brünnler 2009; Bull 1992; Kashima 1994; Poggiolesi 2009), meaning that our logical semantics can be lifted to the language of our proof systems without introducing an extended semantics for nested sequents.

Definition 6 (Formula Interpretation). We define the formula interpretation of a nested sequent $\Gamma \vdash \Delta$ with $\Delta = \Sigma_0, (s_1)[\Sigma_1]_{\pi_1}, \dots, (s_n)[\Sigma_n]_{\pi_n}$ as follows:

$$\iota(\Gamma \vdash \Delta) := \bigwedge \Gamma \rightarrow \bigvee \Sigma_0 \vee \bigvee_{1 \leq i \leq n} \square_{s_i}(\bigvee \Sigma_i)$$

We define $\Gamma \vdash \Delta$ to be valid iff $\iota(\Gamma \vdash \Delta)$ is valid. Also, we note that $\bigwedge \emptyset = \top$ and $\bigvee \emptyset = \perp$, as usual.

A uniform presentation of our nested calculi is given in Fig. 2. We let $\text{NS}(\mathcal{V})$ denote the corresponding nested sequent calculus over a vocabulary \mathcal{V} . Our inference rules make use of the brackets ‘{’ and ‘}’ in the consequent of a nested sequent to indicate that the displayed formula(e) occur in some component. In particular, given a nested sequent $\Gamma \vdash \Delta$, where Δ is of the form $\Sigma_0, (s_1)[\Sigma_1]_{\pi_1}, \dots, (s_1)[\Sigma_i]_{\pi_i}, \dots, (s_n)[\Sigma_n]_{\pi_n}$, the notation $\Gamma \vdash \Delta\{\varphi\}_{\pi_i}$ indicates that φ occurs in Σ_i ; additionally, we use $\Gamma \vdash \Delta\{\varphi\}_{\pi_0}$ to indicate that φ occurs in Σ_0 , i.e. the label π_0 is used to reference the multiset Σ_0 serving as the root of the tree encoded by the consequent.

To make the functionality of each rule in $\text{NS}(\mathcal{V})$ precise, we explicitly state the operation performed by each rule. With the exception of the premise-free (*id*) rule, we explain for each rule how the *premise(s)* (the nested sequent(s) occurring above the horizontal inference line) are obtained from the *conclusion* (the nested sequent occurring below the horizontal inference line). This explanation is consistent with how the rules are applied (bottom-up) during proof-search as described in the following section. Also, in accordance with standard proof-theoretic terminology (Buss 1998; Takeuti 2013), we refer to the formula that is explicitly displayed in the conclusion of a rule as *principal*, and indicate the principal formulae in our explanation of the rules below to make this precise for the reader.

- (*id*) A nested sequent is *initial*, and may be used to begin a derivation, so long as some component contains both p and $\neg p$ (the principal formulae);
- (\vee) If a component Σ_i of the conclusion contains $\varphi \vee \psi$ (the principal formula), then adding φ and ψ to Σ_i yields the premise;

- (\wedge) If a component Σ_i of the conclusion contains $\varphi \wedge \psi$ (the principal formula), then adding φ to Σ_i yields the left premise and adding ψ to Σ_i yields the right premise;
- (\square_s) For any $s \in \mathcal{S}$, if a component Σ_i of the conclusion contains $\square_s \varphi$ (the principal formula), then appending the consequent Δ with $(s)[\varphi]_{\pi'}$, where π' is fresh (i.e. it does not occur in the conclusion), yields the premise;
- (n_s) For any $s \in \mathcal{S}$, we may append the consequent of the conclusion with $(s)[\emptyset]_{\pi'}$ to obtain the premise so long as π' is fresh;
- (\diamond_s^1) For any $s \in \mathcal{S}$, if a component Σ_i of the conclusion contains $\diamond_s \varphi$ (the principal formula), the consequent contains a nesting $(s')[\Sigma]_{\pi'}$, and $s' \preceq_\Gamma^* s$, then adding φ to the nesting $(s')[\Sigma]_{\pi'}$ yields the premise;
- (\diamond_s^2) For any $s \in \mathcal{S}$, if the consequent contains a nesting of the form $(s')[\diamond_s \varphi, \Sigma]_{\pi'}$ with $\diamond_s \varphi$ the principal formula, and $s' \preceq_\Gamma^* s$, then adding φ to the nesting $(s')[\diamond_s \varphi, \Sigma]_{\pi'}$ yields the premise;
- (\diamond_*) If a component Σ_i of the conclusion contains $\diamond_* \varphi$ (the principal formula), then prepending the consequent Δ with φ (i.e. adding φ to Σ_0) yields the premise.

Example 2. Below, we provide an example of a nested sequent derivation. To minimize the width of the proof, we let φ denote $\diamond_s \diamond_* \neg p \vee \square_{s'} p$.

$$\frac{\frac{\frac{s' \preceq s \vdash \varphi, \diamond_s \diamond_* \neg p, \square_{s'} p, (s')[\diamond_* \neg p, \neg p, p]}{(id)} \quad \frac{s' \preceq s \vdash \varphi, \diamond_s \diamond_* \neg p, \square_{s'} p, (s')[\diamond_* \neg p, p]}{(\diamond_*^2)} \quad \frac{s' \preceq s \vdash \varphi, \diamond_s \diamond_* \neg p, \square_{s'} p, (s')[p]}{(\diamond_s^1)}}{\frac{s' \preceq s \vdash \varphi, \diamond_s \diamond_* \neg p, \square_{s'} p}{(\square_{s'})}} \quad (\vee)}{s' \preceq s \vdash \varphi} \quad (\vee)$$

Observe that (\diamond_s^1) is applicable as $s' \preceq_\Gamma^* s$ holds due to the antecedent, and (\diamond_*^2) is applicable as $s' \preceq_\Gamma^* *$ holds by definition (see Def. 5).

We now prove that our calculi are sound (Theorem 1, building on Lemma 1), that is, that every nested sequent derivable in $\text{NS}(\mathcal{V})$ is valid. We then state our completeness theorem, which is a consequence of the work in Section 4.

Lemma 1. Let $\Gamma \vdash \Delta$ be a sequent, $\mathcal{M} = \langle \Pi, \sigma, \delta \rangle$ be a model with $\pi \in \Pi$, and $s, s' \in \mathcal{S}$. If $\mathcal{M}, \pi \models \bigwedge \Gamma$ and $s' \preceq_\Gamma^* s$, then $\mathcal{M}, \pi \models s' \preceq s$.

Proof. Assume that $\mathcal{M}, \pi \models \bigwedge \Gamma$ and $s' \preceq_\Gamma^* s$ for some $s', s \in \mathcal{S}$. There are four cases to consider:

- (1) $s = *$. The result is immediate, since $\sigma(*) = \Pi$ and therefore $\mathcal{M}, \pi \models s' \preceq *$ for every $s' \in \mathcal{S}$ by Def. 4.
- (2) $s' \preceq s \in \Gamma$. From the assumption that $\mathcal{M}, \pi \models \bigwedge \Gamma$ it follows that $\mathcal{M}, \pi \models s' \preceq s$.
- (3) $s = s'$. Then, it is trivially implied that $\mathcal{M}, \pi \models s' \preceq s$ since $\sigma(s') = \sigma(s) \subseteq \sigma(s)$ by Def. 4.
- (4) There are some $s_1, \dots, s_n \in \mathcal{S}$ such that $s' \preceq s_1 \in \Gamma$, $s_i \preceq s_{i+1} \in \Gamma$ for every $1 \leq i \leq n-1$, and $s_n \preceq s \in \Gamma$, that is, $s' \preceq_\Gamma^* s$ is obtained by transitivity on a path in \mathcal{S} . From this, together with the assumption that $\mathcal{M}, \pi \models \bigwedge \Gamma$, it directly follows that $\mathcal{M}, \pi \models s' \preceq s$ by Def. 4. \square

$$\begin{array}{c}
\frac{}{\Gamma \vdash \Delta\{p, \neg p\}_\pi} (id) \quad \frac{\Gamma \vdash \Delta\{\varphi, \psi\}_\pi}{\Gamma \vdash \Delta\{\varphi \vee \psi\}_\pi} (\vee) \quad \frac{\Gamma \vdash \Delta\{\varphi\}_\pi \quad \Gamma \vdash \Delta\{\psi\}_\pi}{\Gamma \vdash \Delta\{\varphi \wedge \psi\}_\pi} (\wedge) \\
\frac{\Gamma \vdash \Delta\{\Box_s \varphi\}_\pi, (s)[\varphi]_{\pi'}}{\Gamma \vdash \Delta\{\Box_s \varphi\}_\pi} (\Box_s)^\dagger_1 \quad \frac{\Gamma \vdash \Delta, (s)[\emptyset]_{\pi'}}{\Gamma \vdash \Delta} (n_s)^\dagger_1 \\
\frac{\Gamma \vdash \Delta\{\Diamond_s \varphi\}_\pi, (s')[\Sigma, \varphi]_{\pi'}}{\Gamma \vdash \Delta\{\Diamond_s \varphi\}_\pi, (s')[\Sigma]_{\pi'}} (\Diamond_s^1)^\dagger_2 \quad \frac{\Gamma \vdash \Delta, (s')[\Diamond_s \varphi, \varphi, \Sigma]_\pi}{\Gamma \vdash \Delta, (s')[\Diamond_s \varphi, \Sigma]_\pi} (\Diamond_s^2)^\dagger_2 \quad \frac{\Gamma \vdash \varphi, \Delta\{\Diamond_* \varphi\}_\pi}{\Gamma \vdash \Delta\{\Diamond_* \varphi\}_\pi} (\Diamond_*)
\end{array}$$

Figure 2: The nested calculus $\text{NS}(\mathcal{V})$ with $\mathcal{V} = \langle \mathcal{P}, \mathcal{S} \rangle$ a vocabulary. We note that π is permitted to be any label from $\{\pi_i \mid i \in \mathbb{N} \setminus \{0\}\}$ and that $\text{NS}(\mathcal{V})$ contains a copy of (\Box_s) , (n_s) , (\Diamond_s^1) , and (\Diamond_s^2) for each $s \in \mathcal{S}$. The side condition \dagger_1 stipulates that the rule is applicable only if the label π' is fresh and \dagger_2 stipulates that the rule is applicable only if $s' \preceq_\Gamma^* s$.

Theorem 1 (Soundness). *If $\Gamma \vdash \Delta$ is derivable in $\text{NS}(\mathcal{V})$, then $\Gamma \vdash \Delta$ is valid.*

Proof. We prove the result by induction on the number of inferences in a given derivation, and assume that Δ is of the form $\Sigma_0, (s_1)[\Sigma_1], \dots, (s_n)[\Sigma_n]$.

Base case. In the base case, our derivation consists of a single application of the (id) rule. Hence,

$$\iota(\Gamma \vdash \Delta\{p, \neg p\}_{\pi_i}) := \bigwedge \Gamma \rightarrow \bigvee \Sigma_0 \vee \bigvee_{1 \leq i \leq n} \Box_{s_i}(\bigvee \Sigma_i)$$

where $p, \neg p \in \Sigma_i$, for some $0 \leq i \leq n$. Regardless, the consequent of the implication above will be satisfied in any model \mathcal{M} , implying that the above implication is valid.

Inductive step. We make a case distinction based on the last rule applied, and show that if the conclusion of the rule is invalid, then at least one of the premises of the rule is invalid, that is to say, we show by contraposition that if the premise(s) is (are) valid, then the conclusion is valid. We only show the (\Box_s) and (\Diamond_s^1) cases as the remaining cases are argued in a similar manner.

(\Box_s) . We assume that Σ_1 is of the form $\Box_s \varphi, \Sigma'_1$ with $\Box_s \varphi$ principal; all remaining cases are similar. Furthermore, let us suppose that $\iota(\Gamma \vdash \Delta\{\Box_s \varphi\}_{\pi_1}) :=$

$$\bigwedge \Gamma \rightarrow \bigvee \Sigma_0 \vee \Box_{s_1}(\Box_s \varphi \vee \bigvee \Sigma'_1) \vee \bigvee_{2 \leq i \leq n} \Box_{s_i}(\bigvee \Sigma_i)$$

is invalid. Then there is a standpoint model $\mathcal{M} := \langle \Pi, \sigma, \delta \rangle$ with a precisification π such that $\mathcal{M}, \pi \not\models \Box_{s_1}(\Box_s \varphi \vee \bigvee \Sigma'_1)$. Hence, there exists a precisification $\pi' \in \sigma(s_1)$ such that $\mathcal{M}, \pi' \not\models \Box_s \varphi$, implying that there exists a precisification $\pi'' \in \sigma(s)$ such that $\mathcal{M}, \pi'' \not\models \varphi$. It thus follows that $\mathcal{M}, \pi \not\models \Box_s \varphi$, showing that the premise of (\Box_s) is invalid.

(\Diamond_s^1) . Suppose that Σ_1 is of the form $\Diamond_s \varphi, \Sigma'_1$ with $\Diamond_s \varphi$ principal; all remaining cases are argued in a similar fashion. Assume that $s' \preceq_\Gamma^* s$ holds and that the following is invalid:

$$\iota(\Gamma \vdash \Delta\{\Diamond_s \varphi\}_\pi, (s')[\Sigma]_{\pi'}) := \bigwedge \Gamma \rightarrow \bigvee \Sigma_0 \vee \Box_{s_1}(\Diamond_s \varphi \vee \bigvee \Sigma'_1) \vee \left(\bigvee_{2 \leq i \leq n} \Box_{s_i}(\bigvee \Sigma_i) \right) \vee \Box_{s'}(\bigvee \Sigma)$$

Therefore, there is a standpoint model $\mathcal{M} := \langle \Pi, \sigma, \delta \rangle$ with $\pi \in \Pi$ such that $\mathcal{M}, \pi \not\models \Box_{s'}(\bigvee \Sigma)$ and $\mathcal{M}, \pi \not\models \Box_{s_1}(\Diamond_s \varphi \vee \bigvee \Sigma'_1)$, and such that $\sigma(s') \subseteq \sigma(s)$, by $s' \preceq_\Gamma^* s$

and Lem. 1. This entails that there exists a precisification $\pi' \in \sigma(s')$ such that $\mathcal{M}, \pi' \not\models \bigvee \Sigma$, and that there exists a precisification $\pi_1 \in \sigma(s_1)$ such that $\mathcal{M}, \pi_1 \not\models \Diamond_s \varphi$. The latter further implies that for every precisification in $\sigma(s)$, and thus for π' (since $\pi' \in \sigma(s') \subseteq \sigma(s)$), that $\mathcal{M}, \pi' \not\models \varphi$. Thus, the premise has been shown invalid. \square

Theorem 2 (Completeness). *If $\Gamma \vdash \Delta$ is valid, then $\Gamma \vdash \Delta$ is provable in $\text{NS}(\mathcal{V})$.*

Proof. The theorem follows from the correct (Theorem 3) and terminating (Theorem 4) proof-search procedure given in the subsequent section (Section 4). \square

4 Automating Standpoint Logic via Proof-Search

We now employ our nested calculi in an algorithm that decides the validity of formulae for propositional standpoint logics. In particular, we design a proof-search algorithm (see Alg. 1 below) which takes a vocabulary \mathcal{V} as a parameter and bottom-up applies rules from $\text{NS}(\mathcal{V})$ in attempt to construct a proof of a given input sequent $\Gamma \vdash \varphi$. We may make the simplifying assumption that our proof-search algorithm only receives inputs of the form $\Gamma \vdash \varphi$ as any nested sequent $\Gamma \vdash \Delta$ with $\Delta = \Sigma_0, (s_1)[\Sigma_1], \dots, (s_n)[\Sigma_n]$ is valid iff $\iota(\Gamma \vdash \Delta)$ is valid iff $\Gamma \vdash \varphi$ with $\varphi = \bigvee \Sigma_0 \vee \Box_{s_1}(\bigvee \Sigma_1) \vee \dots \vee \Box_{s_n}(\bigvee \Sigma_n)$ is valid.

To decrease the complexity of proof-search and obtain (worst-case) complexity-optimality, we introduce a new technique in the context of nested sequents which we refer to as *coloring*. In essence, given the input $\Gamma \vdash \varphi$, the first step of proof-search guesses a *proper coloring* of the formula φ , that is, it labels the formula's subformulae with either an *active* label \circ or an *inactive* label \bullet in a particular manner. Recall that, due to the (\wedge) rule, a proof in $\text{NS}(\mathcal{V})$ has the structure of a binary tree, thus giving rise to the possibility that proof-search is exponential; therefore, our proof-search algorithm uses the aforementioned labels to only generate a single path in this binary tree relative to each coloring, which yields a worst-case complexity-optimal proof-search procedure in CoNP (for the validity problem of $\mathbb{S}(\mathcal{V})$).

Definition 7 (Coloring). *We define a colored formula to be a formula generated via the following grammar in BNF:*

$$\tilde{\varphi} ::= p^* \mid \neg p^* \mid (\tilde{\varphi} \vee \tilde{\varphi})^* \mid (\tilde{\varphi} \wedge \tilde{\varphi})^* \mid (\Diamond_s \tilde{\varphi})^* \mid (\Box_s \tilde{\varphi})^*$$

Algorithm 1: Prove_V

Input: A Nested Sequent: $\Gamma \vdash \varphi$ **Output:** A Boolean: True, False

- 1 Choose a proper coloring $\tilde{\varphi}$ of φ ;
 - 2 **return** ProofSearch_V($\Gamma \vdash \tilde{\varphi}$);
-

with $*$ \in $\{\circ, \bullet\}$. For any colored formula $\tilde{\varphi}$, we let φ be the formula in \mathcal{L}_V obtained by removing all labels \circ and \bullet from $\tilde{\varphi}$. A formula $\tilde{\varphi}$ is properly colored iff $\tilde{\varphi} = f_\circ(\varphi)$, where the non-deterministic coloring function f_\circ and f_\bullet are defined accordingly with $*$ \in $\{\circ, \bullet\}$:

- $f_*(p) = p^*$
- $f_*(\neg p) = \neg p^*$
- $f_*(\varphi \vee \psi) = (f_*(\varphi) \vee f_*(\psi))^*$
- $f_\circ(\varphi \wedge \psi) \in \{(f_\circ(\varphi) \wedge f_\bullet(\psi))^\circ, (f_\bullet(\varphi) \wedge f_\circ(\psi))^\circ\}$
- $f_\bullet(\varphi \wedge \psi) = (f_\bullet(\varphi) \wedge f_\bullet(\psi))^\bullet$
- $f_*(\diamond_s \varphi) = (\diamond_s f_*(\varphi))^*$
- $f_*(\square_s \varphi) = (\square_s f_*(\varphi))^*$

We define $\text{pcs}(\varphi)$ to be the set of all proper colorings of φ , and define a colored nested sequent to be a nested sequent that uses colored formulae as opposed to formulae from \mathcal{L}_V .

We now stipulate our *saturation conditions*. When such conditions are unsatisfied during proof-search it signals that certain inference rules still need to be applied bottom-up. Alternatively, once all such conditions are satisfied this signals that proof-search ought to terminate.

Definition 8 (Saturation Conditions). A colored nested sequent $\Gamma \vdash \Sigma_0, (s_1)[\Sigma_1]_{\pi_1}, \dots, (s_n)[\Sigma_n]_{\pi_n}$ is saturated iff for every $i \in \{0, \dots, n\}$ it satisfies the following conditions:

- id* If $p^\circ \in \Sigma_i$, then $\neg p^\circ \notin \Sigma_i$;
- \vee if $(\tilde{\varphi} \vee \tilde{\psi})^\circ \in \Sigma_i$, then $\tilde{\varphi}^\circ, \tilde{\psi}^\circ \in \Sigma_i$;
- \wedge if $(\tilde{\varphi} \wedge \tilde{\psi})^\circ \in \Sigma_i$, then either $\tilde{\varphi}^\circ \in \Sigma_i$ or $\tilde{\psi}^\circ \in \Sigma_i$;
- \diamond_s if $(\diamond_s \tilde{\varphi})^\circ \in \Sigma_i$ and $s' \preceq_\Gamma^* s$, then for each $j \in \{1, \dots, n\}$ such that $s_j = s'$, $\tilde{\varphi}^\circ \in \Sigma_j$;
- \diamond_* if $(\diamond_* \tilde{\varphi})^\circ \in \Sigma_i$, then $\tilde{\varphi}^\circ \in \Sigma_0$;
- \square_s if $(\square_s \tilde{\varphi})^\circ \in \Sigma_i$, then for some $j \in \{1, \dots, n\}$, $s_j = s$, and $\tilde{\varphi}^\circ \in \Sigma_j$;
- n_s for each $s \in \mathcal{S}$, there exists a $j \in \{1, \dots, n\}$ such that $s_j = s$.

Let us comment on the functionality of our (non-deterministic) proof-search algorithm Prove_V (Alg. 1), which takes ProofSearch_V (Alg. 2) as a subroutine. (NB. Alg. 2 is split between this page and the next due to its length.) As mentioned above, given an input $\Gamma \vdash \varphi$, the algorithm Prove_V guesses a proper coloring $\tilde{\varphi}$ of φ , and then returns the value of ProofSearch_V($\Gamma \vdash \tilde{\varphi}$). We note that ProofSearch_V applies the rules from NS(\mathcal{V}) in a bottom-up manner (each corresponding to a recursive call of the algorithm with the exception of (*id*)). The application of each rule is as follows: (*id*) corresponds to lines 1–3, (\vee) to lines 7–11, (\wedge) to lines 12–16 and 17–21, that respectively yields the left and right premises of (\wedge), (\diamond_s^1) and (\diamond_s^2) to lines

Algorithm 2: ProofSearch_V (Part I)

Input: A Colored Nested Sequent: $\Phi := \Gamma \vdash \Sigma_0, (s_1)[\Sigma_1]_{\pi_1}, \dots, (s_n)[\Sigma_n]_{\pi_n}$ **Output:** A Boolean: True, False

- 1 **if** for some $0 \leq i \leq n$, $p^\circ, \neg p^\circ \in \Sigma_i$ **then**
 - 2 | **return** True;
 - 3 **end**
 - 4 **if** Σ is saturated **then**
 - 5 | **return** False;
 - 6 **end**
 - 7 **if** for some $0 \leq i \leq n$, $(\tilde{\varphi} \vee \tilde{\psi})^\circ \in \Sigma_i$, but $\tilde{\varphi}, \tilde{\psi} \notin \Sigma_i$ **then**
 - 8 | Let $\Sigma'_i := \Sigma_i, \tilde{\varphi}, \tilde{\psi}$;
 - 9 | Let $\Phi' := \Gamma \vdash \Sigma_0, \dots, (s_i)[\Sigma'_i]_{\pi_i}, \dots, (s_n)[\Sigma_n]_{\pi_n}$;
// Replace Σ_i by Σ'_i to obtain Φ' .
 - 10 | **return** Prove_V(Φ');
 - 11 **end**
 - 12 **if** for some $0 \leq i \leq n$, $(\tilde{\varphi}^\circ \wedge \tilde{\psi}^\circ)^\circ \in \Sigma_i$, but $\tilde{\varphi}^\circ \notin \Sigma_i$ **then**
 - 13 | Let $\Sigma'_i := \Sigma_i, \tilde{\varphi}^\circ$;
 - 14 | Let $\Phi' := \Gamma \vdash \Sigma_0, \dots, (s_i)[\Sigma'_i]_{\pi_i}, \dots, (s_n)[\Sigma_n]_{\pi_n}$;
// Replace Σ_i by Σ'_i to obtain Φ' .
 - 15 | **return** Prove_V(Φ');
 - 16 **end**
 - 17 **if** for some $0 \leq i \leq n$, $(\tilde{\varphi}^\bullet \wedge \tilde{\psi}^\circ)^\circ \in \Sigma_i$, but $\tilde{\psi}^\circ \notin \Sigma_i$ **then**
 - 18 | Let $\Sigma'_i := \Sigma_i, \tilde{\psi}^\circ$;
 - 19 | Let $\Phi' := \Gamma \vdash \Sigma_0, \dots, (s_i)[\Sigma'_i]_{\pi_i}, \dots, (s_n)[\Sigma_n]_{\pi_n}$;
// Replace Σ_i by Σ'_i to obtain Φ' .
 - 20 | **return** Prove_V(Φ');
 - 21 **end**
-

22–25, (\diamond_*) to lines 26–29, (\square_s) to lines 30–33, and (n_s) to lines 34–37.

Moreover, ProofSearch_V contrasts with typical proof-search algorithms in that it utilizes the active and inactive labels \circ and \bullet in $\tilde{\varphi}$ to guide its computation and only constructs a *single thread* of the proof.¹ In other words, if a nested sequent $\Gamma \vdash \varphi$ is derivable in NS(\mathcal{V}), then the sequent has a proof in NS(\mathcal{V}) such that ProofSearch_V generates each thread of the proof relative to each proper coloring of φ ; as argued in the lemma below, all such threads may be ‘zipped’ together to reconstruct a full proof of $\Gamma \vdash \varphi$ in NS(\mathcal{V}). In this way, our proof-search algorithm may be used to construct certificates witnessing the validity (by means of a proof in NS(\mathcal{V})) or invalidity (by means of a countermodel) of any input $\Gamma \vdash \varphi$ (see Thm. 3 below for details).

Lemma 2. Let $\Gamma \vdash \varphi$ be a sequent and $\text{pcs}(\varphi)$ the (finite) set of proper colorings of φ . If $\text{ProofSearch}_V(\Gamma \vdash \tilde{\varphi}) = \text{True}$ for all $\tilde{\varphi} \in \text{pcs}(\varphi)$, then there is a proof of $\Gamma \vdash \varphi$ in NS(\mathcal{V}).

Proof. Assume that $\text{ProofSearch}_V(\Gamma \vdash \tilde{\varphi}) = \text{True}$ for ev-

¹A *thread* in a proof is defined in the usual fashion as a maximal path of sequents from the conclusion of the proof to an initial sequent (cf. (Takeuti 2013, p. 14)).

ery $\tilde{\varphi} \in \text{pcs}(\varphi)$, and the following thread of colored nested sequents is generated during its execution:

$$T(\tilde{\varphi}) := \Gamma \vdash \Delta_0, \dots, \Gamma \vdash \Delta_h$$

such that $\Delta_0 = \tilde{\varphi}$ and $\Gamma \vdash \Delta_h$ is an instance of *(id)* (by lines 1–3). Let \mathcal{T} be the set of all such threads.

For a thread $\Gamma \vdash \Delta_0, \dots, \Gamma \vdash \Delta_h \in \mathcal{T}$ and some $0 \leq k \leq h$, the colored nested sequent $\Gamma \vdash \Delta_k$ is:

- *left conjunctive* iff $\Gamma \vdash \Delta_{k+1}$ is obtained from $\Gamma \vdash \Delta_k$ by applying (\wedge) yielding the left premise (lines 12–16), and
- *right conjunctive* iff $\Gamma \vdash \Delta_{k+1}$ is obtained from $\Gamma \vdash \Delta_k$ by applying (\wedge) yielding the right premise (lines 17–21).

A colored nested sequent is *conjunctive* iff it is left or right conjunctive. We now explain how our threads may be transformed into a proof of $\Gamma \vdash \varphi$ in $\text{NS}(\mathcal{V})$.

We assume w.l.o.g. that the initial segments of all threads up to and including the first conjunctive sequent $\Gamma \vdash \Delta_k$ are identical (i.e. we assume that the subroutine $\text{ProofSearch}_{\mathcal{V}}$ executes in a deterministic fashion). Hence, we may form the ‘pseudo-derivation’ shown below left, by making use of the first k sequents of any given thread, where the rules $(r_1), \dots, (r_{k-1})$ are determined on the basis of which lines of $\text{ProofSearch}_{\mathcal{V}}$ were executed.

$$\left| \frac{\frac{\Gamma \vdash \Delta_k}{\vdots} (r_{k-1})}{\Gamma \vdash \varphi} (r_1) \quad \frac{\frac{\Gamma \vdash \Delta_{k+1}^l \quad \Gamma \vdash \Delta_{k+1}^r}{\Gamma \vdash \Delta_k} (\wedge)}{\vdots} (r_1) \right|$$

Let us define $T^k(\tilde{\varphi}) := \Gamma \vdash \Delta_{k+1}, \dots, \Gamma \vdash \Delta_h$ to be the tail of a thread $T(\tilde{\varphi})$ starting from $k+1$. Since the k^{th} colored nested sequent of every thread is conjunctive, we may generate two sets of threads \mathcal{T}_L and \mathcal{T}_R from \mathcal{T} :

$$\mathcal{T}_L := \{T^k(\tilde{\varphi}) \mid \Gamma \vdash \Delta_k \in T(\tilde{\varphi}) \text{ is left conjunctive}\}$$

$$\mathcal{T}_R := \{T^k(\tilde{\varphi}) \mid \Gamma \vdash \Delta_k \in T(\tilde{\varphi}) \text{ is right conjunctive}\}$$

We now extend the ‘pseudo-derivation’ shown above left, with a bottom-up application of (\wedge) to obtain the ‘pseudo-derivation’ shown above right, where $\Gamma \vdash \Delta_{k+1}^l$ and $\Gamma \vdash \Delta_{k+1}^r$ are the elements of \mathcal{T}_L and \mathcal{T}_R occurring directly after $\Gamma \vdash \Delta_k$ in each thread, respectively. By successively repeating the above described process over \mathcal{T}_L and \mathcal{T}_R , a proof in $\text{NS}(\mathcal{V})$ will eventually be built above $\Gamma \vdash \Delta_{k+1}^l$ and $\Gamma \vdash \Delta_{k+1}^r$, giving a proof of $\Gamma \vdash \varphi$ in $\text{NS}(\mathcal{V})$. \square

Example 3. We provide an example of the proof construction process of $\emptyset \vdash \varphi$ with $\varphi := \Box_s(p \vee (\neg p \wedge \neg p)) \wedge \Diamond_s(q \vee \neg q)$ to illustrate the procedure given in Lem. 2 above. First, observe that φ has three proper colorings, given below:

$$\tilde{\varphi}^1 := ((\Box_s(p^\circ \vee (\neg p^\bullet \wedge \neg p^\bullet)^\circ)^\circ) \wedge (\Diamond_s(q^\bullet \vee \neg q^\bullet)^\bullet)^\bullet)^\bullet$$

$$\tilde{\varphi}^2 := ((\Box_s(p^\circ \vee (\neg p^\bullet \wedge \neg p^\circ)^\circ)^\circ) \wedge (\Diamond_s(q^\bullet \vee \neg q^\bullet)^\bullet)^\bullet)^\circ$$

$$\tilde{\varphi}^3 := ((\Box_s(p^\bullet \vee (\neg p^\bullet \wedge \neg p^\bullet)^\bullet)^\bullet) \wedge (\Diamond_s(q^\circ \vee \neg q^\circ)^\circ)^\circ)^\circ$$

Each proper coloring $\tilde{\varphi}^i$ with $i \in \{1, 2, 3\}$ gives rise to a corresponding thread $T(i)$ (shown at the top of Fig. 3) when $\text{ProofSearch}_{\mathcal{V}}(\emptyset \vdash \tilde{\varphi}^i)$ is run. Note that in the figure we have omitted the active and inactive labels, as well as labels

Algorithm 2: ProofSearch $_{\mathcal{V}}$ (Part II)

```

22 if for some  $0 \leq i \leq n$ ,  $(\Diamond_s \tilde{\varphi})^\circ \in \Sigma_i$  and  $s' \preceq_\Gamma^* s$ ,
    and for some  $1 \leq j \leq n$ ,  $s_j = s'$ , but  $\tilde{\varphi} \notin \Sigma_j$  then
23   Let  $\Phi' := \Gamma \vdash$ 
      $\Sigma_0, \dots, (s_j)[\Sigma_j]_{\pi_j}, \dots, (s_n)[\Sigma_n]_{\pi_n}$ ;
     // Add  $\tilde{\varphi}$  to the  $j^{\text{th}}$  nesting to
     obtain  $\Phi'$ .
24   return  $\text{Prove}_{\mathcal{V}}(\Phi')$ ;
25 end
26 if for some  $0 \leq i \leq n$ ,  $(\Diamond_* \tilde{\varphi})^\circ \in \Sigma_i$ , but  $\tilde{\varphi} \notin \Sigma_0$ 
    then
27   Let
      $\Phi' := \Gamma \vdash \tilde{\varphi}, \Sigma_0, (s_1)[\Sigma_1]_{\pi_1}, \dots, (s_n)[\Sigma_n]_{\pi_n}$ ;
     // Add  $\tilde{\varphi}$  to the  $0^{\text{th}}$  component to
     obtain  $\Phi'$ .
28   return  $\text{Prove}_{\mathcal{V}}(\Phi')$ ;
29 end
30 if for some  $0 \leq i \leq n$ ,  $(\Box_s \tilde{\varphi})^\circ \in \Sigma_i$ , but for each
     $1 \leq j \leq n$  such that  $s_j = s$ ,  $\tilde{\varphi} \notin \Sigma_j$  then
31   Let  $\Sigma' := \Gamma \vdash \Sigma_0, \dots, (s_n)[\Sigma_n]_{\pi_n}, (s)[\tilde{\varphi}]_{\pi_{n+1}}$ ;
     // Append  $(s)[\tilde{\varphi}]_{\pi_{n+1}}$  to obtain  $\Phi'$ 
     with  $\pi_{n+1}$  fresh.
32   return  $\text{Prove}_{\mathcal{V}}(\Phi')$ ;
33 end
34 if for some  $s \in \mathcal{S}$  there does not exist a  $1 \leq j \leq n$ 
    such that  $s_j = s$  then
35   Let  $\Phi' := \Gamma \vdash \Sigma_0, \dots, (s_n)[\Sigma_n]_{\pi_n}, (s)[\emptyset]_{\pi_{n+1}}$ ;
     // Append  $(s)[\emptyset]_{\pi_{n+1}}$  to obtain  $\Phi'$ 
     with  $\pi_{n+1}$  fresh.
36   return  $\text{Prove}_{\mathcal{V}}(\Phi')$ ;
37 end

```

of the form π_i , from each thread $T(i)$ to improve readability, and each application of (\wedge) is emphasised with a dashed inference line. By making use of the proof construction process described in Lem. 2, the three threads $T(1)$, $T(2)$, and $T(3)$ can be fused together to generate a proof of $\emptyset \vdash \varphi$ shown at the bottom of Fig. 3.

Theorem 3 (Correctness). *Let $\Gamma \vdash \varphi$ be a sequent.*

1. *If $\text{ProofSearch}_{\mathcal{V}}(\Gamma \vdash \tilde{\varphi}) = \text{True}$ for all proper colorings of φ , then a proof in $\text{NS}(\mathcal{V})$ may be constructed witnessing that $\Gamma \vdash \varphi$ is valid.*
2. *If $\text{ProofSearch}_{\mathcal{V}}(\Gamma \vdash \tilde{\varphi}) = \text{False}$ for some proper coloring of φ , then a counter-model may be constructed witnessing that $\Gamma \vdash \varphi$ is invalid.*

Proof. The first claim follows by Lem. 2 and the soundness of each nested calculus (see Thm. 1); therefore, we focus on the second claim. Suppose that $\text{ProofSearch}_{\mathcal{V}}(\Gamma \vdash \tilde{\varphi}) = \text{False}$ for some chosen proper coloring of φ . Then, $\text{ProofSearch}_{\mathcal{V}}$ generates a saturated nested sequent $\Gamma \vdash \Delta$ with Δ of the form $\Sigma_0, (s_1)[\Sigma_1]_{\pi_1}, \dots, (s_n)[\Sigma_n]_{\pi_n}$. We will use $\Gamma \vdash \Delta$ to construct a counter-model for $\Gamma \vdash \varphi$, thus proving the second claim. Let us define $\mathcal{M} := \langle \Pi, \sigma, \delta \rangle$ as follows:

We now show that $\text{ProofSearch}_{\mathcal{V}}$ (and hence $\text{Prove}_{\mathcal{V}}$) terminates after at most polynomially many (bottom-up) rule applications in the *size* of the input sequent. Supposing that $\Phi := \Gamma \vdash \varphi$ is our input, the size of Φ is defined to be $|\Phi| := |\mathcal{S}| + |\varphi|$. In other words, the size of a sequent Φ is the sum of the cardinality of the set \mathcal{S} of standpoints and the size of φ . The reason the size of a sequent incorporates a measure on the set \mathcal{S} of standpoints from the associated vocabulary \mathcal{V} as opposed to a measure on the set Γ of sharpening statements is that Γ only plays a role in bottom-up applications of (\diamond_s) , which is bounded in part by the cardinality of \mathcal{S} and in part by the number of \square_s modalities occurring in φ , as explained in the proof of Thm. 4 below. That is to say, it is primarily \mathcal{S} and φ that play a fundamental role in the number of inferences (each corresponding to a recursive call) in a thread generated by $\text{Prove}_{\mathcal{V}}$.

Theorem 4 (Termination). *Let $\Phi := \Gamma \vdash \varphi$ be a sequent. Then, the number of recursive calls throughout the computation of $\text{ProofSearch}_{\mathcal{V}}(\Gamma \vdash \varphi)$, and thus $\text{Prove}_{\mathcal{V}}(\Gamma \vdash \varphi)$, is bounded above by a polynomial*

$$p(|\Phi|) = \mathcal{O}(|\Phi|^2).$$

Proof. Let $\Phi := \Gamma \vdash \varphi$ be a nested sequent, and N_{\oplus} be the number of occurrences of the connectives $\{\vee, \wedge\} \cup \{\diamond_s \mid s \in \mathcal{S}\}$ in φ . By the saturation conditions (Def. 8), we know that for each $s \in \mathcal{S}$, the rule (\square_s) will be applied bottom-up at most one time for each occurrence of \square_s in φ , which are bounded by $|\varphi|$. Also, (n_s) will be applied at most one time for each $s \in \mathcal{S}$. Since (\square_s) and (n_s) are the only rules that introduce nestings, then the number of components (i.e. the nestings plus the root) throughout the course of proof-search is bounded by:

$$K := 1 + |\mathcal{S}| + |\varphi|$$

For each occurrence of \vee , \wedge , and \diamond_s in φ (with $s \in \mathcal{S}$), we know by the saturation conditions that (\vee) , (\wedge) , (\diamond_s) , and (\diamond_s) can be applied a maximum number of K times during proof search. Then, since $N_{\vee} + N_{\wedge} + \sum_{s \in \mathcal{S}} N_{\diamond_s} \leq |\varphi|$, the number of recursive calls (i.e. bottom-up applications of rules) during proof-search is bounded by $N := |\varphi| \cdot K$. Finally, $|\mathcal{S}|, |\varphi| \leq |\Phi|$ holds trivially, implying:

$$N \leq |\Phi| \cdot (1 + |\Phi| + |\Phi|)$$

Therefore, it follows that a polynomial $p(|\Phi|) = \mathcal{O}(|\Phi|^2)$ bounds the number of recursive calls of $\text{Prove}_{\mathcal{V}}(\Phi)$. \square

Corollary 1. *Let \mathcal{V} be a vocabulary. Then,*

1. $\mathbb{S}(\mathcal{V})$ is decidable;
2. $\mathbb{S}(\mathcal{V})$ has the finite model property;
3. $\text{Prove}_{\mathcal{V}}$ is worst-case complexity-optimal, deciding the validity problem for $\mathbb{S}(\mathcal{V})$ in CoNP;
4. The validity problem for $\mathbb{S}(\mathcal{V})$ is CoNP-complete.

Proof. Statements 1 and 2 follow from the fact that $\text{Prove}_{\mathcal{V}}$ is a correct (Thm. 3) and terminating (Thm. 4) decision procedure for $\mathbb{S}(\mathcal{V})$ that, in particular, returns a finite counter-model when the input is invalid.

To show statement 3, i.e. that $\text{Prove}_{\mathcal{V}}$ is in CoNP and is worst-case complexity-optimal, we first note that

$\text{Prove}_{\mathcal{V}}$ is a non-deterministic algorithm that takes a sequent $\Phi := \Gamma \vdash \varphi$ as input, guesses a proper coloring of φ , and then constructs a thread that is polynomial in the size of its input. To see this, note that by Thm. 4 the number of (bottom-up) rule applications (i.e. the length of the thread) is bounded by a polynomial $p(|\Phi|) = \mathcal{O}(|\Phi|^2)$. Moreover, since any sequent generated during proof-search can have at most $K \leq 1 + |\mathcal{S}| + |\varphi|$ many components (as stated in the proof of Thm. 4 above), each of which can only be inhabited by at most $|\text{sufo}(\varphi)| = |\varphi|$ many formulae, it follows that each nested sequent in the thread has a size bounded by $\mathcal{O}(|\Phi|^2)$ since $|\mathcal{S}|, |\varphi| \leq |\Phi|$. Therefore, taking the functionality of $\text{Prove}_{\mathcal{V}}$ into account, one can see that if $\text{Prove}_{\mathcal{V}}(\Phi) = \text{False}$, then the corresponding thread is generated in polynomial time and its size is bounded above by a polynomial $q(|\Phi|) = \mathcal{O}(|\Phi|^4)$. Additionally, observe that $\text{Prove}_{\mathcal{V}}$ is worst-case complexity-optimal as the validity problem for classical propositional logic is CoNP-complete, and can be solved by $\text{Prove}_{\mathcal{V}}$ as (id) , (\vee) , and (\wedge) form a sound and complete proof system for propositional logic (cf. (Lyon 2021a, App. B)). Last, statement 4 is an immediate consequence of statement 3. \square

5 Conclusion and Future Work

In this paper, we introduced and employed nested sequent systems to automate reasoning with propositional standpoint logics. To obtain worst-case complexity-optimal proof-search, we presented a novel proof-search technique, referred to as *coloring*, whereby the subformulae of an input formula are non-deterministically colored with (in)active labels, yielding partial proofs (i.e. *threads*) of the input. By means of our technique, we designed a non-deterministic proof-search algorithm deciding the validity of standpoint implications in CoNP, showing how certain threads could be transformed into a counter-model for an invalid input, and how all threads could be transformed into a proof for a valid input. The attainment of these “certificates” from proof-search serve as explanations for the (in)validity of standpoint formulae, thus motivating our proof-theoretic approach.

For future work, we aim to extend our nested systems and proof-search algorithm to cover (i) first-order standpoint logics that (ii) incorporate complex standpoints, which have interesting applications in knowledge integration scenarios. Regarding point (i), placing standpoint logic on a first-order base, rather than a propositional base, increases the applicability of standpoint logic along with its expressivity to better match that of contemporary knowledge representation languages. Our focus in this area is to provide results that can then be extrapolated to widely used decidable fragments of FOL. Regarding point (ii), we note that the set-theoretic interpretation of standpoints permits the definition of complex standpoints built atop atomic ones; e.g. union $s_1 \cup s_2$ (integrating knowledge from multiple perspectives), intersection $s_1 \cap s_2$ (expressing the knowledge jointly shared between multiple perspectives), and difference $s_1 \setminus s_2$ (yielding the sharpening of s_1 by ignoring all precisifications of s_2). Beyond providing nested systems for more expressive formulations of standpoint logic, we also aim to write and evaluate theorem provers based on our nested calculi.

References

- Belnap, N. D. 1982. Display logic. *Journal of philosophical logic* 11(4):375–417.
- Brünnler, K. 2009. Deep sequent systems for modal logic. *Arch. Math. Log.* 48(6):551–577.
- Bull, R. A. 1992. Cut elimination for propositional dynamic logic without *. *Z. Math. Logik Grundlag. Math.* 38(2):85–100.
- Buss, S. R. 1998. An introduction to proof theory. *Handbook of proof theory* 137:1–78.
- Dyckhoff, R. 1992. Contraction-free sequent calculi for intuitionistic logic. *The Journal of Symbolic Logic* 57(3):795–807.
- Fine, K. 1975. Vagueness, truth and logic. *Synthese* 30(3-4):265–300.
- Fitting, M., and Kuznets, R. 2015. Modal interpolation via nested sequents. *Annals of pure and applied logic* 166(3):274–305.
- Gentzen, G. 1935a. Untersuchungen über das logische Schließen. i. *Mathematische Zeitschrift* 39(1):176–210.
- Gentzen, G. 1935b. Untersuchungen über das logische Schließen. ii. *Mathematische Zeitschrift* 39(1):405–431.
- Gómez Álvarez, L., and Bennett, B. 2018. Dealing with conceptual indeterminacy: A framework based on supervaluation semantics. In *Joint Proceedings of MedRACER and WOMoCoE*. CEUR.
- Gómez Álvarez, L., and Rudolph, S. 2021. Standpoint logic: Multi-perspective knowledge representation. In Neuhaus, F., and Brodaric, B., eds., *Proceedings of the 12th International Conference on Formal Ontology in Information Systems*, volume 344 of *FAIA*, 3–17. IOS Press.
- Kashima, R. 1994. Cut-free sequent calculi for some tense logics. *Studia Logica* 53(1):119–135.
- Keefe, R., and Smith, P. 1997. *Theories of vagueness*. The MIT press.
- Lyon, T., and van Berkel, K. 2019. Automating agential reasoning: Proof-calculi and syntactic decidability for STIT logics. In Baldoni, M.; Dastani, M.; Liao, B.; Sakurai, Y.; and Zalila Wenkstern, R., eds., *Proceedings of the 22nd International Conference on Principles and Practice of Multi-Agent Systems*, volume 11873, 202–218. Springer.
- Lyon, T.; Tiu, A.; Goré, R.; and Clouston, R. 2020. Syntactic interpolation for tense logics and bi-intuitionistic logic via nested sequents. In Fernández, M., and Muscholl, A., eds., *Proceedings of the 28th Annual Conference on Computer Science Logic*, volume 152 of *LIPICs*, 28:1–28:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Lyon, T. 2020. On the correspondence between nested calculi and semantic systems for intuitionistic logics. *Journal of Logic and Computation* 31(1):213–265.
- Lyon, T. 2021a. *Refining Labelled Systems for Modal and Constructive Logics with Applications*. Ph.D. Dissertation, Technische Universität Wien.
- Lyon, T. S. 2021b. Nested sequents for intuitionistic modal logics via structural refinement. In Das, A., and Negri, S., eds., *Automated Reasoning with Analytic Tableaux and Related Methods*, 409–427. Cham: Springer.
- Pesquita, C.; Faria, D.; Santos, E.; and Couto, F. M. 2013. To repair or not to repair: Reconciling correctness and coherence in ontology reference alignments. In *Proceedings of the 8th International Conference on Ontology Matching*, volume 1111, 13–24. CEUR.
- Poggioli, F. 2009. The method of tree-hypersequents for modal propositional logic. In Makinson, D.; Malinowski, J.; and Wansing, H., eds., *Towards Mathematical Philosophy*, volume 28 of *Trends in logic*. Springer. 31–51.
- Simpson, A. K. 1994. *The proof theory and semantics of intuitionistic modal logic*. Ph.D. Dissertation, University of Edinburgh. College of Science and Engineering. School of Informatics.
- Slaney, J. K. 1997. Minlog: A minimal logic theorem prover. In McCune, W., ed., *Proceedings of the 14th International Conference on Automated Deduction*, volume 1249 of *Lecture Notes in Computer Science*, 268–271. Springer.
- Takeuti, G. 2013. *Proof theory*, volume 81. Courier Corporation.
- Tiu, A.; Ivanovski, E.; and Goré, R. 2012. Grammar logics in nested sequent calculus: Proof theory and decision procedures. In Bolander, T.; Brünnler, T.; Ghilardi, S.; and Moss, L. S., eds., *Advances in Modal Logic 9, papers from the ninth conference on Advances in Modal Logic*, 516–537. College Publications.
- Wansing, H. 2002. Sequent systems for modal logics. In Gabbay, D. M., and Guenther, F., eds., *Handbook of Philosophical Logic: Volume 8*. Springer. 61–145.