

Proof Theory and Decision Procedures for Deontic STIT Logics

Tim S. Lyon

*Technische Universität Dresden,
Nöthnitzerstraße 46, 01187 Dresden, Germany*

TIMOTHY_STEPHEN.LYON@TU-DRESDEN.DE

Kees van Berkel

*TU Wien, Institute for Logic and Computation
Favoritenstraße 9, 1040, Vienna, Austria*

KEES@LOGIC.AT

Abstract

This paper provides a set of cut-free complete sequent-style calculi for deontic STIT (‘See To It That’) logics used to formally reason about choice-making, obligations, and norms in a multi-agent setting. We leverage these calculi to write a proof-search algorithm deciding deontic, multi-agent STIT logics with (un)limited choice and introduce a loop-checking mechanism to ensure the termination of the algorithm. Despite the acknowledged potential for deontic reasoning in the context of autonomous, multi-agent scenarios, this work is the first to provide a syntactic decision procedure for this class of logics. Our proof-search procedure is designed to provide verifiable witnesses/certificates of the (in)validity of formulae, which permits an analysis of the (non)theoremhood of formulae and act as explanations thereof. We show how the proof system and decision algorithm can be used to automate normative reasoning tasks such as *duty checking* (viz. determining an agent’s obligations relative to a given knowledge base), *compliance checking* (viz. determining if a choice, considered by an agent as potential conduct, complies with the given knowledge base), and *joint fulfillment checking* (viz. determining whether under a specified factual context an agent can jointly fulfill all their duties).

1. Introduction

The logic of ‘Seeing To It That’, commonly referred to as STIT, is a formal framework developed by Belnap, Perloff, and Xu (2001) for the analysis of agential choice in multi-agent settings. Next to its philosophical merits, the formalism has acquired a central position in the investigation of agency within the domain of artificial intelligence (AI). The framework has been adopted and applied to temporal (Ciuni & Lorini, 2017) and epistemic (Broersen, 2011) reasoning, as well as to the investigation of legal concepts in AI (Lorini & Sartor, 2015). Most notably, STIT has been employed for the analysis of *normative choice*, viz. multi-agent choice on the basis of obligations, prohibitions, and permissions. Logics dealing with such concepts are commonly referred to as *deontic*, i.e. relating to duty. In fact, the potential of applying STIT to the deontic setting was recognized from the outset (Bartha, 1993; Horty & Belnap, 1995). More recently, the potential of deontic STIT for reasoning in the context of autonomous vehicles has been identified and explored (Arkoudas, Bringsjord, & Bello, 2005; Shea-Blymyer & Abbas, 2021).

The work by Horty (2001) set the stage for the investigation of deontic STIT as a study in its own right. Horty’s deontic STIT continues to receive considerable attention in the

literature, and it is often recognized that the framework could serve as a suitable basis for developing automated reasoning tools within the context of agent-based, normative reasoning. Despite an increasing demand for automated normative reasoning tools for artificial agents (Arkoudas et al., 2005; van Berkel & Lyon, 2019a; Shea-Blymyer & Abbas, 2020), such tools have yet to be satisfactorily developed for modal logics of agency. Although Arkoudas et al. (2005) obtained preliminary results on automated reasoning for deontic STIT, three essential components were left unaddressed: correct termination of the implementation, constructive generation of explicit proofs, and explainability in the case of failed proof-search. In response, this article provides reasoning tools for deontic STIT logics, which address all three components. In particular, we focus on Horty’s (2001) deontic STIT logic for *dominance ought*, which deals with a utilitarian evaluation of agent-dependent obligations. The well-grounded philosophical motivation of dominance ought made the framework the most extensively studied deontic STIT logic with a wide variety of applications.

Murakami (2004) provided a sound and complete axiomatization for the logic of dominance ought and van Berkel and Lyon (2019b) showed how utilitarian STIT models—those initially provided for dominance ought by Horty (2001)—can be transformed into equivalent relational models. The first upshot of the latter result is that we can safely deal with deontic STIT in the more modular relational setting, omitting reference to utilities. The second upshot is that the alternative semantics facilitates the development of an effective *proof-theory* for deontic STIT, i.e. the formulation of a set of rules (called a *calculus* or *proof system*) capable of deriving all and only the theorems of a particular deontic STIT logic, and which may be leveraged for applications or to study (meta-)logical properties. Preliminary works in the proof-theory of (deontic) STIT logics, provided in (Wansing, 2006; van Berkel & Lyon, 2019a; Lyon & van Berkel, 2019; Negri & Pavlović, 2020; van Berkel & Lyon, 2021; Lyon, 2021), show the potential of proof theory in relation to automated reasoning for deontic STIT logics, which is the primary concern of this article.

Our work is of a twofold nature, reflected by the following two aims. The first aim is our central aim and constitutes the main contribution of this article:

(AIM I) Supply automated reasoning procedures for the deontic STIT logic of dominance ought (and variations thereof) which produce certificates of the (in)validity of input formulae.

As mentioned above, the automation of STIT reasoning is motivated by the need for explicit normative reasoning in agentic settings. For this reason, aim I explicates the need for certificates, i.e. the reasoning procedure must generate explicit proofs of successful proof-search and counter-models in the case of failed proof-search. Such certificates can be seen as explicit *reasons* that explain an input’s (in)validity. This relates to our second aim:

(AIM II) Illustrate how the obtained calculi and proof-search algorithm can be utilized for agent-based normative reasoning tasks.

We demonstrate how the developed proof systems and accompanying proof-search algorithm can be employed for reasoning tasks concerning obligations and choices. The constructed proofs and counter-models will serve as certificates that provide reasons for an agent’s obligations and choices. The upshot of this approach, is that such certificates can be

accessed, analyzed, and assessed, which are important aspects of more transparent and explainable normative reasoning in AI.

Let us discuss these two aims in more detail.

1.1 Aim I: Proof Theory

Modern proof theory has its roots in the work of Gerhard Gentzen (1935a, 1935b) who introduced the *sequent calculus* framework for classical and intuitionistic logic. The characteristic feature of the sequent calculus framework is the utilization of *sequents* in defining inference rules and deriving theorems (sequents are formulae of the form $\varphi_1, \dots, \varphi_n \Rightarrow \psi_1, \dots, \psi_k$, interpreted as stating that if all formulae φ_i in the antecedent hold, then some formula ψ_j in the conclusion holds). The framework proved to be well-suited for the construction of *analytic calculi* (i.e. calculi that possess the *subformula property*), meaning that every formula used to *reach* the conclusion of a proof, *occurs* as a subformula in the conclusion of the proof. The analyticity of a calculus is highly valuable in designing proof-search algorithms that take a formula as input and attempt to construct a proof by applying inference rules in reverse (typically allowing for a counter-model to be constructed if a proof is not found).

Despite its success and utility, it was found that the sequent calculus framework is often not expressive enough to construct analytic calculi for many logics of interest, e.g. see (Wansing, 2002). In response, a diverse number of proof-theoretic formalisms extending Gentzen’s sequent calculus framework were introduced, allowing for the construction of analytic calculi for larger classes of logics. Such formalisms include (prefixed) tableaux (Fitting, 1972, 2014), display calculi (Belnap, 1982; Wansing, 1994), hypersequent calculi (Avron, 1996; Lahav, 2013), labeled calculi (Gabbay, 1996; Simpson, 1994; Viganò, 2000), and nested calculi (Bull, 1992; Kashima, 1994; Brünnler, 2009; Poggiolesi, 2009).

In this paper, we utilize proof systems for deontic STIT logics within the *labeled sequent* formalism (Simpson, 1994; Gabbay, 1996; Viganò, 2000; Hein, 2005; Negri, 2005). This formalism is particularly suitable for handling modal logics, since it allows for semantic information to be explicitly represented in the syntax of sequents. Its name refers to the use of labels, representing worlds in a relational model, that prefix formulae. The proof systems used in this paper first appeared in the PhD thesis of Lyon (2021) and are extensions/variants of the labeled calculi for (deontic) STIT logics provided by both authors in a sequence of papers (van Berkel & Lyon, 2019a; Lyon & van Berkel, 2019; van Berkel & Lyon, 2021). In these works, the problem of designing proof-search algorithms for deontic STIT logics was left open; this problem will be addressed in this paper.

We adopt the labeled formalism since it offers a variety of advantages. First, building calculi within the labeled formalism is straightforward as calculi are easily obtained by transforming the semantic clauses and frame properties associated with the relational semantics of a logic into inference rules (Simpson, 1994; Schmidt & Tishkovsky, 2011). Additionally, the labeled formalism provides uniform and modular presentations for many classes of logics, where the addition, deletion, or modification of inference rules from one calculus yields a calculus for another logic within the considered class (Simpson, 1994; Viganò, 2000). This modularity proves useful in our setting as the addition/deletion of inference rules transforms our proof systems into sound and complete proof systems for various (non-)deontic STIT logics, including the non-deontic traditional STIT logics developed by Belnap, Perloff, and

Xu (2001). Consequently, all of our results straightforwardly transfer to the non-deontic, multi-agent setting. Last, labeled calculi typically exhibit favorable proof-theoretic properties such as the height-preserving invertibility of rules and cut admissibility (Simpson, 1994; Viganò, 2000; Negri, 2005). We leverage such properties to successfully realize aim I.

1.2 Aim II: Normative Reasoning Tasks and Explanations

The constructive stance taken in this work is deliberate. As our proof-search algorithm also supports counter-model extraction, we automate the construction of explicit proofs of derivability and underivability. Our approach, thus, yields insight into the reasons for a formula’s (non-)theoremhood, providing information that explains logical inferences relative to agents, choices, and obligations. This is essential for the secondary aim of our work.

Explanation is a critically important topic in AI (Biran & Cotton, 2017; Miller, Howe, & Sonenberg, 2017; Gunning, Stefik, Choi, Miller, Stumpf, & Yang, 2019; Miller, 2019). It allows us to understand and evaluate why intelligent systems—such as artificial agents—produce particular decisions, predictions, and actions. An often observed consequence of improved explainability is that it increases our trust in such systems. Especially in light of the rapid increase of autonomous intelligent systems, we also find an increasing demand for formal approaches to explainable agency, which is the ability of intelligent agents to explain their choices and actions (Langley, Meadows, Sridharan, & Choi, 2017). Similarly, the demand for explanatory methods for normative reasoning has increased (van Berkel & Straßer, 2022). In particular, the work by van Berkel and Straßer (2024) proposes dialogue models to construct, what are called, deontic explanations for defeasible normative reasoning. The deontic formalisms in these works do not contain explicit concepts of agency.

Explanation is strongly connected to notions such as justification and interpretability (Biran & Cotton, 2017). Following Ye and Johnson (1995), a ‘justification’ denotes a type of explanation that gives an “explicit description of the causal argument or rationale behind each inferential step taken” (p. 158)¹ and ‘interpretability’ signifies, in the words of Biran and Cotton (2017) that a system’s “operations can be understood by a human, either through introspection or through a produced explanation” (p. 8). Hence, it can be seen that constructive proof-search procedures have explanatory value by yielding justifications in terms of proofs—i.e. transparent representations of a step-by-step inferential procedure—and facilitating interpretability via explicit inferential operations that can be inspected.

Logical reasoning is suitable for transparent reasoning tasks, avoiding the opaqueness issues of black box reasoning as often ascribed to machine learning procedures (Gunning et al., 2019). Although such logical methods might not be optimal for real-time reasoning in which decisions must be made fast, they are particularly useful for transparent reasoning procedures for which we can formally specify various reasoning tasks and facilitate explanation in the way defined above. For instance, in this paper, we illustrate how the proof theory of deontic STIT can generate certificates of potential (non-)compliance of agents’ behavior with respect to a set of obligations. By providing *procedural* proof-search algorithms with automated counter-model extraction, we address three core demands for automated

1. This is a common definition in AI. In ethics, however, ‘explanation’ and ‘justification’ are often considered distinct, namely, the former providing reasons for why something happened in the way it happened and the latter providing reasons for why it is good that it happened (Alvarez, 2017).

agent-based normative reasoning: correctness and termination, transparency of reasoning, and explainability in terms of justification and interpretability.

We also apply our deontic STIT calculi to address three normative reasoning tasks: First, our formalism can be applied to determine an agent’s obligations given a certain situation in which the agent resides. That is, given a situation—i.e. a given knowledge base consisting of obligations and facts—the agent can check whether they are bound by certain obligations. We refer to this application as *duty checking*. Second, the calculi can be used to determine whether a certain choice, considered by the agent as potential conduct, complies with the given knowledge base of obligations and facts. We refer to this second application as *compliance checking*. Third, the calculi can be employed to check whether under a specified factual context an agent can (still) jointly fulfill all of their obligations. We call this application *joint fulfillment checking*. These three applications of the framework demonstrate three conceptually different normative reasoning tasks. Furthermore, we discuss the formal similarities between these applications from the perspective of proof theory.

1.3 Related Work

There have been previous attempts to automate or mechanize reasoning with (deontic) STIT logics. An important first step was made by Arkoudas et al. (2005) who encoded a natural deduction system for Horty’s deontic STIT in Athena, allowing for the formal verification of proofs, but omitting the specification of a terminating proof-search algorithm. Sheablymyer and Abbas (2021) wrote a model-checking algorithm for a variant of Horty’s deontic STIT combined with CTL*, which takes a model and obligation as input and determines whether the obligation holds. This significantly differs from our algorithm which takes an arbitrary formula as input and decides the formula, outputting either a proof or counter-model thereof. Proof-theoretic approaches to formal reasoning with non-deontic STIT logics include the tableau systems of Wansing (2006) and the labeled systems of Negri and Pavlović (2020). The former work leaves the specification of a tableau-based decision algorithm open, while the latter work takes a similar approach to the authors’ work in this and previous papers (Lyon & van Berkel, 2019), providing a labeled calculus for non-deontic STIT logic and attempting to secure decidability via terminating proof-search. Dalmonte et al. (2021) develop proof-search with counter-model extraction for a class of non-STIT deontic logics.

As the work of Negri and Pavlović (2020) shares similarities with ours, the authors would like to emphasize the main differences between the labeled calculi applied in this paper and those defined by Negri and Pavlović (2020) for non-deontic multi-agent STIT logic: (i) We consider a more general class of STIT calculi, which includes the deontic extension of traditional multi-agent STIT (Horty, 2001; Murakami, 2004) with and without the STIT principle for limiting the amount of choices available to agents (Belnap et al., 2001, Ch. 17). (ii) In contrast to (Negri & Pavlović, 2020), our calculi adopt a *relational* semantics for STIT, forgoing the more involved—yet redundant (Herzig & Schwarzentruher, 2008)—branching-time structures in the context of atemporal STIT logics. As a consequence of using relational semantics, our calculi are syntactically simpler and more economical which, for instance, significantly decreases the number of rules in the calculi. (iii) We implement an explicit loop-checking mechanism for our proof-search algorithm to ensure correct and terminating proof-search for (deontic) multi-agent STIT logics (with a limited as well as

unlimited choice constraint). At the end of Section 3, we provide an argument as to why proof-search algorithms for (non-deontic) STIT logics without loop-checking are susceptible to problematic cases (i.e. termination may fail; e.g. in Negri and Pavlović (2020)).

Last, (non-)deontic STIT logics were shown decidable via the finite model property by Belnap et al. (2001) and Murakami (2004), respectively. Balbiani, Herzig, and Troquard (2008) studied the complexity of deciding traditional STIT logics, finding that single-agent STIT logic is NP-complete and multi-agent STIT is NEXPTIME-complete. These approaches are model-theoretic, whereas our work is the first to establish the decidability of (non-)deontic STIT logics *syntactically* by means of terminating proof-search.

1.4 Outline of the Paper

Section 2 introduces deontic STIT logics and Section 3 introduces their cut-free labeled sequent calculi. Section 4 is dedicated to a discussion of our proof-search algorithm, and the proposed loop-checking mechanism, which lets us decide (non-)deontic multi-agent STIT logics with (un)limited choice. In Section 5, we apply proof-search to the three agent-based normative reasoning tasks. We conclude and discuss future work in Section 6.

2. Logical Preliminaries

Deontic STIT logics enable reasoning about agents’ choices and obligations in multi-agent scenarios by employing a variety of modal operators. In this section, we discuss the interpretation of our deontic STIT languages and provide a formal semantics for its formulae.

Choices. Different choices may be available to different agents at different moments in time. The characteristic feature of traditional STIT logic is the use of an instantaneous *choice* operator $[i]$ for each agent i , which informally expresses that “agent i sees to it that” (some proposition holds). The operator is instantaneous in the sense that choice refers to what an agent can directly see to at a given moment. In a multi-agent world, a single agent cannot uniquely determine the future by acting. For instance, when I decide to go to a concert, it may be that my friends join me, but also they may stay at home. Nevertheless, if I see to it that I go to the concert, this excludes a future continuation where I stayed at home. Hence, what an agent can do via exercising choice is to constrain or limit the possible courses of events; the modal operator $[i]$ models this idea.

Settledness. Certain states of affairs cannot be altered by any of the agents’ (joint) choices at a given moment. Such states of affairs are *settled true* for that moment. Basic STIT logic includes a *settledness* operator \square . For instance, that it is currently settled true that it is Tuesday, means that no choice is available to any of the agents to see to it that today is *not* Tuesday. In such cases, we sometimes say that what is settled true is realized independently of any of the agents’ choices. The settledness operator plays an essential role in characterizing the relations between different choices of agents. Following Horty and Belnap (1995), we focus on single-moment scenarios.

Obligations. In the context of STIT, obligations prescribe certain choices over others and are captured by an agent-specific *deontic choice* operator \otimes_i for each agent i . We interpret \otimes_i as “agent i ought to see to it that” (some proposition holds) (Horty, 2001). To illustrate, the formula $\otimes_i \mathbf{concert}$ is informally read as “agent i ought to see to it that she attends the concert” (e.g. because she made a promise to a friend).

For each of the above modal operators (viz. $[i]$, \Box , and \otimes_i , respectively), we include the *duals* (viz. $\langle i \rangle$, \Diamond , and \ominus_i). We read $\langle i \rangle \varphi$ as ‘the state of affairs φ may result from a choice made by agent i ’ and take $\Diamond \varphi$ to express that ‘the state of affairs φ is currently possible’. We read $\ominus_i \varphi$ as ‘agent i is permitted to see to it that the state of affairs φ holds’ (cf. permission as the dual of obligation (Hilpinen & McNamara, 2013)). Beyond the aforementioned modalities, our languages also include disjunction \vee , conjunction \wedge , and classical negation. The multi-agent language \mathcal{L}_n with $n \in \mathbb{N}$ is defined as follows:

Definition 1 (The Language \mathcal{L}_n). *Let $Ag := \{0, 1, \dots, n\} \subset \mathbb{N}$ be a finite set of agent labels and let $Var := \{p_0, p_1, p_2, \dots\}$ be a denumerable set of propositional variables. The language \mathcal{L}_n is defined via the following grammar in BNF:*

$$\varphi ::= p \mid \neg p \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \Box \varphi \mid \Diamond \varphi \mid [i] \varphi \mid \langle i \rangle \varphi \mid \otimes_i \varphi \mid \ominus_i \varphi$$

where $i \in Ag$ and $p \in Var$.

Our formulae in \mathcal{L}_n are in negation normal form, which allows us to simplify the sequents employed in our calculi and reduce the number of rules. This means that we restrict applications of negation to propositional variables. Then, we define the negation $\neg \varphi$ of an arbitrary formula $\varphi \in \mathcal{L}_n$ to be the formula where every operator \vee , \wedge , \Box , \Diamond , $[i]$, $\langle i \rangle$, \otimes_i , \ominus_i is replaced with its dual \wedge , \vee , \Diamond , \Box , $\langle i \rangle$, $[i]$, \ominus_i , \otimes_i (respectively), every propositional variable p is replaced with its negation $\neg p$, and $\neg p$ is replaced with its positive version p . We define $\varphi \rightarrow \psi := \neg \varphi \vee \psi$, $\varphi \leftrightarrow \psi := (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$, $\top := p \vee \neg p$, and $\perp := p \wedge \neg p$. In what follows, we refer to the *complexity* of a formula in the usual way, corresponding to the number of symbols in a formula.

To illustrate \mathcal{L}_n , let Jade (j) be an agent cycling to her office in London. Then, the formula $\otimes_j \text{left_jade} \wedge \Diamond \text{left_jade} \wedge [j] \neg \text{left_jade}$ describes that ‘‘Jade is obliged to cycle on the left-hand side of the road, it is realizable that she cycles on the left side, and she chooses not to do so,’’ in other words, Jade sees to it that her obligation is violated.

We interpret \mathcal{L}_n formulae over special types of relational models, called DS_n^k -models. The parameter n refers to the number of agents involved and k to the maximal number of choices for each agent at each moment, and which imposes no maximum when $k = 0$. Since our logic concerns instantaneous (i.e. atemporal) choice-making, it suffices to consider single-moment frameworks (Balbiani et al., 2008), forgoing the use of traditional branching time structures often employed in atemporal STIT logics (Belnap et al., 2001) (the STIT language \mathcal{L}_n is not expressive enough to reason about branching time structures).

Definition 2 (DS_n^k -frames, -models). *Let $n \in \mathbb{N}$ and for each agent label $i \in Ag$, let us define $R_{[i]}(w) := \{v \in W \mid (w, v) \in R_{[i]}\}$. A DS_n^k -frame is defined to be a tuple of the form $F := (W, \{R_{[i]} \mid i \in Ag\}, \{I_{\otimes_i} \mid i \in Ag\})$ with W a non-empty set of worlds w, v, u, \dots and:*

- (C1) For all $i \in Ag$, $R_{[i]} \subseteq W \times W$ is an equivalence relation.
- (C2) For all $u_1, \dots, u_n \in W$, $\bigcap_{i \in Ag} R_{[i]}(u_i) \neq \emptyset$.
- (C3) If $k > 0$, then for all $i \in Ag$ and $w_0, w_1, \dots, w_k \in W$,

$$\bigvee_{0 \leq m \leq k-1, m+1 \leq j \leq k} R_{[i]} w_m w_j.$$

- (D1) For all $i \in Ag$, $I_{\otimes_i} \subseteq W$.
 (D2) For all $i \in Ag$, $I_{\otimes_i} \neq \emptyset$.
 (D3) For all $i \in Ag$ and $w, v \in W$, if $w \in I_{\otimes_i}$ and $v \in R_{[i]}(w)$, then $v \in I_{\otimes_i}$.

A DS_n^k -model is a tuple $M := (F, V)$ where F is a frame and $V: Var \rightarrow \mathcal{P}(W)$ is a valuation function mapping propositional variables to subsets of W .

Following Definition 2, a DS_n^k -model consists of a set W representing a single moment in time at which the agents from Ag are making decisions and where each world in W designates a different configuration of states of affairs that could hold at that moment. Conditions (Ci) and (Di) refer, respectively, to the *choice properties* and *deontic properties*. Condition (C1) partitions W into *choice-cells* (i.e. equivalence classes) for each $i \in Ag$, where each choice-cell represents a choice available to agent i at the considered moment. The partition imposed on W must satisfy the *independence of agents* condition, and may satisfy a *limited choice* condition (Belnap et al., 2001). Independence of agents, expressed by condition (C2), stipulates that no agent can keep another agent from exercising an available choice at the moment of choice, i.e. regardless of the choices made by the agents, some set of possible states of affairs ensues. Formally, (C2) ensures that any combination of choices made by the agents is consistent, i.e. if we select a choice-cell for each agent, their intersection is non-empty. The limited choice condition, captured by (C3), stipulates that each agent has a maximum of k choices available to choose from at each moment and is only enforced when $k > 0$ (i.e. imposing no limitation on the number of choices when $k = 0$). Formally, (C3) states that if the limit is k , then for $k + 1$ many worlds at least two worlds must be in the same choice-cell.² The condition (D1) ensures that all deontically optimal worlds at the present moment for any agent i (which are those worlds contained in I_{\otimes_i}) are realizable. Condition (D2) ensures that, for each agent i , at least one deontically optimal world exists at a moment, and (D3) states that if a deontically optimal world exists in a choice, then every world in that choice is deontically optimal. In other words, (D1)–(D3) ensure the existence of an obligatory choice for each agent i at a given moment. Last, each DS_n^k -model contains a *valuation function* V mapping propositional variables to sets of worlds.

Definition 3 (Semantics). *Let M be a DS_n^k -model and let $w \in W$ of M . The satisfaction of a formula $\varphi \in \mathcal{L}_n$ in M at w is recursively defined accordingly:*

1. $M, w \Vdash p$ iff $w \in V(p)$
2. $M, w \Vdash \neg p$ iff $w \notin V(p)$
3. $M, w \Vdash \varphi \wedge \psi$ iff $M, w \Vdash \varphi$ and $M, w \Vdash \psi$
4. $M, w \Vdash \varphi \vee \psi$ iff $M, w \Vdash \varphi$ or $M, w \Vdash \psi$
5. $M, w \Vdash \Box\varphi$ iff $\forall u \in W, M, u \Vdash \varphi$

2. Belnap et al. (2001, Ch. 9) argue that one may want to limit the amount of choices due to the possibility of ‘busy choosers’, which are agents making infinitely many choices in a finite period. For some STIT logics, busy choosers have undesirable effects, e.g. where refraining from refraining does not equal doing.

6. $M, w \Vdash \Diamond\varphi$ iff $\exists u \in W, M, u \Vdash \varphi$
7. $M, w \Vdash [i]\varphi$ iff $\forall u \in R_{[i]}(w), M, u \Vdash \varphi$
8. $M, w \Vdash \langle i \rangle\varphi$ iff $\exists u \in R_{[i]}(w), M, u \Vdash \varphi$
9. $M, w \Vdash \otimes_i\varphi$ iff $\forall u \in I_{\otimes_i}, M, u \Vdash \varphi$
10. $M, w \Vdash \ominus_i\varphi$ iff $\exists u \in I_{\otimes_i}, M, u \Vdash \varphi$

A formula φ is DS_n^k -valid (written $\Vdash_{\text{DS}_n^k} \varphi$) iff for every DS_n^k -model M , it is satisfied at every world in the domain W of M . We define the logic DS_n^k as the set of DS_n^k -valid formulae.

The semantic evaluation of the propositional connectives implies that our logic extends classical propositional logic. We note in passing that sound and complete axiomatizations of DS_n^k with respect to the above semantic characterization are provided by Lyon (2021).

Let us consider an example. Suppose there are two agents, Jade and Kai (where $Ag := \{j, k\}$), who are cycling toward each other on a two lane road. Both have two choices: keep cycling left or keep cycling right. Let `left_jade` and `right_jade`, and `left_kai` and `right_kai` denote the states of affairs in which Jade and Kai cycle left and right, respectively. In the graphical illustration of this model (see Figure 2), the set $W = \{w_1, w_2, w_3, w_4\}$ represents a moment in time with four possible situations, e.g. w_1 denotes the situation in which both Jade and Kai cycle on the left. The model in Figure 2 shows that Jade has two choices, one consisting of the worlds w_1 and w_3 , and one consisting of w_2 and w_4 , whereas Kai has one choice consisting of w_1 and w_2 , and one consisting of w_3 and w_4 . Graphically, the horizontal choice-cells outlined with dotted lines ‘ $\cdot \cdot \cdot$ ’ represent Kai’s choices and the vertical choice-cells outlined with dashed lines ‘ $- - -$ ’ represent Jade’s choices. There are two cases in which a collision between Jade and Kai can be avoided: when both keep cycling left and when both keep cycling right. These two scenarios are represented by the intersecting choices leading to w_1 , respectively w_4 in Figure 2. A collision `coll` holds at those intersecting choices where one of the two cycles left and the other right.³ In Figure 2 the propositional variable `left_jade` is mapped to the set $\{w_1, w_3\}$, meaning that in these possible worlds Jade cycles left. The negation $\neg p$ of a propositional variable p holds at those worlds where p does *not* hold (e.g. $\neg\text{coll}$ holds at w_1 and w_3 in our example model).

Although there are two outcomes in the cycling scenario of Figure 2 in which a collision is avoided, we assume that UK traffic law obliges cyclists to always cycle on the left-hand side of the road. In Figure 2, Jade’s choice consisting of w_1 and w_3 is obligatory for Jade and Kai’s choice consisting of w_1 and w_2 is obligatory for Kai. Thus, based on the assumed traffic laws, if both Jade and Kai commit to their obligatory choices, they jointly see to it that a collision is avoided. In Figure 2, obligatory choices are shaded. In relation to this, due to fact that DS_n^k -models ensure that at least one obligatory choice exists for each agent $i \in Ag$, each agent’s set of obligations is consistent, i.e. no agent is obliged to see to it that φ and $\neg\varphi$ for some proposition φ . It can be straightforwardly checked that the example DS_n^k -model in Figure 2 satisfies all properties from Definition 2.

3. Avoiding a collision is expressed as: $\Box(\neg\text{coll} \rightarrow ((\text{right_jade} \wedge \text{right_kai}) \vee (\text{left_jade} \wedge \text{left_kai})))$.




Figure 1: An illustration of the DS_n -model $M := (F, V)$ built atop the DS_n -frame $F := (W, \{R_{[j]}, R_{[y]}\}, \{I_{\otimes_j}, I_{\otimes_y}\})$ such that $W := \{w_1, w_2, w_3, w_4\}$, $R' := \{(w_i, w_i) \mid 1 \leq i \leq 4\}$, Jade's relation $R_{[j]} := \{(w_1, w_3), (w_3, w_1), (w_2, w_4), (w_4, w_2)\} \cup R'$, Kai's relation $R_{[k]} := \{(w_1, w_2), (w_2, w_1), (w_3, w_4), (w_4, w_3)\} \cup R'$, $I_{\otimes_j} := \{w_1, w_3\}$, $I_{\otimes_k} := \{w_1, w_2\}$, and $V := \{(\text{left_jade}, \{w_1, w_3\}), (\text{right_jade}, \{w_2, w_4\}), (\text{left_kai}, \{w_1, w_2\}), (\text{right_kai}, \{w_3, w_4\}), (\text{coll}, \{w_2, w_3\})\}$.

3. Proof Theory for Deontic STIT Logics

We present a class of labeled sequent calculi, which generalize Gentzen-style sequent systems by including semantic information directly in the syntax of sequents. The calculi are variations of the calculi given by van Berkel and Lyon (2021), and were first defined by Lyon (2021). The syntactic expressions in our calculi, called *labeled sequents*—or *sequents* for short—encode semantic information by prefixing formulae with labels (representing worlds in a relational model) and include relational atoms (encoding the accessibility relations).

Definition 4 (Sequent). *Let $Ag = \{0, 1, \dots, n\}$ be a finite set of agent labels and let $Lab := \{w, u, v, \dots\}$ be a denumerable set of labels. A sequent is defined to be an expression of the form $\Lambda := \mathcal{R} \Rightarrow \Gamma$, where \mathcal{R} is a (potentially empty) finite set of relational atoms of the form $R_{[i]}wu$ and $I_{\otimes_i}w$, and Γ is a (potentially empty) finite set of labeled formulae of the form $w : \varphi$, where i ranges over Ag , φ ranges over \mathcal{L}_n , and w, u range over Lab .*

We use $\mathcal{R}, \mathcal{Q}, \dots$ (potentially annotated) to denote sets of relational atoms, and use Γ, Δ, \dots (potentially annotated) to denote sets of labeled formulae. We often use $\Lambda := \mathcal{R} \Rightarrow \Gamma$ to denote an entire sequent, where we refer to \mathcal{R} as the *antecedent* of Λ , and Γ as the *consequent* of Λ . Intuitively, a sequent $\mathcal{R} \Rightarrow \Gamma$ is interpreted as stating that if all relational atoms in \mathcal{R} hold, then some labeled formula $w : \varphi \in \Gamma$ holds (see Definition 6 for the semantic interpretation of a sequent). We use the notation $Lab(\mathcal{R})$, $Lab(\Gamma)$, and $Lab(\Lambda)$ to represent the set of labels contained in \mathcal{R} , Γ , and Λ , respectively; e.g. $Lab(R_{[2]}wv \Rightarrow u : p) = \{w, u, v\}$. Furthermore, for a label w and set of labeled formulae Γ , we define $\Gamma \upharpoonright w := \{\varphi \mid w : \varphi\}$ to be the set of formulae prefixed with the label w .

$$\begin{array}{c}
 \frac{}{\mathcal{R} \Rightarrow w : p, w : \neg p, \Gamma} \text{(id)} \quad \frac{\mathcal{R} \Rightarrow w : \varphi, \Gamma \quad \mathcal{R} \Rightarrow w : \psi, \Gamma}{\mathcal{R} \Rightarrow w : \varphi \wedge \psi, \Gamma} (\wedge) \\
 \\
 \frac{\mathcal{R} \Rightarrow w : \varphi, w : \psi, \Gamma}{\mathcal{R} \Rightarrow w : \varphi \vee \psi, \Gamma} (\vee) \quad \frac{\mathcal{R}, R_{[1]}w_1u, \dots, R_{[n]}w_nu \Rightarrow \Gamma}{\mathcal{R} \Rightarrow \Gamma} \text{(IOA)}^\dagger \\
 \\
 \frac{\mathcal{R} \Rightarrow u : \varphi, \Gamma}{\mathcal{R} \Rightarrow w : \Box \varphi, \Gamma} (\Box)^\dagger \quad \frac{\mathcal{R} \Rightarrow w : \Diamond \varphi, u : \varphi, \Gamma}{\mathcal{R} \Rightarrow w : \Diamond \varphi, \Gamma} (\Diamond) \quad \frac{\mathcal{R}, R_{[i]}ww \Rightarrow \Gamma}{\mathcal{R} \Rightarrow \Gamma} \text{(Ref}_i\text{)} \\
 \\
 \frac{\mathcal{R}, I_{\otimes_i}u \Rightarrow w : \ominus_i \varphi, u : \varphi, \Gamma}{\mathcal{R}, I_{\otimes_i}u \Rightarrow w : \ominus_i \varphi, \Gamma} (\ominus_i) \quad \frac{\mathcal{R}, R_{[i]}wu \Rightarrow w : \langle i \rangle \varphi, u : \varphi, \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow w : \langle i \rangle \varphi, \Gamma} (\langle i \rangle) \\
 \\
 \frac{\mathcal{R}, R_{[i]}wu \Rightarrow u : \varphi, \Gamma}{\mathcal{R} \Rightarrow w : [i]\varphi, \Gamma} ([i])^\dagger \quad \frac{\mathcal{R}, I_{\otimes_i}u \Rightarrow u : \varphi, \Gamma}{\mathcal{R} \Rightarrow w : \otimes_i \varphi, \Gamma} (\otimes_i)^\dagger \quad \frac{\mathcal{R}, I_{\otimes_i}u \Rightarrow \Gamma}{\mathcal{R} \Rightarrow \Gamma} \text{(D2}_i\text{)}^\dagger \\
 \\
 \frac{\mathcal{R}, R_{[i]}wu, R_{[i]}wv, R_{[i]}wv \Rightarrow \Gamma}{\mathcal{R}, R_{[i]}wu, R_{[i]}wv \Rightarrow \Gamma} \text{(Euc}_i\text{)} \quad \frac{\mathcal{R}, R_{[i]}wu, I_{\otimes_i}w, I_{\otimes_i}u \Rightarrow \Gamma}{\mathcal{R}, R_{[i]}wu, I_{\otimes_i}w \Rightarrow \Gamma} \text{(D3}_i\text{)} \\
 \\
 \frac{\left\{ \mathcal{R}, R_{[i]}w_mw_j \Rightarrow \Gamma \mid 0 \leq m \leq k-1, m+1 \leq j \leq k \right\}}{\mathcal{R} \Rightarrow \Gamma} \text{(APC}_i^k\text{)}
 \end{array}$$

Figure 2: The G3DS_n^k calculi (Lyon, 2021), where $|Ag| = n \in \mathbb{N}$ and $k \in \mathbb{N}$. The superscript \dagger on (\Box) , $([i])$, (\otimes_i) , (IOA) , and (D2_i) indicates that u is a eigenvariable, i.e. it does not occur in the conclusion. We have $(\langle i \rangle)$, $([i])$, (\ominus_i) , (\otimes_i) , (Ref_i) , (Euc_i) , (APC_i^k) , (D2_i) , and (D3_i) rules for each $i \in Ag$. (APC_i^k) is omitted when $k = 0$.

Definition 6 (Sequent Semantics). *Let $M = (W, \{R_{[i]} \mid i \in Ag\}, \{I_{\otimes_i} \mid i \in Ag\}, V)$ be a DS_n^k -model with $I : \text{Lab} \mapsto W$ an interpretation function mapping labels to worlds. The satisfaction of a relational atom $R_{[i]}wu$ or $I_{\otimes_i}w$ (written $M, I \models R_{[i]}wu$, respectively $M, I \models I_{\otimes_i}w$) and a labeled formula $w : \varphi$ (written $M, I \models w : \varphi$) is defined as:*

- $M, I \models R_{[i]}wu$ iff $(w^I, u^I) \in R_{[i]}$
- $M, I \models I_{\otimes_i}w$ iff $w^I \in I_{\otimes_i}$
- $M, I \models w : \varphi$ iff $M, w^I \Vdash \varphi$

A labeled sequent $\Lambda := \mathcal{R} \Rightarrow \Gamma$ is satisfied in M with I (written, $M, I \models \Lambda$) iff if $M, I \models R_{[i]}wu$ and $M, I \models I_{\otimes_i}w$ for all $R_{[i]}wu, I_{\otimes_i}w \in \mathcal{R}$, then $M, I \models w : \varphi$ for some $w : \varphi \in \Gamma$. We say that a labeled sequent Λ is falsified in M with I iff $M, I \not\models \Lambda$, that is, Λ is not satisfied by M with I . Last, a labeled sequent Λ is DS_n^k -valid (written $\models \Lambda$) iff it is satisfiable in every DS_n^k model M with every interpretation function I . We say that a labeled sequent Λ is DS_n^k -invalid iff $\not\models \Lambda$, i.e. Λ is not DS_n^k -valid.

$$\frac{\mathcal{R} \Rightarrow \Gamma}{\mathcal{R}(w/u) \Rightarrow \Gamma(w/u)} \text{ (Sub)} \quad \frac{\mathcal{R} \Rightarrow \Gamma}{\mathcal{R}, \mathcal{R}' \Rightarrow \Gamma, \Gamma'} \text{ (Wk)} \quad \frac{\mathcal{R} \Rightarrow w : \varphi, \Gamma \quad \mathcal{R} \Rightarrow w : \neg\varphi, \Gamma}{\mathcal{R} \Rightarrow \Gamma} \text{ (Cut)}$$

Figure 3: The set StrR of structural rules consists of the rules above.

As shown in Theorem 7 below, the G3DS_n^k calculi for deontic STIT logics possess fundamental properties such as the *height-preserving (hp) admissibility* of label substitutions (Sub) and weakenings (Wk) (see Figure 3). Moreover, all rules of the calculi are *height-preserving invertible*, and each calculus admits a syntactic proof of cut admissibility. We briefly clarify the meaning of these properties, which will recur throughout the course of this paper, and after, list the properties possessed by each G3DS_n^k calculus.

Let the rule (r) be of the form:

$$\frac{\Lambda_1 \quad \dots \quad \Lambda_n}{\Lambda} (r)$$

We say that the rule (r) is *(hp)-admissible* in a calculus *iff* if $\Lambda_1, \dots, \Lambda_n$ have proofs (with heights h_1, \dots, h_n , respectively), then Λ has a proof (with a height $h \leq \max\{h_1, \dots, h_n\}$). We let the *i-inverse* of (r) , written (\hat{r}_i) , be the rule obtained by taking the conclusion of (r) to be the premise of (\hat{r}_i) and the i^{th} premise of (r) to be the conclusion of (\hat{r}_i) . We say that an n -ary rule (r) is *(hp)-invertible* in a calculus *iff* (\hat{r}_i) is (hp)-admissible for each $1 \leq i \leq n$. Last, a *label substitution* (w/u) is defined in the usual way, that is, for a labeled sequent Λ , $\Lambda(w/u)$ is obtained by replacing each occurrence of u by w in Λ . For example, $(R_{[i]}wu, I_{\otimes_i}u \Rightarrow u : \varphi)(v/u) = (R_{[i]}wv, I_{\otimes_i}v \Rightarrow v : \varphi)$.

Admissibility and invertibility properties serve a twofold purpose: first, such properties are useful in establishing the completeness of the G3DS_n^k calculi (Theorem 7), and second, they are useful in showing the decidability of our logics via proof-search (Section 4).

Theorem 7 (G3DS_n^k Properties (Lyon, 2021, Ch. 3.4)). *Let $|Ag| = n \in \mathbb{N}$ and $k \in \mathbb{N}$.*

1. For all $\varphi \in \mathcal{L}_n$, $\text{G3DS}_n^k \vdash \mathcal{R} \Rightarrow w : \varphi, w : \neg\varphi, \Gamma$;
2. The (Sub) and (Wk) rules are *hp-admissible* in G3DS_n^k ;
3. All rules in G3DS_n^k are *hp-invertible*;
4. The rule (Cut) is *admissible* in G3DS_n^k ;
5. G3DS_n^k is *sound*, i.e. if $\vdash_{\text{G3DS}_n^k} \Lambda$, then $\models \Lambda$;
6. G3DS_n^k is *complete*, i.e. $\Vdash_{\text{DS}_n^k} \varphi$, then $\vdash_{\text{G3DS}_n^k} w : \varphi$.

4. Proof-search and Decidability

We now put the G3DS_n^k calculi to use and demonstrate that each system serves as a basis for automated reasoning with the corresponding logic DS_n^k . In particular, we design a proof-search procedure, referred to as Prove_n^k (see Algorithm 1 below), that decides the validity

of each formula $\varphi \in \mathcal{L}_n$, and which additionally provides witnesses for the answers it yields. That is to say, Prove_n^k bottom-up applies rules from G3DS_n^k to the input formula, and when a proof is found (i.e. proof-search succeeds), it follows that the input formula is valid; when a proof is not found (i.e. proof-search fails), we show that a counter-model witnessing the invalidity of the formula can be extracted.

An interesting feature of our proof-search algorithm is the inclusion of a novel loop-checking mechanism, which ensures that proof-search avoids entering an infinite loop. Although our loop-checking mechanism is sophisticated, it is necessary. In fact, at the end of the section, we give a counter-example to terminating proof-search in the absence of loop-checking, illustrating the need for such a mechanism.

The first tool we introduce, to be used during proof-search, is the notion of an \mathcal{R}^i -path. Intuitively, an \mathcal{R}^i -path exists between two labels in a set \mathcal{R} of relational atoms when one label can be reached from the other by means of an undirected sequence of relational atoms of the form $R_{[i]}wu$. Semantically, an \mathcal{R}^i -path existing between two labels w and u means that both are interpreted as worlds within the same choice-cell.

Definition 8 (\mathcal{R}^i -Path). *Let $w \sim_i u \in \{R_{[i]}wu, R_{[i]}uw\}$ and let \mathcal{R} be a set of relational atoms. An \mathcal{R}^i -path of relational atoms from a label w to u occurs in \mathcal{R} (written $w \sim_i^{\mathcal{R}} u$) iff $w = u$, $w \sim_i u$, or there exist labels v_i (with $i \in \{0, \dots, n\}$) such that $w \sim_i v_0, \dots, v_n \sim_i u$ occurs in \mathcal{R} .*

By making use of the above definition, it is straightforward to verify the following:

Lemma 9. *If $\Lambda = \mathcal{R} \Rightarrow \Gamma$ is a sequent, then $\sim_i^{\mathcal{R}}$ is an equivalence relation over $\text{Lab}(\Lambda)$.*

The use of \mathcal{R}^i -paths is uniquely beneficial in managing bottom-up rule applications of (IOA) during proof-search. In essence, given a sequent Λ and $n = |Ag|$ many labels $w_1, \dots, w_n \in \text{Lab}(\Lambda)$, we check to see if a label $u \in \text{Lab}(\Lambda)$ exists such that $w_i \sim_i^{\mathcal{R}} u$ for every $i \in \{1, \dots, n\}$, and if not, then (IOA) is applied bottom-up to the sequent during proof-search. This avoids unnecessary applications of the (IOA) rule during proof-search since if such a label $u \in \text{Lab}(\Lambda)$ exists, then u is interpreted as a world in the intersection of the $n = |Ag|$ many choice-cells that respectively contain w_1, \dots, w_n , meaning (IOA) need not be bottom-up applied relative to these labels as they already satisfy **(C2)**.

We introduce a second helpful notion in the management of applications of the (IOA) rule, namely, the notion of a *thread*, which is defined below.

Definition 10 (Thread). *We define a thread from $\mathcal{R}_0 \Rightarrow \Gamma_0$ to $\mathcal{R}_h \Rightarrow \Gamma_h$ to be a sequence $\mathcal{T} = (\mathcal{R}_i \Rightarrow \Gamma_i)_{i=0}^h$ of sequents such that*

1. *for each $i \in \{0, \dots, h-1\}$, each sequent $\mathcal{R}_{i+1} \Rightarrow \Gamma_{i+1}$ is obtained from $\mathcal{R}_i \Rightarrow \Gamma_i$ by a bottom-up application of a rule in G3DS_n^k with the former sequent serving as a premise and the latter serving as the conclusion, and*
2. *if a sequent $\mathcal{R}_{i+1} \Rightarrow \Gamma_{i+1}$ is obtained from $\mathcal{R}_i \Rightarrow \Gamma_i$ by a bottom-up application of (IOA), meaning that $\mathcal{R}_{i+1} = \mathcal{R}_i, R_{[1]}w_1u, \dots, R_{[n]}w_nu$, then none of the labels w_1, \dots, w_n were introduced via a prior application of (IOA).*

Occasionally, we say that a thread \mathcal{T} is from $\mathcal{R}_0 \Rightarrow \Gamma_0$ or to $\mathcal{R}_h \Rightarrow \Gamma_h$ if we only aim to indicate the first or last element of \mathcal{T} . Last, we call h the height of the thread \mathcal{T} and $\mathcal{R}_h \Rightarrow \Gamma_h$ a top sequent.

The first condition in the above definition states that a thread corresponds to a single path in a ‘partial’ G3DS_n^k proof, i.e. the path may not necessarily end at an initial sequent. The second condition stipulates that bottom-up applications of (IOA) are only applied to labels *not* introduced by prior (IOA) applications. We shall see that our proof-search procedure always constructs threads with this second property by definition as this ensures the termination of proof-search (see Lemma 25 and Theorem 26 below).

Rather than considering single bottom-up applications of the (IOA) rule during proof-search, we define an operation, referred to as IoaOp , which sequentially applies (IOA) to certain collections of labels from a given sequent.

Definition 11 ($\text{IOA}^{\mathcal{T}}$ -tuple, $\text{IOA}^{\mathcal{T}}$ -satisfied, IoaOp). Let $\Lambda = \mathcal{R} \Rightarrow \Gamma$ be a sequent, $|Ag| = n$, and $\mathcal{T} = (\mathcal{R}_i \Rightarrow \Gamma_i)_{i=0}^h$ be a thread from $\mathcal{R}_0 \Rightarrow \Gamma_0$ to Λ , i.e. $\mathcal{R} = \mathcal{R}_h$ and $\Gamma = \Gamma_h$.

We define an $\text{IOA}^{\mathcal{T}}$ -label in Λ to be a label introduced by a bottom-up application of (IOA) in \mathcal{T} , and a non- $\text{IOA}^{\mathcal{T}}$ -label to be a label in Λ that is not an $\text{IOA}^{\mathcal{T}}$ -label. A tuple $(w_1, \dots, w_n) \in \text{Lab}(\Lambda)$ of length n is defined to be an $\text{IOA}^{\mathcal{T}}$ -tuple iff for each $i \in \{1, \dots, n\}$, the label w_i is a non- $\text{IOA}^{\mathcal{T}}$ -label.

We say that a tuple $(w_1, \dots, w_n) \in \text{Lab}(\Lambda)^n$ is $\text{IOA}^{\mathcal{T}}$ -satisfied iff there exists an $\text{IOA}^{\mathcal{T}}$ -label $u \in \text{Lab}(\Lambda)$ such that $w_i \sim_i^{\mathcal{R}} u$ for all $i \in Ag$. We say that Λ is $\text{IOA}^{\mathcal{T}}$ -satisfied iff each tuple in $\text{Lab}(\Lambda)^n$ is $\text{IOA}^{\mathcal{T}}$ -satisfied.

Let \mathbf{I} be the set of all $\text{IOA}^{\mathcal{T}}$ -tuples of Λ that are not $\text{IOA}^{\mathcal{T}}$ -satisfied. We define

$$\text{IoaOp}(\Lambda) := \mathcal{R}, \bigcup_{\vec{w} \in \mathbf{I}} R_{[i]} w_i u \Rightarrow \Gamma$$

with u fresh for each $\text{IOA}^{\mathcal{T}}$ -tuple $\vec{w} = (w_1, \dots, w_n) \in \mathbf{I}$.

There are two reasons for the introduction of the above operation: first, after an application of IoaOp we are ensured that if we transform our labeled sequent into a counter-model, the model will satisfy condition (C2) (independence of agents) which is required in proving our proof-search procedure correct (see Lemma 12). Second, this operation limits bottom-up applications of (IOA) to only certain labels, namely, labels that were not introduced by bottom-up applications of (IOA). This ensures that bottom-up applications of (IOA) do not enter into an infinite loop during proof-search, introducing a fresh label ad infinitum.

We may disregard proofs where the ($[i]$) rule is applied to $\text{IOA}^{\mathcal{T}}$ -labels, and so, $\text{IOA}^{\mathcal{T}}$ -labels will never give rise to new relational atoms of the form $R_{[i]}uv$ during proof-search, ultimately facilitating termination. This point is discussed in more detail in Remark 19.

Lemma 12. Let $\mathcal{T} = (\mathcal{R}_i \Rightarrow \Gamma_i)_{i=0}^h$ be a thread where $\Lambda = \mathcal{R}_{h-1} \Rightarrow \Gamma_{h-1}$ and $\Lambda' = \mathcal{R}_h \Rightarrow \Gamma_h = \text{IoaOp}(\Lambda)$. If $(w_1, \dots, w_n) \in \text{Lab}(\Lambda')^n$, then (w_1, \dots, w_n) is $\text{IOA}^{\mathcal{T}}$ -satisfied in Λ' .

Proof. Let $(w_1, \dots, w_n) \in \text{Lab}(\Lambda')$. We have two cases: either (i) $(w_1, \dots, w_n) \in \text{Lab}(\Lambda)^n$, or (ii) $(w_1, \dots, w_n) \notin \text{Lab}(\Lambda)^n$.

(i) Suppose that $(w_1, \dots, w_n) \in \text{Lab}(\Lambda)^n$ and let $\{i_1, \dots, i_m\} \subseteq \{1, \dots, n\}$ such that each w_{i_j} with $j \in \{1, \dots, m\}$ is an $\text{IOA}^{\mathcal{T}}$ -label. Then, since each $\text{IOA}^{\mathcal{T}}$ -label w_{i_j} was introduced via

a bottom-up application of (IOA), there exists a label u_{i_j} such that $R_{[i_j]}u_{i_j}w_{i_j} \in \mathcal{R}_{h-1}$. Let (v_1, \dots, v_n) be the tuple obtained by replacing each occurrence of w_{i_j} in (w_1, \dots, w_n) with u_{i_j} . Observe that (v_1, \dots, v_n) is an $\text{IOA}^{\mathcal{T}}$ -tuple since no v_i with $i \in \{1, \dots, n\}$ is an $\text{IOA}^{\mathcal{T}}$ -label by Definition 10 (i.e. threads only allow bottom-up applications of the (IOA) rule to non- $\text{IOA}^{\mathcal{T}}$ -labels). Then, (v_1, \dots, v_n) is $\text{IOA}^{\mathcal{T}}$ -satisfied in Λ' as IoaOp was applied, implying the existence of a label z such that $v_i \sim_i^{\mathcal{R}_h} z$ for all $i \in \{1, \dots, n\}$. Since $R_{[i_j]}u_{i_j}w_{i_j} \in \mathcal{R}_{h-1} \subseteq \mathcal{R}_h$, we also have that $u_{i_j} \sim_i^{\mathcal{R}_h} w_{i_j}$ for all $j \in \{1, \dots, m\}$, implying that $u_{i_j} \sim_i^{\mathcal{R}_h} z$ for each u_{i_j} by Lemma 9. It therefore follows that (w_1, \dots, w_n) is $\text{IOA}^{\mathcal{T}}$ -satisfied in Λ' .

(ii) Let us assume that $(w_1, \dots, w_n) \notin \text{Lab}(\Lambda)^n$, meaning that the following set

$$N_1 := \{i \in \{1, \dots, n\} \mid w_i \in \text{Lab}(\Lambda') \setminus \text{Lab}(\Lambda)\}$$

is non-empty. We define $N_2 := \{1, \dots, n\} \setminus N_1$, and note that N_1 contains the indices of all fresh labels w_i introduced by IoaOp to Λ' and N_2 contains the indices of all labels w_j occurring in Λ . Since each w_i with $i \in N_1$ was introduced by IoaOp , we know that there exists a label $u_i \in \text{Lab}(\Lambda)$ such that $R_{[i]}u_iw_i \in \mathcal{R}_h$, i.e. $w_i \sim_i^{\mathcal{R}_h} u_i$. If we consider the tuple (v_1, \dots, v_n) obtained by replacing each occurrence of w_i in (w_1, \dots, w_n) with each such u_i , then $(v_1, \dots, v_n) \in \text{Lab}(\Lambda)$, implying that (v_1, \dots, v_n) is $\text{IOA}^{\mathcal{T}}$ -satisfied by case (i) above. Hence, there exists a label z such that $v_i \sim_i^{\mathcal{R}} z$ for each $i \in \{1, \dots, n\}$. For any label w_i in (w_1, \dots, w_n) with $i \in N_1$, we have that $w_i \sim_i^{\mathcal{R}_h} u_i$ and $u_i \sim_i^{\mathcal{R}_h} z$, meaning that $w_i \sim_i^{\mathcal{R}} z$ for each $i \in N_1$ by Lemma 9. Since $w_j \sim_j^{\mathcal{R}_h} z$ holds for each $j \in N_2$ as well, we have that (w_1, \dots, w_n) is indeed $\text{IOA}^{\mathcal{T}}$ -satisfied in Λ' . \square

We now introduce and prove the admissibility of a selection of inference rules (shown in Figure 4). The (\square^*) , (\otimes_i^*) , and $([i]^*)$ rules are particularly helpful in controlling how \square , \otimes_i , and $[i]$ formulae are processed during proof-search. The (\square^*) and (\otimes_i^*) rules allow for the label of the principal formula to be changed, and the $([i]^*)$ rule allows for the principal formula $u : [i]\varphi$ to be ‘unpacked’ at a label w so long as a relational atom of the form $R_{[i]}wu$ is present in the sequent. The $(\langle i \rangle^*)$ rule plays a crucial role in ensuring the correctness of our proof-search algorithm (see Theorem 23, the $\langle i \rangle\xi$ case, case (ii)^{*} of the inductive step on p. 862). Last, the (Sym_i) rule is helpful in showing the admissibility of (Tra_i) , which is used in establishing the admissibility of $(\langle i \rangle^*)$.

Such rules will be used to change the functionality of proof-search so that labels are only introduced via the (\square^*) , (\otimes_i^*) , $([i])$, and $([i]^*)$ rules in a “treelike manner” as witnessed by Lemma 15 below and which gives rise to a loop-checking mechanism (see Definition 16) employed in proof-search to ensure termination. We come back to this later, in Remark 19.

Lemma 13. *All rules in Figure 4 are admissible in G3DS_n^k .*

Proof. We first show the admissibility of the rules (Sym_i) , (Tra_i) , (\square^*) , and (\otimes_i^*) . The first two rules leverage the hp-admissibility of (Wk) and the latter two rules rely on the hp-invertibility of (\square) and (\otimes_i) .

$$\frac{\frac{\mathcal{R}, R_{[i]}wu, R_{[i]}uw \Rightarrow \Gamma}{\mathcal{R}, R_{[i]}ww, R_{[i]}wu, R_{[i]}uw \Rightarrow \Gamma} (\text{Wk})}{\frac{\mathcal{R}, R_{[i]}ww, R_{[i]}wu \Rightarrow \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow \Gamma} (\text{Ref}_i)} (\text{Euc}_i) \quad \frac{\mathcal{R} \Rightarrow w : \square\varphi, \Gamma}{\mathcal{R} \Rightarrow v : \varphi, \Gamma} \text{Theorem 7-3}}{\mathcal{R} \Rightarrow u : \square\varphi, \Gamma} (\square)$$

$$\begin{array}{c}
 \frac{\mathcal{R}, R_{[i]}wu, R_{[i]}uw \Rightarrow \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow \Gamma} \text{ (Sym}_i) \quad \frac{\mathcal{R}, R_{[i]}wu, R_{[i]}uv, R_{[i]}wv \Rightarrow \Gamma}{\mathcal{R}, R_{[i]}wu, R_{[i]}wv \Rightarrow \Gamma} \text{ (Tra}_i) \\
 \\
 \frac{\mathcal{R}, R_{[i]}wu, R_{[i]}wv \Rightarrow u : [i]\varphi, v : \varphi, \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow u : [i]\varphi, \Gamma} ([i]^*)^\dagger \quad \frac{\mathcal{R} \Rightarrow w : \Box\varphi, \Gamma}{\mathcal{R} \Rightarrow u : \Box\varphi, \Gamma} (\Box^*) \\
 \\
 \frac{\mathcal{R}, R_{[i]}wu \Rightarrow w : \langle i \rangle\varphi, u : \langle i \rangle\varphi, u : \varphi, \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow w : \langle i \rangle\varphi, \Gamma} (\langle i \rangle^*) \quad \frac{\mathcal{R} \Rightarrow w : \otimes_i\varphi, \Gamma}{\mathcal{R} \Rightarrow u : \otimes_i\varphi, \Gamma} (\otimes_i^*)
 \end{array}$$

 Figure 4: Admissible rules. The side condition \dagger stipulates that v must be an eigenvariable.

$$\begin{array}{c}
 \frac{\mathcal{R}, R_{[i]}wu, R_{[i]}uv, R_{[i]}wv \Rightarrow \Gamma}{\mathcal{R}, R_{[i]}wu, R_{[i]}uw, R_{[i]}uv, R_{[i]}wv \Rightarrow \Gamma} \text{ (Wk)} \quad \frac{\mathcal{R} \Rightarrow w : \otimes_i\varphi, \Gamma}{\mathcal{R}, I_{\otimes_i}v \Rightarrow v : \varphi, \Gamma} \text{ Theorem 7-3} \\
 \frac{\mathcal{R}, R_{[i]}wu, R_{[i]}uw, R_{[i]}uv \Rightarrow \Gamma}{\mathcal{R}, R_{[i]}wu, R_{[i]}uv \Rightarrow \Gamma} \text{ (Euc}_i) \quad \frac{\mathcal{R} \Rightarrow u : \otimes_i\varphi, \Gamma}{\mathcal{R} \Rightarrow u : \otimes_i\varphi, \Gamma} (\otimes_i) \\
 \text{ (Sym}_i)
 \end{array}$$

To show the admissibility of $([i]^*)$, we invoke the hp-admissibility of (Wk) as well as (Cut) admissibility (Theorem 7).

$$\begin{array}{c}
 \Pi_1 = \left\{ \frac{\frac{\mathcal{R}, R_{[i]}wu, R_{[i]}wv \Rightarrow u : \langle i \rangle[i]\varphi, u : \langle i \rangle\neg\varphi, u : [i]\varphi, w : [i]\varphi, \Gamma}{\mathcal{R}, R_{[i]}wu, R_{[i]}wv \Rightarrow u : \langle i \rangle[i]\varphi, u : \langle i \rangle\neg\varphi, w : [i]\varphi, \Gamma} \text{ (Wk)} \quad \frac{\mathcal{R}, R_{[i]}wu, R_{[i]}wv \Rightarrow u : \langle i \rangle[i]\varphi, u : \langle i \rangle\neg\varphi, u : [i]\varphi, w : [i]\varphi, \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow u : \langle i \rangle[i]\varphi, u : \langle i \rangle\neg\varphi, w : [i]\varphi, \Gamma} \text{ (Ref}_i)}{\mathcal{R}, R_{[i]}wu \Rightarrow u : \langle i \rangle[i]\varphi, u : \langle i \rangle\neg\varphi, w : [i]\varphi, \Gamma} \text{ (Cut)} \right. \\
 \\
 \Pi_2 = \left\{ \frac{\frac{\mathcal{R}, R_{[i]}wu, R_{[i]}wv \Rightarrow u : [i]\varphi, v : \varphi, \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow w : [i]\varphi, u : [i]\varphi, \Gamma} ([i]) \quad \frac{\mathcal{R}, R_{[i]}wu \Rightarrow w : [i]\varphi, u : [i]\varphi, u : \langle i \rangle[i]\varphi, \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow w : [i]\varphi, u : \langle i \rangle[i]\varphi, \Gamma} \text{ (Wk)} \quad \frac{\mathcal{R}, R_{[i]}wu \Rightarrow w : [i]\varphi, u : [i]\varphi, u : \langle i \rangle[i]\varphi, \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow u : \langle i \rangle[i]\varphi, u : [i]\varphi, \Gamma} \text{ (Wk)}}{\mathcal{R}, R_{[i]}wu \Rightarrow u : \langle i \rangle[i]\varphi, u : [i]\varphi, \Gamma} \text{ (Cut)} \right. \\
 \\
 \Pi_3 = \left\{ \frac{\frac{\mathcal{R}, R_{[i]}wu, R_{[i]}uz, R_{[i]}uv, R_{[i]}zv \Rightarrow z : \langle i \rangle\neg\varphi, v : \varphi, v : \neg\varphi, \Gamma}{\mathcal{R}, R_{[i]}wu, R_{[i]}uz, R_{[i]}uv, R_{[i]}zv \Rightarrow z : \langle i \rangle\neg\varphi, v : \varphi, \Gamma} \text{ (Euc}_i)}{\mathcal{R}, R_{[i]}wu, R_{[i]}uz, R_{[i]}uv \Rightarrow z : \langle i \rangle\neg\varphi, v : \varphi, \Gamma} ([i]) \cdot 2 \quad \frac{\mathcal{R}, R_{[i]}wu, R_{[i]}uz, R_{[i]}uv, R_{[i]}zv \Rightarrow z : \langle i \rangle\neg\varphi, v : \varphi, v : \neg\varphi, \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow u : [i]\langle i \rangle\neg\varphi, u : [i]\varphi, \Gamma} \text{ (Wk)}}{\mathcal{R}, R_{[i]}wu \Rightarrow u : [i]\langle i \rangle\neg\varphi, u : [i]\varphi, \Gamma} \text{ (Cut)} \right. \\
 \\
 \frac{\Pi_3 \quad \Pi_2}{\mathcal{R}, R_{[i]}wu \Rightarrow u : [i]\varphi, \Gamma} \text{ (Cut)}
 \end{array}$$

The admissibility of $(\langle i \rangle^*)$ follows from the hp-admissibility of (Wk) and cut admissibility.

$$\Pi_1 = \left\{ \frac{\frac{\mathcal{R}, R_{[i]}wu \Rightarrow w : \langle i \rangle \varphi, u : \langle i \rangle \varphi, u : \varphi, \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow w : \langle i \rangle \varphi, u : \langle i \rangle \varphi, \Gamma} (\langle i \rangle)}{\frac{\mathcal{R}, R_{[i]}wu \Rightarrow w : \langle i \rangle \varphi, w : \langle i \rangle \langle i \rangle \varphi, u : \langle i \rangle \varphi, \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow w : \langle i \rangle \varphi, w : \langle i \rangle \langle i \rangle \varphi, \Gamma} (\langle i \rangle)} (\text{Wk}) \right.$$

$$\Pi_2 = \left\{ \frac{\frac{\frac{\mathcal{R}, R_{[i]}wu, R_{[i]}wv', R_{[i]}wu', R_{[i]}u'v' \Rightarrow v' : \neg \varphi, v' : \varphi, w : \langle i \rangle \varphi, \Gamma}{\mathcal{R}, R_{[i]}wu, R_{[i]}wv', R_{[i]}wu', R_{[i]}u'v' \Rightarrow v' : \neg \varphi, w : \langle i \rangle \varphi, \Gamma} (\text{Tra}_i)}{\frac{\mathcal{R}, R_{[i]}wu, R_{[i]}wu', R_{[i]}u'v' \Rightarrow v' : \neg \varphi, w : \langle i \rangle \varphi, \Gamma}{\mathcal{R}, R_{[i]}wu \Rightarrow w : [i][i] \neg \varphi, w : \langle i \rangle \varphi, \Gamma} ([i]) \cdot 2} (\text{Theorem 7}) \right.$$

$$\frac{\frac{\Pi_1}{\mathcal{R} \Rightarrow w : \langle i \rangle \varphi, \Gamma} \quad \frac{\Pi_2}{\mathcal{R} \Rightarrow w : \langle i \rangle \varphi, \Gamma}}{\mathcal{R} \Rightarrow w : \langle i \rangle \varphi, \Gamma} (\text{Cut})$$

□

We now introduce the fundamental component of our loop-checking mechanism, i.e. the notion of a *generation tree*. A generation tree is a graph that tracks how certain labels are introduced within a thread, and can be used to check if a label w has an ancestor u associated with the same set of formulae (in which case we say that w is *blocked*; see Definition 16 below). In such a case, the $([i])$ rule need not be applied to the leaf of the path containing w and u in the tree. As discussed in Lemma 25 below, this loop-checking mechanism bounds the depth of the generation tree, ultimately securing terminating proof-search.

Definition 14 (Generation Tree). *Let $\mathcal{T} = (\mathcal{R}_i \Rightarrow \Gamma_i)_{i=0}^h$ be a thread from $\emptyset \Rightarrow w : \varphi$ to Λ . We define the generation tree $G_\Lambda^\mathcal{T} = (V, E)$ as follows:*

- $w \in V$;
- if (\square) is applied bottom-up in \mathcal{T} with $w : \square \varphi$ principal and $u : \varphi$ auxiliary, then add u to V and (w, u) to E ;
- if (\otimes_i) is applied bottom-up in \mathcal{T} with $w : \otimes_i \varphi$ principal, and $I_{\otimes_i} u$ and $u : \varphi$ auxiliary, then add u to V and (w, u) to E ;
- if $(D2_i)$ is applied bottom-up in \mathcal{T} with $I_{\otimes_i} u$ auxiliary, add u to V and (w, u) to E ;
- if $([i])$ is applied bottom-up in \mathcal{T} with $u : [i] \varphi$ principal and $v : \varphi$ auxiliary, then add v to V and (u, v) to E ;
- if $([i]^*)$ is applied bottom-up in \mathcal{T} with $R_{[i]}uv$ and $v : [i] \varphi$ principal, and $R_{[i]}uz$ and $z : \varphi$ auxiliary, then add z to V and (u, z) to E .⁴

Note that in the (\square) and (\otimes_i) cases, the edge (w, u) added to E always ensures that u is a child of the root w . The significance of this is explained later on in Remark 19.

Lemma 15. *Let $\mathcal{T} = (\mathcal{R}_i \Rightarrow \Gamma_i)_{i=0}^h$ be a thread from $\emptyset \Rightarrow w : \varphi$. The generation tree $G_\Lambda^\mathcal{T} = (V, E)$ is a tree.*

4. We note that u must already occur in V , so we are indeed permitted to add (u, z) to E (see Lemma 15).

Proof. We prove the result by induction on the height of the thread \mathcal{T} .

Base case. If $h = 0$, then $\mathcal{T} = \emptyset \Rightarrow w : \varphi$, implying that $G_\Lambda^\mathcal{T} = (V, E)$ with $V = \{w\}$ and $E = \emptyset$, which is a tree.

Inductive step. Let $\mathcal{T} = (\mathcal{R}_i \Rightarrow \Gamma_i)_{i=0}^{h+1}$ be a thread to Λ with $\mathcal{T}' = (\mathcal{R}_i \Rightarrow \Gamma_i)_{i=0}^h$ a thread to Λ' . By IH, we know that $G_{\Lambda'}^{\mathcal{T}'} = (V', E')$ is a tree. If Λ' was obtained from Λ by a rule other than (\Box) , (\otimes_i) , $(D2_i)$, $([i])$, or $([i]^*)$ then $G_\Lambda^\mathcal{T} = G_{\Lambda'}^{\mathcal{T}'}$ by Definition 14. Alternatively, if Λ' was obtained from Λ by (\Box) , (\otimes_i) , $(D2_i)$, or $([i])$, then due to the eigenvariable condition imposed on each rule, $G_\Lambda^\mathcal{T} = (V, E)$ where $V = V' \cup \{v\}$ and $E = E' \cup \{(u, v)\}$ with $u \in V'$ and v fresh. Last, concerning the $([i]^*)$ case, recall that the (IOA) rule is never applied to an $\text{IOA}^\mathcal{T}$ -label in a thread \mathcal{T} by Definition 10 and 11. As such, if $([i]^*)$ is bottom-up applied to the sequent $\mathcal{R}, R_{[i]}uv \Rightarrow v : [i]\varphi, \Gamma$ to obtain the sequent $\mathcal{R}, R_{[i]}uv, R_{[i]}uz \Rightarrow v : [i]\varphi, z : \varphi, \Gamma$, then u cannot be an $\text{IOA}^\mathcal{T}$ -label, meaning, either $u = w$ or the label was introduced by a prior bottom-up application of (\Box) , (\otimes_i) , $(D2_i)$, $([i])$, or $([i]^*)$. Therefore, $u \in V'$ and $G_\Lambda^\mathcal{T} = (V, E)$ where $V = V' \cup \{z\}$ and $E = E' \cup \{(u, z)\}$ with z fresh. Hence, in the (\Box) , (\otimes_i) , $(D2_i)$, $([i])$, and $([i]^*)$ cases, $G_\Lambda^\mathcal{T}$ is nothing more than $G_{\Lambda'}^{\mathcal{T}'}$ with a new edge protruding from a vertex to a fresh node, showing that $G_\Lambda^\mathcal{T}$ is indeed a tree. \square

We adopt a notion of blocking in our setting, used to bound the number of labels introduced during proof-search. It is based on the work of Horrocks and Sattler (2004), who introduced loop-checking by means of (in)direct blocking in tableaux for description logics, and Tiu, Ianovski, and Goré (2012), who applied a loop-checking mechanism for grammar logics in the context of nested sequent calculi. Both methods work in fundamentally the same way: in the former, a tableau is constructed, which is a tree whose nodes are sets of description logic formulae, and in the latter, a proof of nested sequents is constructed, where each nested sequent is a tree whose nodes are sets of grammar logic formulae. In both cases, the trees ‘grow’ throughout proof-search and when two nodes are encountered along the same branch with the same set of formulae, one stops growing the branch, thus limiting the size of the tableau or nested sequent. In our setting, although our labeled sequents do not form a tree, we may leverage generation trees to employ a similar mechanism.

Definition 16 (Blocking). *Let $\Lambda = \mathcal{R} \Rightarrow \Gamma$ be a sequent and $\mathcal{T} = (\mathcal{R}_i \Rightarrow \Gamma_i)_{i=0}^h$ be a thread from $\emptyset \Rightarrow w : \varphi$ to Λ . We say that u is directly blocked by its proper ancestor v in the generation tree $G_\Lambda^\mathcal{T}$ iff (i) v is not the root of $G_\Lambda^\mathcal{T}$, (ii) $I_{\otimes_i}u \in \mathcal{R}$ iff $I_{\otimes_i}v \in \mathcal{R}$, for all $i \in \text{Ag}$, and (iii) $\Gamma \upharpoonright u = \Gamma \upharpoonright v$. We refer to such a u that is directly blocked as a loop node, and to such a v as a loop ancestor. We say that u is indirectly blocked iff it has a proper ancestor in the generation tree $G_\Lambda^\mathcal{T}$ that is directly blocked. Last, we say that u is blocked in Λ iff it is directly or indirectly blocked in the generation tree $G_\Lambda^\mathcal{T}$, and that u is unblocked otherwise.*

Remark 17. *By the definition of a generation tree (Definition 14), no label introduced via a bottom-up application of (IOA) (i.e. an $\text{IOA}^\mathcal{T}$ -label) will occur as a node in the generation tree. Hence, all $\text{IOA}^\mathcal{T}$ -labels will be considered unblocked by Definition 16 above.*

Next, we introduce *saturation conditions*, used to determine when the remaining rules of G3DS_n^k are no longer applicable during proof-search and ‘enough’ information has been introduced to ensure that a counter-model can be extracted in the case of failure. Our conditions are motivated by conditions enforced in proof-search algorithms for description logics (Horrocks & Sattler, 2004) and multi-modal grammar logics (Tiu et al., 2012).

Definition 18 (Saturation Conditions). *Let \mathcal{T} be a thread to Λ . We define the saturation conditions accordingly:*

- (**C_{id}**) *for each $w \in \text{Lab}(\Lambda)$, if $w : \varphi \in \Gamma$, then $w : \neg\varphi \notin \Gamma$;*
- (**C_v**) *for each $w \in \text{Lab}(\Lambda)$, if $w : \varphi \vee \psi \in \Gamma$, then $w : \varphi, w : \psi \in \Gamma$;*
- (**C_^**) *for each $w \in \text{Lab}(\Lambda)$, if $w : \varphi \wedge \psi \in \Gamma$, then either $w : \varphi$ or $w : \psi \in \Gamma$;*
- (**C_◇**) *for each $w, u \in \text{Lab}(\Lambda)$, if $w : \diamond\varphi \in \Gamma$, then $u : \varphi \in \Gamma$;*
- (**C_□**) *for each $w \in \text{Lab}(\Lambda)$, if $w : \square\varphi \in \Gamma$, then there exists an unblocked u in the generation tree $G_\Lambda^\mathcal{T}$ such that $u : \varphi \in \Gamma$;*
- (**C_{\langle i \rangle}**) *for each $w, u \in \text{Lab}(\Lambda)$ and $i \in \text{Ag}$, if $w : \langle i \rangle\varphi \in \Gamma$ and $R_{[i]}wu \in \mathcal{R}$, then $u : \varphi, u : \langle i \rangle\varphi \in \Gamma$;*
- (**C_[i]**) *for each $w \in \text{Lab}(\Lambda)$ and $i \in \text{Ag}$, if $w : [i]\varphi \in \Gamma$ with w unblocked in the generation tree $G_\Lambda^\mathcal{T}$, then there exists a $u \in \text{Lab}(\Lambda)$ such that $R_{[i]}wu \in \mathcal{R}$ and $u : \varphi \in \Gamma$;*
- (**C_{ref}**) *for each $w \in \text{Lab}(\Lambda)$, $R_{[i]}ww \in \mathcal{R}$;*
- (**C_{euc}**) *for each $w, u, v \in \text{Lab}(\Lambda)$, if $R_{[i]}wu, R_{[i]}wv \in \mathcal{R}$, then $R_{[i]}uv \in \mathcal{R}$;*
- (**C_{⊖_i}**) *for each $w, u \in \text{Lab}(\Lambda)$ and $i \in \text{Ag}$, if $w : \ominus_i\varphi \in \Gamma$ and $I_{\otimes_i}u \in \mathcal{R}$, then $u : \varphi \in \Gamma$;*
- (**C_{⊗_i}**) *for each $w \in \text{Lab}(\Lambda)$ and $i \in \text{Ag}$, if $w : \ominus_i\varphi \in \Gamma$, then there exists an unblocked u in the generation tree $G_\Lambda^\mathcal{T}$ such that $I_{\otimes_i}u \in \mathcal{R}$ and $u : \varphi \in \Gamma$;*
- (**C_{D2}**) *for each $i \in \text{Ag}$ there exists an unblocked w in the generation tree $G_\Lambda^\mathcal{T}$ such that $I_{\otimes_i}w \in \mathcal{R}$;*
- (**C_{D3}**) *for each $w, u \in \text{Lab}(\Lambda)$ and $i \in \text{Ag}$, if $I_{\otimes_i}w, R_{[i]}wu \in \mathcal{R}$, then $I_{\otimes_i}u \in \mathcal{R}$;*
- (**C_{APC}**) *for each $w_0, \dots, w_k \in \text{Lab}(\Lambda)$ and $i \in \text{Ag}$, $R_{[i]}w_jw_m \in \mathcal{R}$ for some $j, m \in \{0, \dots, k\}$ with $j \neq m$.*

We refer to Λ as *stable* iff it satisfies all above conditions and every $\text{IOA}^\mathcal{T}$ -tuple of the sequent is $\text{IOA}^\mathcal{T}$ -satisfied.

Our proof-search algorithm Prove_n^k is shown in Algorithm 1 below, and decides if a formula is DS_n^k -valid, meaning we take n and k to be fixed, input parameters. Prove_n^k is split into three parts due to its length. We note that the algorithm applies rules from G3DS_n^k and admissible rules from Figure 4 bottom-up on an input of the form $\emptyset \Rightarrow w_0 : \varphi$ with $\varphi \in \mathcal{L}_n$. Each recursive call of the algorithm corresponds to a bottom-up application of a rule, and at any given step, Prove_n^k has generated a finite number of threads from $\emptyset \Rightarrow w_0 : \varphi$ that form a ‘partial proof’ of the input. As we will see below, if the input is valid, then eventually a proof in G3DS_n^k may be extracted from successful proof-search, and if the input is invalid, then a counter-model may be extracted from failed proof-search. We also note that if $k = 0$, i.e. we do not bound the number of choices available to an agent, then we omit lines 15–20 and do not enforce the saturation condition (**C_{APC}**).

The saturation conditions, along with $\text{IOA}^{\mathcal{T}}$ -satisfiability, are used to guide the computation of Prove_n^k . Such conditions serve a twofold role: first, they are used to determine the (in)applicability of a rule during proof-search, thus letting us control when rules are bottom-up applied, and yielding termination of the algorithm (Theorem 26). Second, such conditions ensure that enough information has entered into a sequent during proof-search to be able to extract a counter-model in the case where the input formula is invalid (Theorem 23). In particular, if a stable sequent is encountered during proof-search, then it can be transformed into a finite DS_n^k -model (see Definition 20 and Lemma 22 below). We remark that it is decidable to check if a sequent (which is a finite object) is stable.

The proof-search strategy of our algorithm Prove_n^k is, intuitively, as follows: First, Prove_n^k is split into three parts, namely, lines 5-20, lines 21-40, and lines 41-60. Lines 5-20 bottom-up apply the inference rules (Ref_i) , (Euc_i) , (D3_i) , and (APC_i^k) , which only affect relational atoms. Lines 21-40 bottom-up apply the inference rules (\vee) , (\wedge) , (\diamond) , $([i]^*)$, and (\ominus_i) , which only affect labeled formulae. All rules in lines 5-20 and 21-40 are *non-generating* in the sense that they do not introduce new labels. Conversely, lines 41-60 apply the *generating* rules (\square) , (\otimes_i) , (D2_i) , $([i])$, $([i]^*)$, and IoaOp (i.e. a sequence of (IOA) applications) that bottom-up introduce new, fresh labels to sequents. We note that (\square) and (\otimes_i) are preceded by applications of (\square^*) and (\otimes_i^*) , respectively, as discussed in more detail in the remark below. Since all rules of G3DS_n^k are hp-invertible (Theorem 7-3), we can bottom-up apply our rules in any order. However, we find it conceptually easier to prove facts about Prove_n^k by grouping rules of a similar functionality together.

Remark 19. *As it is important for understanding Prove_n^k , we discuss our use of the admissible rules (\square^*) , (\otimes_i^*) , and $([i]^*)$. Lines 41-43 bottom-up apply (\square^*) followed by (\square) and lines 44-47 bottom-up apply (\otimes_i^*) followed by (\otimes_i) . In each case, the label w of $w : \square\psi$ and $w : \otimes_i\psi$ is changed (via the (\square^*) and (\otimes_i^*) rules) to the ‘source’ label w_0 before bottom-up applying (\square) and (\otimes_i) , respectively. (NB. We always assume that the input to Prove_n^k is a sequent of the form $\emptyset \Rightarrow w_0 : \varphi$, meaning w_0 will be the root of every generation tree constructed throughout proof-search.) This has the effect that each label introduced by the bottom-up application of the subsequent (\square) or (\otimes_i) is a child of the root of the generation tree. Observe that by the definition of a blocked node (Definition 16), that a node in the generation tree cannot be blocked by the root of the generation tree. This means that fresh labels generated by (\square) and (\otimes_i) are unblocked by definition, which is crucial for constructing a counter-model when proof-search fails (see Definition 20 and Lemma 22 below) as our counter-model is built using only unblocked labels. Hence, labels generated by bottom-up applications of (\square) and (\otimes_i) will always exist within our counter-model, which will be vital in proving the correctness of Prove_n^k (see Theorem 23 below).*

Lines 52-57 bottom-up apply $([i])$ to unblocked non- $\text{IOA}^{\mathcal{T}}$ -labels and $([i]^)$ to $\text{IOA}^{\mathcal{T}}$ -labels. Applying $([i])$ to only unblocked non- $\text{IOA}^{\mathcal{T}}$ -labels limits the growth of sequents during proof-search since repeated bottom-up applications of $([i])$ will eventually produce blocked non- $\text{IOA}^{\mathcal{T}}$ -labels as discussed in the proof of termination (see Theorem 26 below). Our reason for the use of $([i]^*)$ however, is that since if a sequent $\mathcal{R}, R_{[i]}wu \Rightarrow \Gamma$ is encountered during proof-search with $u : [i]\varphi \in \Gamma$ and u an $\text{IOA}^{\mathcal{T}}$ -label, then applying $([i]^*)$ lets us ‘unpack’ $u : [i]\varphi$ and obtain $\mathcal{R}, R_{[i]}wu, R_{[i]}v \Rightarrow \Gamma, v : \varphi$ with v fresh rather than $\mathcal{R}, R_{[i]}wu, R_{[i]}uv \Rightarrow \Gamma, v : \varphi$. This prevents chains of relational atoms from being introduced that protrude outward from*

Algorithm 1: Prove_n^k

Input: A Labeled Sequent: $\Lambda = \mathcal{R} \Rightarrow \Gamma$

Output: A Boolean: True, False

```

1 if  $w : \varphi, w : \neg\varphi \in \Gamma$  then
2   | return True;
3 if  $\mathcal{R} \Rightarrow \Gamma$  is stable then
4   | return False;
5 if for some  $w \in \text{Lab}(\Lambda), R_{[i]}ww \notin \mathcal{R}$  then
6   | Let  $\mathcal{R}' := R_{[i]}ww, \mathcal{R}$ ;
7   | return  $\text{Prove}_n^k(\mathcal{R}' \Rightarrow \Gamma)$ 
8 if for some  $w, u, v \in \text{Lab}(\Lambda), R_{[i]}wu, R_{[i]}wv \in \mathcal{R}$ , but  $R_{[i]}uv \notin \mathcal{R}$  then
9   | Let  $\mathcal{R}' := R_{[i]}uv, \mathcal{R}$ ;
10  | return  $\text{Prove}_n^k(\mathcal{R}' \Rightarrow \Gamma)$ ;
11 if for some  $w, u \in \text{Lab}(\Lambda), I_{\otimes_i}w, R_{[i]}wu \in \mathcal{R}$ , but  $I_{\otimes_i}u \notin \mathcal{R}$  then
12  | Let  $\mathcal{R}' := I_{\otimes_i}u, \mathcal{R}$ ;
13  | return  $\text{Prove}_n^k(\mathcal{R}' \Rightarrow \Gamma)$ ;
14 if for some  $w_0, \dots, w_k \in \text{Lab}(\Lambda), R_{[i]}w_jw_m \notin \mathcal{R}$  for each  $j, m \in \{0, \dots, k\}$  with
    |  $j \neq m$  then
15  | Let  $\mathcal{R}_{m,j} := R_{[i]}w_mw_j, \mathcal{R}$  (with  $0 \leq m \leq k-1$  and  $m+1 \leq j \leq k$ );
16  | if  $\text{Prove}_n^k(\mathcal{R}_{m,j} \Rightarrow \Gamma) = \text{False}$  for some  $m$  and  $j$  then
17  |   | return False;
18  | else
19  |   | return True;
20  | end

```

IOA^T-labels. In fact, none of the generating rules are ever applied to an IOA^T-label, meaning such labels can be ignored in our generation tree, and thus do not give rise to infinite loops during proof-search, as they never give rise to fresh labels during proof-search.

We now explain how a counter-model (which we call a *stability model*) can be obtained from failed proof-search, and then show the correctness of the proof-search algorithm.

Definition 20 (Stability Model). *Let \mathcal{T} be a thread from $\emptyset \Rightarrow w_0 : \varphi$ to $\Lambda := \mathcal{R} \Rightarrow \Gamma$ generated by $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi)$. We define the stability model relative to Λ to be $M^\Lambda := (W^\Lambda, \{R_{[i]}^\Lambda \mid i \in \text{Ag}\}, \{I_{\otimes_i}^\Lambda \mid i \in \text{Ag}\}, V^\Lambda)$ such that*

- W^Λ is the set of unblocked labels in Λ ;
- $R_{[i]}^\Lambda$ is the Euclidean closure of the set $S_{[i]}^\Lambda$, where $(w, u) \in S_{[i]}^\Lambda$ iff (i) $R_{[i]}wu \in \mathcal{R}$ with w and u both unblocked in G_Λ^T , or (ii) $R_{[i]}wv \in \mathcal{R}$ with w unblocked, and v directly blocked by u in G_Λ^T .⁵

5. Since Λ is stable, we know it satisfies the saturation condition (**C_{ref}**), meaning $(w, w) \in S_{[i]}^\Lambda$ for each $w \in W^\Lambda$. Hence, taking the Euclidean closure of $S_{[i]}^\Lambda$ implies that $R_{[i]}^\Lambda$ will be both reflexive and Euclidean, i.e. it will be an equivalence relation.

Algorithm 1: Prove_n^k (Continued)

```

21 if there exists a  $w : \varphi \vee \psi \in \Gamma$ , but either  $w : \varphi \notin \Gamma$  or  $w : \psi \notin \Gamma$  then
22   | Let  $\Gamma' := w : \varphi, w : \psi, \Gamma$ ;
23   | return Provenk( $\mathcal{R} \Rightarrow \Gamma'$ );
24 if there exists a  $w : \varphi \wedge \psi \in \Gamma$ , but  $w : \varphi, w : \psi \notin \Gamma$  then
25   | Let  $\Gamma_1 := w : \varphi, \Gamma$ ;
26   | Let  $\Gamma_2 := w : \psi, \Gamma$ ;
27   | if Provenk( $\mathcal{R} \Rightarrow \Gamma_i$ ) = False for some  $i \in \{1, 2\}$  then
28   |   | return False;
29   | else
30   |   | return True;
31   | end
32 if  $w : \diamond\varphi \in \Gamma$ , but  $u : \varphi \notin \Gamma$  for some  $u \in \text{Lab}(\Lambda)$  then
33   | Let  $\Gamma' := u : \varphi, \Gamma$ ;
34   | return Provenk( $\mathcal{R} \Rightarrow \Gamma'$ );
35 if  $w : \langle i \rangle\varphi \in \Gamma$ , but  $u : \varphi, u : \langle i \rangle\varphi \notin \Gamma$  for some  $u \in \text{Lab}(\Lambda)$  such that  $R_{[i]}wu \in \mathcal{R}$ 
    then
36   | Let  $\Gamma' := u : \varphi, u : \langle i \rangle\varphi, \Gamma$ ;
37   | return Provenk( $\mathcal{R} \Rightarrow \Gamma'$ );
38 if  $w : \ominus_i\varphi \in \Gamma$ , but  $u : \varphi \notin \Gamma$  for some  $u \in \text{Lab}(\Lambda)$  such that  $I_{\otimes_i}u \in \mathcal{R}$  then
39   | Let  $\Gamma' := u : \varphi, \Gamma$ ;
40   | return Provenk( $\mathcal{R} \Rightarrow \Gamma'$ );
    
```

- $I_{\otimes_i}^\Lambda = \{w \mid I_{\otimes_i}w \in \mathcal{R} \text{ and } w \text{ is unblocked.}\};$
- $w \in V^\Lambda(p)$ iff $w : \neg p \in \Gamma$ and w is unblocked.

Lemma 21. Let $\Lambda := \mathcal{R} \Rightarrow \Gamma$ be a stable sequent with $w, u \in W^\Lambda$. If $w \sim_i^{\mathcal{R}} u$, then $R_{[i]}^\Lambda wu$.

Proof. If $w \sim_i^{\mathcal{R}} u$, then since Λ is stable, and satisfies both **(C_{ref})** and **(C_{euc})**, it follows that $R_{[i]}^\Lambda wu$. \square

We now show that any stability model is indeed a DS_n^k-model of Definition 2.

Lemma 22. Let \mathcal{T} be a thread from $\emptyset \Rightarrow w_0 : \varphi$ to the stable sequent $\Lambda := \mathcal{R} \Rightarrow \Gamma$ generated by Prove_n^k($\emptyset \Rightarrow w_0 : \varphi$) = **false**. Then, the stability model M^Λ is a finite DS_n^k-model.

Proof. First, since Λ is a stable sequent generated from Prove_n^k($\emptyset \Rightarrow w_0 : \varphi$), $w_0 \in W^\Lambda$ ensures its non-emptiness. We show that the model satisfies **(C1)–(C3)** and **(D1)–(D3)**.

- (C1)** To prove that **(C1)** holds, we show that for all $i \in \text{Ag}$, $R_{[i]}^\Lambda \subseteq W^\Lambda \times W^\Lambda$ is an equivalence relation. By Definition 20, we know that $R_{[i]}^\Lambda$ is the Euclidean closure of the $S_{[i]}^\Lambda$ relation. Since Λ is stable, we know that it satisfies the saturation condition **(C_{ref})**, implying that $R_{[i]}^\Lambda$ is reflexive and Euclidean, and so, is an equivalence relation.

Algorithm 1: Prove_n^k (Continued)

```

41 if  $w : \Box\varphi \in \Gamma$ , but  $u : \varphi \notin \Gamma$  for all  $u \in \text{Lab}(\Lambda)$  then
42   | Replace  $w : \Box\varphi \in \Gamma$  with  $w_0 : \Box\varphi$  and let  $\Gamma' := v : \varphi, \Gamma$  with  $v$  fresh;
43   | return Provenk( $\mathcal{R} \Rightarrow \Gamma'$ );
44 if  $w : \otimes_i\varphi \in \Gamma$ , but  $u : \varphi \notin \Gamma$  for all  $u \in \text{Lab}(\Lambda)$  with  $I_{\otimes_i}u \in \mathcal{R}$  then
45   | Replace  $w : \otimes_i\varphi \in \Gamma$  with  $w_0 : \otimes_i\varphi$ ;
46   | Let  $\mathcal{R}' := I_{\otimes_i}v, \mathcal{R}$  and  $\Gamma' := v : \varphi, \Gamma$  with  $v$  fresh;
47   | return Provenk( $\mathcal{R}' \Rightarrow \Gamma'$ );
48 if for all  $w \in \text{Lab}(\Lambda)$ ,  $I_{\otimes_i}w \notin \mathcal{R}$  then
49   | Let  $\mathcal{R}' := I_{\otimes_i}u, \mathcal{R}$  with  $u$  fresh;
50   | return Provenk( $\mathcal{R}' \Rightarrow \Gamma$ );
51 if  $w : [i]\varphi \in \Gamma$  with  $w$  unblocked in  $G_\Lambda^\top$ , but  $u : \varphi \notin \Gamma$  for all  $u \in \text{Lab}(\Lambda)$  such that
     $R_{[i]}wu \in \mathcal{R}$  then
52   | if  $w$  is not an IOAT-label then
53     | Let  $\mathcal{R}' := R_{[i]}wv, \mathcal{R}$  and  $\Gamma' := v : \varphi, \Gamma$  with  $v$  fresh;
54     | return Provenk( $\mathcal{R}' \Rightarrow \Gamma'$ );
55   | if  $w$  is an IOAT-label then
56     | // As  $w$  is an IOAT-label,  $\mathcal{R}$  is of the form  $R_{[i]}vw, \mathcal{R}'$ .
57     | Let  $\mathcal{R}'' := R_{[i]}vz, R_{[i]}vw, \mathcal{R}'$  and  $\Gamma' := z : \varphi, \Gamma$  with  $z$  fresh;
58     | return Provenk( $\mathcal{R}'' \Rightarrow \Gamma'$ );
59 if for some non-IOAT-labels  $w_1, \dots, w_n \in \text{Lab}(\Lambda)$ , there does not exist an IOAT-label
     $u \in \text{Lab}(\Lambda)$  such that  $w_i \sim_i^{\mathcal{R}} u$  for each  $i \in \text{Ag}$  then
60   | Let  $\Lambda' := \text{IoaOp}(\Lambda)$ ;
61   | return Provenk( $\Lambda'$ );

```

(C2) Suppose that $w_1, \dots, w_n \in W^\Lambda$. Since Λ is saturated, we know that it is IOA^T-satisfied, meaning that an IOA^T-label $u \in \text{Lab}(\Lambda)$ exists such that $w_i \sim_i^{\mathcal{R}} u$ for each $i \in \text{Ag}$. By Remark 17, we know that u is unblocked, and so, $u \in W^\Lambda$. By the (C_{ref}) and (C_{auc}) saturation conditions, it follows from the fact that $w_i \sim_i^{\mathcal{R}} u$ for each $i \in \text{Ag}$ that $R_{[i]}w_i u \in \mathcal{R}$ for each $i \in \text{Ag}$. Since $w_1, \dots, w_n \in W^\Lambda$, each label is unblocked by Definition 20, and as mentioned above, u is unblocked as well, implying that $R_{[i]}^\Lambda w_i u$ for each $i \in \text{Ag}$. Thus, (C2) is satisfied in this case.

(C3) Suppose that $w_0, \dots, w_k \in W^\Lambda$ and let $i \in \text{Ag}$. Since Λ is stable and satisfies (C_{APC}), we know that $R_{[i]}w_j w_m \in \mathcal{R}$ holds for some $j, m \in \{0, \dots, k\}$. By definition, $R_{[i]}^\Lambda w_j w_m$ holds, implying the desired result.

(D1) Follows immediately from the definition of $I_{\otimes_i}^\Lambda$.

(D2) Let $i \in \text{Ag}$. Since Λ is stable, we know that Λ satisfies (C_{D2}), meaning that there exists an unblocked $w \in W^\Lambda$ such that $I_{\otimes_i}w \in \mathcal{R}$. Hence, by the definition of $I_{\otimes_i}^\Lambda$, $w \in I_{\otimes_i}^\Lambda$, implying that $I_{\otimes_i}^\Lambda$ is non-empty.

(D3) Let $i \in Ag$, $w, u \in W^\Lambda$, and suppose that $w \in I_{\otimes_i}^\Lambda$ and $R_{[i]}^\Lambda wu$. Let us define $v \sim_i^S z$ iff (v, z) or $(z, v) \in S_{[i]}^\Lambda$. (NB. See Definition 20 above for the definition of $S_{[i]}^\Lambda$.) We prove the claim by induction on the minimal length of an undirected path $w \sim_i^S v_1, \dots, v_n \sim_i^S u$ from w to u in $S_{[i]}^\Lambda$ and omit the case when the path is of length zero (i.e. $w = u$) as the case is trivial.

Base case. Suppose that $(w, u) \in S_{[i]}^\Lambda$; the case when $(u, w) \in S_{[i]}^\Lambda$ is shown similarly. We have two cases to consider: either (i) $R_{[i]}^\Lambda wu \in \mathcal{R}$, or (ii) $R_{[i]}^\Lambda wu \notin \mathcal{R}$.

(i) If $R_{[i]}^\Lambda wu \in \mathcal{R}$, then since $w \in I_{\otimes_i}^\Lambda$, we know that $I_{\otimes_i} w \in \mathcal{R}$ by Definition 20. By the saturation condition (**C_{D3}**), it follows that $I_{\otimes_i} u \in \mathcal{R}$, and since $u \in W^\Lambda$, we know that u is unblocked, implying that $u \in I_{\otimes_i}^\Lambda$.

(ii) If $R_{[i]}^\Lambda wu \notin \mathcal{R}$, then $(w, u) \in S_{[i]}^\Lambda$ must hold due to clause (ii) of the definition of $S_{[i]}^\Lambda$ (see Definition 20 above). Therefore, there is a loop node v with u its loop ancestor such that $R_{[i]}^\Lambda wv \in \mathcal{R}$. Since $w \in I_{\otimes_i}^\Lambda$, we know that $I_{\otimes_i} w \in \mathcal{R}$, implying that $I_{\otimes_i} v \in \mathcal{R}$ by the saturation condition (**C_{D3}**). By Definition 16, it follows that $I_{\otimes_i} u \in \mathcal{R}$ since u is the loop ancestor of v . Since u is unblocked, we have that $u \in I_{\otimes_i}^\Lambda$.

Inductive step. Let $w \sim_i^S v_1, \dots, v_n \sim_i^S u$ be a minimal undirected path between w and u whose length is $n + 1$. By IH, we know that $v_n \in I_{\otimes_i}^\Lambda$, implying that $I_{\otimes_i} v_n \in \mathcal{R}$ since v_n is unblocked. (NB. Every label participating in $S_{[i]}^\Lambda$ and $R_{[i]}^\Lambda$ is unblocked by definition.) Suppose that $(u, v_n) \in S_{[i]}^\Lambda$; the case when $(v_n, u) \in S_{[i]}^\Lambda$ is shown similarly. We have two cases to consider: either (i) $R_{[i]}^\Lambda uv_n \in \mathcal{R}$, or (ii) $R_{[i]}^\Lambda uv_n \notin \mathcal{R}$.

(i) If $R_{[i]}^\Lambda uv_n \in \mathcal{R}$, then by conditions (**C_{ref}**) and (**C_{auc}**), we know $R_{[i]}^\Lambda v_n u \in \mathcal{R}$. We may conclude that $u \in I_{\otimes_i}^\Lambda$ by an argument similar to case (i) of the base case above.

(ii) Suppose that $R_{[i]}^\Lambda uv_n \notin \mathcal{R}$. Then, $(u, v_n) \in S_{[i]}^\Lambda$ must hold due to clause (ii) of the definition of $S_{[i]}^\Lambda$, implying the existence of a loop node z with v_n its loop ancestor such that $R_{[i]}^\Lambda uz \in \mathcal{R}$. Since $I_{\otimes_i} v_n \in \mathcal{R}$ and v_n is the loop ancestor of z , we have that $I_{\otimes_i} z \in \mathcal{R}$. By the fact that $R_{[i]}^\Lambda uz \in \mathcal{R}$ and by saturation conditions (**C_{ref}**) and (**C_{auc}**), we know that $R_{[i]}^\Lambda zu \in \mathcal{R}$, showing that $I_{\otimes_i} u \in \mathcal{R}$. It follows that $u \in I_{\otimes_i}^\Lambda$.

Last, we note that by condition (**C_{id}**), the valuation V^Λ is well-defined, and since M^Λ was extracted from Λ , which is a finite object, M^Λ will be finite. \square

We are now in the position to prove correctness.

Theorem 23 (Correctness). *Let $n, k \in \mathbb{N}$ and $\varphi \in \mathcal{L}_n$.*

- (i) *If $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi) = \mathbf{true}$, then there is a proof of $\emptyset \Rightarrow w_0 : \varphi$ in G3DS_n^k witnessing that φ is DS_n^k -valid.*
- (ii) *If $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi) = \mathbf{false}$, then there is finite DS_n^k -model witnessing that φ is DS_n^k -invalid.*

Proof. (i) Every instruction in Prove_n^k can be seen as a bottom-up application of a rule in G3DS_n^k or of an admissible rule from Figure 4 with **IoOp** equivalent to a series of bottom-up (IOA) applications. We can invoke Theorem 7-1 to obtain proofs of all top sequents.

(ii) Let us suppose that $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi) = \text{false}$. It follows that $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi)$ generated a thread $\mathcal{T} = (\mathcal{R}_i \Rightarrow \Gamma_i)_{i=0}^k$ from $\emptyset \Rightarrow w_0 : \varphi$ to $\Lambda := \mathcal{R} \Rightarrow \Gamma$ with Λ stable. We may construct the stability model M^Λ as specified in Definition 20. By Lemma 22, we know that M^Λ is a finite DS_n^k -model. Let us now prove that for any $u \in W^\Lambda$ with $u : \psi \in \Gamma$, $M^\Lambda, u \not\models \psi$. We show this by induction on the complexity of ψ and note that since $w_0 : \varphi \in \Gamma$, the above claim implies that $M^\Lambda, w_0 \not\models \varphi$, meaning that $w_0 : \varphi$ is not provable in G3DS_n^k by soundness (Theorem 7). We only consider the inductive step when ψ is of the form $\langle i \rangle \xi$, $\ominus_i \xi$, $[i] \xi$, and $\otimes_i \xi$ since all remaining cases are simple or similar.

$\langle i \rangle \xi$. Suppose $u : \langle i \rangle \xi \in \Gamma$ and assume that for some arbitrary $v \in W^\Lambda$, $R_{[i]}^\Lambda uv$. We aim to show that $M^\Lambda, v \Vdash \xi$, and have two cases to consider: either (i) $R_{[i]} uv \in \mathcal{R}$, or (ii) $R_{[i]} uv \notin \mathcal{R}$.

(i) If $R_{[i]} uv \in \mathcal{R}$, then since Λ is stable, we know that $v : \xi \in \Gamma$. Hence, by IH and the definition of $R_{[i]}^\Lambda$ we have that $M^\Lambda, v \Vdash \xi$.

(ii) Suppose that $R_{[i]} uv \notin \mathcal{R}$. We define $u \sim_i^S v$ iff (u, v) or $(v, u) \in S_{[i]}^\Lambda$, and prove that $v : \xi, v : \langle i \rangle \xi \in \Gamma$ by induction on the minimal length of an undirected path $u \sim_i^S z_1, \dots, z_n \sim_i^S v$ from u to v in $S_{[i]}^\Lambda$. (NB. See Definition 20 above for the definition of $S_{[i]}^\Lambda$.) Note that the path cannot be of length 0, since then $u = v$, implying that $R_{[i]} uu \notin \mathcal{R}$ by our assumption, which contradicts the fact that $R_{[i]} uu \in \mathcal{R}$ by the saturation condition (\mathbf{C}_{ref}).

Base case. Suppose that $(u, v) \in S_{[i]}^\Lambda$; the case when $(v, u) \in S_{[i]}^\Lambda$ is shown similarly. Then, $(u, v) \in S_{[i]}^\Lambda$ must hold due to clause (ii) of the definition of $S_{[i]}^\Lambda$ (Definition 20). Hence, there is a loop node z with v its loop ancestor such that $R_{[i]} uz \in \mathcal{R}$. Since $u : \langle i \rangle \xi \in \Gamma$, we know that $z : \xi, z : \langle i \rangle \xi \in \Gamma$, implying that $v : \xi, v : \langle i \rangle \xi \in \Gamma$ as v is the loop ancestor of z .

Inductive step. Let $u \sim_i^S z_1, \dots, z_n \sim_i^S v$ be a minimal undirected path between u and v whose length is $n + 1$. By IH, we know that $z_n : \xi, z_n : \langle i \rangle \xi \in \Gamma$. Suppose that $(v, z_n) \in S_{[i]}^\Lambda$; the case when $(z_n, v) \in S_{[i]}^\Lambda$ is shown similarly. We have two cases to consider: either (i) $R_{[i]} vz_n \in \mathcal{R}$, or (ii) $R_{[i]} vz_n \notin \mathcal{R}$.

(i*) If $R_{[i]} vz_n \in \mathcal{R}$, then by saturation conditions (\mathbf{C}_{ref}) and (\mathbf{C}_{euc}), we know $R_{[i]} z_n v \in \mathcal{R}$. We may conclude that $v : \xi, v : \langle i \rangle \xi \in \Gamma$ since Λ satisfies the saturation condition ($\mathbf{C}_{(i)}$).

(ii*) Suppose that $R_{[i]} vz_n \notin \mathcal{R}$. Then, $(v, z_n) \in S_{[i]}^\Lambda$ must hold due to clause (ii) of the definition of $S_{[i]}^\Lambda$, implying the existence of a loop node z with z_n its loop ancestor such that $R_{[i]} vz \in \mathcal{R}$. Since z_n is the loop ancestor of z , we have that $z : \xi, z : \langle i \rangle \xi \in \Gamma$. By the fact that $R_{[i]} vz \in \mathcal{R}$ and by saturation conditions (\mathbf{C}_{ref}) and (\mathbf{C}_{euc}), we know that $v : \xi, v : \langle i \rangle \xi \in \Gamma$.⁶

Hence, regardless of the length of a minimal path between u and v , we have that $v : \xi, v : \langle i \rangle \xi \in \Gamma$, showing that $M^\Lambda, v \Vdash \xi$ by IH. As v is arbitrary, may conclude that $M^\Lambda, u \not\models \langle i \rangle \xi$.

$\ominus_i \xi$. Suppose $u : \ominus_i \xi \in \Gamma$. Since Λ is stable we know that $v : \xi \in \Gamma$ for all v such that $I_{\otimes_i} v \in \mathcal{R}$. Hence, by IH and the definition of $I_{\otimes_i}^\Lambda$ we have that $M^\Lambda, v \not\models \xi$ for all $v \in W^\Lambda$ such that $I_{\otimes_i}^\Lambda v$, implying $M^\Lambda, u \not\models \ominus_i \xi$.

$[i] \xi$. Suppose $u : [i] \xi \in \Gamma$. Since Λ is stable, we know that there exists a $v \in \text{Lab}(\Lambda)$ such that $R_{[i]} uv \in \mathcal{R}$ and $v : \xi \in \Gamma$. There are two cases: (i) v is a loop node or (ii) v is not.

6. This case of our proof demonstrates the utility of the $(\langle i \rangle^*)$ rule. Since $z : \langle i \rangle \xi$ occurs in Γ , and $R_{[i]} vz \in \mathcal{R}$ implies $R_{[i]} zv \in \mathcal{R}$ by the saturation conditions (\mathbf{C}_{ref}) and (\mathbf{C}_{euc}), we know that $v : \xi, v : \langle i \rangle \xi \in \Gamma$. If we utilized the $(\langle i \rangle)$ rule rather than $(\langle i \rangle^*)$, we would not immediately have that $z : \langle i \rangle \xi \in \Gamma$.

(i) If v is a loop node, then there exists a loop ancestor z (which is therefore unblocked) such that $z : \xi \in \Gamma$ and $R_{[z]}^\Lambda uz$. Hence, by IH, we have that for some $z \in W^\Lambda$, $M^\Lambda, z \not\models \xi$.

(ii) By IH, the definition of W^Λ , and the definition of $R_{[z]}^\Lambda$, it follows that there exists a $v \in W^\Lambda$ such that $R_{[z]}^\Lambda uv$ and $M^\Lambda, v \not\models \xi$.

Hence, $M^\Lambda, u \not\models [z]\xi$.

$\otimes_i \xi$. Suppose $u : \otimes_i \xi \in \Gamma$. Since Λ is stable, we know that there exists an unblocked $v \in Lab(\Lambda)$ such that $I_{\otimes_i} v \in \mathcal{R}$ and $v : \xi \in \Gamma$. By IH, the definition of W^Λ , and the definition of $I_{\otimes_i}^\Lambda$, it follows that there exists a $v \in I_{\otimes_i}^\Lambda$ such that $I_{\otimes_i}^\Lambda v$ and $M^\Lambda, v \not\models \xi$. Hence, $M^\Lambda, u \not\models \otimes_i \xi$. \square

Let us now argue that Prove_n^k terminates regardless of its input $\emptyset \Rightarrow w_0 : \varphi$ for $\varphi \in \mathcal{L}_n$. We first define the set of subformulae of a given formula φ (Definition 24), which due to its finiteness limits the formulae that can occur during proof-search. We then prove that only a finite number of labels can be created during proof-search, implying that only a finite number of sequents can be generated, and yielding the termination of Prove_n^k .

Definition 24 (Subformula). *Let φ be a formula in \mathcal{L}_n . We define the set $\text{suf}(\varphi)$ of subformulae of φ recursively as follows:*

- $\text{suf}(\varphi) = \{\varphi\}$ for $\varphi \in \{p, \neg p \mid p \in \text{Var}\}$;
- $\text{suf}(\nabla\psi) = \{\nabla\psi\} \cup \text{suf}(\psi)$ for $\nabla \in \{\Box, \Diamond\} \cup \{\otimes_i, [i], \ominus_i, \langle i \rangle \mid i \in \text{Ag}\}$;
- $\text{suf}(\psi \circ \chi) = \text{suf}(\psi) \cup \text{suf}(\chi)$ for $\circ \in \{\wedge, \vee\}$.

Lemma 25. *Let $n, k \in \mathbb{N}$ and $\varphi \in \mathcal{L}_n$.*

1. *For each $\Box\psi \in \text{suf}(\varphi)$, (\Box^*) and (\Box) are applied bottom-up at most once in any thread during the computation of $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi)$;*
2. *For each $\otimes_i\psi \in \text{suf}(\varphi)$, (\otimes_i^*) and (\otimes_i) are applied bottom-up at most once in any thread during the computation of $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi)$;*
3. *For each $i \in \text{Ag}$, $(D2_i)$ is applied bottom-up at most once in any thread during the computation of $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi)$;*
4. *$([i])$ and $([i]^*)$ will be applied only a finite number of times in any thread of the computation of $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi)$;*
5. *IoaOp will be applied only a finite number of times in any thread of the computation of $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi)$.*

Proof. Claims 1 and 2 are straightforward since once (\Box) and (\otimes_i) are applied bottom-up (which is preceded by an application of (\Box^*) and (\otimes_i^*) , respectively, as discussed in Remark 19) with $w_0 : \Box\psi$ or $w_0 : \otimes_i\psi$ principal, then the saturation conditions (\mathbf{C}_\Box) and (\mathbf{C}_{\otimes_i}) will hold for *any* labeled formula of the form $u : \Box\psi$ or $u : \otimes_i\psi$ for the remainder of proof-search. Similarly, once $(D2_i)$ is applied bottom-up for an agent $i \in \text{Ag}$, it never need be applied for that agent again as (\mathbf{C}_{D2}) will continue to be satisfied for the remainder of proof-search. We therefore dedicate the remainder of the proof to showing claims 4 and 5.

4. Observe that any generation tree defined relative to a thread \mathcal{T} generated during the computation of $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi)$ has a finite branching factor bounded by the number of agents plus the number of \square , \otimes_i , and $[i]$ subformulae of φ . Moreover, since each label u can only be associated with formulae from $\{I_{\otimes_i}u \mid i \in \text{Ag}\} \cup \text{sufo}(\varphi)$, which is a finite set, we know the depth of the generation tree is bounded by a finite number as eventually a loop node must occur along any given path in the generation tree. Hence, since each bottom-up application of $([i])$ and $([i]^*)$ corresponds to an edge in a generation tree, we have that $([i])$ and $([i]^*)$ can be bottom-up applied only a finite number of times in any thread.

5. Note that only IoaOp and bottom-up applications of (\square) , (\otimes_i) , (D2_i) , $([i])$, and $([i]^*)$ add fresh labels during proof-search. By claims 1–4 above, we know that if we consider any thread of proof-search then eventually all possible labels that could be added via (\square) , (\otimes_i) , (D2_i) , $([i])$, and $([i]^*)$ will have been added. After such a point, IoaOp will eventually be applied if the top sequent of a thread is not yet $\text{IOA}^{\mathcal{T}}$ -satisfied, yielding an $\text{IOA}^{\mathcal{T}}$ -satisfied sequent Λ due to Lemma 12. Since no other bottom-up applications of rules introduce labels that could violate the $\text{IOA}^{\mathcal{T}}$ -satisfiability of a sequent, all sequents generated after Λ will be $\text{IOA}^{\mathcal{T}}$ -satisfied as well, meaning that IoaOp will never be applied again. Hence, IoaOp will only be applied a finite number of times in any thread during proof-search. \square

Theorem 26 (Termination). *For each $n, k \in \mathbb{N}$ and $\varphi \in \mathcal{L}_n$, $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi)$ terminates.*

Proof. By Lemma 25 above we know that only a finite number of labels will be introduced during the computation of $\text{Prove}_n^k(\emptyset \Rightarrow w_0 : \varphi)$. Furthermore, each label will only be associated with a finite number of formulae from $\text{sufo}(\varphi)$ and a finite number of relational atoms. Hence, only a finite number of possible sequents can be generated during proof-search. Since each recursive call of proof-search, i.e. bottom-up application of a rule, creates a strictly larger sequent (as each recursive call introduces at least one new relational atom or labeled formula to a sequent), Prove_n^k never generates the same sequent twice in any given thread, implying that the algorithm will eventually terminate due to the finite space of sequents that can be generated. \square

Last, as a consequence of the above we obtain decidability and the finite model property.

Corollary 27 (Decidability and FMP). *For each $n, k \in \mathbb{N}$, the logic DS_n^k is decidable and has the finite model property.*

4.1 Necessity of Loop-Checking in Proof-Search

We show that proof-search is not guaranteed to terminate without loop-checking, thus motivating its introduction and use. We provide an example of a (non-deontic) STIT formula $\diamond[1]p \vee \diamond[2]q$, such that proof-search in G3DS_2^2 does not terminate if we disregard loop-checking. In order to improve readability, we represent labeled sequents as labeled graphs, and explain how these graphs are expanded during proof-search:

Definition 28 (Graph of Labeled Sequent). *Let $\Lambda = \mathcal{R} \Rightarrow \Gamma$ be a labeled sequent. We define its corresponding graph to be the triple $G(\Lambda) = (V, E, L)$ such that $V = \text{Lab}(\Lambda)$, $E = \{(w, u, i) \mid R_{[i]}wu \in \mathcal{R}\}$, and $L(w) = \Gamma \upharpoonright w$.*




Figure 5: Three ‘snapshots’ of proof-search without loop-checking, showing that proof-search will continue ad infinitum if loop-checking is not enforced.

Now, let us suppose that our proof-search algorithm omits loop-checking, meaning that lines 51–57 (see p. 860) are removed from Prove_n^k and replaced with the following lines of pseudo-code that simply apply the $([i])$ rule bottom-up when the condition is met.

Algorithm 2: Applying the $([i])$ Rule without Loop-checking

- 1 **if** $w : [i]\varphi \in \Gamma$, but $u : \varphi \notin \Gamma$ for all $u \in \text{Lab}(\Lambda)$ such that $R_{[i]}wu \in \mathcal{R}$ **then**
 - 2 Let $\mathcal{R}' := R_{[i]}wv, \mathcal{R}$ and $\Gamma' := v : \varphi, \Gamma$ with v fresh;
 - 3 **return** $\text{Prove}_n^k(\mathcal{R}' \Rightarrow \Gamma')$;
 - 4 **end**
-

Now, suppose our proof-search algorithm takes $\Rightarrow w : \Diamond[1]p \vee \Diamond[2]q$ as an input. Then, the derivation shown below will initially be constructed yielding the labeled node w shown at the top of the graph in Figure 5 (which corresponds to the top sequent in the derivation). The loops/relational atoms $R_{[1]}ww, R_{[2]}ww$ have been omitted in the graph. For simplicity, we always omit the presentation of loops in Figure 5 and assume their presence.

$$\begin{array}{c}
 \frac{R_{[1]}ww, R_{[2]}ww \Rightarrow w : \diamond[1]p, \diamond[2]q, w : [1]p, w : [2]q}{R_{[1]}ww, R_{[2]}ww \Rightarrow w : \diamond[1]p, \diamond[2]q, w : [1]p} (\diamond) \\
 \frac{R_{[1]}ww, R_{[2]}ww \Rightarrow w : \diamond[1]p, \diamond[2]q}{R_{[1]}ww, R_{[2]}ww \Rightarrow w : \diamond[1]p \vee \diamond[2]q} (\vee) \\
 \frac{R_{[1]}ww, R_{[2]}ww \Rightarrow w : \diamond[1]p \vee \diamond[2]q}{R_{[1]}ww \Rightarrow w : \diamond[1]p \vee \diamond[2]q} (\text{Ref}_2) \\
 \frac{R_{[1]}ww \Rightarrow w : \diamond[1]p \vee \diamond[2]q}{\Rightarrow w : \diamond[1]p \vee \diamond[2]q} (\text{Ref}_1)
 \end{array}$$

As the top sequent in the derivation above is not stable, our algorithm will continue proof-search, thus yielding the derivation shown below, whose top sequent corresponds to the graph shown in ‘snapshot 1’, i.e. the dashed rectangle S_1 in Figure 5, though without the horizontal 1-edge between u_0 and u_1 . We refer to this graph as G_1 . To improve readability, and due to the size of the labeled sequents occurring in the derivation below, we have included ellipses to indicate the side formulae that are inherited from lower sequents. In Figure 5, we have opted to denote pairs of relational atoms that form a 2-cycle (e.g. $R_{[1]}wu_0, R_{[1]}u_0w$ and $R_{[1]}wu_1, R_{[2]}u_1w$) as a single, undirected edge to simplify presentation.

$$\begin{array}{c}
 \frac{R_{[1]}wu_1, \dots \Rightarrow u_1 : q, u_1 : [1]p, u_1 : [2]q, \dots}{R_{[1]}wu_1, \dots \Rightarrow u_1 : q, u_1 : [1]p, \dots} (\diamond) \\
 \frac{R_{[1]}wu_1, \dots \Rightarrow u_1 : q, u_1 : [1]p, \dots}{R_{[1]}wu_1, R_{[2]}u_1w, \dots \Rightarrow u_1 : q, \dots, w : [2]q} (\diamond) \\
 \frac{R_{[1]}wu_1, R_{[2]}u_1w, \dots \Rightarrow u_1 : q, \dots, w : [2]q}{R_{[1]}wu_1, R_{[1]}u_1u_1, R_{[2]}u_1u_1, \dots \Rightarrow u_1 : q, \dots, w : [2]q} (\text{Euc}_2) \\
 \frac{R_{[1]}wu_1, R_{[1]}u_1u_1, R_{[2]}u_1u_1, \dots \Rightarrow u_1 : q, \dots, w : [2]q}{R_{[1]}wu_1, R_{[1]}u_1u_1, \dots \Rightarrow u_1 : q, \dots, w : [2]q} (\text{Ref}_2) \\
 \frac{R_{[1]}wu_1, R_{[1]}u_1u_1, \dots \Rightarrow u_1 : q, \dots, w : [2]q}{R_{[2]}wu_1, \dots \Rightarrow u_1 : q, \dots, w : [2]q} (\text{Ref}_1) \\
 \frac{R_{[2]}wu_1, \dots \Rightarrow u_1 : q, \dots, w : [2]q}{R_{[1]}wu_0, \dots \Rightarrow u_0 : p, u_0 : [1]p, u_0 : [2]q, \dots, w : [2]q} ([2]) \\
 \frac{R_{[1]}wu_0, \dots \Rightarrow u_0 : p, u_0 : [1]p, u_0 : [2]q, \dots, w : [2]q}{R_{[1]}wu_0, \dots \Rightarrow u_0 : p, u_0 : [1]p, \dots, w : [2]q} (\diamond) \\
 \frac{R_{[1]}wu_0, \dots \Rightarrow u_0 : p, u_0 : [1]p, \dots, w : [2]q}{R_{[1]}wu_0, R_{[1]}u_0w, \dots \Rightarrow u_0 : p, \dots, w : [2]q} (\diamond) \\
 \frac{R_{[1]}wu_0, R_{[1]}u_0w, \dots \Rightarrow u_0 : p, \dots, w : [2]q}{R_{[1]}wu_0, R_{[1]}u_0u_0, R_{[2]}u_1u_1, \dots \Rightarrow u_0 : p, \dots, w : [2]q} (\text{Euc}_1) \\
 \frac{R_{[1]}wu_0, R_{[1]}u_0u_0, R_{[2]}u_1u_1, \dots \Rightarrow u_0 : p, \dots, w : [2]q}{R_{[1]}wu_0, R_{[1]}u_0u_0, \dots \Rightarrow u_0 : p, \dots, w : [2]q} (\text{Ref}_2) \\
 \frac{R_{[1]}wu_0, R_{[1]}u_0u_0, \dots \Rightarrow u_0 : p, \dots, w : [2]q}{R_{[1]}wu_0, \dots \Rightarrow u_0 : p, \dots, w : [2]q} (\text{Ref}_1) \\
 \frac{R_{[1]}wu_0, \dots \Rightarrow u_0 : p, \dots, w : [2]q}{R_{[1]}ww, R_{[2]}ww \Rightarrow w : \diamond[1]p, \diamond[2]q, w : [1]p, w : [2]q} ([1])
 \end{array}$$

One can see that $(\mathbf{C}_{[1]})$ and $(\mathbf{C}_{[2]})$ are unsatisfied in G_1 due to the occurrence of $[1]p$ at u_1 and $[2]q$ at u_0 , respectively. We assume w.l.o.g. that proof-search applies $([2])$ bottom-up, creating a new label/vertex u_2 emanating from u_0 , which can be seen in Figure 5. Proof-search will then apply (Ref_1) and (Ref_2) , introducing loops at u_2 , followed by an application of (Euc_2) , introducing a 2-edge between u_2 and u_0 . Recall that $k = 2$, that is, each agent has a maximum of two choices. Since no relational atom of the form $R_{[1]}vz$ with $v, z \in \{u_0, u_1, u_2\}$ and $v \neq z$ exists within our labeled sequent/graph constructed thus far, i.e. there is no 1-edge between u_0 , u_1 , or u_2 , the $(\mathbf{C}_{\text{APC}})$ saturation condition is not satisfied. Hence, proof-search will apply (APC_1^2) bottom-up, yielding a premise with $R_{[1]}u_1u_0$, followed by an application of (Euc_1) , introducing $R_{[1]}u_0u_1$. After applying the (\diamond) rule bottom-up twice, introducing the formulae $[1]p$ and $[2]q$ at u_2 , we obtain the labeled sequent/graph shown in ‘snapshot 2’ in Figure 5, i.e. the graph in the dashed rectangle S_2 , though without the diagonal 2-edge between u_1 and u_2 . We refer to this graph as G_2 .

Observe that the $(\mathbf{C}_{[1]})$ saturation condition is not satisfied due to the occurrence of $[1]p$ at u_2 in G_2 . Hence, the algorithm will apply the $([1])$ rule bottom up, introducing the label u_3 protruding from u_2 , followed by applications of (\mathbf{Ref}_1) , (\mathbf{Ref}_2) , and (\mathbf{Euc}_1) . One can see that (\mathbf{CAPC}) is not satisfied since a 2-edge does not exist between u_1 , u_2 , or u_3 , and so, (\mathbf{APC}_2^2) will be applied bottom-up. This rule application, in conjunction with an application of (\mathbf{Euc}_2) , introduces $R_{[2]u_1u_2}$ and $R_{[2]u_2u_1}$, which are represented as an undirected 2-edge between u_1 and u_2 . After two applications of the (\diamond) rule, we obtain the labeled sequent/graph in ‘snapshot 3’ in Figure 5, i.e. the graph in the dashed rectangle S_3 , although without the diagonal 1-edge between u_1 and u_3 . We call this graph G_3 .

The $(\mathbf{C}_{[1]})$ saturation condition is not satisfied in G_3 due to the occurrence of $[2]q$ at u_3 . As a consequence, our proof-search algorithm will apply the $([2])$ rule bottom up, introducing the label u_4 protruding from u_3 , followed by applications of (\mathbf{Ref}_1) , (\mathbf{Ref}_2) , and (\mathbf{Euc}_2) . At this stage, (\mathbf{CAPC}) is not satisfied in the graph since a 1-edge does not exist between u_1 , u_3 , or u_4 , and so, (\mathbf{APC}_1^2) will be applied bottom-up. This rule application yields a premise that, in conjunction with an application of (\mathbf{Euc}_1) , will introduce $R_{[1]u_1u_3}$ and $R_{[1]u_3u_1}$ (represented as the undirected 1-edge between u_1 and u_3). After two applications of the (\diamond) rule, we obtain the labeled sequent/graph shown in Figure 5. By repeating the above pattern, one can see that the left path in the labeled sequent/graph will continue to grow through alternating bottom-up applications of the $([1])$ and $([2])$ rules, i.e. proof-search does not terminate for the formula $\diamond[1]p \vee \diamond[2]q$ in the absence of loop-checking.

We emphasize that, since the chosen formula is \ominus_i - and \otimes_i -free, the above example shows the importance of a loop-checking mechanism for (non-deontic) STIT as well.⁷

5. Applications: Duty, Compliance, and Joint Fulfillment Checking

Proof-search algorithms make a logic suitable for reasoning tasks. In this last section, we discuss three tasks of particular interest to agent-based normative reasoning:

Duty Checking: Determine an agent’s obligations relative to a given knowledge base.

Compliance Checking: Determine if a choice, considered by an agent as potential conduct, complies with the given knowledge base.

Joint Fulfillment Checking: Determine whether under a specified factual context an agent can (still) jointly fulfill all their duties.

The referred knowledge base is a set consisting of obligations and facts and, in what follows, we distinguish between the norm base and the factual context of the knowledge base. We provide an example for each of these applications. For the sake of readability, we consider a single-agent setting in all cases, but the approach generalizes to multi-agent reasoning. In considering the first two tasks, we assume that the given knowledge base is consistent, unless stated otherwise (we note that consistency can also be determined through proof-search).

7. Although the proof-search algorithm for non-deontic STIT logic by Negri and Pavlović (2020) is susceptible to the above counter-example, we conjecture that a modified version of our loop-checking mechanism can be used to ensure termination in their algorithm.

As argued for in Section 1, the derivations and (counter-)models obtained through our applications have explanatory value⁸ since both provide justifications of (non-)theoremhood by representing a constructive step-by-step reasoning process interpretable by the explainee.

5.1 Duty Checking

Consider the following scenario: Yara borrowed a hammer from a friend (Kai) in order to repair a leaking shed. Furthermore, Yara promised to return the hammer to Kai before noon. Let \mathbf{n} be the proposition “Yara arrives at Kai’s place before noon (to return the hammer).” We say that (due to the promise) Yara is under the obligation to return the hammer before noon, i.e. $\otimes_y \mathbf{n}$. We assume that Yara is free to see to it that she arrives before noon or not, represented as $\diamond[y] \mathbf{n} \wedge \diamond[y] \neg \mathbf{n}$. Now, Yara knows that Kai lives around the corner and considers her means of going there, e.g. she could go by foot or take the car. Suppose Yara entertains the idea to go by foot, i.e. $\diamond[y] \mathbf{f}$. Moreover, let us assume that Yara knows, as a fact, that walking will ensure that she arrives by noon, i.e. $\Box(\mathbf{f} \rightarrow \mathbf{n})$.

The knowledge base for this scenario is the collection

$$\Sigma := \{\otimes_y \mathbf{n}, \diamond[y] \mathbf{n}, \diamond[y] \neg \mathbf{n}, \diamond[y] \mathbf{f}, \diamond[y] \neg \mathbf{f}, \Box(\mathbf{f} \rightarrow \mathbf{n})\},$$

though we simplify it to $\Sigma := \{\otimes_y \mathbf{n}, \diamond[y] \neg \mathbf{n}, \diamond[y] \mathbf{f}, \Box(\mathbf{f} \rightarrow \mathbf{n})\}$ as $\diamond[y] \mathbf{n}$ and $\diamond[y] \neg \mathbf{f}$ are implied by the other formulae. Since Yara is obliged to meet Kai, Yara wonders whether she has the obligation to go by foot, i.e. $\otimes_y \mathbf{f}$. Accordingly, we check whether the formula

$$\varphi := (\otimes_y \mathbf{n} \wedge \diamond[y] \neg \mathbf{n} \wedge \diamond[y] \mathbf{f} \wedge \Box(\mathbf{f} \rightarrow \mathbf{n})) \rightarrow \otimes_y \mathbf{f}$$

is valid, i.e. we can check whether $w : \varphi$ is derivable in G3DS_1^0 , where $Ag = \{y\}$ and the limited choice parameter $k = 0$ since we have not enforced an *a priori* limit on Yara’s number of choices. Hence, we check whether the duty $\otimes_y \mathbf{f}$ is implied by the context Σ .

To test the validity of φ , let us run our proof-search algorithm $\text{Prove}_1^0(\emptyset \Rightarrow w : \varphi)$. Doing so, one finds that φ is invalid as the stable sequent $\Lambda = \mathcal{R} \Rightarrow \Gamma_0, \Gamma_1, \Gamma_2, \Gamma_3$ is generated via proof-search, where each set in Λ is as follows:

$$\mathcal{R} := R_{[y]} ww, R_{[y]} uu, R_{[y]} vv, R_{[y]} zz, I_{\otimes_y} z;$$

$$\Gamma_0 := w : \ominus_y \neg \mathbf{n}, w : \Box \langle y \rangle \mathbf{n}, w : \Box \langle y \rangle \neg \mathbf{f}, w : \diamond(\mathbf{f} \wedge \neg \mathbf{n}), w : \mathbf{f} \wedge \neg \mathbf{n}, w : \mathbf{f}, w : \otimes_y \mathbf{f};$$

$$\Gamma_1 := u : \langle y \rangle \mathbf{n}, u : \mathbf{n}, u : \mathbf{f} \wedge \neg \mathbf{n}, u : \mathbf{f};$$

$$\Gamma_2 := v : \langle y \rangle \neg \mathbf{f}, v : \neg \mathbf{f}, v : \mathbf{f} \wedge \neg \mathbf{n}, v : \neg \mathbf{n};$$

$$\Gamma_3 := z : \mathbf{f}, z : \mathbf{f} \wedge \neg \mathbf{n}, z : \neg \mathbf{n}.$$

As explained in the previous section, we may transform Λ into a counter-model M^Λ for φ . We define the stability model $M := (W, R_{[y]}, I_{\otimes_y}, V)$ in accordance with Definition 20, where $W = \{w, u, v, z\}$, $R_{[y]} = \{(w, w), (u, u), (v, v), (z, z)\}$, $I_{\otimes_y} = \{z\}$, and $V(\mathbf{n}) = \{w, v, z\}$ and $V(\mathbf{f}) = \{v\}$. A graphical representation of M^Λ is presented in Figure 6.

8. For some purposes, formal proofs may be less explanatory than informal ones containing reasoning gaps. This is especially true when the intended explainee is a layperson (Miller, 2019). For experts, derivations and interpretable models serve as appropriate explanations; cf. (Biran & Cotton, 2017).

5.2 Compliance Checking

To determine whether a considered choice complies with a given knowledge base, consisting of a norm base and a factual context, we must determine whether the considered choice does not conflict with an obligation entailed by the norm base under the given factual context (recall that we assume the knowledge base to be consistent in this scenario).

We emphasize that checking whether an agent’s conduct is *consistent* with a given knowledge base does not define compliance checking. Consider the simple example where Jade is obliged to drive on the left-hand side of the road, i.e. $\otimes_j \text{left_jade}$ and where she considers driving on the right-hand side instead, i.e. $[j] \text{right_jade}$. Furthermore, assume that driving left and right are mutually exclusive, i.e. $\Box(\text{left_jade} \rightarrow \neg \text{right_jade})$. Such a scenario can be consistently modeled by a single-agent two-choice moment with one choice—the obligatory one—representing Jade driving on the left and another choice with Jade cycling on the right. The scenario is consistent, but there is no compliance: Jade violates her obligation by driving on the right side of the road. A choice is not compliant whenever its performance entails a norm violation. Therefore, we check compliance by determining whether there is no explicit obligation to the contrary of what the agent considers as potential conduct.

Reconsider the scenario where Yara promised Kai to return the borrowed hammer before noon, i.e. $\otimes_y \mathbf{n}$. She considers going there by foot or by car and both are available choices, i.e. $\Diamond[y] \mathbf{f} \wedge \Diamond[y] \text{car}$. Yara wonders whether taking the car—i.e. $[y] \text{car}$ —complies with her promise. The knowledge base formalizing this situation is the following set of formulae:

$$\Sigma := \{ \otimes_y \mathbf{n}, \Diamond[y] \mathbf{f} \wedge \Diamond[y] \text{car} \}$$

To see if the choice $[y] \text{car}$ complies with Σ we need to check whether the formula

$$(\otimes_y \mathbf{n} \wedge \Diamond[y] \mathbf{f} \wedge \Diamond[y] \text{car}) \rightarrow \otimes_y \neg \text{car}$$

is derivable. If not, then this means there is no implicit obligation for Yara to not take the car and, so, taking the car complies with the current normative situation expressed by Σ .

Clearly, $[y] \text{car}$ complies due to the syntactic independence with $\otimes_y \mathbf{n}$. As can be seen, compliance checking is a form of duty checking, where it checks for a duty to the contrary of what the agent considers doing. Since we already discussed successful proof-search and counter-model extraction in the context of duty checking, our discussion above suffices, and we omit demonstration of failed proof-search.

5.3 Joint Fulfillment Checking

Last, we are interested in checking whether a consistent norm base together with a consistent factual context preserve consistency when put together. If the resulting set of formulae is consistent this tells us that given the context at hand the agents can jointly satisfy all of their obligations. If the resulting set is inconsistent this means that the factual context does not allow for jointly fulfilling all duties. We observe that, when taken together with a specified factual context, a norm base may imply obligations that are not implied when considering the norm base in isolation. Hence, in contrast to the first two reasoning tasks, this particular task is about logical consistency.

The following example illustrates this third reasoning task. Suppose that Yara promised her friend Lisa to pick up a package. That is, Yara has the obligation $\otimes_y \mathbf{p}$, where \mathbf{p} is the proposition “the package is picked up.” Let us assume that Yara knows the post-office closes at noon and realizes that if she goes to pick up the package, she will not be able to return the borrowed hammer to Kai at noon, i.e. $\Box([y]\mathbf{p} \rightarrow \neg[y]\mathbf{n})$. We may formalize the situation as follows:

$$\{\diamond[y]\mathbf{n}, \diamond[y]\neg\mathbf{n}, \diamond[y]\mathbf{p}, \diamond[y]\neg\mathbf{p}, \Box([y]\mathbf{n} \rightarrow \neg[y]\mathbf{p}), \otimes_y \mathbf{n}, \otimes_y \mathbf{p}\}$$

The first four formulae describe Yara’s available choices, namely, Yara has the choice to either return or not return the hammer to Kai before noon, and the choice to pick up or not pick up the package for Lisa. The fifth formula expresses that the choices available to Yara are mutually exclusive. The last two formulae encode the obligations applicable to Yara. Are Yara’s obligations consistent with the facts of the situation?

We can utilize our proof-search algorithm to derive $(\otimes_y \mathbf{n} \wedge \otimes_y \mathbf{p} \wedge \Box([y]\mathbf{p} \rightarrow \neg[y]\mathbf{n})) \rightarrow \perp$ in $\mathbf{G3DS}_1^0$, where our set of agents is $Ag = \{y\}$ and $k = 0$. By applying the hp-admissibility of (Wk) to the derivation given below, we can indeed show that the above norm base is inconsistent, that is, the following formula is a theorem in \mathbf{DS}_1^0 :

$$(\diamond[y]\mathbf{n} \wedge \diamond[y]\neg\mathbf{n} \wedge \diamond[y]\mathbf{p} \wedge \diamond[y]\neg\mathbf{p} \wedge \Box([y]\mathbf{n} \rightarrow \neg[y]\mathbf{p}) \wedge \otimes_y \mathbf{n} \wedge \otimes_y \mathbf{p}) \rightarrow \perp$$

Due to the size of the proof of $(\otimes_y \mathbf{n} \wedge \otimes_y \mathbf{p} \wedge \Box([y]\mathbf{p} \rightarrow \neg[y]\mathbf{n})) \rightarrow \perp$, we split the proof into various sections and present each below. We denote the first proof shown below as Π_0 . The sub-derivation Π' constituting the right branch of Π_0 is obtained in a similar way to the left branch.

$$\frac{\frac{\frac{\frac{I_{\otimes_y} v, I_{\otimes_y} z, \dots \Rightarrow \dots, v : [y]\mathbf{p}, z : \mathbf{p}, z : \neg\mathbf{p}}{\quad} (\text{id})}{I_{\otimes_y} v, I_{\otimes_y} z, \dots \Rightarrow \dots, v : [y]\mathbf{p}, z : \mathbf{p}} (\ominus_y)}{I_{\otimes_y} v, R_{[y]}vv, R_{[y]}vz, R_{[y]}zv, R_{[y]}zz, \dots \Rightarrow \dots, v : [y]\mathbf{p}, z : \mathbf{p}} (\text{D3}_y)}{I_{\otimes_y} v, R_{[y]}vv, R_{[y]}vz, \dots \Rightarrow \dots, v : [y]\mathbf{p}, z : \mathbf{p}} (\text{Ref}_y), (\text{Euc}_y)}{\frac{I_{\otimes_y} v, R_{[y]}vv, R_{[y]}vz, \dots \Rightarrow \dots, v : [y]\mathbf{p}}{I_{\otimes_y} v, R_{[y]}vv, R_{[y]}vz, \dots \Rightarrow \dots, v : [y]\mathbf{p}} ([y])} (\text{Ref}_y), (\text{Euc}_y)}{\frac{I_{\otimes_y} v, R_{[y]}vv, \dots \Rightarrow \dots, v : [y]\mathbf{p} \wedge [y]\mathbf{n}}{I_{\otimes_y} v, R_{[y]}vv, \dots \Rightarrow \dots, w : [y]\mathbf{p}, u : \mathbf{p}} (\diamond)} (\text{Ref}_y), (\text{Euc}_y)}{\frac{I_{\otimes_y} v, \dots \Rightarrow \dots, w : [y]\mathbf{p}, u : \mathbf{p}}{R_{[y]}ww, R_{[y]}wu, R_{[y]}uw, R_{[y]}uu, I_{\otimes_y} v \Rightarrow \dots, w : [y]\mathbf{p}, u : \mathbf{p}} (\text{D2}_y)} (\text{Ref}_y), (\text{Euc}_y)}{\frac{R_{[y]}ww, R_{[y]}wu \Rightarrow \dots, w : [y]\mathbf{p}, u : \mathbf{p}}{R_{[y]}ww \Rightarrow \dots, w : [y]\mathbf{p}} ([y])} (\text{Ref}_y), (\text{Euc}_y)}{\frac{I_{\otimes_y} v, I_{\otimes_y} z, \dots \Rightarrow \dots, v : [y]\mathbf{p}, z : \mathbf{p}}{I_{\otimes_y} v, R_{[y]}vv, \dots \Rightarrow \dots, v : [y]\mathbf{n}} (\Pi')} (\wedge)} (\wedge)$$

The derivation of $R_{[y]}ww \Rightarrow \dots, w : [y]\mathbf{n}$, referred to as Π_1 , is similar to the derivation Π_0 applying the same rules bottom-up in the same order. By making use of Π_0 and Π_1 , we can complete our proof as follows:

$$\frac{\frac{\frac{\frac{\frac{\frac{\Pi_0}{R_{[y]}ww \Rightarrow w : \ominus_y \neg\mathbf{n}, w : \ominus_y \neg\mathbf{p}, w : \diamond([y]\mathbf{p} \wedge [y]\mathbf{n}), w : [y]\mathbf{p} \wedge [y]\mathbf{n}, w : \perp}}{\quad} (\wedge)}{R_{[y]}ww \Rightarrow w : \ominus_y \neg\mathbf{n}, w : \ominus_y \neg\mathbf{p}, w : \diamond([y]\mathbf{p} \wedge [y]\mathbf{n}), w : \perp}}{\quad} (\diamond)}{R_{[y]}ww \Rightarrow w : \ominus_y \neg\mathbf{n} \vee \ominus_y \neg\mathbf{p} \vee \diamond([y]\mathbf{p} \wedge [y]\mathbf{n}) \vee \perp}} (\vee) \times 3}{\frac{R_{[y]}ww \Rightarrow w : \ominus_y \neg\mathbf{n} \vee \ominus_y \neg\mathbf{p} \vee \diamond([y]\mathbf{p} \wedge [y]\mathbf{n}) \vee \perp}{\Rightarrow w : \ominus_y \neg\mathbf{n} \vee \ominus_y \neg\mathbf{p} \vee \diamond([y]\mathbf{p} \wedge [y]\mathbf{n}) \vee \perp}} (\text{Ref}_y)}{\Rightarrow w : (\otimes_y \mathbf{n} \wedge \otimes_y \mathbf{p} \wedge \Box([y]\mathbf{p} \rightarrow \neg[y]\mathbf{n})) \rightarrow \perp} =$$

The above proof illustrates how our proof-search algorithm may be used to check the consistency of knowledge bases including obligations, possibly relative to a factual context.

As a final reflection, we observe that, from an abstract point of view, these specific reasoning tasks can be reduced to variations of validity checking. That is, they all use proof-search with automated counter-model extraction to check the validity of a specified formula. Nevertheless, these three different applications conceptually highlight different reasoning tasks with varying deontic value. If the agent’s considered behavior is not compliant we know that there is a conflicting obligation implied by the norm base. If the agent’s considered behavior is compliant this tells us that the action is permissible, but it does not tell us anything about the deontic value of the action (the action may be merely compatible with the agent’s obligatory actions). The third application is also different, checking whether the obligations of a norm base can be consistently fulfilled in a specified factual context irrespective of any considered actions.

6. Conclusion

In this paper, we demonstrated how to automate normative reasoning for deontic STIT logics by means of proof-search algorithms founded upon proof systems developed in (Lyon, 2021) and inspired by the proof systems for related logics introduced in (van Berkel & Lyon, 2019a, 2021; Lyon & van Berkel, 2019). In doing so, this work is the first to provide terminating proof-search procedures for deontic as well as non-deontic multi-agent STIT logics with (un)limited choice. The key to obtaining terminating proof-search was the introduction of a loop-checking mechanism, which ensured the introduction of only finitely many labels during proof-search. As argued for in Section 4, such a loop-checking mechanism is needed to effectively handle the general case of STIT logics with unlimited and limited choice axioms.

The present work provides a promising outlook for future research. For instance, the proof theory of (and automated reasoning with) extensions of deontic STIT logics remains to be investigated: e.g. temporal (van Berkel & Lyon, 2019b; Ciuni & Lorini, 2017) and epistemic extensions (Broersen, 2011). In particular, temporal extensions of deontic STIT logics are of interest in the context of verification of obligations for autonomous vehicles (Shea-Blymyer & Abbas, 2021). Decidability remains an open question for various temporal extensions of deontic STIT (Ciuni & Lorini, 2017). Additionally, in relation to normative reasoning tasks, an interesting direction would be to extend deontic STIT logic with conditional obligations (Horty, 2001, Ch. 5). Such obligations are often used to consistently model challenging contrary-to-duty scenarios which are scenarios in which a violation—i.e. $[i]\varphi \wedge \otimes_i \neg\varphi$ —has occurred and the agent needs to find out her obligations *given* the sub-ideality of the situation (Hilpinen & McNamara, 2013).

From a technical point of view, we believe that an important step in automated agential normative reasoning can be taken by writing and evaluating theorem provers based on our proof-search algorithms for the logics in this paper. Due to the modularity of our approach such results may be adapted to extensions of the deontic STIT logics discussed here. On a related note, derivations generated by our proof-search algorithm are not necessarily minimal and shorter proofs may exist. Optimization procedures may serve as a fruitful future research direction.

In the context of explainable normative reasoning, we point out that this work provides only a promising first step in showing how logical methods can be harnessed in the context of STIT and Explainable AI. As argued, constructive proof-search procedures that yield counter-models in case of failed proof-search have the advantage of providing a justification of a formula’s (non-)theoremhood and make the reasons for the conclusion interpretable by providing a step-by-step reasoning procedure. Such explanations take experts instead of laypersons as the intended explainee. However, certain refinements can be made to increase the perspicuity of our proofs and counter-models. For instance, relevance conditions can be imposed on our proof calculi; when we derive the obligation $\otimes_y \mathbf{f}$ from the knowledge base $\{\otimes_y \mathbf{n}, \Box(\mathbf{f} \rightarrow \mathbf{n}), \Box(\mathbf{n} \rightarrow \mathbf{f})\}$ the premise $\Box(\mathbf{f} \rightarrow \mathbf{n})$ is not strictly needed in the derivation and thus can be removed from the proof as it is irrelevant in explaining the derived duty.

Acknowledgments

This research was supported by the European Research Council (ERC) Consolidator Grant 771779 (DeciGUT) and the Austrian Science Fund (FWF) project W1255-N23.

References

- Alvarez, M. (2017). Reasons for Action: Justification, Motivation, Explanation. In Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017 edition). Metaphysics Research Lab, Stanford University.
- Arkoudas, K., Bringsjord, S., & Bello, P. (2005). Toward ethical robots via mechanized deontic logic. In *AAAI Fall Symposium on Machine Ethics*, pp. 17–23. The AAAI Press Menlo Park, CA.
- Avron, A. (1996). The method of hypersequents in the proof theory of propositional non-classical logics. In Hodges, W., Hyland, M., Steinhorn, C., & Truss, J. (Eds.), *From Foundations to Applications: European Logic Colloquium*, p. 1–32. Clarendon Press.
- Balbani, P., Herzig, A., & Troquard, N. (2008). Alternative axiomatics and complexity of deliberative STIT theories. *Journal of Philosophical Logic*, 37(4), 387–406.
- Bartha, P. (1993). Conditional obligation, deontic paradoxes, and the logic of agency. *Annals of Mathematics and Artificial Intelligence*, 9(1-2), 1–23.
- Belnap, N. (1982). Display logic. *Journal of Philosophical Logic*, 11(4), 375–417.
- Belnap, N., Perloff, M., & Xu, M. (2001). *Facing the future: agents and choices in our indeterminist world*. Oxford University Press, Oxford.
- van Berkel, K., & Lyon, T. (2019a). Cut-free calculi and relational semantics for temporal STIT logics. In Calimeri, F., Leone, N., & Manna, M. (Eds.), *European Conference on Logics in Artificial Intelligence (JELIA)*, pp. 803–819, Cham. Springer.
- van Berkel, K., & Lyon, T. (2019b). A neutral temporal deontic STIT logic. In *International Workshop on Logic, Rationality and Interaction*, pp. 340–354. Springer.
- van Berkel, K., & Lyon, T. (2021). The varieties of ought-implies-can and deontic STIT logic. In Liu, F., Marra, A., Portner, P., & Putte, F. V. D. (Eds.), *Deontic Logic and*

- Normative Systems: 15th International Conference (DEON2020/2021, Munich)*, pp. 57–76. College Publications.
- van Berkel, K., & Straßer, C. (2022). Reasoning with and about norms in logical argumentation. In Toni, F., Polberg, S., Booth, R., Caminada, M., & Kido, H. (Eds.), *Frontiers in Artificial Intelligence and Applications: Computational Models of Argument, proceedings (COMMA22)*, Vol. 353, pp. 332 – 343. IOS press.
- van Berkel, K., & Straßer, C. (2024). Towards deontic explanations through dialogue. In Čyras, K., Kampik, T., Cocarascu, O., & Rago, A. (Eds.), *2nd International Workshop on Argumentation for eXplainable AI (ArgXAI)*. CEUR Workshop Proceedings (To Appear).
- Biran, O., & Cotton, C. (2017). Explanation and justification in machine learning: A survey. In *IJCAI-17 Workshop on Explainable AI (XAI)*, pp. 8–13.
- Broersen, J. (2011). Deontic epistemic STIT logic distinguishing modes of mens rea. *Journal of Applied Logic*, 9(2), 137–152.
- Brünnler, K. (2009). Deep sequent systems for modal logic. *Archive for Mathematical Logic*, 48(6), 551–577.
- Bull, R. A. (1992). Cut elimination for propositional dynamic logic without *. *Mathematical Logic Quarterly*, 38(1), 85–100.
- Ciuni, R., & Lorini, E. (2017). Comparing semantics for temporal STIT logic. *Logique et Analyse*, 61(243), 299–339.
- Dalmonte, T., Lellmann, B., Olivetti, N., & Pimentel, E. (2021). Hypersequent calculi for non-normal modal and deontic logics: countermodels and optimal complexity. *Journal of Logic and Computation*, 31(1), 67–111.
- Fitting, M. (1972). Tableau methods of proof for modal logics.. *Notre Dame Journal of Formal Logic*, 13(2), 237–247.
- Fitting, M. (2014). Nested sequents for intuitionistic logics. *Notre Dame Journal of Formal Logic*, 55(1), 41–61.
- Gabbay, D. M. (1996). *Labelled deductive systems*, Vol. 33 of *Oxford Logic Guides*. Clarendon Press/Oxford Science Publications.
- 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.
- Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., & Yang, G.-Z. (2019). Xai—explainable artificial intelligence. *Science Robotics*, 4(37), eaay7120.
- Hein, R. (2005). Geometric theories and proof theory of modal logic. Master’s thesis, Technische Universität Dresden.
- Herzig, A., & Schwarzentruher, F. (2008). Properties of logics of individual and group agency. In Areces, C., & Goldblatt, R. (Eds.), *Proceedings of the 7th Conference on Advances in Modal Logic (AIML)*, Vol. 7, pp. 133–149. College Publications.

- Hilpinen, R., & McNamara, P. (2013). Deontic logic: A historical survey and introduction. In *Handbook of Deontic Logic and Normative Systems*, Vol. 1, pp. 3–136. College Publications Milton Keynes.
- Horrocks, I., & Sattler, U. (2004). Decidability of SHIQ with complex role inclusion axioms. *Artificial Intelligence*, 160(1-2), 79–104.
- Horty, J. F. (2001). *Agency and deontic logic*. Oxford University Press.
- Horty, J. F., & Belnap, N. (1995). The deliberative STIT: A study of action, omission, ability, and obligation. *Journal of Philosophical Logic*, 24(6), 583–644.
- Kashima, R. (1994). Cut-free sequent calculi for some tense logics. *Studia Logica*, 53(1), 119–135.
- Lahav, O. (2013). From frame properties to hypersequent rules in modal logics. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pp. 408–417.
- Langley, P., Meadows, B., Sridharan, M., & Choi, D. (2017). Explainable agency for intelligent autonomous systems. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI'17*, p. 4762–4763. AAAI Press.
- Lorini, E., & Sartor, G. (2015). Influence and responsibility: A logical analysis. In Rotolo, A. (Ed.), *Frontiers in Artificial Intelligence and Applications: Legal Knowledge and Information Systems (JURIX 2015)*, Vol. 279, pp. 51–60.
- Lyon, T. (2021). *Refining Labelled Systems for Modal and Constructive Logics with Applications*. Ph.D. thesis, Technische Universität Wien.
- Lyon, T., & van Berkel, K. (2019). Automating agential reasoning: Proof-calculi and syntactic decidability for STIT logics. In Baldoni, M., Dastani, M., Liao, B., Sakurai, Y., & Zalila Wenkstern, R. (Eds.), *Principles and Practice of Multi-Agent Systems - 22nd International Conference (PRIMA)*, Vol. 11873 of LNCS, pp. 202–218. Springer.
- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence*, 267, 1–38.
- Miller, T., Howe, P., & Sonenberg, L. (2017). Explainable AI: Beware of inmates running the asylum or: How I learnt to stop worrying and love the social and behavioural sciences. *arXiv preprint arXiv:1712.00547*, abs/1712.00547.
- Murakami, Y. (2004). Utilitarian deontic logic. In Schmidt, R. A., Pratt-Hartmann, I., Reynolds, M., & Wansing, H. (Eds.), *Proceedings of the 5th Conference on Advances in Modal Logic (AIML)*, pp. 211–230.
- Negri, S. (2005). Proof analysis in modal logic. *Journal of Philosophical Logic*, 34, 507.
- Negri, S., & Pavlović, E. (2020). Proof-theoretic analysis of the logics of agency: The deliberative STIT. *Studia Logica*, 109, 1–35.
- Poggiolesi, F. (2009). The method of tree-hypersequents for modal propositional logic. In Makinson, D., Malinowski, J., & Wansing, H. (Eds.), *Towards Mathematical Philosophy*, Vol. 28 of *Trends in logic*, pp. 31–51. Springer.
- Schmidt, R. A., & Tishkovsky, D. (2011). Automated synthesis of tableau calculi. *Logical Methods in Computer Science*, 7(2).

- Shea-Blymyer, C., & Abbas, H. (2020). A deontic logic analysis of autonomous systems' safety. In *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control*, pp. 1–11, New York. Association for Computing Machinery.
- Shea-Blymyer, C., & Abbas, H. (2021). Algorithmic ethics: Formalization and verification of autonomous vehicle obligations. *ACM Transactions on Cyber-Physical Systems (TCPS)*, 5(4), 1–25.
- Simpson, A. K. (1994). *The proof theory and semantics of intuitionistic modal logic*. Ph.D. thesis, University of Edinburgh.
- Tiu, A., Ianovski, E., & Goré, R. (2012). Grammar logics in nested sequent calculus: Proof theory and decision procedures. In *Proceedings of the 9th Conference on Advances in Modal Logic (AIML)*, pp. 516–537.
- Viganò, L. (2000). *Labelled Non-Classical Logics*. Springer Science & Business Media.
- Wansing, H. (1994). Sequent calculi for normal modal propositional logics. *Journal of Logic and Computation*, 4(2), 125–142.
- Wansing, H. (2002). Sequent systems for modal logics. In Gabbay, D. M., & Guentner, F. (Eds.), *Handbook of Philosophical Logic: Volume 8*, pp. 61–145. Springer, Dordrecht.
- Wansing, H. (2006). Tableaux for multi-agent deliberative-STIT logic.. In *Proceedings of the 6th Conference on Advances in Modal Logic (AIML)*, pp. 503–520.
- Ye, L. R., & Johnson, P. E. (1995). The impact of explanation facilities on user acceptance of expert systems advice. *Mis Quarterly*, 19, 157–172.