SMT-Based Deontic Reasoning for Åqvist Logics

Christian Köll, Agata Ciabattoni (✉), Dmitry Rozplokhas

Institute of Logic and Computation, TU Wien
ch.koell@gmail.com,agata@logic.at,dmitry@logic.at

Abstract. Building on the small-model constructions for Åqvist's deontic logics introduced in [24], we present Deo-SMT, an SMT-based reasoner implemented in Z3 for checking validity and generating countermodels. Deo-SMT covers all four of Åqvist's logics (**E**, **F**, **F+(CM)**, **G**) and provides countermodel visualizations as text, matrices, and directed graphs. Our tool outperforms the existing Isabelle/HOL approach, while providing a lightweight and accessible interface for normative reasoning.

1 Introduction

Normative reasoning plays a central role in fields such as law, artificial intelligence, and ethics, where decisions must respect explicit or implicit norms. Norms prescribe, permit, or forbid certain actions or states of affairs, and are captured in terms of obligations, permissions, and prohibitions. Deontic logics provide formal tools for representing and reasoning about them [12,13]. Broadly speaking, these logics fall into two families: norm-based systems, which treat norms as inference rules, and preference-based systems, which interpret normative claims through an ordering of possible worlds. Among the most prominent representatives of the preference-based family are Åqvist's logics [1] **E**, **F**, and **G**, long recognized as foundational frameworks for modeling conditional obligation, denoted as $\bigcirc(\phi \mid \psi)$ (read as "ϕ is obligatory, given ψ"). Originally introduced to model Contrary-to-Duty (CTD) obligations, which are obligations that arise when other obligations are violated, Åqvist's systems interpret deontic modalities using a possible-worlds semantics with a preference (betterness) relation. **F+(CM)** [20] is a more recent extension of **F** that incorporates cautious monotonicity, a well-known principle from non-monotonic reasoning [11]; **F+(CM)** can be viewed as Preferential Conditional Logic $\mathbb{PCL}$ [5] supplemented with reflexivity and the absoluteness axiom, ensuring that the ranking is not world-relative; the shallow fragment of **F+(CM)** coincides with the KLM logic P [17]. **E**, **F**, **F+(CM)**, and **G** form a modular hierarchy of increasing deductive strength. At the top of the hierarchy, **G** coincides with $\mathbb{VTA}$ [18,14], one of Lewis's conditional logics, and its shallow fragment is the KLM logic R [17].

Åqvist systems have been extensively studied from semantic, axiomatic, and proof theory perspectives, e.g., [21,6,8,7]. However, these works do not directly address practical questions such as whether, given a large body of norms and a set of factual circumstances, a particular obligation applies, fails to apply, or conflicts with others. Nor do they provide practical mechanisms for assessing the

consistency of real-world normative systems, for example, determining whether provisions of the General Data Protection Regulation (GDPR) are consistent or contradictory.

As automated decision-making becomes more widespread, the need for tools that can reason with these logics has become increasingly pressing. A SAT-based tool for I/O logics [19], a prominent approach to norm-based reasoning, proposed in [26]. For preference-based systems, the available tools are based on embeddings into Higher Order Logic [23,22,3,2]; they support versatile experimentation but remain inefficient for object-level tasks like proof search and countermodel generation, since they encode the logic's semantics inside expressive, and potentially undecidable, higher-order frameworks.

In this paper, we build on the LEGO small-model constructions of [24] to introduce Deo-SMT, an efficient reasoning tool for all four propositional Åqvist systems. Following the ideas of Friedman and Halpern [10], the constructions assemble polynomial-size models from elementary components: chains, antichains, and cliques of worlds drawn from an arbitrary model. This yields NP-completeness for the logics and enables an encoding into classical propositional logic. Using these small-model constructions (LEGO for **E**, **F**, and **F+(CM)**, and the construction of [10] for **G** in its $\mathbb{VTA}$ reformulation), we design SMT-based decision procedures in Z3 for both validity checking and countermodel generation. Deo-SMT supports all four logics in Åqvist's family and provides countermodel visualizations in textual, matrix, and graph formats.

Empirically, Deo-SMT drastically outperforms existing approaches, achieving an average speedup factor of $\times 31.8$ compared to the Isabelle/HOL encodings of Parent and Benzmüller [23,22,3,2]. This gain in efficiency reflects a different design trade-off: Isabelle/HOL-based approaches provide the strong guarantees of mechanically verified higher-order reasoning, whereas Deo-SMT prioritizes efficient object-level reasoning through SMT-based deduction and countermodel generation, together with an accessible interface for practical normative reasoning.

By combining the theoretical uniformity of the LEGO constructions with the efficiency of SMT solvers, our tool offers the first lightweight, scalable decision procedure for the full family of propositional Åqvist's preference-based deontic logics, making automated normative reasoning significantly more accessible.

Deo-SMT is available at https://lego.guest.logic.at/.

2 Åqvist's Systems in a Nutshell

Åqvist's systems are a family of modal logics for modeling conditional obligations. Building on the preference-based approach to conditional obligations proposed by [16], Åqvist introduced a hierarchy of bi-modal logics, combining the monadic alethic modality $\Box$ of **S5** and the dyadic deontic modality $\bigcirc$. The semantics of these systems is based on the following variant of preference models [18,25].

Definition 1. *A* preference model *is a triple* $\langle W, \succeq, \mathbb{V}\rangle$ *where* W *is a (non-empty) set of worlds, the preference relation* $\succeq$ *is a binary relation on* W*, and*

$\mathbb{V}\colon \mathbb{P} \to 2^W$ *is a valuation function. We denote the set of worlds of* M *by* $W(M)$*, and write* $w \succ u$ *to denote that both* $w \succeq u$ *and* $u \not\succeq w$*.*

When evaluating a modal formula w.r.t. a preference model, $\Box$ is interpreted as a universal modality and $\bigcirc(\varphi \mid \psi)$ is true when φ holds in all *best* (i.e. $\succ$-maximal) worlds among the worlds satisfying ψ:

$$\begin{array}{ll} M, w \models \Box\varphi & \textit{iff} \quad M, u \models \varphi \ \ \forall u \in W(M) \\ M, w \models \bigcirc(\varphi \mid \psi) & \textit{iff} \quad M, u \models \varphi \ \ \forall u \in best_M(\psi) \\ & \textit{where } best_M(\psi) = \{u : M, u \models \psi \ \textit{and} \ \forall u' \succ u : M, u' \not\models \psi\} \end{array}$$

Definition 2. (Limit conditions) *M is* limited *if for any formula φ satisfied in some world of M, $best_M(\varphi) \neq \emptyset$. M is* smooth *if for any φ and any world w satisfying φ, either $w \in best_M(\varphi)$ or there is $u \in best_M(\varphi)$ such that $u \succ w$.*

The weakest Åqvist logic, **E**, imposes no conditions on $\succeq$; **F** imposes only *limitedness* to rule out contradictory obligations; and the strongest system, **G**, assumes that $\succeq$ is a smooth preorder, as is standard for counterfactuals in the Lewisian tradition. Hence **G** coincides with Lewis' $\mathbb{VTA}$ [18], while **E** and **F** do not coincide with any known counterfactual logics due to the possible non-transitivity of $\succeq$. **F** + (**CM**) is a further addition to the Åqvist family that drops the totalness requirement from **G**, and thus lies strictly between **F** and **G** in deductive strength; in fact, it coincides with $\mathbb{PCLTA}$, a known extension of $\mathbb{PCL}$ [15]. All four Åqvist systems enjoy the finite model property [21]. Moreover, as shown in [24], for finite-model characterizations the limit conditions can be replaced by more natural frame conditions: for **F** + (**CM**) and **G**, smoothness is entailed by finiteness and transitivity, and for **F**, limitedness can be replaced by acyclicity. The table 1 summarizes the semantic characterizations for the systems.

Åqvist's logic	Original characterization [21]	Finite-model characterization [24]
E	arbitrary models	finite models
F	limited models	finite, acyclic models
F + (**CM**)	smooth, transitive models	finite, transitive models
G	smooth, transitive, total models	finite, transitive, total models

Table 1. Semantical characterizations of Åqvist's logics.

To assess the computational complexity of conditional logics, [10] introduced *small model constructions* showing that every satisfiable formula is satisfied by some model of bounded (polynomial) size w.r.t. the size of the formula. In particular, they gave constructions for $\mathbb{VTA}$ (i.e., **G**) and $\mathbb{PCLTA}$ (i.e., **F** + (**CM**)) with, respectively, a linear and a quadratic number of worlds. [24] extended this approach to all Åqvist logics, providing alternative constructions with a quadratic number of worlds. Our Deo-SMT will rely on these bounds.

3 Tool Description

The structure of Deo-SMT is in Figure 1.

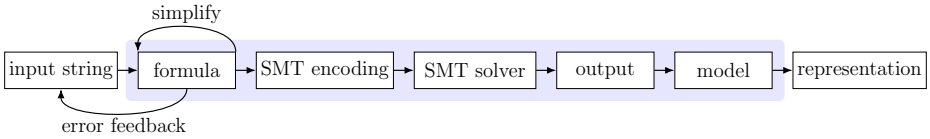


Fig. 1. Tool Structure

Parsing The input text is translated into a logic formula using a custom parser for the obligation syntax, supporting infix notation with symbolic operators (e.g. `&`, `/`, `~`); syntactic errors are reported with their locations.

Formula Simplification The parsed formula is simplified to minimize its size.

Model Encoding The simplified formula is encoded into classical propositional logic. This encoding depends on the selected Åqvist logic and is then passed to the SMT-solver Z3 [9].

Minimal countermodel extraction If the formula is not valid, a countermodel of minimal size is extracted from the model returned by Z3.

Countermodel representation The extracted countermodel is presented in a human-readable form. Different representations are used.

Formula Simplification. It reduces encoding size and eliminates semantic redundancies, improving performance and clarifying the formula structure. In particular, reducing modal operators lowers the upper bound on the number of worlds needed for a countermodel, achieving the largest reduction in encoding size. The syntactic and semantic simplification rules are applied bottom-up so that subformulas are simplified before their parents, avoiding missed opportunities in nested contexts. For example, $\neg\Box\top \wedge (\top \to \Box\Box\phi) \rightsquigarrow \bot \wedge \Box\phi \rightsquigarrow \bot$.

Model Encoding The encoding into classical propositional logic, follows the approach in [24]. For any formula φ and Åqvist logic $\mathcal{L}$ define a propositional formula $F_{\mathcal{L}}(\varphi)$, which is valid iff φ is valid in $\mathcal{L}$. $F_{\mathcal{L}}(\varphi)$ encodes a countermodel M for φ with worlds $\{w_1 \dots w_{N(\varphi)}\}$, where $N(\varphi)$ is the size bound derived from the small model constructions (see below). The encoding uses the following variables.

- $p_{i,j}$ for $1 \le i, j \le N(\varphi)$ encoding $w_i \succeq w_j$
- v_i^{ψ} for $1 \le i \le N(\varphi)$ and $\psi \in \mathit{Sub}\mathcal{F}(\varphi)$, where $\mathit{Sub}\mathcal{F}(\varphi)$ denotes the subformulas of φ, encoding $M, w_i \models \psi$

If ψ is not a propositional variable, v_i^{ψ} depends on the v-variables for the immediate subformulas of ψ. The dependences are encoded by the following set of propositional equivalences:

$$
\begin{aligned}
C^{ev}(\varphi) = \quad & \{ \quad v_i^{\neg\psi} \Leftrightarrow (\neg v_i^{\psi}) \quad | \quad (\neg\psi) \in \mathit{Sub}\mathcal{F}(\varphi) \quad \}_{1\le i\le N(\varphi)} \\
\cup \{ \quad & v_i^{\psi_1\wedge\psi_2} \Leftrightarrow (v_i^{\psi_1} \wedge v_i^{\psi_2}) \quad | \quad (\psi_1 \wedge \psi_2) \in \mathit{Sub}\mathcal{F}(\varphi) \quad \}_{1\le i\le N(\varphi)} \\
\cup \{ \quad & v_i^{\psi_1\vee\psi_2} \Leftrightarrow (v_i^{\psi_1} \vee v_i^{\psi_2}) \quad | \quad (\psi_1 \vee \psi_2) \in \mathit{Sub}\mathcal{F}(\varphi) \quad \}_{1\le i\le N(\varphi)} \\
\cup \{ \quad & v_i^{\psi_1\to\psi_2} \Leftrightarrow (v_i^{\psi_1} \to v_i^{\psi_2}) \quad | \quad (\psi_1 \to \psi_2) \in \mathit{Sub}\mathcal{F}(\varphi) \quad \}_{1\le i\le N(\varphi)} \\
\cup \{ \quad & v_i^{\Box\beta} \Leftrightarrow (\textstyle\bigwedge_{j=1}^{N(\varphi)} v_j^{\beta}) \quad | \quad (\Box\beta) \in \mathit{Sub}\mathcal{F}(\varphi) \quad \}_{1\le i\le N(\varphi)} \\
\cup \{ \quad & v_i^{\bigcirc(\gamma \mid \alpha)} \Leftrightarrow (\textstyle\bigwedge_{j=1}^{N(\varphi)} (v_j^{\gamma} \vee \neg v_j^{\alpha} \vee (\bigvee_{t=1}^{N(\varphi)} (p_{t,j} \wedge \neg p_{j,t} \wedge v_t^{\alpha}))) \\
& \qquad\qquad\qquad\qquad | \quad \bigcirc(\gamma \mid \alpha) \in \mathit{Sub}\mathcal{F}(\varphi) \quad \}_{1\le i\le N(\varphi)}
\end{aligned}
$$

To encode transitivity and totality of $\succeq$, we use the following sets of formulas:

$$\mathcal{C}^{trans}(\varphi) = \{p_{i,j} \wedge p_{j,k} \Rightarrow p_{i,k}\}_{1 \leq i,j,k \leq N(\varphi)} \quad \mathcal{C}^{total}(\varphi) = \{p_{i,j} \vee p_{j,i}\}_{1 \leq i,j \leq N(\varphi)}$$

The relation $\succeq$ is acyclic if there exists a relation $\succeq^t$ that is transitive, irreflexive and contains $\succeq$. Using additional variables $t_{i,j}$ for $1 \leq i,j \leq N(\varphi)$ encoding the fact $w_i \succeq^t w_j$, we encode acyclicity with the following set of formulas.

$$\begin{aligned} C^{acyclic}(\varphi) = \quad & \{\ (t_{i,j} \wedge t_{j,k}) \Rightarrow t_{i,k}\ \}_{1 \leq i,j,k \leq N(\varphi)} \\ \cup\ & \{\ \neg t_{i,i}\ \}_{1 \leq i \leq N(\varphi)} \\ \cup\ & \{\ p_{i,j} \Rightarrow t_{i,j}\ \}_{1 \leq i,j \leq N(\varphi)} \end{aligned}$$

The encoding is then passed to the Z3 solver that determines whether a countermodel exists for the formula, thereby showing that it is invalid.

Size Bound Calculation The required number of worlds in a countermodel forms the basis for encoding formulas into propositional logic. The calculations of the size bound for each logic are based on the small model constructions in [24]. We introduce the notation: $\text{Nec}(\varphi) = \{\,\Box\phi \in \mathit{Sub}\mathcal{F}(\varphi)\,\}$, $\text{Obl}(\varphi) = \{\,\bigcirc(\psi \mid \phi) \in \mathit{Sub}\mathcal{F}(\varphi)\,\}$, $\mathit{Cond}(\varphi) = \{\alpha \mid \bigcirc(\gamma \mid \alpha) \in \mathit{Sub}\mathcal{F}(\varphi)\}$. Without determining for each subformula $\Box\phi$ and $\bigcirc(\psi \mid \phi)$ whether it is satisfied, we use the following upper bound on the number of falsifying worlds $\text{Fal}(\varphi, M)$.

$|\text{Fal}(\varphi, M)| \leq 1 + |\text{Nec}(\varphi)| + |\text{Obl}(\varphi)|$

The small model constructions yield upper bounds on the size of countermodels in each logic, given by the following formulas.

$$\begin{aligned} \mathit{sizeBound}_E &= 3 \cdot |\mathit{Cond}(\varphi)| + |\mathit{Fal}(\varphi, M)| \cdot (|\mathit{Cond}(\varphi)| + 1) \\ \mathit{sizeBound}_F &= |\mathit{Cond}(\varphi)| + |\mathit{Fal}(\varphi, M)| \cdot (|\mathit{Cond}(\varphi)| + 1) \\ \mathit{sizeBound}_{F+(CM)} &= |\mathit{Fal}(\varphi, M)| \cdot (|\mathit{Cond}(\varphi)| + 1) \end{aligned}$$

For the logic $\mathbb{VTA}$ (i.e. **G**), Friedman and Halpern [10] use a different approach for constructing a countermodel. They take one world from the original model for each modality. This results in a linear size bound of the model and the following formula, which we use in our tool: $\mathit{sizeBound}_G = 1 + |\text{Nec}(\varphi)| + |\text{Obl}(\varphi)|$.

Finding a Minimal Countermodel When reasoning about a formula, we call Z3 multiple times. To establish validity, we first attempt to construct a countermodel using the maximum number of worlds allowed by the encoding.

To encode the countermodels for all of Åqvist's logics, we combine the sets of formulas in the following way , adding $\neg v_1^\varphi$ to impose that φ is falsified in the model in the world #1.

$$\begin{aligned} F_{\mathbf{E}}(\varphi) &= \neg v_1^\varphi \wedge \bigwedge \mathcal{C}^{ev}(\varphi) \\ F_{\mathbf{F}}(\varphi) &= \neg v_1^\varphi \wedge \bigwedge \mathcal{C}^{ev}(\varphi) \wedge \bigwedge \mathcal{C}^{acyclic}(\varphi) \\ F_{\mathbf{F+(CM)}}(\varphi) &= \neg v_1^\varphi \wedge \bigwedge \mathcal{C}^{ev}(\varphi) \wedge \bigwedge \mathcal{C}^{trans}(\varphi) \\ F_{\mathbf{G}}(\varphi) &= \neg v_1^\varphi \wedge \bigwedge \mathcal{C}^{ev}(\varphi) \wedge \bigwedge \mathcal{C}^{trans}(\varphi) \wedge \bigwedge \mathcal{C}^{total}(\varphi) \end{aligned}$$

If no countermodel exists at this upper bound, then no countermodel exists at all and the formula is valid. Since most formulas can be falsified with few worlds, we perform encodings and SMT calls with an increasing number of worlds, starting from 1, until a minimal model is found.

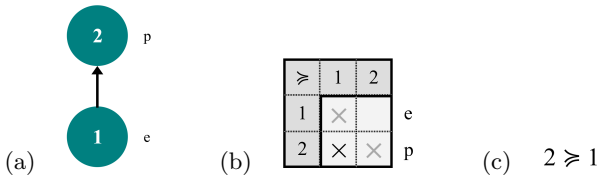


Fig. 2. Representations of the GDPR model. (a) Graph. (b) Matrix. (c) Text.

4 Countermodel Representation

To illustrate how Deo-SMT presents countermodels, following [4], we use the norms below from the GDPR (Regulation EU 2016/679) as a case study (we denote "processed lawfully" by p and "erase the data" by e):

1. Personal data shall be processed lawfully, i.e. $\bigcirc(\mathrm{p} \mid \top)$
2. If the personal data have been processed unlawfully, the controller has the obligation to erase the personal data in question $\bigcirc(\mathrm{e} \mid \neg\mathrm{p})$
3. It is obligatory to keep the personal data provided that it is processed lawfully $\bigcirc(\neg\mathrm{e} \mid \mathrm{p})$

Assume now that some data in the context of such an agreement has been processed unlawfully, i.e. ¬p. Norms 1 and 2 form a standard Contrary-To-Duty pattern, and norms 1-3 combined with the fact ¬p are inconsistent in the best known deontic logic: Standard Deontic Logic [27] (aka the normal modal logic KD). By contrast, in Åqvist systems 1-4 are consistent. E.g., in the logic **E**, the tool produces a two-world model $M = \langle W, \succeq, \mathbb{V}\rangle$ in which all four norms hold simultaneously. The formal definition of this model (see Definition 1) is: $W = \{w_1, w_2\}$, $\succeq = \{(w_2, w_1)\}$, and $\mathbb{V}(w_1) = \{\mathrm{p}\}$ $\mathbb{V}(w_2) = \{\mathrm{e}\}$.

Deo-SMT visualizes the model as (a) a directed graph, (b) a matrix, and (c) text only, as demonstrated in Fig. 2.

Directed graph. In the graph view, worlds are nodes labeled by their number (i.e., node i representing world w_i), and an edge from i to j indicates $w_j \succeq w_i$. For larger models, and especially those with a dense preference relation, graphs are hard to read. For small models (up to three worlds), crossing-free layouts are used, either linear or triangular, with better worlds placed higher when possible.

Matrix. Countermodels can also be shown as a matrix with rows and columns indexed by worlds. An entry in column i and row j, indicates that $w_i \succeq w_j$, i.e. that world w_i is preferred over world w_j. We improve the readability of the matrix representation, especially for bigger models, by an interactive user interface. Indeed, hovering over a matrix entry or world number highlights the corresponding better and worse worlds and the true variables in the corresponding worlds.

Text representation. The preference relation can also be presented textually by listing preference statements in formal notation. Instead of listing each pair separately, related statements are combined into compact expressions, e.g., merging $w_1 \succeq w_2$ and $w_3 \succeq w_1$ into $w_3 \succeq w_1, w_2$.

To construct these compact relations, worlds with identical sets of better or inferior worlds are grouped together, and their relations are listed jointly.

For efficiency, the sizes of these sets are tracked to quickly identify and merge matching groups.

Remark 1. The most suitable representation for a countermodel depends on its size. For models with up to three worlds, all three representations are shown. For larger models with more than three worlds, the graph representation is omitted. The text representation of the preference relation remains useful for quickly identifying the most and least preferred worlds of a model.

5 Comparison to Existing Tools

We compare Deo-SMT to the existing automated approach by Parent and Benzmüller [23,22,3,2] using Isabelle/HOL. Their work embeds Åqvist's logics into higher-order logic via a shallow semantic embedding, enabling both automated and interactive reasoning about the logics and their properties. Isabelle/HOL, an interactive theorem prover based on classical higher-order logic, provides a suitable meta-logic for formalizing such complex systems. Their approach leverages this embedding to express and automatically verify properties and correspondences using higher-order reasoning.

The comparison below is carried out with respect to the following criteria:

- **Goals:** Conceptually, the two approaches pursue complementary goals. The Isabelle/HOL-based framework is designed for formal verification, correspondence theory, and meta-logical analysis of Åqvist's logics, whereas Deo-SMT targets efficient automated deduction, countermodel generation, and user-oriented feedback.
- **Supported logics:** Our tool uniformly supports all Åqvist's logics: **E**, **F**, **F+(CM)**, **G**. In contrast, the Isabelle/HOL approach encodes each logic separately, each requiring a dedicated formalization.
- **Performance:** Deo-SMT uses finite model construction via propositional encoding and SAT-based solving. We tested both tools for the logic **E**, since most Isabelle/HOL implementations target it, on a benchmark of 10,460 formulas which are all possible formulas up to operator nesting depth at most 2, modulo equivalence under variable renaming, associativity and commutativity. On this benchmark, Deo-SMT reduces average runtime by a factor of $\times 31.8$ compared to the Isabelle/HOL-based solver of Parent and Benzmüller [22], when performing validity checking and countermodel construction under the same formula-dependent size bounds.
- **Simplification:** Deo-SMT simplifies the input formula, which improves readability and also computational performance.
- **Model readability:** Deo-SMT provides countermodels in three human-readable formats (graphical, matrix, and text), while Isabelle/HOL uses Nitpick with lambda terms and explicit relation assignments.
- **Minimal models:** Our tool automatically finds a minimal model.
- **Input format:** Deo-SMT operates directly on formulas written in the logic's syntax. By contrast, Isabelle/HOL requires familiarity with formal proof environments and function syntax.

Acknowledgements. The work is partially supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No 101034440, the FWF project I 6372-N (Grant-DOI 10.55776/I6372), and the WWTF project ICT22-023 (Grant ID: 10.47379/ICT22023).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Åqvist, L.: Deontic logic. In: Gabbay, D., Guenthner, F. (eds.) Handbook of Philosophical Logic: Volume II, pp. 605–714. Springer, Dordrecht (1984)
2. Benzmüller, C., Farjami, A., Parent, X.: Åqvist's dyadic deontic logic E in HOL. FLAP **6**(5), 733–754 (2019)
3. Benzmüller, C., Parent, X., van der Torre, L.W.N.: Designing normative theories for ethical and legal reasoning: Logikey framework, methodology, and tool support. Artif. Intell. **287**, 103348 (2020)
4. Benzmüller, C., Parent, X., van der Torre, L.W.N.: A deontic logic reasoning infrastructure. In: Manea, F., Miller, R.G., Nowotka, D. (eds.) Sailing Routes in the World of Computation - CiE 2018. Lecture Notes in Computer Science, vol. 10936, pp. 60–69. Springer (2018)
5. Burgess, J.: Quick completeness proofs for some logics of conditionals. Notre Dame Journal of Formal Logic **22**(3), 76–84 (1981)
6. Ciabattoni, A., Olivetti, N., Parent, X.: Dyadic obligations: proofs and countermodels via hypersequents. In: Proceedings of PRIMA 2022. pp. 54–71. Springer (2022)
7. Ciabattoni, A., Rozplokhas, D., Tesi, M.: GL-based calculi for pcl and its deontic cousin. In: Proceedings of JELIA 2025 (2025)
8. Ciabattoni, A., Tesi, M.: Sequents vs hypersequents for Åqvist systems. In: Proceedings of IJCAR 2024. pp. 176–195. Springer (2024)
9. De Moura, L., Bjørner, N.: Z3: An efficient smt solver. In: International conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 337–340. Springer (2008)
10. Friedman, N., Halpern, J.Y.: On the complexity of conditional logics. In: Proceedings of KR 1994. pp. 202–213. Elsevier (1994)
11. Gabbay, D.M.: Theoretical foundations for non-monotonic reasoning in expert systems. In: Apt, K.R. (ed.) Logics and Models of Concurrent Systems. pp. 439–457. Springer Berlin Heidelberg, Berlin, Heidelberg (1985)
12. Gabbay, D.M., Horty, J., Parent, X., van der Meyden, R., van der Torre, L. (eds.): Handbook of Deontic Logic and Normative Systems. College Publications (2013)
13. Gabbay, D.M., Horty, J.F., Parent, X., van der Meyden, R., van der Torre, L. (eds.): Handbook of Deontic Logic and Normative Systems, Volume 2. College Publications, London (2021)
14. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Analytic tableaux calculi for KLM logics of nonmonotonic reasoning. ACM Trans. Comput. Log. **10**(3), 18:1–18:47 (2009)
15. Girlando, M.: On the Proof Theory of Conditional Logics. Theses, Aix-Marseille Universite; Helsinki University (Feb 2019)

16. Hansson, B.: An analysis of some deontic logics. In: Deontic logic: Introductory and systematic readings, pp. 121–147. Springer (1971)
17. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. Artif. Intell. **44**(1-2), 167–207 (1990)
18. Lewis, D.: Counterfactuals, blackwells (1973)
19. Makinson, D., van der Torre, L.W.N.: Input/output logics. J. Philos. Log. **29**(4), 383–408 (2000)
20. Parent, X.: Maximality vs. optimality in dyadic deontic logic. Journal of Philosophical Logic **43**, 1101–1128 (2014)
21. Parent, X.: Preference semantics for Hansson-type dyadic deontic logic: a survey of results. In: Gabbay, D., Horty, J., Parent, X., van der Torre, L., van der Meyden, R. (eds.) Handbook of Deontic Logic and Normative Systems (vol. 2), pp. 7–70. College Publications, London (2021)
22. Parent, X., Benzmüller, C.: Automated verification of deontic correspondences in isabelle/hol-first results. In: Proceedings of Automated Reasoning in Quantified Non-Classical Logics, 4th International Workshop (associated with FLoC and IJCAR 2022). pp. 92–108 (2022)
23. Parent, X., Benzmüller, C.: Conditional normative reasoning as a fragment of HOL. J. Appl. Non Class. Logics **34**(4), 561–592 (2024)
24. Rozplokhas, D.: Lego-like small model constructions for åqvist's logics. In: Ciabattoni, A., Gabelaia, D., Sedlár, I. (eds.) Advances in Modal Logic, pp. 631–651. College Publications (2024)
25. Shoham, Y.: A semantical approach to nonmonotic logics. In: Proceedings of the Symposium on Logic in Computer Science (LICS '87). pp. 275–279. IEEE Computer Society (1987)
26. Steen, A.: A reduction of input/output logics to sat. Journal of Applied Non-Classical Logics pp. 1–34 (04 2026)
27. Von Wright, G.H.: Deontic logic. Mind **60**(237), 1–15 (1951)