

Secondary Publication



Lawniczak, Lara; Benzmüller, Christoph

Logical Modalities within the European AI Act : An Analysis

Date of secondary publication: 06.02.2026

Version of Record (Published Version), Conferenceobject

Persistent identifier: urn:nbn:de:bvb:473-irb-112989x

Primary publication

Lawniczak, Lara; Benzmüller, Christoph (2025): Logical Modalities within the European AI Act : An Analysis, in: Juliano Maranhão (Ed.), ICAIL '25: Proceedings of the Twentieth International Conference on Artificial Intelligence and Law, ACM, pp. 348–353, doi: 10.1145/3769126.3769202.

Legal Notice

This work is protected by copyright and/or the indication of a licence. You are free to use this work in any way permitted by the copyright and/or the licence that applies to your usage. For other uses, you must obtain permission from the rights-holders.

This document is made available under a Creative Commons license.



The license information is available online:

<https://creativecommons.org/licenses/by/4.0/legalcode>

Logical Modalities within the European AI Act: An Analysis

Lara Lawniczak
University of Bamberg
Bamberg, Germany
lara.lawniczak@gmx.de

Christoph Benz Müller
AI Systems Engineering
University of Bamberg
Bamberg, Germany
Department of Mathematics and Computer Science
Freie Universität Berlin
Berlin, Germany
christoph.benzmueller@uni-bamberg.de

Abstract

The paper presents a comprehensive analysis of the European AI Act in terms of its logical modalities, with the aim of preparing its formal representation, for example, within the logic-pluralistic Knowledge Engineering Framework and Methodology (LogiKey). LogiKey develops computational tools for normative reasoning based on formal methods, employing Higher-Order Logic (HOL) as a unifying meta-logic to integrate diverse logics through shallow semantic embeddings. This integration is facilitated by Isabelle/HOL, a proof assistant tool equipped with several automated theorem provers. The modalities within the AI Act and the logics suitable for their representation are discussed. For a selection of these logics, embeddings in HOL are created, which are then used to encode sample paragraphs. Initial experiments evaluate the suitability of these embeddings for automated reasoning, and highlight key challenges on the way to more robust reasoning capabilities.

CCS Concepts

• **Computing methodologies** → **Knowledge representation and reasoning**; • **Applied computing** → **Law**; • **Theory of computation** → *Higher order logic*; *Automated reasoning*.

Keywords

European AI Act, Modalities, Non-classical logics, Higher-order Logic, Proof Assistants, Automated Reasoning

ACM Reference Format:

Lara Lawniczak and Christoph Benz Müller. 2025. Logical Modalities within the European AI Act: An Analysis. In *20th International Conference on Artificial Intelligence and Law (ICAIL 2025)*, June 16–20, 2025, Chicago, IL, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3769126.3769202>

1 Introduction

In the last decade, the representation of, and the reasoning with, legal information has gained growing attention in the AI and Law community. Alongside established approaches like Akoma Ntoso [29] and LegalRuleML [3, 30], research is currently being invested in the exploration and application of the logic-pluralistic knowledge engineering framework and methodology LogiKey [8], which

develops computational tools for normative reasoning based on formal methods.

An interesting and relevant use case for the LogiKey framework is the European AI Act [16], which came into force in August 2024. The AI Act classifies AI systems into three risk levels, and assigns different rules and regulations to each category. If successful, the logical formalisation of key concepts of the AI Act within the LogiKey framework could, for example, prepare the way for sophisticated automated or interactive compliance checks and support in various ways the technical enforcement of the new regulation. It is not only the AI Act itself that is of interest, but also the specific follow-up regulations and standardisations that will be triggered by the AI Act, which will then impose specific regulatory constraints on concrete applications of AI systems.

Before adequate logical representations can be developed, it is essential to identify the modalities that occur in the AI Act and the logics, or combinations of logics, suited to adequately represent these. This paper addresses this challenge as its main contribution. In the long run, however, it is rather the concrete *instantiations* of the abstract legislation of the AI Act, i.e. the more concrete laws that are expected to follow from the adoption of the AI Act in the following years, that are relevant and interesting for formalisation. We expect that the findings on the identified modalities will (to a large extent) overlap and apply to them as well.

The paper is organized as follows. Section 2 lists and discusses the modalities found in the relevant parts of the AI Act. The findings of this section are relevant beyond the LogiKey context. Section 3 explores potential logic systems for representing these modalities. Section 4 briefly presents preliminary results on embedding selected logics in the LogiKey framework and using them to formalize some exemplary sections of the AI Act. Section 5 concludes the paper.

2 Modalities in the AI Act

This section presents the modalities (and other challenges) identified in the AI Act and illustrates them with examples. In order to find all the modalities in the document, it was read several times, with different foci of interest, and visualized with the help of mind maps and tables.



This work is licensed under a Creative Commons Attribution 4.0 International License. *ICAIL 2025, Chicago, IL, USA*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1939-4/25/06

<https://doi.org/10.1145/3769126.3769202>

2.1 Obligations, Prohibitions, and Permissions

The AI Act is a European piece of legislation intended to regulate the use of AI systems. In it *Obligations* are expressed with the word "shall",¹ as in this sentence from Article 8:

"High-risk AI systems shall comply with the requirements established in this section." [16, §8(1)]

The negated shall, "shall not", is employed to express *Prohibitions*. An example is:

"Where an importer has sufficient reason to consider that a high-risk AI system is not in conformity with this Regulation, or is falsified, or accompanied by falsified documentation, it shall not place the system on the market until it has been brought into conformity." [16, §23(2)]

For the expression of *Permissions*, two strategies can be identified in the AI Act: They are introduced by either "may" or "is/are empowered to".

2.2 Contrary-to-Duty-Obligations

A special type of obligations that occurs within the AI Act deserves attention: *contrary-to-duty obligations* (CTDs) are obligations that arise only if a primary obligation is not fulfilled, meaning that it is "conditional on [...] violating [a] primary obligation" [27].

An example of a CTD in the AI Act can be found in Article 20:

"Providers of high-risk AI systems which consider or have reason to consider that a high-risk AI system that they have placed on the market or put into service is not in conformity with this Regulation shall immediately take the necessary corrective actions to bring that system into conformity, to withdraw it, to disable it, or to recall it, as appropriate." [16, §20(1)]

It has been stated before that high-risk systems must comply with the requirements [16, §16(1)]. This is the primary obligation in the given context. However, obligations are not always fulfilled; they may also be violated. In the event of such a violation, the obligation to take corrective action becomes relevant. This is a CTD, applicable only when the primary obligation has been violated. CTD situations are often represented in a typical structure following Chisholm [14]. The discussed example then looks as follows:

- (1) It ought to be that high-risk AI systems comply with the requirements in the regulation.
- (2) It ought to be that if a high-risk AI system does not comply with the requirements, providers take corrective actions.
- (3) If a system complies with the requirements, the provider must not take any corrective actions.²
- (4) Concrete Situation: The system does not comply with the requirements.

Notice that some of these obligations in this example are agentive for the provider. The aspects of agency and agentive obligations are

¹Within legal acts such as the AI Act, the word "shall" can also have performative or constitutive functions instead of deontic ones. However, the primary function of "shall" in the AI Act, is the introduction of deontic obligations.

²It has been brought to our attention that an interpretation as "need not" might be more appropriate. We are thankful for this remark and aim to uncover the differences between these interpretations in the future.

discussed separately in Section 3. First, two special kinds of CTDs are considered.

2.2.1 Contrary-to-Duty-Obligations Involving Multiple Agents. Some CTDs within the AI Act relate to multiple agents, which complicates matters, as illustrated by the following example from §36:

"If the notifying authority comes to the conclusion that the notified body no longer meets the requirements laid down in Article 31 or that it is failing to fulfil its obligations, it shall restrict, suspend or withdraw the designation as appropriate, depending on the seriousness of the failure to meet those requirements or fulfil those obligations." [16, §36(4)]

For simplicity, the notion of an agent belief, expressed in the words *comes to the conclusion that*, is ignored (since beliefs are the subject of Section 4). Also, disregard the temporal notion contained in *no longer*, since CTDs involving temporality are discussed in Section 2.2.2.

The focus here lies on the involvement of two distinct agents: The notifying authority and the notified body. The situation is as follows: The notified body has several primary obligations that are specified within the AI Act [16, §31]. If it violates these obligations, another obligation arises. However, this obligation is agentive for the notifying authority, not for the notified body that violated the primary obligation.

2.2.2 Contrary-to-Duty-Obligations Involving Temporality. There are cases in the AI Act where CTDs occur in combination with temporal constraints, stating that a certain property used to be fulfilled but *no longer* is. An example is:

"Where a notifying authority has sufficient reason to consider that a notified body no longer meets the requirements laid down in Article 31, or that it is failing to fulfil its obligations, the notifying authority shall without delay investigate the matter with the utmost diligence." [16, §36(4)]

To understand the CTD, §30 must be taken into account. There it is stated that the notifying authority "may notify only conformity assessment bodies which have satisfied the requirements laid down in Article 31" [16, §30(1)]. Unfortunately, this example does not involve only temporality but also multiple agents as discussed in the Section 2.2.1. To focus on only temporality, we reformulate the regulation as follows:

Only conformity assessment bodies that fulfill the requirements in Article 31 before the notification may be notified (adapted from Article 30). If a conformity assessment body that was notified (= notified body) no longer meets the requirements laid down in Article 31, the matter shall be investigated further with the utmost diligence (adapted from Article 36).

Now it becomes visible how this example is different from a typical CTD: Technically speaking, the notified body is only obligated to fulfill the requirements in Article 31 at one point in time *before* being notified. Hence, the fact that a notified body *no longer* fulfills the requirements in Article 31 does not equal a violation of the previous obligation. The obligation to further investigate the matter

then is just a normal obligation, not a CTD, and could be expressed as such. In this case, the logic representing this situation must not only provide adequate obligation operators, but must also be able to express temporality in order to properly capture the notion of *before* (as discussed in Section 2.5).

Another possibility is disregarding the temporality here and translating the example to match the usual CTD structure:

- (1) It ought to be that notified bodies fulfill the requirements in Article 31.
- (2) It ought to be that if a notified body does not comply with the requirements in Article 31, further investigations are started.
- (3) If a notified body does comply with the requirements in Article 31, no further investigation must be started.
- (4) Concrete Situation: A notified body does not fulfill the requirements in Article 31.

Reformulated as above, the example can be treated like a typical CTD.

2.3 Agency and Agentive Obligations

Since different agents interact with an AI system during its lifetime, it is plausible that these agents have different duties towards the system based on their relation to it. Such obligations then are not general but agentive: They only hold for a specific type of agent, e.g. for providers or importers. For instance, the sentence

"Providers of high-risk AI systems shall (...) have a quality management system in place which complies with Article 17" [16, §16(c)]

is such an agentive obligation because it expresses what the provider ought to do. Agentive obligations are different from general obligations, and must be distinguishable.

2.4 Beliefs of Agents

Apart from agentive obligations, another modality relating to agency is crucial in the AI Act: Agent beliefs. For example, Article 20 states:

"Providers of high-risk AI systems which consider or have reason to consider that a high-risk AI system that they have placed on the market or put into service is not in conformity with this Regulation shall immediately take the necessary corrective actions to bring that system into conformity, to withdraw it, to disable it, or to recall it, as appropriate." [16, §20(1)]

Here, the provider's belief that a high-risk AI system is not in conformity with the regulation is relevant, not whether the system is, in fact, not conforming. Beliefs can be either right or wrong, and it is essential to differentiate them from facts.

2.5 Temporality and Temporal Notions

Time is an important concept within the AI Act, with many articles specifying obligations that must be fulfilled before certain events occur. For instance, Article 23 states:

"Before placing a high-risk AI system on the market, importers shall ensure that the system is in conformity with this Regulation by verifying that (...) "[16, §23(1)]

This obligation for importers is tied to a specific period and point in time; it must be fulfilled *before* the system is placed on the market.

Additionally, some sentences contain hidden temporal notions, such as this phrase in Article 9: "When implementing the risk management system as provided for in paragraphs 1 to 7, providers shall..." [16, §9(9)]. The sentence implies simultaneity which becomes obvious if it is rephrased as follows: "While implementing the risk management system, providers are obligated to (...)"

2.6 Exceptions

Within the AI Act, the concept of exceptions from general rules emerges as a recurring theme. Consider this sentence from Article 5:

"The following AI practices shall be prohibited: (...) the use of 'real-time' remote biometric identification systems in publicly accessible spaces for the purposes of law enforcement, unless and in so far as such use is strictly necessary for one of the following objectives:..." [16, §5(1)]

This sentence outlines a general rule – the prohibition of certain AI practices – followed by exceptions based on specific objectives.

2.7 Fuzziness

Fuzzy statements can be found in many places within the AI Act. They can be identified by their vagueness and usually involve statements that are not true or false, but true or false to a certain degree. The following expressions are examples of such fuzziness:

"unless and in as far as such use is strictly necessary for one of the following objectives" [16, §5(1)], "to the extent to which..." [16, §7], and "at a minimum" [16, §11(1)].

All these statements express vague notions or involve degrees of truth or necessity.

2.8 Combinations of Modalities

The modalities of the AI Act do not appear in isolation. Some of the examples discussed in Section 2.2 have already shown that two or more modalities are often combined within a single sentence. While this was ignored in the previous discussion for the sake of simplicity, it cannot be ignored when constructing a logical system intended to cover the entire AI Act, or larger coherent parts of it. Such a system must not only accurately represent all the relevant modalities individually, but also capture what happens when they are combined.

3 Which Logics?

In this section we discuss logics that seem suitable to facilitate a representation of the AI Act in HOL.

Standard Deontic Logic (SDL) is the most studied system of deontic logic. It contains a monadic deontic obligation operator depending on an accessibility relation and can successfully express and reason with obligations, permissions, and prohibitions [27]. However, SDL reaches its limits and leads to inconsistencies in CTD scenarios [18].

A logic that enables expressing and reasoning with CTDs is Dyadic Deontic Logic (DDL) by Carmo and Jones [12, 13]. It differentiates between ideal and actual obligations and introduces a conditional obligation operator, thereby enabling the expression of CTDs that arise when ideal obligations have been violated. Related to DDL is Åqvist’s system E for conditional obligation [2].

To represent agentive obligations, a modal logic of agency is needed. While there are many ideas for formulating such a logic [1, 11, 17, 23, 39, 40], the most prominent theory of agency is a branch called Seeing-To-It-That (STIT) logic that originated in the works of Belnap [4] and Belnap and Perloff [5]. STIT theory introduces a STIT operator of the form $stit_a F$, which expresses that an agent a sees to it that F holds. Given that agency appears in the AI Act mainly via agentive obligations, a suitable agency logic must be able to formulate ought-to-do statements for specific agents. Several authors have researched this issue and presented their approaches, which could be apt to help represent the AI Act [21, 22, 28, 38, 41].

Moving to the representation of agent beliefs, the field called epistemic logic becomes interesting. Epistemic logic allows for the exploration of different ideas of knowledge and beliefs of agents and their logical relations [35]. Popular epistemic logics include modal logics S4 and S5 based on possible world semantics, and providing modal operators representing what agents *know* and *believe*.

For the representation of temporal constructs, we may turn to temporal logic, a class of modal logics that can formally represent and reason with temporality [19]. A popular system is called Tense Logic (TL) and dates back to the work of Prior [26, 32–34]. TL treats propositions as tensed and introduces temporal operators into the language, thereby allowing for the representation of temporal relations like the ones in the AI Act.

Non-monotonic logic provides a way to deal with exceptions from general rules. It is designed to handle defeasible inference, enabling reasoners to retract conclusions when warranted by additional information, such as the presence of conditions allowing for an exception [37].

Finally, a logic that can do justice to vague and fuzzy statements is called fuzzy logic. Fuzzy logic assigns a degree of truth to a proposition that is expressed on a scale within the real unit interval $[0, 1]$, with 0 equal to total falsehood and one equal to total truth [15]. From the point of view of LogiKey, however, the fuzzy logic variants of Hájek [20] are particularly interesting, since they can be seen as generalized multi-value logics and as such are more amenable to formalization; in fact, [20] provides a good blue-print for doing so in LogiKey extending the formalization idea presented in [36].

4 Initial Experiments

In this section we give a brief summary of the findings and results from the experiments on embedding some of the logics using the LogiKey approach and using these logics to represent parts of the AI Act; for details see [24]. Experiments have so far focused on only a selection of the identified modalities, namely obligations, prohibitions, permissions, CTDs, agency, and agentive obligations.

LogiKey leverages higher-order logic (HOL) as a unifying meta-logic to enable the modelling of diverse object logics via shallow semantic embeddings [6]. This integration is facilitated, for example, within the higher-order proof assistant tool *Isabelle/HOL*, which

comes with state-of-the-art automated theorem provers (ATP), satisfiability modulo theories (SMT) solvers, and model finders; however, other proof assistant systems, such as Lean or Coq, or the TPTP infrastructure, could be employed instead.

To represent obligations (as well as permissions and prohibitions), SDL was identified as a suitable logic in Section 3. A trustworthy embedding of SDL in Isabelle/HOL using the LogiKey approach already exists [7] and has been used to represent parts from the AI Act, particularly Article 5 on prohibited AI systems. This was successful: The paragraphs could be adequately represented, and both Sledgehammer (an automated theorem proving tool integrated with Isabelle/HOL) and Nitpick [10] (a model finder integrated with Isabelle/HOL) could reason with them correctly [24]. Sledgehammer was used to automatically check the validity of example judgements, and Nitpick to find countermodels for invalid judgements and also to prove the consistency of whole contexts.

Like SDL, DDL already has a faithful embedding in Isabelle/HOL [9]. Using this embedding it was possible to represent selected CTDs from the AI Act without problems. The formalization of the selected examples from the AI Act was straightforward, and Nitpick and Sledgehammer could work with them successfully [24]. Since DDL can express everything SDL can and more [18], the DDL embedding is suitable for representing obligations, prohibitions, permissions, and CTDs. Encodings of Article 5 and selected CTDs from the AI Act in DDL can be found in the appendix of the preprint paper [25].³ Further work includes similar experiments adapting an existing embedding of Åqvist’s system E in Isabelle/HOL [31].

Additionally, logics for the representation of agency and agentive obligations were explored. In particular, an embedding into HOL of Temporal Deontic STIT Logic [38], a logic appropriate for representing both actions of agents and agentive obligations, was created. While most axioms postulated in [38] were provable via Sledgehammer and none were refuted, Nitpick could not find a model confirming the consistency of the embedding. This indicates that TDS logic has (presumably) only infinite models, making it unsuitable for model finding/checking in the context of any formalization of AI Act [24] with existing tools.

As an alternative to TDS, the authors explored the possibility of extending DDL with additional operators for the representation of agency and agentive obligations. Two different variants were introduced: One representing agent types as constants, with separate accessibility relations and obligation operators for each agent, and one representing agent types as function types, with a general accessibility relation and a general obligation operator taking an agent as an input parameter. Both variants introduced a STIT-like operator as a constant and equipped it with minimal axiomatization, namely to ensure that anything an agent sees to actually holds [24].

While the first variant is more restrictive, with agentive obligations always applying to a generic agent type without the possibility of specifying further constraints or relations, the second variant is more flexible, allowing groupings based on characteristics of agent types or relations and formulating obligations for specific subsets of agents. Unfortunately, the second variant performs worse in

³The Isabelle/HOL source files can also be downloaded here: <http://logikey.org/tree/master/2025-ICAIL-Data>.

other aspects: It struggles with the representation of CTDs, and Nitpick can not find a model for cardinalities of i (domain of possible worlds) bigger than 1. For the first variant, a model is found up to a cardinality of $i=2$, and CTDs can be represented without problems. For both variants, all but one of the prominent lemmas for DDL as studied by Benzmüller et al. could be proven [24]. Whether either of the variants can be adapted or simplified to perform better remains an open question. The embedding of TDS and the two extensions of DDL can be found in the appendix of [25].

The experiments have shown that while elementary concepts such as obligations and CTDs can be formalized and reasoned with correctly in HOL, the representation of complex constructs such as agency requires more sophisticated logics. Reasoning can be challenging for model finders and automated theorem provers, and further experiments need to be conducted to explore feasibility and automation performance as the studied contexts grow.

5 Conclusion

An extensive analysis of the AI Act was carried out, which led to the identification of different contained modalities that pose a challenge to attempts to formalise and reason with it: obligations, permissions, prohibitions, CTDs, agency and agentive obligations, beliefs of agents, temporality and temporal notions, fuzziness, and exceptions. Suggestions for logics suitable to represent these modalities were presented, and results of some first experiments with these logics and examples from the AI Act were reported. As of now, experiments have only been conducted for the following modalities (and some combinations of): obligations, permissions, prohibitions, CTDs, agency, and agentive obligations.

While the presented project is not yet completed, the reported material provides a starting point and references for future research. In the LogiKey context, the challenge is to create embeddings of the suggested logics, choose and provide suitable logic combinations, and experiment with increasingly larger parts of the AI Act.

Based on the information gathered here, the possibility of representing legal information in HOL and using it in automated reasoning should be further explored and improved.

Acknowledgments

We are grateful to the anonymous reviewers for their helpful feedback and comments.

References

- [1] A. R. Anderson. 1962. Logic, Norms, and Roles. *Ratio (Misc.)* 4, 1 (1962), 36–49.
- [2] Lennart Åqvist. 2002. Deontic Logic. In *Handbook of Philosophical Logic: Volume 8*, D. M. Gabbay and F. Guenther (Eds.). Springer Netherlands, Dordrecht, 147–264. https://doi.org/10.1007/978-94-010-0387-2_3
- [3] Tara Athan, Guido Governatori, Monica Palmirani, Adrian Paschke, and Adam Wyner. 2015. LegalRuleML: Design Principles and Foundations. In *Reasoning Web. Web Logic Rules*, Wolfgang Faber and Adrian Paschke (Eds.), Vol. 9203. Springer, Cham, 151–188. https://doi.org/10.1007/978-3-319-21768-0_6
- [4] Nuel Belnap. 1991. Backwards and Forwards in the Modal Logic of Agency. *Philosophy and Phenomenological Research* 51, 4 (Dec. 1991), 777. <https://doi.org/10.2307/2108182>
- [5] Nuel Belnap and Michael Perloff. 1988. Seeing to It That: A Canonical Form for Agentives. *Theoria* 54, 3 (1988), 175–199. <https://doi.org/10.1111/j.1755-2567.1988.tb00717.x>
- [6] Christoph Benzmüller. 2019. Universal (Meta-)Logical Reasoning: Recent Successes. *Science of Computer Programming* 172 (2019), 48–62. <https://doi.org/10.1016/j.scico.2018.10.008>
- [7] Christoph Benzmüller and Xavier Parent. 2018. First Experiments with a Flexible Infrastructure for Normative Reasoning. <https://doi.org/10.48550/ARXIV.1804.02929>
- [8] Christoph Benzmüller, Xavier Parent, and Leendert van der Torre. 2020. Designing Normative Theories for Ethical and Legal Reasoning: LogiKey Framework, Methodology, and Tool Support. *Artificial Intelligence* 287 (2020), 103348. <https://doi.org/10.1016/j.artint.2020.103348>
- [9] Christoph Benzmüller, Ali Farjami, and Xavier Parent. 2022. Dyadic Deontic Logic in HOL: Faithful Embedding and Meta-Theoretical Experiments. In *New Developments in Legal Reasoning and Logic*, Shahid Rahman, Matthias Armgardt, and Hans Christian Nordtveit Kvernesen (Eds.), Vol. 23. Springer International Publishing, Cham, 353–377. https://doi.org/10.1007/978-3-030-70084-3_14
- [10] Jasmin Christian Blanchette and Tobias Nipkow. 2010. Nitpick: A Counterexample Generator for Higher-Order Logic Based on a Relational Model Finder. In *ITP 2010 (Lecture Notes in Computer Science, Vol. 6172)*, Matt Kaufmann and Lawrence C. Paulson (Eds.), Springer, 131–146. https://doi.org/10.1007/978-3-642-14052-5_11
- [11] Mark A. Brown. 1988. On the logic of ability. *Journal of Philosophical Logic* 17, 1 (Feb. 1988), 1–26. <https://doi.org/10.1007/BF00249673>
- [12] José Carmo and Andrew J. I. Jones. 2002. Deontic Logic and Contrary-to-Duties. In *Handbook of Philosophical Logic: Volume 8*, D. M. Gabbay and F. Guenther (Eds.). Springer, Dordrecht, 265–343. https://doi.org/10.1007/978-94-010-0387-2_4
- [13] José Carmo and Andrew J. I. Jones. 2022. Carmo and Jones’ logic for contrary-to-duty obligations revised. *J. Log. Comput.* 32, 7 (2022), 1352–1364. <https://doi.org/10.1093/LOGCOM/EXAC026>
- [14] R. M. Chisholm. 1963. Contrary-to-Duty Imperatives and Deontic Logic. *Analysis* 24, 2 (1963), 33–36. <https://doi.org/10.1093/analys/24.2.33>
- [15] Petr Cintula, Christian G. Fermüller, and Carles Noguera. 2023. Fuzzy Logic. In *The Stanford Encyclopedia of Philosophy* (Summer 2023 ed.), Edward N. Zalta and Uri Nodelman (Eds.). Metaphysics Research Lab, Stanford University, Stanford, CA, USA.
- [16] EU commission. 2024. AI Act. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- [17] Frederic B. Fitch. 1963. A Logical Analysis of Some Value Concepts. *The Journal of Symbolic Logic* 28, 2 (1963), 135–142. <http://www.jstor.org/stable/2271594>
- [18] D. Gabbay, J. Horty, X. Parent, R. van der Meyden, and L. van der Torre (Eds.). 2013. *Handbook of Deontic Logic and Normative Systems*. Vol. 1. College Publications, London, UK.
- [19] Valentin Goranko and Antje Rumberg. 2023. Temporal Logic. In *The Stanford Encyclopedia of Philosophy* (Fall 2023 ed.), Edward N. Zalta and Uri Nodelman (Eds.). Metaphysics Research Lab, Stanford University, Stanford, CA, USA.
- [20] Petr Hájek. 1998. *Metamathematics of Fuzzy Logic*. Trends in Logic, Vol. 4. Kluwer, Dordrecht. <https://doi.org/10.1007/978-94-011-5300-3>
- [21] John Horty. 2001. *Agency and Deontic Logic*. Oxford University Press, New York.
- [22] John F. Horty and Nuel Belnap. 1995. The Deliberative Stit: A Study of Action, Omission, Ability, and Obligation. *Journal of Philosophical Logic* 24, 6 (1995), 583–644. <http://www.jstor.org/stable/30227231>
- [23] Stig Kanger. 1972. Law and logic. *Theoria* 38, 3 (1972), 105–132. <https://doi.org/10.1111/j.1755-2567.1972.tb00928.x>
- [24] Lara Lawnczak. 2024. *Towards Automated Ethical Reasoning: Representing the AI Act in Higher-Order-Logic*. Master’s thesis. Otto-Friedrichs-Universität Bamberg.
- [25] Lara Lawnczak and Christoph Benzmüller. 2025. Logical Modalities within the European AI Act: An Analysis. <https://doi.org/10.48550/arXiv.2501.19112> arXiv:2501.19112 [cs.AI]
- [26] Czesław Lejewski. 1959. Time and Modality, by A. N. Prior, Clarendon Press: Oxford University Press, 1957. Pp. Viii + 148. *Philosophy* 34, 128 (1959), 56–. <https://doi.org/10.1017/s0031819100029776>
- [27] Paul McNamara and Frederik Van De Putte. 2022. Deontic Logic. In *The Stanford Encyclopedia of Philosophy* (Fall 2022 ed.), Edward N. Zalta and Uri Nodelman (Eds.). Metaphysics Research Lab, Stanford University, Stanford, CA, USA.
- [28] Y. Murakami. 1998. Utilitarian Deontic Logic. In *Advances in Modal Logic*, Marcus Kracht, Maarten de Rijke, Heinrich Wansing, and Michael Zakharyashev (Eds.). CSLI Publications, Stanford, CA, USA, 211–230.
- [29] OASIS Open. 2018. *Akoma Ntoso Version 1.0*. OASIS. <https://www.oasis-open.org/2018/09/11/akoma-ntoso-v1-0-akn-oasis-standard-published/>
- [30] Monica Palmirani, Guido Governatori, Antonino Rotolo, Said Tabet, Harold Boley, and Adrian Paschke. 2011. LegalRuleML: XML-based rules and norms. In *RuleML 2011*, Frank Olken, Monica Palmirani, and Davide Sottara (Eds.), LNCS, Vol. 7018. Springer, Berlin, Heidelberg. <https://doi.org/10.1007/978-3-642-24908-2>
- [31] Xavier Parent and Christoph Benzmüller. 2024. Conditional Normative Reasoning as a Fragment of HOL. *Journal of Applied Non-Classical Logics* 34, 4 (2024), 561–592. <https://doi.org/10.1080/11663081.2024.2386917>
- [32] A. N. Prior. 1959. Thank Goodness That’s Over. *Philosophy* 34, 128 (1959), 12–17. <https://doi.org/10.1017/s0031819100029685>
- [33] Arthur N. Prior. 1967. *Past, Present and Future*. Clarendon P., Oxford.
- [34] Arthur N. Prior. 1968. *Papers on Time and Tense*. Oxford University Press, UK, Oxford, UK.
- [35] Rasmus Rendsvig, John Symons, and Yanjing Wang. 2024. Epistemic Logic. In *The Stanford Encyclopedia of Philosophy* (Summer 2024 ed.), Edward N. Zalta and Uri Nodelman (Eds.). Metaphysics Research Lab, Stanford University, Stanford,

- CA, USA.
- [36] Alexander Steen and Christoph Benzmüller. 2016. Sweet SIXTEEN: Automation via Embedding into Classical Higher-Order Logic. *Logic and Logical Philosophy* 25 (2016), 535–554. <https://doi.org/10.12775/LLP.2016.021>
- [37] Christian Strasser and G. Aldo Antonelli. 2019. Non-monotonic Logic. In *The Stanford Encyclopedia of Philosophy* (summer 2019 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University, Stanford, CA, USA.
- [38] Kees van Berkel and Tim Lyon. 2019. A Neutral Temporal Deontic STIT Logic. In *LORI 2019*, Patrick Blackburn, Emiliano Lorini, and Meiyun Guo (Eds.). LNCS, Vol. 11813. Springer, 340–354. https://doi.org/10.1007/978-3-662-60292-8_25
- [39] Georg Henrik Von Wright. 1963. *Norm and Action: A Logical Enquiry*. Routledge and Kegan Paul, New York, NY, USA.
- [40] Georg Henrik Von Wright. 1981. On the Logic of Norms and Actions. In *New Studies in Deontic Logic: Norms, Actions, and the Foundations of Ethics*, Risto Hilpinen (Ed.). Springer Netherlands, Dordrecht, 3–35. https://doi.org/10.1007/978-94-009-8484-4_1
- [41] Ming Xu. 2015. Combinations of Stit with Ought and Know. *Journal of Philosophical Logic* 44, 6 (2015), 851–877. <https://doi.org/10.1007/s10992-015-9365-7>