

# Ontologies, Arguments, and Large-Language Models

John Beverley<sup>1,2</sup>, Francesco Franda<sup>2,3</sup>, Hedi Karray<sup>2,4</sup>, Dan Maxwell<sup>5</sup>, Carter Benson<sup>1,2</sup>, and Barry Smith<sup>1,2</sup>

<sup>1</sup> SUNY University at Buffalo, USA

<sup>2</sup> National Center for Ontological Research, USA

<sup>3</sup> University of Neuchâtel, Switzerland

<sup>4</sup> University of Toulouse, France

<sup>5</sup> KadSci LLC, USA

**Abstract.** The explosion of interest in large language models (LLMs) has been accompanied by concerns over the extent to which generated outputs can be trusted, owing to the prevalence of bias, hallucinations, and so forth. Accordingly, there is a growing interest in knowledge representation (KR) tools – ontologies and knowledge graphs, in particular - to make LLMs more trustworthy. This rests on the long history of KR in constructing human-comprehensible justification for model outputs as well as traceability concerning the impact of evidence on other evidence. Understanding the nature of arguments and argumentation is critical to justification and traceability, especially when LLM output conflicts with what is expected by users. The central contribution of this article is to extend the Arguments Ontology (ARGO) - an ontology of terms and relations specific to the domain of argumentation and evidence broadly construed - into the space of LLM outputs in the interest of promoting justification and traceability. We outline a strategy for creating ARGO-based ‘blueprints’ to help LLM users explore justifications for outputs. We conclude by describing critical applications at the intersection of LLM and knowledge representation research.

**Keywords:** Ontology · Arguments · Semantic Reasoning · Large Language Models

## 1 Introduction

The explosion of interest in large language models (LLMs) has been accompanied by concerns over the extent to which generated outputs can be trusted, owing largely to the prevalence of hallucinations [3] and bias [51]. Each may be at least partially addressed by intervention at various points of LLM development, such as by pre-training on vetted data [12] and human-in-the-loop reinforcement [13]. A natural strategy for addressing hallucinations post-inferencing involves **fact-checking**, the process of evaluating whether claims asserted to be true, are indeed true [9]. Traditionally, this process involves identifying asserted claims,

relevant evidence or counter-evidence, and delivering a verdict based on these results [10].

Work exploring fact-checking [6] with respect to LLMs has focused on the outputs of domain-specific prompts, for example climate change [11], disease [7], or Twitter [8], though recent work has sought to expand such scope [4]. In naive models, given claim and evidence pairs, a tag is generated indicating whether the evidence supports or undermines the claim. More sophisticated models attempt to provide missing evidence for or against claims [9] or correct claims based on existing evidence [7]. Regardless of the strategies adopted, there is a clear reliance on the relationship between claims and evidence, a relationship often characterized in terms of *arguments*.

In this respect, fact-checking research for LLMs dovetails with traditional applications of *ontologies* - controlled vocabularies of terms representing entities and relationships among them - which have been identified as useful for enhancing LLM accuracy [1][2]. Ontologies often serve as *lingua francas* to promote semantic interoperability across otherwise disparate datasets, doing so by providing explicit machine-understandable schema for domains of interest. Similarly, ontologies often serve as *lingua francas* to promote what we might call *human interoperability* across the otherwise disparate vocabularies of people, facilitating human understanding in a machine-understandable manner [14][43]. Ontologies intersect with fact-checking research on LLMs through providing *justification* - human comprehensible arguments in favor of or against model outputs - and *traceability* - human comprehensible arguments concerning impacts of evidence on other evidence.

Little research has been conducted at the intersection of LLMs, fact-checking, and ontology engineering. In what follows, we aim to remedy this gap by connecting an ontology of arguments and an ontology of LLM biases, to lay foundations for research on how ontology representations and tooling can promote and support justification and traceability for fact-checking strategies.

## 2 Adequacy Constraints for Argument Ontologies

### 2.1 Hallmarks of Arguments

Claims are often asserted through sentences, syntactic patterns of characters used to convey content, for example "Snow is white" expresses that snow is white. Sentences are not identical to the contents expressed by them: "Schnee ist weiß" expresses the same content as "Snow is white". Additionally, sentences and content may be used in different ways. One might assert "Snow is white" or ask whether "Snow is white", the former being an assertion and the latter a question.

Arguments can be described as collections of content used to support the truth of other content. This is in keeping with the arguments being, fundamentally, vehicles for persuasion:

1. If Susan leaves work early, she will go home and to the gym.

2. SUPPOSE Susan leaves work early.
3. Hence, Susan will go home and to the gym.
4. Hence, Susan will go home.
5. Hence, if Susan leaves work early then Susan will go home.

Lines 1 and 2 are intended to support line 3, which supports 4, while lines 2 and 4 support 5. In this example, line 1 is used to assert the corresponding content is true. Line 2 reflects a supposition that the corresponding content is true in the interest of drawing out consequences. Lines 3 follows from 1 and 2; 4 from 5, and 5 from 2 and 4. As illustrated, arguments may be complex, with claims depending on other claims in support of an overall conclusion. This example also highlights how sentence contents occupy distinct roles in the context of arguments. Lines 1 and 2 being premises of the argument, lines 3 and 4 being sub-conclusions, and line 5 being the main conclusion. This distinction is crucial for identifying textbook question-begging arguments:

1. Ghosts exist.
2. Hence, Ghosts exist.

Where the same sentence content plays distinct argument roles.

Each of the preceding examples illustrates arguments using content that may be true or false, but arguments in general need not be so restricted [44]:

1. Hold the door if you want to keep your job!
2. You want to keep your job.
3. Hence, hold the door!

Lines 1 and 3 are expressed by commands, the content of which is not typically taken to be true or false [37][41]. Any general ontology of arguments should then allow for representing arguments containing contents that are neither true nor false.

## 2.2 Ontology Best Practices

A well-designed ontology should satisfy certain adequacy constraints that reflect the main goal of ontology development, which is interoperability of heterogeneous data and information systems [27][15]. Ontologies should be accurate, adaptable, consistent, and provide clear annotations [20][21]. With respect to **accuracy**, an ontology should accurately represent entities and relationships within its stated scope. With respect to **adaptability**, any ontology should be designed to be reused by and reuse from other ontologies to the extent possible. With respect to **clarity**, terms in ontologies should be given unambiguous, clear, labels, alternative labels, definitions, and definition sources, to promote understanding across a variety of potential stakeholders. With respect to **consistency**, ontologies should be logically consistent, as demonstrated by an automated reasoner such as those readable by Web Ontology Language (OWL) based reasoners [25][26].

These best practices are borne out of years of ontology development by a variety of users, and indeed are codified as principles in large ontology development communities, such as the Open Biological and Biomedical (OBO) Foundry [32] and the Industrial Ontologies Foundry (IOF) [33]. To encourage best practices, further principles for ontology design have emerged from these communities, such as having ontologies within their respective communities extend from a top-level architecture: Basic Formal Ontology (BFO) [30][28]. BFO - an ISO/IEC 21838:2 top-level ontology standard [34] - contains high-level general terms such as *object* and *process*. Extending from a common top-level architecture promotes semantic interoperability by ensuring that no matter how far ontologies extend towards specific domain content, e.g. *electrons*, *tables*, *whales*, they nevertheless share a common formal language. Extending from BFO indeed promotes explainability and traceability, as all extended ontology terms should be accompanied by definitions [31] following the scheme: **A is a B that Cs**, where "A" is a subclass of "B" and differentiated from other subclasses of B by "C" [15]. For example, a **disorder** is a clinically abnormal part of an extended organism.

Altogether, these observations suggest the following constraints should be respected for an ontology of arguments, the ontology should:

- Contain clear labels, definitions following the "A is a B that Cs" scheme, and other disambiguating annotations. [**Clarity**]
- Extend from a top-level architecture, such as BFO. [**Adaptability**]
- Distinguish sentences from their contents and sentence contents from particular roles within an argument. [**Accuracy**]
- Allow for representation of multiple sentence and content types in arguments. [**Accuracy**]
- Depict arguments as consisting of premises or suppositions supporting one or more conclusions. [**Accuracy**]
- Be represented in OWL2 and verified for logical consistency using associated reasoners. [**Consistency**]

In the next section, we examine existing ontologies of arguments or evidence, noting where they fall short of one or more adequacy constraints.

### 2.3 Existing Ontologies of Arguments or Evidence

Of the ontologies dealing with argumentation and evidence we reviewed, none adequately reflect hallmarks of arguments and none satisfy all of the preceding criteria for ontology development. A common issue, exhibited for instance by the Legal Knowledge Interchange Format (LKIF) Core Ontology [35], is the conflation of arguments and evidence. LKIF views arguments 'as reasons expressed through a medium'. This overlooks the complexity of arguments involving reasoning chains and multiple inferences.

The Argument Interchange Format (AIF) does not make such a conflation, [46] but instead conflates the content of arguments with what the arguments are about [36]. The Argument Model Ontology (AMO), based on the influential

Toulmin model of argumentation [48], distinguishes claims, evidence, warrant, rebuttals, and so on, but does not distinguish the *contents* of sentences comprising arguments from the *roles* played by such contents within an argument [45]. The Semantic Science Integrated Ontology (SIO) provides a canonical treatment of arguments, validity, soundness, and so forth, but defines the contents of sentences as sentences “expressing something true or false”, conflating contents with sentences themselves [38]. Perhaps most worrisome, however, is that SIO appears to have taken labels and definitions from BFO, such as *site* and *process* but coined new unique identifiers of these terms rather than use those of BFO, a practice that is in direct conflict with the goals of semantic interoperability.

Most of the preceding do not adopt a top-level architecture, though there are argument ontologies that do, such as the OBO Foundry Evidence Ontology (EO) [49]. While EO adopts BFO as a top-level, it nevertheless does not align with it, as EO houses its terms outside the BFO hierarchy, though terms such as “evidence” are defined as clearly falling under BFO’s root class. Additionally, the scope of EO is restricted to the biological domain. On the other hand, while the Explanation Ontology (EXO) adopts a top-level architecture, namely SIO, it inherits the serious issues exhibited by that import [50].

### 3 ARGO

The Arguments Ontology (ARGO) is a small ontology designed to satisfy the constraints identified in the previous section. ARGO accordingly extends from BFO, leveraging resources from other BFO-conformant ontologies such as the Information Artifact Ontology (IAO) [36], an extension of BFO designed to represent information and information bearers.

ARGO distinguishes sentences and sentence content. The ARGO class **expression** consists of patterns of character shapes in a language, such as the string of characters comprising this clause. **sentence** is a subclass of **expression**, instances of which satisfy some conventional rules of grammar. Both are distinct from the class **statement**, which represents the contents of sentences, e.g. the **sentence** “Susan is happy” expresses the **statement** that Susan is happy, which is plausibly the meaning of the sentence and what one would believe by it. **Statements** are a subclass of the IAO **information content entity**, roughly, entities that are about things in the world, such as the content of a book or the information encoded in a docx file on a hard drive. **Information content entities** are identical across bearers, e.g. the content of a given PDF may exist across distinct laptops. Consequently, **statements** may have identical instances across bearers. Two observant friends of Susan, for instance, may both believe Susan is happy, each expressing this by uttering “Susan is happy”.

When a statement is a constituent of an argument, it is typically as either a premise or conclusion; a given statement may serve as a conclusion in one argument and a premise in another. Moreover, premises and conclusions are always used in arguments, whereas statements need not be. Premises are linked to conclusions insofar as they are offered as support or evidence for conclusions in

arguments. Plausibly, this link between premises and a conclusion is an action—a passing from some collection of statements to another statement because one believes the latter is justified, supported, or entailed by the former statements. We reflect this link between premises and conclusions by defining a class **act of inferring**. A premise is a statement that stands in a particular relation to an argument as a result of being the input of an **act of inferring**; a conclusion is a **statement** that stands in a particular relation to an argument as a result of being the output of an **act of inferring**. The relations ‘has input’ and ‘has output’ are reused from the Common Core Ontologies, and defined roughly in such a way that a **premise** is input to an **act of inferring** and a **conclusion** is the output of such an act; we then say that a given argument **has premise** some **premise**, and similarly **has conclusion** some **conclusion**.

Premises are affirmed in arguments; suppositions are accepted in arguments. To suppose that a statement is true, is often done for the sake of some further inferential goal, e.g. hypothetical deliberation, indirect reasoning, and so on. We capture these distinctions in terms of differing acts, namely, an **act of affirming** in which an agent believes a **statement** is true or false based on evidence, and an **act of accepting** in which an agent entertains a **statement** as true or false independent of belief or evidence. Both may be inputs to **acts of inferring**, which links premises and suppositions to conclusions. Regarding conclusions, they may be affirmed or accepted:

1. If Susan leaves work early, she will go home and to the gym.
2. SUPPOSE Susan leaves work early.
3. Hence, Susan will go home and to the gym.
4. Hence, Susan will go home.
5. Hence, if Susan leaves work early then Susan will go home.

Here, 1 is affirmed while 2 is accepted. Lines 3, 4, and 5 are conclusions; line 3 is inferred based on a combination of an affirmed line and an accepted line, which suggests it is itself accepted. Accordingly, 4 is accepted as it is based on the accepted 3, while 5 is best described in terms of an agent affirming the connection between what is supposed and a consequence of it, which is not itself a case of affirming supposed content.

Every **subconclusion** stands in a ‘subconclusion in’ relation to a **complex argument**, requiring the **statement** is, first, an affirmed or accepted input in an **act of inferring** in an **argument** and, second, affirmed or accepted output in an **act of inferring** in an **argument** distinct from the first, where third, both **arguments** are parts of the **complex argument** to which the **statement** stands in the ‘subconclusion in’ relation.

**Arguments** are ordered collections of **statements** involving **premises**, **suppositions**, and a single **conclusion**. In turn, **premises** are logically equivalent to **statements** that are affirmed inputs of **acts of inferring**, where **conclusions** are equivalent to **statements** that are affirmed or accepted outputs of **acts of inferring**. **Suppositions** are **statements** that are accepted inputs of **acts of inferring**, whereas **subconclusions** are **statements** that are af-

firmed or accepted inputs and outputs of distinct **acts of inferring** in distinct **arguments** that are proper parts of a **complex argument**.

There are many different purposes one might have in constructing an argument. The paradigm case involves arguing, where an individual provides an argument with the intent of convincing others that the conclusion of the argument is true. We characterize this process as an **act of arguing**. One can argue successfully or unsuccessfully, but one cannot argue without intending to convince one’s audience of some conclusion; even when arguing against a given conclusion, you are still arguing in favor of *some* conclusion. Of course, one may have no intention to convince others of some conclusion; one may be creating an argument for the purpose of interpretation or to anticipate what an opponent might say during a debate. In such cases, one is not arguing; rather, one is merely creating arguments, which we characterize as a process of **act of argument creation**. An **act of arguing** may have an **act of argument creation** as process part, if in the process of arguing one creates an argument. Creating an argument involves a series of steps, at least one of which is an **act of inferring**. An **act of argument creation** is related to the **argument** in that it creates by the ‘is created by’ relation. In Figure 1 we have a complex argument that is cre-

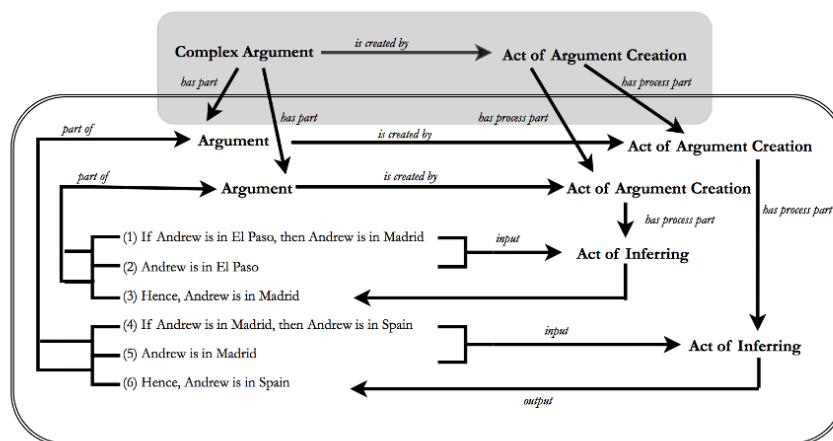


Fig. 1. Complex Argument

ated by an **act of argument creation** which has two proper **act of argument creation** parts, each of which has an **act of inferring** part, respectively. The **statements** exhibited by lines 1 and 2 are inputs to an **act of inferring**, just as lines 4 and 5 are inputs to another such act. Each has as its output a respective conclusion **statement**, where each **statement** is a part of some **argument**, which in turn is proper part of the overall **complex argument**. Importantly,

note that 3 and 5 exhibit the same **statement** used differently across two simple arguments that make up a complex argument. These observations support defining **complex argument** defined as an **argument** with at least one proper **argument** part and which has only **argument** parts.

We have thus far discussed relationships among **statements** within **arguments** but now turn to **statements** between **arguments**, i.e. textitcounterarguments. A counterargument must stand in a certain relation to another argument. For example, argument A may be a counterargument to B if B has a conclusion that stands in a relation of contradiction to the conclusion of argument A. More often, cases of counterarguments undermine, but do not contradict, some part of another argument. For example, argument A may have some counterargument B if the conclusion of B raises concerns for the justification of one or more premises of A. To capture such generality, ARGO adopts an 'opposes' relation, which holds between premises, suppositions, conclusions, or subconclusions across arguments. Sub-relations include: negates, contradicts, and undermines. Similarly, a 'supports' relation is introduced to illustrate potentially favorable evidence. These are, of course, coarse-grained. A fuller treatment will introduce degrees of support and opposition.

## 4 Explainable, Traceable, LLMs

Ontology artifacts are often written in the Web Ontology Language (OWL), which supports automated model and proof creation. Standard ontology reasoners, such as Hermit [52] and [53] are used to provide explanations for conclusions drawn from data sets, in a manner align closely with common sense reasoning; additionally, ontologies are leveraged to generate assumptions needed to support a given claim represented in an ontology. In each case, ontologies provide human-understandable explanations for output as well as deductive traces useful for exploring putative commitments that led to a certain output.

We envision leveraging ARGO to address hallucinations the realm of LLM fact-checking. Our first step involves the identification - using ARGO as a guide - of claims extracted from LLM output, tagging identified collections of claims as exhibiting arguments, having premises, suppositions, and conclusions. Automating the tagging using ARGO will involve standard NLP pre-processing, tokenization and named entity recognition to identify potential claims and related entities. When a given **statement** found in gathered evidence contradicts another **statement** found in an argument, we annotate this relationship. Leveraging the distinctions drawn between, say, premises and suppositions, we can also assess the extent to which evidence opposes some **statement** is damaging to an argument. Finding counterevidence for a premise may undermine an entire argument; counterevidence for a supposition may undermine only the sub-argument leveraging the supposition.

OWL2 reasoners may, moreover, be used to facilitate the identification of conflicts across arguments. Suppose **statement** A conflicts with **statement** B. A will likely be in conflict with B - to some extent - regardless of what



arguments in which one finds A (affected by whether A is a supposition, premise, or conclusion). OWL2 reasoners can be used to track such relationships across argument contexts, thereby providing an avenue for traceability.

Core to the blueprint application of ARGO is the generation of explanations for verdicts made about identified claims. ARGO-based blueprints foster explainability through the use of well-curated definitions for argument components and inferences which use them. These explanations will offer valuable insights into the argumentative structure of LLM outputs. Integration of ARGO with LLMs for fact-checking represents a significant advancement in enhancing reliability and trustworthiness. Our blueprint strategy offers a systematic approach to fact-checking that is grounded in clear definitions, logical consistency, and adaptability to various domains.

## References

1. Pan, S., Luo, L., Wang, Y., Chen, C., Wang, J., Wu, X.: Unifying large language models and knowledge graphs: A roadmap. arXiv preprint arXiv:306.08302 (2023)
2. Sequeda, J., Allemang, D., Jacob, B.: A Benchmark to Understand the Role of Knowledge Graphs on Large Language Model’s Accuracy for Question Answering on Enterprise SQL Databases. arXiv preprint arXiv:2311.07509 (2023)
3. Tang, L., Sun, Z., Idnay, B. et al.: Evaluating large language models on medical evidence summarization. *npj Digital Medicine*, 6, 158 (2023).
4. Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., Liu, T.: A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. arXiv preprint arXiv:2311.05232 (2023)
5. Chen, J., Kim, G., Sriram, A., Durrett, G., Choi, E.: Complex Claim Verification with Evidence Retrieved in the Wild. arXiv preprint arXiv:2305.11859 (2023)
6. Gao, M., Ruan, J., Sun, R., Yin, X., Yang, S., Wan, X.: Human-like Summarization Evaluation with ChatGPT. ArXiv preprint, abs/2304.02554 (2023)
7. James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. FEVER: a Large-scale Dataset for Fact Extraction and VERification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. <https://aclanthology.org/N18-1074> DOI: 10.18653/v1/N18-1074.
8. Nicolas Prolochs. Community-Based Fact-Checking on Twitter’s Birdwatch Platform. *ArXiv*, abs/2104.07175, 2022.
9. Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. A Survey on Automated Fact-Checking. *Transactions of the Association for Computational Linguistics*, 10:178–206, 2022.
10. Dustin Wright and Isabelle Augenstein. Claim Check-Worthiness Detection as Positive Unlabelled Learning. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 476–488, Online, 2020. Association for Computational Linguistics.
11. Thomas Diggelmann, Jordan L. Boyd-Graber, Jannis Bulian, Massimiliano Ciaramita, and Markus Leippold. CLIMATE-FEVER: A dataset for verification of real-world climate claims. *CoRR*, abs/2012.00614, 2020.

12. Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. Siren’s song in the AI ocean: A survey on hallucination in large language models. *CoRR*, abs/2309.01219, 2023b.
13. Yuxia Wang, Revanth Gangi Reddy, Zain Muhammad Mujahid, Arnav Arora, Aleksandr Rubashevskii, Jiahui Geng, Osama Mohammed Afzal, Liangming Pan, Nadav Borenstein, Aditya Pillai, Isabelle Augenstein, Iryna Gurevych, and Preslav Nakov. Factcheck-GPT: End-to-End Fine-Grained Document-Level Fact-Checking and Correction of LLM Output. *arXiv preprint arXiv:2311.09000*, 2023.
14. Roberto Confalonieri and Giancarlo Guizzardi. On the Multiple Roles of Ontologies in Explainable AI. *ArXiv*, abs/2311.04778, 2023.
15. Arp, Robert and Spear, Andrew D. and Smith, Barry. *Building Ontologies with Basic Formal Ontology*. MIT Press, 2015.
16. Keet, Maria C. *Introduction to Ontology Engineering*. College Publications, 2020.
17. Uschold, Mike and Grüninger, Michael. “Ontologies: Principles, Methods and Applications.” *The Knowledge Engineering Review*, vol. 11, no. 2, pp. 93–136, 1996.
18. Raad, Joe and Cruz, Christophe. “A survey on ontology evaluation methods”, *Proceedings of the International Conference on Knowledge Engineering and Ontology Development*, part of the 7th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management, 2018.
19. Kulvatunyou, B. and Drobnjakovic, M. and Ameri, F. and Will, C. and Smith, B. “The Industrial Ontologies Foundry (IOF) Core Ontology”, *Formal Ontologies Meet Industry (FOMI) 2022*, Tarbes, FR, 2022.
20. Vrandečić, Denny. “Ontology Evaluation.” in: *Handbook on Ontologies*, 200, pp. 293–313.
21. Hogan, Aidan et al. “An empirical survey of linked data conformance.”, *Journal of Web Semantics*, 14, 2012, pp. 14-44.
22. Zaveri, Amrapali et al. “Quality assessment for linked data: A survey”, *Semantic Web*, 7.1, 2016, pp. 63-93.
23. Neuhaus, Fabian and Hastings, Janna. “Ontology is Consensus Creation, Not (Merely) Representation”, *Applied Ontology*, 17, 2022.
24. Fox, Mark and Grüninger, Michael. “The Role of Competency Questions in Enterprise Engineering”, *Proceedings of the IFIP WG5.7 Workshop on Benchmarking - Theory and Practice*, 1994.
25. Hermit, an OWL 2 reasoner, 2023. URL: <http://www.hermit-reasoner.com/>
26. Pellet, an OWL 2 reasoner, 2023. URL: <https://github.com/stardog-union/pellet>
27. Abukwaik, H. et al. “Interoperability-related architectural problems and solutions in information systems: a scoping study.” In: *Software architecture. ECSA 2014. Lecture notes in computer science*, vol. 8627. Springer, Cham, 2014.
28. Tchouanguem, Justine Flore et al. “BFO-based ontology enhancement to promote interoperability in BIM.” *Applied Ontology*, 16, 4, 2021, pp. 453–479.
29. Bandrowski, A. et al. “The Ontology for Biomedical Investigations.” *PLoS ONE*, 11 (4), 2016, e0154556. doi: 10.1371/journal.pone.0154556.
30. Otte, N.J. and Beverley, J. and Ruttenberg, A. “BFO: Basic Formal Ontology.” *Applied Ontology*, 2022, pp. 17-43.
31. Seppälä, Selja and Ruttenberg, Alan and Schreiber, Yonatan and Smith, Barry. “Definitions in ontologies.” *Cahiers de Lexicologie*, 109 (2), 2016, pp. 175-207.
32. Smith, B. et al. “The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration.” *Nat Biotechnol*, 25, 2007, pp. 1251-1255. doi:10.1038/nbt1346.

33. Kulvatunyou, Boonserm; Wallace, Evan; Kiritsis, Dimitris; Smith, Barry; Will, Chris (2018). “The Industrial Ontologies Foundry Proof-of-Concept Project.” NIST.
34. ISO/IEC. “Information technology — Top-level ontologies (TLO) — Part 2: Basic Formal Ontology (BFO).” 2021.
35. Breuker, Joost, Hoekstra, Rinke, and Boer, Alexander (eds). *OWL Ontology of Basic Legal Concepts (LKIF-Core)*. Estrella, University of Amsterdam, 2006. Available at: [www.estrellaproject.org/doc/D1.4-OWL-Ontology-of-Basic-Legal-Concepts.pdf](http://www.estrellaproject.org/doc/D1.4-OWL-Ontology-of-Basic-Legal-Concepts.pdf).
36. Ceusters, Werner and Smith, Barry. “Aboutness: Towards Foundations for the Information Artifact Ontology.” In: *Proceedings of the Sixth International Conference on Biomedical Ontology (ICBO)* CEUR Vol. 1515, 2015, pp. 1-5.
37. Clark-Younger, H. “Imperatives and the More Generalized Tarski Thesis.” *Thought: A Journal of Philosophy*, 3, 4, 2015, pp. 314-320.
38. Dumontier, Michel et al. “The Semantiscience Integrated Ontology (SIO) for Biomedical Knowledge Discovery.” *Journal of Biomedical Semantics*, 5:14, 2014.
39. van Eemeren, Frans H. and Grootendorst, Rob. *Argumentation, Communication, and Fallacies: A Pragma-Dialectical Perspective*. Routledge, 1992.
40. Hovda, Paul. “What is Classical Mereology?” *Journal of Philosophical Logic*, 38, 2009, pp. 55-82.
41. Jørgensen, Jørgen. “Imperatives and logic.” *Erkenntnis*, 7, 1938, pp. 288-98.
42. McCullough, Laurence B., Coverdale, John H., and Chervenak, Frank A. “Argument-Based Medical Ethics: A Formal Tool for Critically Appraising the Normative Medical Ethics Literature.” *American Journal of Obstetrics and Gynecology*, 191, no. 4, 2004, pp. 1097-102.
43. Menzel, Christopher. “Knowledge Representation, the World Wide Web, and the Evolution of Logic.” *Synthese*, 182, 2011, pp. 269-295.
44. Parsons, Josh. “Command and Consequence.” *Philosophical Studies*, 164, 1, 2013, pp. 69-92.
45. Peroni, Silvio and Vitali, Fabio. “The Arguments Model Ontology (AMO).” Published April 5, 2011. Accessed August 23, 2017 at: <http://www.essepuntato.it/lode/http://www.essepuntato.it/2011/02/argumentmodel>.
46. Rahwan, Iyad and Banihashemi, Bitan. “Arguments in OWL: A Progress Report.” In: *Computational Models of Argument: Proceedings of COMMA 2008*, T
47. Tuomela, Raimo. “Belief versus Acceptance.” *Philosophical Explorations*, 3, 2, 2000, pp. 122-137.
48. Toulmin, Stephen. *The Uses of Argument*, 2nd edition (1st edition 1958). Cambridge University Press, 2003.
49. Chibucos, M.C., Siegele, D.A., Hu, J.C., Giglio, M. “The Evidence and Conclusion Ontology (ECO): Supporting GO Annotations.” *Methods in Molecular Biology*, vol. 1446, 2017, pp. 245-259. PMID: 27812948; PMCID: PMC6377151.
50. Chari, Shruthi et al. “Explanation Ontology: A General-Purpose, Semantic Representation for Supporting User-Centered Explanations.”
51. Koo, Ryan et al. “Benchmarking Cognitive Biases in Large Language Models as Evaluators.” ArXiv abs/2309.17012 (2023): n. pag.
52. Shearer R, Motik B, Horrocks I. (2008) HermiT: A Highly-Efficient OWL Reasoner. OWLED.
53. Evren Sirin BP, Bernardo Grau C, Kalyanpur Aditya, Yarden Katz. (2007) Pellet: A practical OWL-DL reasoner Web Semantics: Science, Services and Agents on the World Wide Web. 5(2):51-3.