ARGO: Arguments Ontology

John Beverley [a1], J. Neil Otte[b], Francesco Franda[a], Jean-Baptiste Guillon[c], Brian Donohue[a], Alan Ruttenberg[a], and Yonatan Schreiber [a2]

Abstract: Although the last decade has seen a proliferation of ontological approaches to arguments, many of them employ ad hoc solutions to representing arguments, lack interoperability with other ontologies, or cover arguments only as part of a broader approach to evidence. To provide a better ontological representation of arguments, we present the Arguments Ontology (ArgO), a small ontology for arguments that is designed to be imported and easily extended by researchers who work in different upper-level ontology frameworks, different logics, and different approaches to argument evaluation. Unlike most ontological approaches to arguments, ArgO utilizes Basic Formal Ontology (BFO) as an upper-level ontology, and may be used alongside other commonly used ontologies in the BFO framework, including both the Information Artifact Ontology (IAO), and the Information Entity Ontology (INFO). Critically, our proposal is principled, based on rigorous definitions and formal axioms out of which characterizations of arguments naturally fall. It is our hope that ArgO may assist researchers in many projects, including: integrating heterogeneous sources of evidence, structuring the content of semantic wikis, and enhancing semantic reasoning.

Keywords: Ontology, Arguments, Argumentation Theory, Basic Formal Ontology, Semantic Reasoning

Introduction

Information systems in many domains—including, law, medicine, historical records, and humanities journals—structure and store data concerning arguments. Arguments are critical to understanding the motivations of and disputes among agents, as well as evaluating the evidence for and against various claims. For instance, take the following argument:

> Whereas Iraq continues to aid and harbor other international terrorist organizations, the Congress of the United States is warranted in granting support for the President to use the Armed Forces of the United States to defend the United States against the continuing threat posed by Iraq.[3]

This passage is a modification of the joint resolution authorizing the use of force by the United States during the second Iraq war. The original passage provides many other reasons that attempt to justify and lend support for Congress's decision. In its modified form here, we see only one reason, presented in the first sentence, whereas the final sentence contains the conclusion of the argument. We can look at this passage as a

---

[1]Corresponding author: John Beverley, johnbeve@buffalo.edu.
[2]*a* SUNY University at Buffalo; *b* Johns Hopkins University Applied Physics Lab; *c* Collège de France
[3]Adapted from a statement of the Office of the Press Secretary of the United States (2002).

mere set of sentences, but this is uninformative—what is necessary to grasp is the structure of the argument itself, that there is a direction of the intended inference whereby we move from the premises to the conclusion. Once we understand this structure, we can see that this passage is not merely a report or a description, but a concerted effort to have an audience accept the conclusion.

The rationale for the second Iraq war was one of the most widely debated issues of the last decade, and the sources of evidence concerning this debate are scattered in many places. In this regard, it is not so different from many other projects in science, where aggregating evidence from disparate sources and identifying their evidential relations to a conclusion is a massive investigative undertaking. Representing such complicated sources of evidence is one task for formal representations of arguments. A second task is to structure a procedure of evaluation. This allows a formal treatment of arguments to aid investigation, guiding an investigative process through a series of steps. Such tools are widely used in education assessment[4], as well as in research and medical ethics[5], where argument schemes are used to guide a participant through a process of evaluation and documentation.

To begin putting these pieces together with the aid of machines, we need to understand texts not only as consisting of word tokens and characters, but also as being about things in the world, and as expressing content that bears relationships of logical entailment and support. In information science, this work is performed by ontologies, critical tools for integrating data across different languages, servers, and conceptual frameworks. Following Arp et al. (2015), we take an ontology to be a representation of reality comprised of a taxonomy as a proper part, whose representations are intended to designate some combination of universals, defined classes, and certain relations between them.[6] Ontologies are routinely implemented in a decidable formal language, such as the Web Ontology Language (OWL), allowing them to be processed by machines. In the last decade, many ontologies for representing arguments have been developed with the goal of structuring, integrating, and reasoning over represented chains of inference.

A well-designed ontology requires adherence to certain methodological requirements. For instance, entities in an ontology should have clear annotations, including definitions for each class and relation. Ontology reasoning resources require precise formal representations of data, often as logically defined classes organized in a taxonomic structure of class and subclass relations.[7] Using Aristotelian, or genus-species, definitions aids in reasoning tasks, for where a class (A) is formally defined as a subclass of (B) with differentia (C), then any instance of (A) may be inferred to be an instance of (B) that is distinguished from other instances of B by (C). Moreover, precise

---

[4]Rapanta et al. (2016)

[5]McCullough et al. (2004)

[6]Arp et al. (2015), p. 181

[7]Ontology reasoners (Fact++, Hermit) are based in a decidable fragment of first-order logic. In particular, of this family of description logics, these reasoners employ SROIQ, an expressive fragment of first-order logic restricted to classes and binary relations, the latter having further sub-relation restrictions. SROIQ is the foundation of the widely-used web ontology language OWL2, found in the ontology editor Protégé.

formal representations should be compatible with, where possible, at least one widely-used upper-level ontology. This is because ontologies were designed, in part, to solve the problem of data silos, stockpiles of data coded in parochial languages inaccessible to other semantic technologies.[8] Upper-level ontologies provide a lingua franca on which consistently extended domain-level ontologies may be based, thereby improving accessibility among data sets in various domains.

Of the ontologies of arguments we reviewed, most do not use upper-level ontologies. Moreover, most existing ontologies seem narrowly focused, reflecting the concerns of a researcher's data set, or the particular end-goals of a single project. This conclusion was evidenced, for example, both in a frequent lack of clear definitions, as well as a tendency to conflate the physical bearers of information (e.g. patterns of ink on paper) with information (e.g. a proposition) itself. Thus, these ontologies often fail to meet a series of criteria necessary for creating an ontology of arguments that may be re-purposed by multiple researchers who may use the same ontology for different projects.

Few ontological approaches to arguments also employ a realist methodology, according to which thought, experience, and knowledge are characteristically (if also partially and fallibly) about reality, as opposed to about concepts or ideas. This realist approach to ontology design is a hallmark of the BFO framework[9], and we have found it is critical to sustainable efforts in the field. For the present discussion, this should be understood to mean that we take arguments to be *bona fide* parts of the world, independent of ideas or perspectives on them—though representing these is certainly an important goal as well.

In addition to methodological requirements, the particular domain of arguments suggests adequacy constraints. For example, arguments are often expressed through combinations of sentences, but we should not confuse arguments with collections of sentences. Sentences, at first approximation, are repeated patterns of characters in a language used to convey meanings or contents. For example, the English sentence "Snow is white" may be used by a speaker to express that *Snow is white*. However, arguments are neither sentences nor collections of sentences *per se*. Rather, arguments are collections of the contents expressed by such sentences. For just as the German sentence "Schnee ist weiß" may express the same content as the English sentence "Snow is white", the same argument may be expressed using different sentences in different languages.[10]

Relatedly, sentences and sentence contents should be distinguished from illocutionary
force, for instance, where a sentence is an assertion, command, or question. The same sentence may be used to express content with different force, as when the sentence "You will water my plants" may be used as an assertion, a command, or question. In

[8]See Smith (2012), Menzel, (2011), and Gruninger et al. (2014).

[9]For discussion and motivation, see Arp et al. (2015).

[10]This is also true within a language. This point may also be made by examples of sentences that include indexical elements. For instance, the sentence "You are the most beautiful girl in the room" will express different contents across contexts of utterance.

each case, the sentence and much of what it is about is the same, but what is conveyed may be distinct in terms of force.

In many cases, arguments are taken to be composed of the contents of declarative sentences, which are often called 'propositions'. Nevertheless, an ontology of arguments should remain open to arguments involving content conveyed with other speech acts. For instance, commands are typically expressed in imperative sentences. For example[11]:

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

The first and third sentences are imperatives that express content with the force of a command. Arguments of this kind are the subject of the field of imperative logic[12], and since an ontology of arguments should characterize what is true of all arguments, the components of arguments should not be identified with the contents of declaratives alone.

In remaining open to various logics, an ontology of arguments should be ecumenical with respect to logical form and rules of inference as well. For instance, adequate characterizations of imperative and interrogative content suggest logics more permissive than classical will be needed.[13] Also, arguments are frequently expressed in sentences employing modal terms, including obligation, permissibility, necessity, and possibility—formal representations of which typically go beyond classical logic. Consider, too, that arguments in some domains may be better suited by treatments in non-classical logics, for instance, arguments in mathematics may be better represented within intuitionist frameworks. An adequate ontology of arguments will permit users the flexibility to choose which logic best characterizes a domain of interest.

An ontology of arguments should also distinguish the function of sentence contents within an argument from the contents themselves; that is, the ontology should distinguish the meanings of sentences from their functions as premises, conclusions, and the like, within an argument. In a typical case, an argument consists of some number of premises and a single conclusion that follows from, is supported by, or is entailed by the premises. In some cases, an argument may also involve a supposition used to support the conclusion in some manner, as in the following case:

(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.

[11]Throughout our examples, we present arguments in standard form, where the parts of an argument are numbered to reflect the order of a chain of inference, and not the order of presentation in a text.
[12]For background, see Jørgensen, (1938), Vranas (2008), and Mastop (2011).
[13]See Parsons (2013) and Clark-Younger (2014).

Here, we have an argument consisting of a supposition on line (2), with consequences extending to line (4). Given (1) and supposing (2), (5) intuitively follows.

In other cases, an argument may have a single premise and a conclusion, where the premise and conclusion share the same content. For example, the following question-begging argument plausibly represents the same content bearing two different functions within the argument:

(1) God exists.
(2) Hence, God exists.

An ontology of arguments should permit distinguishing the mere repetition of sentence content in an argument, from that content's function within an argument as, for instance, a premise or conclusion.

An argument ontology meeting such adequacy constraints provides a firm foundation for computational reasoning and gains the advantages of upper-level compatibility, while remaining flexible enough to approach a wide range of related phenomena and purposes. In summary, we list the methodological and adequacy constraints of the domain as follows:

Methodological Constraints

The preceding remarks motivate the following methodological constraints on any argument ontology:

1. The ontology should have logically defined classes whose definitions employ an Aristotelian or genus-species form.
2. The ontology should be compatible with a widely used upper-level ontology, such as Basic Formal Ontology (BFO).
3. The ontology and its imports should be logically consistent with respect to a widely used formal language, such as first-order logic.
4. The ontology should distinguish informational entities from their physical bearers (e.g. novels from books).

Adequacy Constraints

Similarly, the following adequacy constraints on any ontology of arguments are motivated:

1. The ontology should depict arguments in a canonical manner, as consisting of reasons for a conclusion.
2. The ontology should distinguish sentences, from the contents or meanings expressed by those sentences, as well as the sentence contents themselves from their particular functions within an argument (e.g. premise).

3. The ontology should be capable—presently or by the extension of its classes—of representing not only the intended structure of an argument, but also the potential weaknesses of an argument either with its present classes or by creating subclasses that extend from existing classes.
4. The ontology should allow—presently or by the extension of its classes—multiple approaches to logic, such that it is not committed to a single logic (e.g. classical logic).
5. The ontology should be capable—presently or by the extension of its classes—of facilitating different approaches to truth, possibility, probability, and the formal properties of arguments, such as entailment and soundness.

Ontological Approaches to Arguments

In our survey of ontologies of arguments, we have not encountered an ontology of arguments that can meet all of these criteria. In particular, there is no other ontology of arguments with which we are familiar that satisfies our second adequacy criterion, as many ontologies simply allow unicode strings to count as parts of arguments. As for the other criteria, many ontologies of arguments succeed in part, but no ontology satisfies them all, and for this reason, they are nearly always incapable of being re-used for many other projects. It is not our goal to provide a comprehensive review of the rich variety of ontological, and more broadly semantic, approaches to arguments, so in what follows we restrict our discussion to a few prominent cases for illustration.

First, some current ontological approaches to arguments can often be concerned with broad approaches to evidence rather than arguments *per se*, and this orientation leads to errors. For example, in the Legal Knowledge Interchange Format (LKIF) Core Ontology, an argument is understood as "a reason that is expressed through some medium."[14] This notion of argument seems more in line with what Duncan Kennedy has called, "argument bites," stereotyped bits of reasoning "that legal reasoners use when the legal issue is one that permits a reference to the policies or purposes or underlying objectives of the legal order, rather than a legal issue that can be satisfactorily resolved through deductive rule application or by reference to binding precedent."[15] In this way, the LKIF characterizes arguments as principles that are appealed to in a variety of legal contexts.

This characterization fails adequacy criterion 1, for even among legal experts, this use of "argument" is peculiar, reflecting the particular concern that legal ontologies have with norms, rules, precedents, and principles of the law and interpretation, rather than more complex arguments involving multiple inferences and chains of reasoning tied together by subconclusions.[16] Arguments are composed of a particular conclusion that is being asserted by the arguer, as well as the particular reasons, suppositions,

---

[14]Breuker et al. (2006), p. 50.
[15]Kennedy (1991), p. 75.
[16]Cf. Baronett, (2012): "legal arguments contain at least one premise and can be appreciated and understood when you are able to grasp the underlying logic."

conjectures, and premises that the arguer takes to speak in favor of, provide support for, or entail the conclusion.

This mistake may also be a consequence of the popular use of argumentation schemes in implementations of computational reasoning. On a first approximation, schemes are stereotypical inference patterns used in both deductive and inductive arguments. Many researchers in artificial intelligence and argument mapping have proposed hierarchies that reflect different approaches to the form of such schemes. One such approach is detailed in Walton (1996), which provides an account of a 'Walton scheme', an account that has been influential for researchers wishing to show the consequences of presumptive reasoning on automated inference. Each Walton scheme includes a conclusion and a set of premises containing meta-linguistic variables. Take, for instance, the following scheme, which Walton et al. (2008) refers to as 'Argument from Cause to Effect'[17]:

> Generally, if A occurs, the B will (might occur).
> In this case, A occurs (might occur).
> Therefore, in this case, B will (might occur).

Each scheme is then accompanied by critical questions, where the critical questions allow users to both identify and anticipate weaknesses in the argument based on the kind of scheme being employed.

Such schemes are useful extensions of argument ontologies[18], but ontologists should be careful to maintain a distinction—if only a practical one—between schemes and arguments.[19] In our survey, we often found that ontologists implemented schemes as classes, and then treated individual arguments as instances of the class Scheme.[20] This is incorrect, since it violates the true path rule according to which everything that is true of instances of the parent class is true of instances of the child of that parent. Every instance of a scheme contains meta-linguistic variables in its component sentences; however, this is only true of some arguments.

Furthermore, one might argue that a scheme has a different function than arguments: a scheme is not created to convince others of the truth of its conclusion; rather, it is intended to function as a clarifying description of the logical structure of a set of arguments. Of course, this point might be debated (e.g. Do some arguments—particularly *reductios*—not also function as debunking descriptions of an opposing argument?). But even if we decide some arguments with variables can function, like

---

[17]Walton et al. (2008) p. 168.

[18]E.g., van Eemeren and Grootendorst (1992) have shown that argumentation schemes are useful in determining whether or not arguments are fallacious.

[19]The AIF ontology accomplishes this distinction, but only by distinguishing within the graph I-nodes that represent the claims of a domain of discourse from S-nodes, which represent patterns of reasoning. However, we find our approach preferable, since what unites the instances of I-nodes and S-nodes is that they are both information content entities—a subclass of generically dependent continuant in BFO. Thus, we allow them to be part of a single taxonomical structure.

[20]Cf. the classification presented in Feng and Hirst (2011).

schemes, as descriptions of other arguments, it is nonetheless not a common trait among all arguments.

These considerations are sufficient for not treating arguments as subclasses of schemes. But are schemes arguments? This is a matter for further theoretical debate, and we find Walton et al. (2008) is unclear on this point.[21] As a practical matter, we prefer to conceive of schemes as descriptions of arguments rather than as arguments themselves, though we acknowledge this issue remains unresolved, and that the same set of sentences may be used to express an argument in one context and exclusively provide a description of some set of arguments in another. We will revisit this issue in section IV.b.

One goal of ontologies of arguments is to be able to reason over an extended web consisting of chains of inferences. Such inference chains may form long and often highly complex links that, in turn, might have different properties of support or force. The Argument Interchange Format (AIF) is one example of such an effort.[22] The AIF is designed to depict not arguments (which the AIF does not define) but instead 'argument entities', which it represents as nodes in a directed graph it calls an "argument network". While the AIF itself demonstrates the utility of programming languages to show the relatedness of arguments, the AIF itself is inconsistent with the use and goals of upper-level ontologies, since it does not distinguish the representation of information from the representation of the world that the information is about within the class structure of its ontology.

Other ontologies of arguments may also define a class of arguments, too broadly, construing arguments "as analogies, as counterevidence, as (rhetorical) questions, as hypotheticals, and even in the form of irony."[23] Of course, arguments are routinely communicated through analogies, rhetorical questions, hypothetical descriptions, and extended forms of irony. But, in many cases, what is pragmatically implicated is an assertion. For example, in certain circumstances a rhetorical question such as "Weren't you at the shopping mart Tuesday" pragmatically implicates the assertion "You were at the shopping mart on Tuesday." Similarly, analogies, when they are arguments, often take the form: S has property P; T is relevantly similar to S; hence, T has property P (or similar property P'). As for irony, take the case of Jonathan Swift's *Modest Proposal*[24]. Swift's work clearly expresses an argument (perhaps several). But what is it about, say, Swift's wickedly ironic *Modest Proposal* that makes the work express an argument? Similarly, what is it about an argument by analogy that separates it from mere analogies (e.g. that ravens are like writing desks)? A natural answer to these questions is structure, and a natural explication of structure in the domain of

---

[21] Although elsewhere in the text Walton et al. (2008) hold that schemes can be recognized as deductive arguments, they begin with the following characterization: "Argumentation schemes are forms of argument (structures of inference) that represent structures of common types of arguments used in everyday discourse, as well as in special contexts like those of legal argumentation and scientific argumentation." p. 1.

[22] Rahwan, I & Banihashemi, B. (2008) discusses the AIF and updates it in their ArgDF ontology, which employs OWL-DL to reason over argument networks.

[23] Cf. Breuker et al. (2006), p. 79.

[24] Swift (1729).

arguments appeals to enthymematic premises, suppositions, and conclusions. Simply permitting analogies as arguments overlooks this distinction.

The model of argument defended by Toulmin[25] has been used in projects such as the Argument Model Ontology (AMO)[26]. According to Toulmin, much of what it is to be an argument varies from field to field, with some features being field-dependent and others being field-invariant. He therefore focuses on an analysis of arguments that emphasizes justification rather than inference, where one begins with a claim and seeks justification for it. Such arguments he calls 'practical arguments'. Toulmin details six components required for their analysis: claims, evidence, warrant, backing, rebuttal, and qualifiers.

Although Toulmin's account of practical arguments has been influential on researchers in argumentation theory, its implementations in ontologies like the AMO fail to clearly demarcate arguments from the background of related sources of evidence and counterevidence. Using Toulmin's categories, the AMO class Argument inherits the following restrictions: a) *has evidence* min. 1, b) *has warrant* min. 1, and c) *has claim* exactly 1. Thus, it both excludes arguments containing only one premise or supposition as well as circular arguments in which a statement is used to infer itself, and also allows for arguments that may have parts other than those included in the Toulmin model, for the AMO does not, in its implementation, provide a closure axiom for its class 'argument', so it is unclear what is excluded from its extension. Also, in failure of our second adequacy criterion, there is no distinction drawn between sentence contents considered independently of any argument, from sentence contents that appear to be doing a specific kind of work within an argument whereupon the AMO could call them warrants, evidences, backings, claims, and so on.

There are no processes in the AMO as well, which leads to an unwelcome proliferation of relations in their OWL implementation. This can happen when it is assumed that predicates in a sentence always express relations, as when the sentence "Barry argues for liberty" appears to show arguing as a way Barry is related to liberty. This is unfortunate, since when relations proliferate in an ontology, they can keep information from being found through other connections that may otherwise be available by using classes. If we instead treated the example with classes, we could say that Barry is an instance of a person, who engages in an instance of an act of arguing, where this act creates an argument, and that argument has, as its conclusion, a statement, where that statement expresses a preference for liberty. Treating the verb in this sentence as an instance of class of activity rather than as a relation also allows us to more easily state when and where this act occurred.

These errors are not made by the Semanticscience Integrated Ontology (SIO), an ontology developed to facilitate biomedical knowledge discovery.[27] The SIO is distinguished among the ontological treatments of arguments we have surveyed by representing arguments in a canonical manner familiar to most readers of introductory logic textbooks. Its subclasses of arguments include valid, sound, deductive, and

[25]Toulmin (2003).
[26]Peroni and Vitali (2011).
[27]Dumontier, et al. (2014).

inductive arguments, and it also includes premises and conclusions among its classes, which facilitates the representation of argument parts. Furthermore, it includes an exclusively binary treatment of truth and falsity, implementing truth-values as subclasses of the class ‘information quality entity’. This allows SIO to be used to represent not only arguments with true premises, but also arguments with false ones, and to distinguish the true from the false. For many purposes, this may be a valuable feature to have. For instance, if statements known to be true or false can be tagged as such, then the ontology may be used to query for all those arguments that rest on false premises, whereupon the evidence for their conclusions may be re-evaluated. In addition, further properties such as soundness—defined as the property of a deductive argument that is both valid and has all true premises—may be inferred by looking for valid arguments with all true premises.

However, there are many parts of SIO that could be improved. First, the ontology itself does not use and is not compatible with common upper-level ontologies or widely used ontologies of information, and it would need to be rebuilt in order to become so. For instance, all information content entities[28] in SIO are a subclass of Object, which is defined as “an entity that is wholly identifiable at any instant of time during which it exists”. This definition is at odds with the BFO definition of Object, which defines an object as a maximal causally unified material entity. Furthermore, the definition of information content entity is ‘an object that requires some background knowledge or procedure to correctly interpret’. This definition is simultaneously too broad and too narrow, as there are some material artifacts (e.g. a computer) that fit this definition but are not themselves information content entities, and there are also information content entities that do not require any special background knowledge or procedure to correctly interpret, apart from the mere knowledge of a language or symbol system, and this does not seem to be the intended meaning.

Second, the SIO asserts that all arguments are a kind of ‘proposition’, where proposition is incorrectly defined as “a sentence expressing something true or false.”[29] It thus fails our first adequacy criterion. Most researchers in logic and the philosophy of language conceive of propositions not as sentences, but rather as the meaning or content declarative sentences express. And even where this distinction is clearly made, the arguments expressed by atomic sentences are not propositions, since an atomic sentence may be a sentence, but because it contains no inference, cannot be an argument. Finally, like AMO, the SIO fails criterion three, for it draws no distinction between the content of a sentence and when that content is used as part of an argument, as say, a premise, subconclusion, or supposition. Users of the ontology are thus required to choose whether they want to count some instance of a proposition as one or all of these, despite it being the case that whether a statement is a premise,

---

[28] The class ‘information content entity’ is broadly used among ontologies that employ the BFO framework. For a review on its characteristics and use, see Ceusters and Smith (2015).

[29] The OWL version of the SIO was accessed April 16, 2017 here: https://raw.githubusercontent.com/micheldumontier/semanticscience/master/ontology/sio/release/sio-release.owl

conclusion, or subconclusion depends entirely on its relation to an argument, and that relation will itself vary across arguments.

ArgO: An Arguments Ontology

In what follows, we introduce the Arguments Ontology (ArgO)[30], a small ontology of arguments designed to meet the constraints we have described and to be readily adopted and extended by others. ArgO was built using Basic Formal Ontology, and is a logical extension of the Information Artifact Ontology (IAO), a widely used information ontology employed in the biomedical domain. Here, we also provide a rigorous axiomatization in first-order logic, as well as some suggestions for natural extensions of ArgO that users may wish to adopt.

Statements, Background Ontological Commitments, and Arguments

As we have said, one must clearly distinguish sentences, sentence content, and force. While allowing for these distinctions, we restrict our discussion here to the content of *declarative sentences*, where by declarative sentence we mean a sentence that may be used to affirm truth or falsity. Such contents are sometimes called *statements*, and are held to be the primary bearers of truth values and the objects of propositional attitudes, such as belief.[31] For example, the sentence “Susan is happy” expresses the statement *Susan is happy*, which may be believed, is truth-apt, and is plausibly the meaning of the declarative sentence quoted. As a mere pattern of characters in a language, a sentence itself is none of these.

We are guided in our treatment of statements by BFO in general, and by the Information Artifact Ontology (IAO) in particular, which is a mid-level extension of BFO designed to represent information and information bearers. Statements fit naturally within the Information Artifact Ontology (IAO) class Information Content Entities, roughly, entities about things in the world. These include the content of sentences in a book, as well as the information encoded in Excel files on a hard drive. With respect to BFO, all information content entities are instances of Generically Dependent Continuant. Generically dependent continuants are those enduring entities lacking temporal parts that are identical across bearers; for example, a PDF may exist across distinct laptops. Hence, we treat statements as, broadly, entities about reality that may have identical instances across bearers. Two observant friends of Susan, for instance, may both believe the same content *Susan is happy*, expressing this content by uttering respectively “Susan is happy”.

Although we do not deal at length with sentences, it is instructive to examine how BFO and IAO provide resources to distinguish sentences from statements. Starting with BFO, generically dependent continuant entities are contrasted with specifically dependent continuants, which are not shareable. Examples include the particular color

---

[30]Project files may be accessed at: https://github.com/johnbeve/Argument-Ontology

[31]See McGrath (2012). In philosophical traditions, the content of declarative sentences is often called a “proposition”. Long-standing disputes over this term make it unwieldy for our purposes; hence, we use the more neutral of the two. We do not intend our use to commit us to *abstracta*, which would conflict with the realism adopted by BFO (see Ceusters and Smith 2015).

of a given apple, or a given sequence of shapes of ink inscribed on a particular page. In every case, specifically dependent continuants depend on some instance of the BFO class Material Entity, which consists of continuant entities that have matter as parts. For example, a particular redness (a specifically dependent continuant) depends on a particular apple (a material entity). Generically dependent continuant entities also bear a relation of dependence. In every case, generically dependent continuants must be concretized in some specifically dependent continuant, and thus indirectly, must have some material entity bearer. For instance, the repeatable generically dependent continuant pattern reflected in the specifically dependent continuant shapes of the material entity ink, is said to be concretized in the shapes, which in turn, depend on the ink (a material entity).

The IAO inherits these distinctions among generically dependent continuants, specifically dependent continuants, and material entities. Where Information Content Entity is an extension of Generically Dependent Continuant, the IAO extends Specifically Dependent Continuant with the class Information Carrier, and Material Entity with the class Information Bearer. Roughly, instances of the class Information Carrier are qualities of instances of the class Information Bearer. Applied to our ink on a page example, the repeatable pattern (an information content entity) is concretized in the shape (an information carrier) that depends on the ink (an information bearer).

ArgO adopts these distinctions, treating sentences as information carriers borne by various information bearers, though we will not have much more to say about these entities. For our purposes, it suffices to observe statements such as *Susan is happy* are information content entities concretized in sentences that are information carriers. There relations are illustrated in Figure 1.

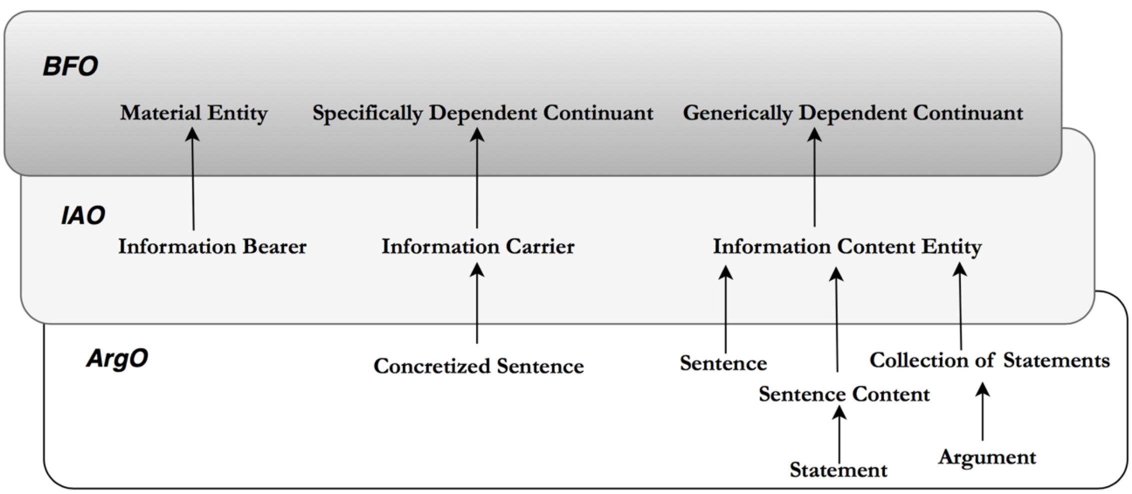

Fig. 1: ArgO classes and their parents in the Information Artifact Ontology (IAO) and Basic Formal Ontology (BFO).

Like statements, we claim arguments are best understood as instances of Information Content Entity, through an intermediary class: Collection of Sentence Contents.[32] The

[32]Cf. The use of "Object" vs "Object Aggregate" in (BFO 2.0). We avoid introducing aggregates along these lines for sentence contents since, among other reasons, aggregate entities must be

class Arguments is thus a subclass of Collection of Sentence Contents, where the sentence contents in that collection are ordered.[33]

Whereas mere instances of Collection of Sentence Contents may have no ordering, every instance of Arguments must be partially ordered. In this regard, arguments are again distinguished from textual entities, such as paragraphs containing multiple sentences in which arguments are expressed, as such sentences may appear in any ordering. We turn next to the ordered components of arguments.

Premises, Conclusions, and Suppositions

Let us set aside sentence contents for the moment and discuss statements. When a statement is a constituent of an argument, it is typically as either a premise or conclusion. Importantly, this is independent of the identity of an instance of Statement. For example, a given statement that serves as the conclusion of an argument, as expressed in line (3) of the following:

(1) If Andrew is in the Prado, then Andrew is in Madrid.
(2) Andrew is in the Prado.
(3) Hence, Andrew is in Madrid.

may serve as a premise in a distinct argument, as expressed in line (2):

(1) If Andrew is in Madrid, then Andrew is in Spain.
(2) Andrew is in Madrid.
(3) Hence, Andrew is in Spain.

A natural first thought then is that *being a premise* and *being a conclusion* are two roles a given statement can bear within an argument.[34] Early in our thinking, we explored using roles to distinguish statements from premises and conclusions by appealing to the BFO class Role. Roles in BFO are realizable entities that may be gained or lost by a bearer based on context, and which are, in every case, correlated with a realized process. For example, a student role may be gained by an individual in an academic context, and is realized in processes, such as studying for an exam and completing coursework.[35] Similarly, premises and conclusions may naturally be understood as statements taking on different roles in particular arguments that are realized in

mereologically disjoint, but a given argument may have repeating statement parts in distinct positions. For instance, in "(1) God exists; (2) Hence, God exists", would not count as an aggregate of statements, since, arguably, the content of (1) overlaps the content of (2). By contrast, our approach treats arguments as wholes that have parts.

[33] Note that we take a partial ordering to also be a kind of ordering. In addition, we leave open whether there are other disjoint siblings of Argument reflecting distinct orderings of sentence contents.

[34] Accordingly, an ontology of arguments should not treat statement as a subclass of premise, supposition, or conclusion. SIO appears to treat *Andrew is in Madrid* not as a type of statement, but as a type of premise. This leaves SIO unable to say what content is common between our sample arguments. More sharply, such a proposal would have difficulty explicating why an argument such as "(1) God exists; (2) Hence, God exists", is a bad argument.

[35] BFO 2.0, pp. 56–57.

respective processes of correlated types, in particular, processes of inference. However, the BFO class Role is a subclass of Specifically Dependent Continuant, where a given instance of Role is borne by a material entity on which this specifically dependent continuant depends. Since the class Specifically Dependent Continuant is disjoint from Generically Dependent Continuant, and because Statement is a subclass of the latter, statements cannot bear roles. Hence, we cannot rely on this initially plausible class to characterize premises and conclusions.

This leaves us still trying to account for the difference between statements *per se* on the one hand, and premises and conclusions on the other. A second difference is that premises and conclusions are always used in arguments, whereas statements need not be. Yet a third difference is that 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. Such an act is naturally characterized as a process in BFO, since an act has temporal parts and is not instantaneous. We reflect this link between premises and conclusions by defining a class Act of Inferring. This class, moreover, reveals the lines along which we may characterize these statements. Roughly, premises, conclusions, and the like, are not entities in their own right, distinct from statements. Rather, 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 taken from the Common Core Ontologies[36], where they are defined as follows:

> *has input* $=_{def}$ a relation between a Continuant and a Processual Entity such that the presence of the Continuant at the beginning of the Processual Entity is a necessary condition for the start of the Processual Entity.
>
> *has output* $=_{def}$ a relation between a Continuant and a Processual Entity such that the presence of the Continuant at the end of the Processual Entity is a necessary condition for the completion of the Processual Entity.

One of the features motivating our original pursuit of roles for premises and conclusions was that roles are associated with processes. Treating premises and conclusions, instead, as defined subclasses of Statement that bear input and output relations to an act of inferring maintains this consideration. All that is further needed is a restriction on premises and conclusions to the effect that they are particular to arguments: this we accomplish by treating premises as mere instances of statements that stand in a 'has premise' relation to some argument (the same for conclusions, *mutatis mutandis*). This restriction mirrors the restriction on roles, which holds that they only inhere in their material bearer in a social context (as when one is a doctor within

---

[36] Accessed 1/1/18: https://github.com/CommonCoreOntology/CommonCoreOntologies

the social context of a hospital). However, with such a restriction in place, we acquire much of what a treatment using roles would otherwise have accomplished.

Suppositions, like premises, provide support for the conclusion of an argument, but are subtly distinct. For whereas premises used in an argument are affirmed by the arguer, suppositions are more often *accepted* rather than affirmed.[37] Affirming a statement is closely tied with believing it. Affirming involves aiming at truth, is typically informed by evidence, and is subject to rationality constraints. For example, sincerely affirming the statement *It is raining outside now* entails the affirming agent believes the statement is true, and suggests one formed this belief based on evidence. Moreover, affirming is minimally restricted in that individuals seem unable to affirm what is clearly impossible, or at least experience significant cognitive resistance when attempting to do so.[38] Premises in a given argument seem best described as affirmed, as they appear to share these characteristics.

In contrast, merely accepting a statement is not so intimately tied to belief, need not aim at truth, may not be informed by evidence, and is not obviously subject to rationality constraints. For example, an individual might plausibly accept *It is raining outside now* while claiming they do not actually believe it, acknowledging the statement is false, but admitting they have no evidence either way. Indeed, one may accept what is clearly impossible with little cognitive resistance, for instance, accepting *It is neither raining outside now nor not raining outside now*. Relatedly, statements are typically accepted for the sake of some further goal, such as continuing a conversation, hypothetical deliberation, or indirect reasoning.

Suppositions in a given argument seem best described as accepted, as they appear to share these characteristics. For suppositions, unlike premises, may be disbelieved, uninfluenced by evidence, known false, known impossible, and are often employed to explore the consequences of hypothetical commitments. We reflect these observations by characterizing a supposition as an accepted but not affirmed statement that is input in an act of inferring that provides support or justification for the conclusion of an argument. Similarly, we augment the previous characterization of premise, where a premise is now an affirmed statement that is input in an act of inferring that provides support or justification for the conclusion of an argument. Moreover, we reflect the distinction between affirming and accepting as one between distinct acts. An act of affirming is a process in which an agent participates, where the agent believes the statement is either true or false based on evidence and within some rationality constraints. An act of accepting is a process in which an agent participates, where the agent entertains a statement as true or false, regardless of belief, evidence, or rationality constraints, as input to an act of inferring.

Are conclusions affirmed or accepted? In many cases, conclusions are clearly affirmed. For example, consider the argument represented in:

---

[37]Cf. Stalnaker (1984), Arcangeli (2016), and Toumela (2000).

[38]Observe, too, the converse relationship between believing and affirming: “I believe it is raining outside but I do not affirm it” suggests one may believe while not affirming, but the infelicity of “I affirm it is raining outside but I do not believe it” suggest affirming requires belief. For our purposes then, if one affirms a statement then they believe it, but not conversely. Additionally, supposing is independent of both believing and affirming.

(1) If Andrew is in the Prado, then Andrew is in Madrid.
(2) Andrew is in the Prado.
(3) Hence, Andrew is in Madrid.

where an agent making this argument affirms (1) and (2) and participates in an act of inferring from the premises to the conclusion. Plausibly in this case, the agent affirms the content of (3) as well. In other cases, it is not so clear whether a conclusion is affirmed or accepted. Consider:

(1) If Andrew is in the Prado, then Andrew is in Madrid and Andrew is in Spain.
(2) SUPPOSE | Andrew is in the Prado.
(3) | Hence, Andrew is in Madrid and Andrew is in Spain.
(4) | Hence, Andrew is in Spain.
(5) Hence, if Andrew is in the Prado, then Andrew is in Spain.

where an agent plausibly participates in several acts of inferring, in particular, from (1) and (2) to (3), from (3) to (4), and finally from (1)-(4) to (5). Here, it seems the respective outputs of these acts of inferring are in some cases accepted and some cases affirmed. For example, an agent affirming the content of (1), accepting the content of (2), and inferring to the content of line (3) is plausibly understood as accepting (3), rather than affirming it, as the agent does not thereby believe (3) is true full-stop. Similarly, for line (4). Rather, what the agent affirms is a connection between what is supposed and associated consequences of the supposition. This is reflected in line (5), the output of (1)-(4), which is best understood as affirmed. Thus, generally speaking, it seems conclusions may be affirmed or accepted, based on whether the input to an act of inferring is a premise or a supposition. With that in mind, we augment our characterization of conclusions as affirmed *or accepted* outputs of an act of inferring thought to be supported or justified by corresponding affirmed or accepted inputs.

This last example suggests arguments may exhibit multiple conclusions. However, many treat a given argument as having only one conclusion (often distinguished as the "main conclusion"), a treatment we adopt as well. We turn next to explicating this feature of ArgO.

Subconclusions and Complex Arguments

Arguments are often complex, involving several premises and perhaps suppositions, and multiple inferential steps taken toward the main conclusion. Consider our initial pair of arguments collapsed into a single argument:

(1) If Andrew is in the Prado, then Andrew is in Madrid.
(2) Andrew is in the Prado.
(3) Hence, Andrew is in Madrid.
(4) If Andrew is in Madrid, then Andrew is in Spain.
(5) Andrew is in Madrid.

(6) Hence, Andrew is in Spain.

There are, at least, two ways to view the argument, as two simple arguments involving two acts of inferring, as depicted in Figure 2:

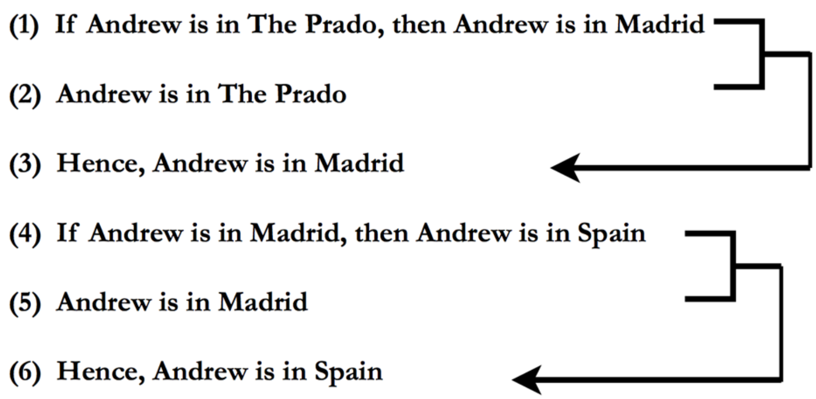

Fig. 2: Arguments Simple-A and Simple-B; arrows indicate distinct acts of inferring

Or as one complex argument, involving a single act of inferring:

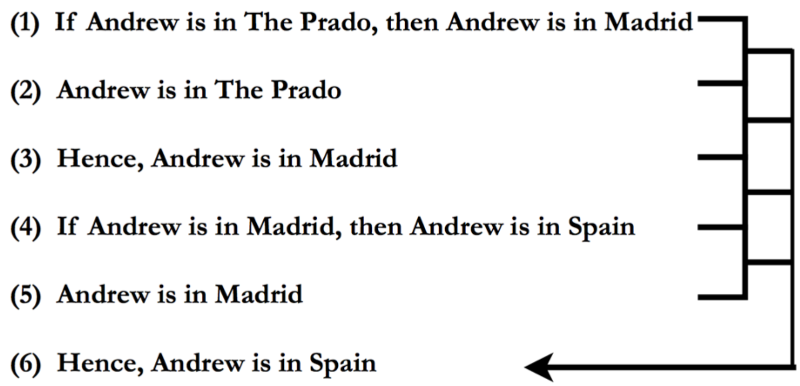

Fig. 3: Argument C; arrow indicates one act of inferring

In Figure 2, we have two acts of inferring, four premises ((1), (2), (4), (5)), and two conclusions: (3) and (6), and we have two distinct arguments, each with their own acts of inferring. Call the argument reflected in (1)-(3), Simple-A, and that reflected in (4)-(6), Simple-B. In Figure 3, we have one act of inferring, five premises (1)-(5), and one conclusion (6). There is one argument here; call it C. Now, Simple-A and Simple-B are clearly components of C, and note that the content of (3) and (5) is the same. However, with respect to Simple-A, this content is the output of an act of inferring, namely, a conclusion, and with respect to Simple-B, this content is input to an act of inferring, specifically, a premise. Similarly, in C, this content is a premise. It is standard to consider the function of (3) in C a *subconclusion* of C, an inferential step toward the main conclusion of the argument, represented by (6). However, treating (3) as a subconclusion of a single argument in this manner overlooks the role (3) plays among the component arguments of C. Subconclusions must be characterized with respect to multiple arguments that stand in a parthood relation. Moreover, connections among component arguments in a complex argument should reflect the role played by subconclusions.

The BFO parthood relation holding between continuants, and in particular, generically dependent continuants, provides resources for making sense of the parthood relations with complex arguments.[39] With BFO parthood, we may define a class relevant to our characterization of subconclusion: Complex Argument. An instance of Complex Argument is an argument with at least one argument as a proper part, and which has only argument as parts. Since the parthood relation is a partial order, characterizing complex arguments in this manner generates ordering among arguments of greater complexity. Thus, a complex argument A may be decomposed into complex parts, B and C, themselves decomposable into complex parts. Ultimately, the decomposition will result in simple argument parts, represented by I, J, E, K, L, G, and H in Figure 4. The directed arrows represent parthood; for example, Argument I is part of Argument D, which is part of Argument B, and so on:

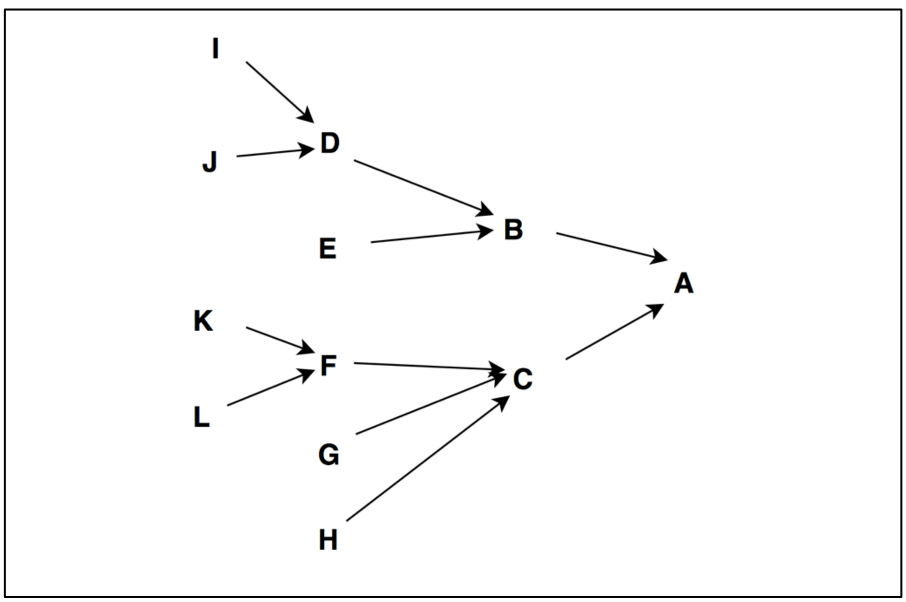

Fig. 4: Partial order over complex argument parts.

[39]More specifically, BFO's mereology is bifurcated, with a ternary, temporally indexed relation governing parts of continuants, and a binary relation governing parts of occurrents. Here, we suppress the ternary index for continuants. Both parthood relations are provably equivalent to the minimal extensional mereology of Simons (1987). Hence, parthood is reflexive, antisymmetric, transitive, ensures weak supplements, and unique fusions for overlap. Proper parthood is easily defined in terms of parthood, e.g. x is proper part of y just in case x is part of y and x≠y, and is irreflexive, asymmetric, and transitive, and hence, i.e. a strict order.

Every subconclusion can then be understood as a statement that bears a 'subconclusion in' relation to a complex argument. This relation requires that 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. This definition of complex argument ensures any instance involves an act of inferring and has at least two arguments as a proper part. Combined with the plausible commitment that each instance of Argument has only *one* conclusion, so that each argument that is part of a complex argument has its own conclusion, as does the complex argument of which they are parts, entails each instance of Complex Argument includes an instance of Subconclusion. Thus, our definition of subconclusion appeals to multiple arguments and the respective functions such statements have within them.

We take care here to emphasize two features of our commitments thus far. First, assuming the following is an argument:

(1) If Andrew is in the Prado, then Andrew is in Madrid.
(2) Andrew is in the Prado.
(3) If Andrew is in Madrid, then Andrew is in Spain.
(4) Hence, Andrew is in Spain.

does not entail it is a *complex* argument. This is because the number of acts of inferring involved in producing this argument will be relative to an interpretation, and we do not wish to provide strict interpretation criteria for users of the ontology—allowing users instead to employ those criteria that suit their projects. Relatedly, it is not obvious that this argument is decomposable into argument parts. One may object, claiming (1) and (2) clearly entail the antecedent of line (3), call it (3*), which together clearly entail (4). Hence, three arguments are exhibited implicitly in (1)-(4) which are intuitively parts of this argument, and which involve a statement as a subconclusion. Hence, (1)-(4) is complex. But this objection misunderstands the task of representing arguments with ontologies. Lines (1)-(4) and the assumed *act of inferring* from (1)-(3) to (4) do not entail the existence of a further *act of inferring*, say, from (1)-(2) to (3*). Regardless of how obvious the step may seem, licensing normative corrections extends beyond our task of representing arguments. Yet without this addition, this argument is not obviously decomposable into argument parts and thus not complex.

Second, subconclusions interact with suppositions in a satisfying way on our proposal. We illustrate by returning to our example argument involving a supposition, which we may characterize as complex:

(1) If Andrew is in the Prado, then Andrew is in Madrid and Andrew is in Spain.
(2) SUPPOSE Andrew is in the Prado.
(3) Hence, Andrew is in Madrid and Andrew is in Spain.
(4) Hence, Andrew is in Spain.

(5) Hence, if Andrew is in the Prado, then Andrew is in Spain.

Here, we have three arguments. First, from (1) and (2) to (3); second, from (3) to (4); third, from (1)-(4) to (5). Call the first, again, Simple-A, the second Simple-B, and the third C. Clearly Simple-A and Simple-B are parts of C. Clearly, too, we have subconclusions on (3) and (4). With respect to Simple-A, (3) is the conclusion, and with respect to Simple-B and C, (3) is a premise. On the other hand, with respect to Simple-B, (4) is the conclusion, and with respect to C, (4) is a premise. Hence, both (3) and (4) are subconclusions. Moreover, with respect to Simple-A, (3) is plausibly understood as accepted, as is (4) with respect to Simple-B, and as are both (3) and (4) with respect to C. Then much like our characterization of conclusions, subconclusions may be accepted as in this complex argument, or affirmed, as in the example with which we began this section.

Revisiting Arguments

We said earlier that the same sentence contents can be a premise in one argument, a conclusion in another, and thus premises and conclusions are best understood as sentence contents that bear relations to particular arguments. This entails that an instance of Sentence Content itself can be identical across arguments, even if its function as a premise or conclusion is not.

This feature satisfies our adequacy constraint 2, and permits tracking of sentence contents across arguments independently of how they are used in arguments. In contrast to premises, conclusions, and suppositions, subconclusions are not relativized to unique arguments, but rather, to combinations of complex arguments and decomposable parts, reflecting the cross-argument axioms defining the relation.

Taking stock, we have arguments as ordered collections of sentence contents involving premises, suppositions, and a single conclusion, and complex arguments as arguments with arguments as proper parts. In turn, premises are equivalent to sentence contents that are affirmed inputs of acts of inferring, whereas conclusions are equivalent to sentence contents that are affirmed or accepted outputs of acts of inferring. Suppositions are sentence contents that are accepted inputs of acts of inferring, whereas subconclusions are sentence contents that are affirmed or accepted inputs and outputs of distinct acts of inferring in distinct arguments that are proper parts of a complex argument. Finally, sentence contents are parts of arguments as well as parts of collections of sentence contents.

Arguing, Creating Arguments, and Chains of Reasoning

Thus far, our characterizations make use of various acts that deserve attention, so we turn now to the process side of ArgO. We have made much use of the act of inferring in our discussion. Yet, this is not the only act of importance for our topic. Consider, 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 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. Of course, one may have no intention to convince others of some conclusion; one may be creating an argument for its own sake, or for the purpose of interpretation and analysis, 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.[40] Then again, it might not. Quite often, when engaging in an act of arguing we merely reuse previously created arguments. When we do so, we engage in an act of arguing that has no act of argument creation as process part.

Because we are interested in the composition of arguments, we take as our focus the act of argument creation. 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 that it creates by the *is created by* relation, which holds between continuant and a process. These relations are depicted in Figure 5.

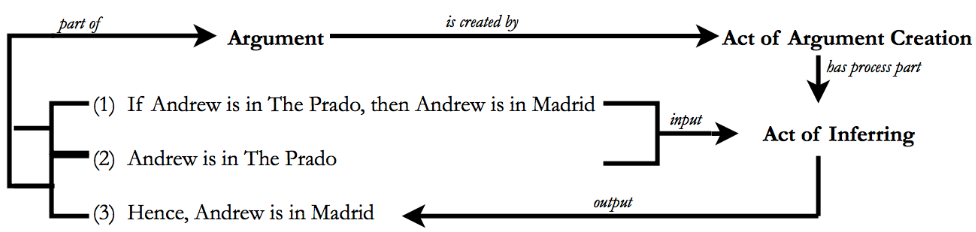

Fig. 5: Simple Act of Argument Creation.

Here, we have a simple argument with premises (1) and (2), conclusion (3), and an act of inferring. The premises and conclusion are instances of Statement that are in different relations to the argument, the super-relation of which is 'part of'. These relations are created in the act of argument creation by the act of inferring.

Complex arguments involve a single act of inferring on one disambiguation and a corresponding single act of argument creation, as illustrated in Figure 6.

---

[40]Here, we rely on the occurrent parthood relation of BFO, where a process part is an occurrent part of some other process.

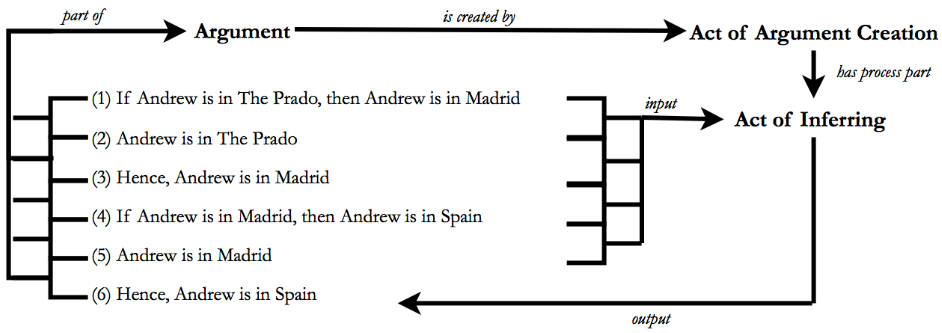

Fig. 6: Complex Argument with one Act of Inferring.

However, complex arguments have parts, each of which involves an *act of inferring* and an *act of argument creation*, illustrated by a disambiguation of our familiar complex argument in Figure 7.

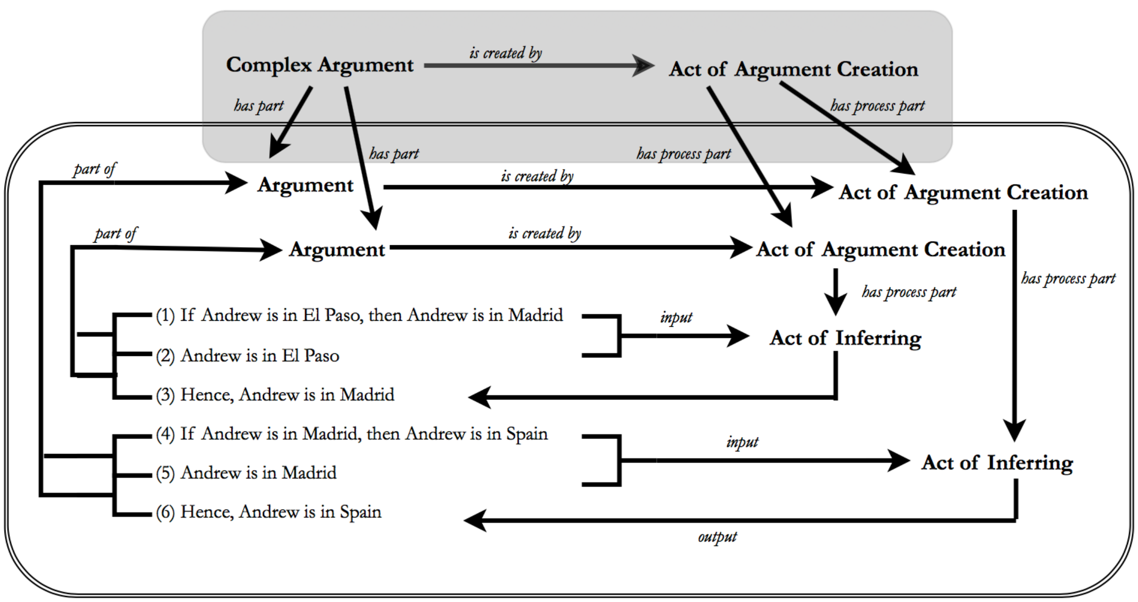

Fig. 7: Complex Argument, emphasis on argument parts.

Here, we have two arguments each with an act of inferring and act of argument creation, and each of which stands in the *is created by* relation to a distinct argument. Since these distinct arguments are parts of a larger complex argument, they are tied to a larger act of argument creation that stands in the *is created by* relation to the complex argument. The complex argument *is created by* an act of argument creation that has, as process parts, two acts of argument creation, paralleling the argument parts that compose the complex argument.

These acts are distinguished from other processes by employing language, whether in thought, speech, or writing, and thus we unify these acts by extending the BFO class Process to Language Act, where a language act is defined as an intentional act involving language. See Figure 8.

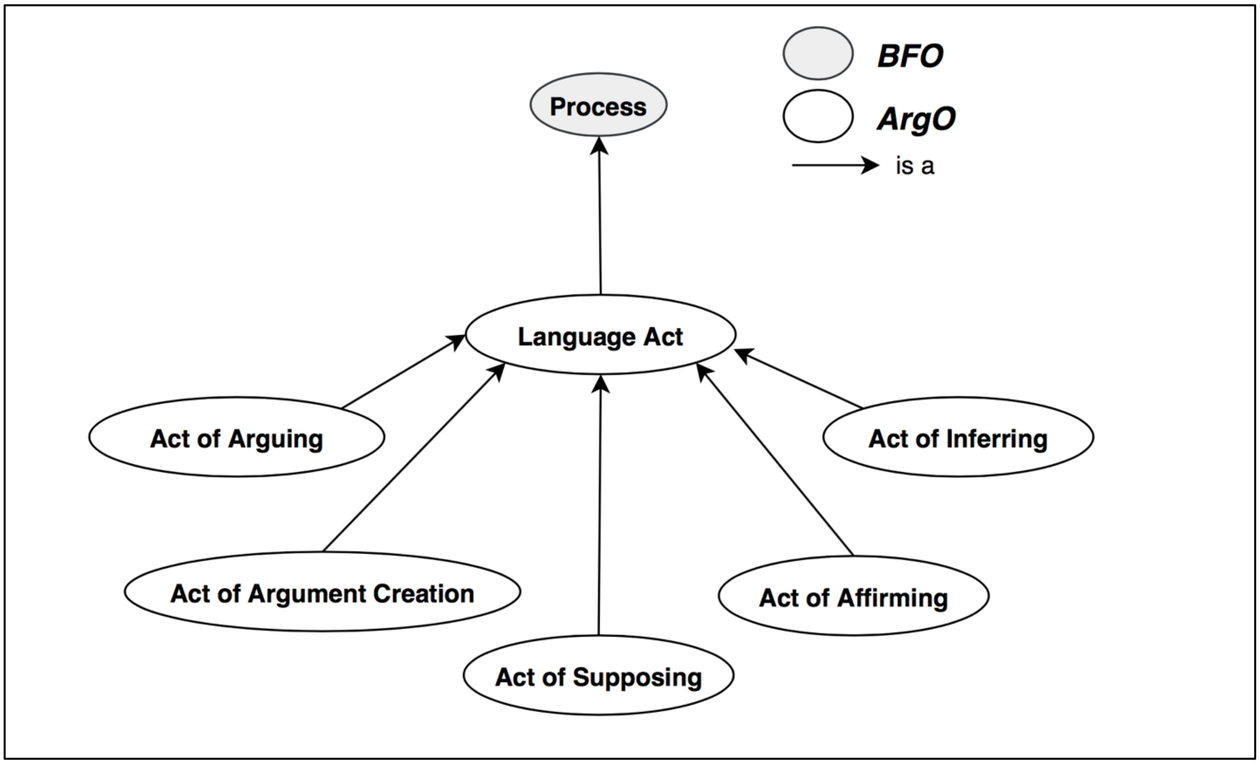

Fig. 8: BFO Process and ArgO Acts relationships.

Just as BFO distinguishes continuants from occurrents, we have taken similar pains here to represent both aspects of an argument. An argument is not a process, but an enduring entity that exists completely at every time during which it exists. In the case of a simple argument, the argument and its parts may be seen in Figure 9, whereas Figure 10 depicts the occurrent side of arguments, where every act of argument creation has, as process part, some act of inferring, and this act of inferring has, in turn, particular functional relationships to sentence contents within the argument itself.

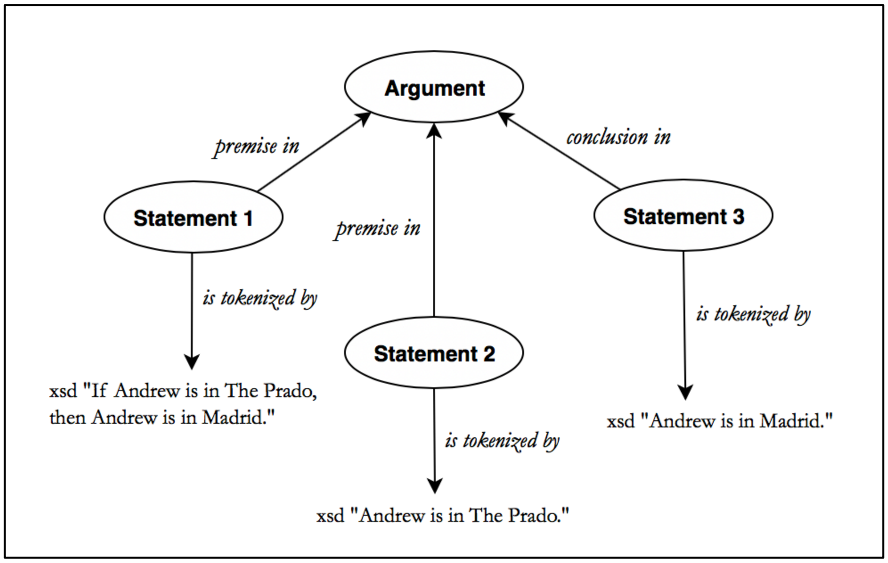

Fig. 9: Simple Argument.

In order to keep the presentation of these graph simple, we adopt here the relation *is tokenized by*, which here holds between an information content entity and a string. In our implementation, however, the full relation follows the practice of the IAO, in having information content entities concretized in information quality entities, which inhere in information bearers, and it is these information bearers that have text values that are strings.

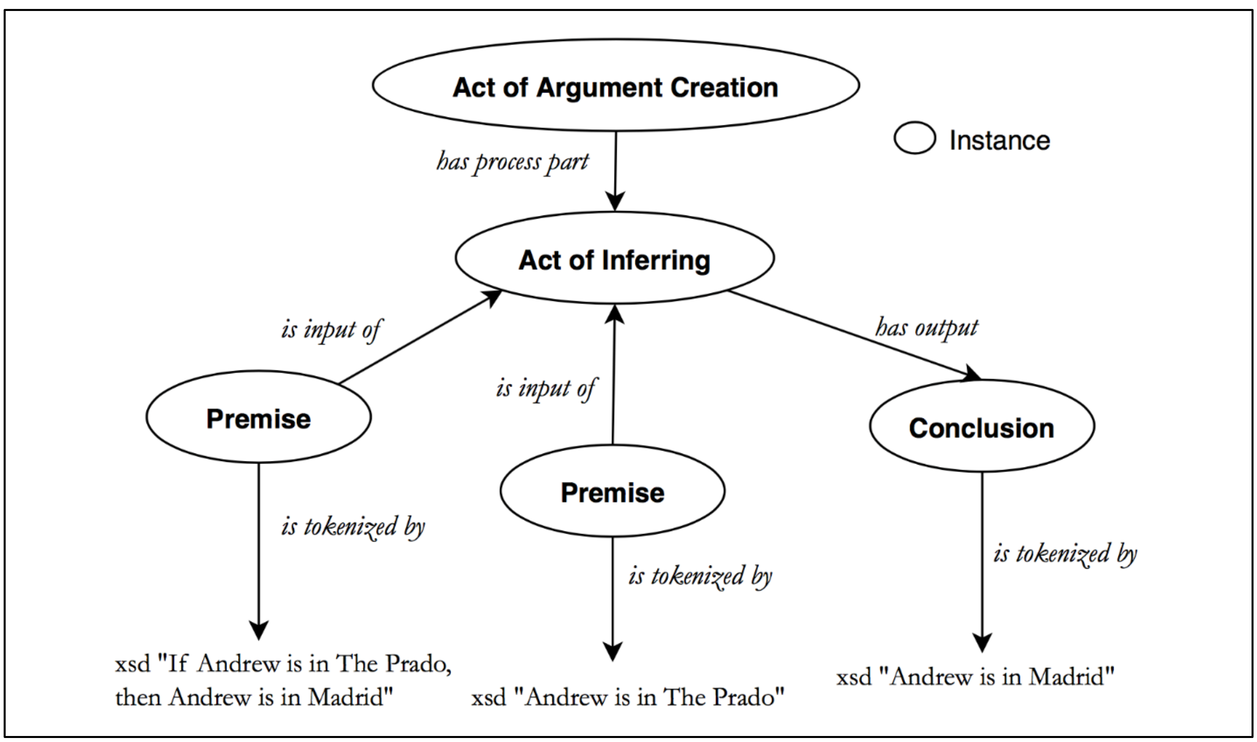

Fig. 10: Act of Argument Creation for a Simple Argument.

Before turning to our axiomatization of ArgO and formal results, let us also look at a complex argument involving a supposition. Consider:

(1) If Andrew is in the Prado, then Andrew is in Madrid and Andrew is in Spain
(2) SUPPOSE Andrew is in the Prado
(3) Hence, Andrew is in Madrid and Andrew is in Spain
(4) Hence, Andrew is in Spain
(5) Hence, if Andrew is in the Prado, then Andrew is in Spain

As in the previous case, we wish to represent both the continuant and occurrent aspects at play in the argument. In Figure 11, we depict the enduring parts of the argument, where line (2) of argument differs from line (1) in being a supposition from which lines (3) and (4) are inferred. However, we represent these subconclusions as the outputs of two distinct acts of inferring, which are themselves process parts of the act of inferring, see in Figure 12, that *has output* the conclusion of the argument.

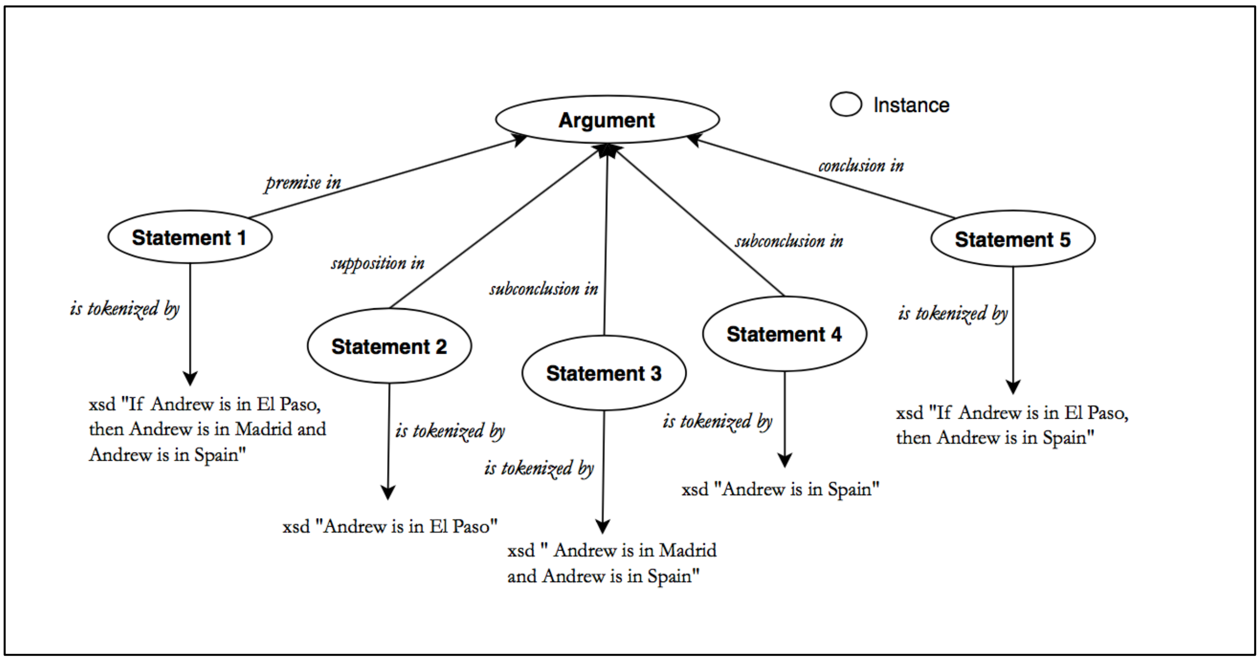

Fig. 11: Complex Argument with Supposition.

It is not necessary to know the precise order or times on which these various cognitive acts occur in order to posit their existence. Frequently, agential reasoning begins with a conclusion and then looks for supporting reasons, while other times a conclusion really is a final output of a process of mental deduction. Still, a natural temporal ordering suggests itself in the presentation of an argument, where the *act of inferring* that outputs the conclusion of the argument precedes any others that output subconclusions in a standard ordering of temporal intervals. A full treatment of a given argument will reflect temporal precedence of cognitive acts; however, we leave this aspect for specific users of our ontology to work out in each case.

We find using processes is a natural way to think about the work of a statement within an argument. In addition, it also allows for a symmetric pattern of ordering between occurrent process and continuant parts, since an act of argument creation will have ordered acts of inferring as process parts, and we can read off of these process part relationships the part relationships that hold among the argument and its parts. This means, for instance, that if a premise entails a subconclusion, and that subconclusion entails another subconclusion, that there is a parallel series of parthood relationships on the process side, such that there are two acts of inferring, both of which are process parts of an act of argument creation, both of which share a statement as a participant.

ArgO and Our Proposed Constraints

We began our discussion by defending a series of methodological and adequacy constraints an ontology of arguments should meet, and we found other ontologies of arguments failed one or more of these conditions. By contrast, ArgO employs only logically defined classes employing a genus-species form of definition that the ontology inherits from retaining compatibility with BFO and IAO, whose classes meet this condition as well. ArgO is also exclusively concerned with processes and information content entities, and thus it maintains the distinction between informational entities, on the one hand, and the material entities that bear them on the other. Furthermore, as we report in the following section, the axioms of ArgO have been proven to be logically consistent. In these regards, ArgO meets the methodological constraints with which we began.

Regarding the adequacy constraints on an ontology of arguments, ArgO begins by characterizing arguments in a familiar way: as reasons for a conclusion, where those reasons include premises, suppositions, and subconclusions. Unlike other approaches, ArgO formally distinguishes sentences from the content of sentences, as well as the function of a statement within an argument whereupon it is called a premise, a conclusion, and so forth. These features satisfy adequacy constraints 1 and 2.

Whereas other treatments of arguments have focused on desiderata for assessment, ArgO is distinguished by focusing on mereological composition, where we are concerned both with the occurrent parts of processes that produce arguments and their components, as well as the continuant parts among the components themselves. Thus, on its own, ArgO does not impose a single logic with regard to rules of inference and composition, nor a series of schemes that may be used to prescribe inference rules or classify arguments, nor even a treatment of the truth and falsity of statements. Yet as we will discuss shortly, ArgO may be easily extended in any of these ways by application ontologies that fulfill specific purposes. Because we have designed ArgO to remain open to such extensions, we find it also satisfies the remaining adequacy constraints 3, 4, and 5.

Axiomatization of ArgO

ArgO has been implemented in first-order logic and OWL 2 DL and each implementation is a conservative extension of the Information Artifact Ontology (IAO), a widely used information ontology employed in the biomedical domain, implementations of which are themselves conservative extensions of implementations of BFO.[41]

Here, we include sample axioms characterizing the classes and relations of ArgO. The reader is directed to the ArgO Github repository for the complete set of axioms.[42] Our language is unsorted first-order logic supplemented with identity, with standard logical connectives. All variables are implicitly bound, and leading universal quantifiers are omitted. We reserve the variable "t" for instances of the BFO class Temporal Region.

We assert axioms characterizing the relationships among the classes of ArgO, the IAO, and BFO, such as:

1. *InformationContentEntity(x) → GenericallyDependentContinuant(x)*
2. *CollectionOfSentenceContents(x) → InformationContentEntity(x)*
3. *Argument(x) → CollectionOfSentenceContents(x)*
4. *ComplexArgument(x) → Argument(x)*
5. *SentenceContent(x) → InformationContentEntity(x)*
6. *Statement(x) → SentenceContent(x)*
7. *LanguageAct(x) → Process(x)*
8. *ActOfArguing(x) → (LanguageAct(x) & ~(ActOfArgumentCreation(x) | ActOfAccepting(x) | ActOfAffirming(x) | ActOfInferring(x)))*
9. *ActOfArgumentCreation(x) → (LanguageAct(x) & ~(ActOfArguing(x) | ActOfAccepting(x) | ActOfAffirming(x) | ActOfInferring(x)))*

And so on. We also assert that all sibling classes in the ontology are disjoint. In addition, we assert axioms characterizing the classes and relations of ArgO in accordance with the definitions we have defended. For example:

10. *Argument(x) → ∃y∃z∃t ((premiseIn(y,x,t) | suppositionIn(y,x,t)) & conclusionIn(z,x,t) & ∀u(conclusionIn(u,x,t) → z=u))*
11. *ComplexArgument(x) → ∃y∃t (Argument(y) & properContinuantPartOfAt(y,x,t)))*
12. *premiseIn(x,y,t) → (SentenceContent(x) & Argument(y) & continuantPartOfAt(x,y,t))*

[41] More specifically, a theory T', or consistent set of first-order axioms, is a (proof-theoretic) conservative extension of theory T if and only if T and T' have overlapping signatures, and any theorem of T is a theorem of T'. In this sense, Argo represented as a theory, i.e. set of consistent first-order axioms, overlaps the signature(s) of IAO (and BFO), and any theorem entailed by IAO (or BFO) is entailed by Argo.

[42] https://github.com/johnbeve/Argument-Ontology

13. *premiseIn(x,y,t) → ∃z∃w∃u(ActOfAffirming(z) & outputOf(x,z,t) & ActOfInferring(w) & inputOf(x,w,t) & ActOfArgumentCreation(u) & isCreatedBy(y,u,t) & properOccurrentPartOf(z,u) & properOccurrentPartOf(w,u))*
14. *conclusionIn(x,y,t) → (SentenceContent(x) & Argument(y) & continuantPartOfAt(x,y,t))*
15. *conclusionIn(x,y,t) → ∃z∃w∃u((ActOfAffirming(z) | ActOfAccepting(z)) & outputOf(x,z,t) & ActOfInferring(w) & outputOf(x,w,t) & ActOfArgumentCreation(u) & properOccurrentPartOf(z,u) & isCreatedBy(y,u,t) & properOccurrentPartOf(w,u))*
16. *subconclusionIn(x,y,t) → (SentenceContent(x) & ComplexArgument(y) & (premiseIn(x,y,t) | suppositionIn(x,y,t)))*
17. *subconclusionIn(x,y,t) → ∃z(Argument(z) & properContinuantPartOfAt(z,y,t) & conclusionIn(x,z,t))*
18. *ActOfArgumentCreation(x) → ∃y∃t (Argument(y) & isCreatedBy(y,x,t))*
19. *ActOfInferring(x) → ∃y∃z∃t (Statement(y) & Statement(z) & inputOf(y,x,t) & outputOf(z,x,t))*
20. *(isCreatedBy(x,y,t) & isCreatedBy(z,y,t))-> x=z)*
21. *(ActOfArgumentCreation(x) & ActOfArgumentCreation(y) & ComplexArgument(z) & isCreatedBy(z,x,t) & properContinuantPartOfAt(w,z,t) & isCreatedBy(w,y,t)) → properOccurrentPartOf(y,x)*

Additionally, we provide sample theorems that follow from our axiomatization. Proofs were generated with the automated theorem proving software Prover9, and can be found in our repository. Results include that the *premiseIn*, *conclusionIn*, *suppositionIn*, and *isCreatedBy* relations are both irreflexive and asymmetric, as well as more interesting theorems such as:

22. *(ActOfInferring(x) & ActOfArgumentCreation(y) & ActOfArgumentCreation(z) & occurrentPartOf(x,y) & occurrentPartOf(y,z)) → occurrentPartOf(x,z)*
23. *(ActOfArgumentCreation(x) & ComplexArgument(y) & isCreatedBy(y,x,t) & Argument(z) & properContinuantPartOfAt(z,y,t) → ∃w(ActOfArgumentCreation(w) & isCreatedBy(z,w,t) & occurrentPartOf(w,x))*

Countermodels were also demonstrated for certain sentences, using the finite model checker Mace4 bundled with Prover9. Moreover, using Mace4, we were also able to demonstrate satisfiability for our axiom set, and thus consistency. More importantly, we were able to demonstrate satisfiability when the axioms are saturated in various ways. For example, a model was generated for a complex argument composed of four statements, complete with premises, suppositions, a subconclusion and conclusion. This model thereby also contained simple arguments, with respective premises, suppositions, and conclusions. Various verified models can be found in the ArgO repository.

Extensions to ArgO

We turn now to detailing how ArgO might be extended. The first extension displays how ArgO may be used to examine weaknesses among arguments. The second displays how ArgO may be used with different logical systems, including: classical, multi-valued, and intuitionistic. Each extension is accompanied by an axiomatization, demonstrating its formal relationship to ArgO.

Extensions: Narrow and Broad Counterarguments

We have so far spoken about inter-relationships among arguments that involve the composition of complex arguments; however, we now turn to intra-relations among arguments, where independent arguments have a variety of possible relationships to one another. Here, we discuss one important family of such relationships, but what we say here may also apply for representations of other relations, such as supporting.

In every case, counterarguments are themselves arguments, but an argument is only a counterargument when it stands in a certain relation to another argument. Such relations come in many different forms. For example, for some argument A there may be some counterargument B if counterargument B has a conclusion that stands in a relation of contradiction to the conclusion of argument A. However, this is a narrow case. More often, cases of counterarguments undermine, but do not contradict, some part of another argument. Such cases hold when, some argument A may have some counterargument B if the conclusion of B raises concerns for the justification of one or more premises of A. Concerns over justification are evidential and not, strictly speaking, logical. Nevertheless, this is a notion of counterargument well worth representing.

Because counterarguments are partially constituted by a wide range of distinct relations, we do not treat them as a class in ArgO. Nevertheless, it is straightforward to use our classes to construct queries to return various kinds of counterargument. We need only extend our proposal slightly to include an *opposes* relation, which holds between premises, suppositions, conclusions, and subconclusions across arguments. This relation is broad enough to permit more refined sub-relations, for example: *negates*, *contradicts*, and *undermines*. Hence, a query to return all of the narrow counterarguments for a given target argument might look for all arguments whose conclusions stand in the *opposes* relation to the conclusion of the target argument. See Figure 12.

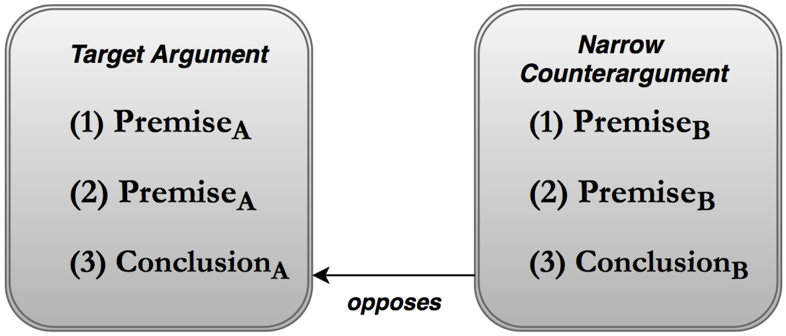

Fig. 2: Narrow Counterargument.

Similarly, a query to return all of the broad counterarguments for a given target argument, might look for all arguments with any premise, supposition, conclusion, or subconclusion that stands in the *opposes* relation to the conclusion of the target argument.

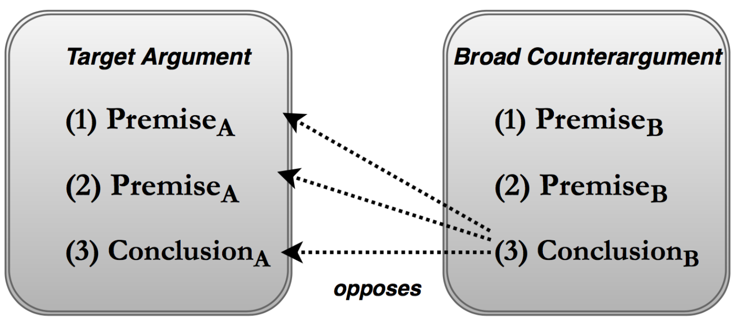

Fig. 3: Broad Counterargument.

See Figure 13. Here, the dotted arrow lines indicate the holding of any one of these opposes relations is sufficient to count an argument as a broad counterargument. Presumably, with the latter query issuing a larger return than the former.

Extensions: Argument Schemes, Truth, Probability, Validity

Many researchers who wish to extend the present ontology will likely want to incorporate a treatment of the formal properties of arguments. In our desire to remain neutral with respect to different approaches, we have not built into our ontology a univocal treatment of this subject. For instance, ArgO allows researchers who wish to use Walton argument schemes to do so by introducing a class of schemes as a subclass of the class information content entity, with each particular scheme serving as an instance of the class. In particular, the Walton argument scheme called “abductive argument” includes cases of arguing from the existence of a data set in a given case to the best explanation of the data set. The scheme ‘abductive argument’ is thus not a class, but an instance, and this instance describes instances of Argument that share the "abductive argument" form.

Whereas most logicians take propositions to be the proper bearers of truth and falsity, because—as we explained in section III.a.—we eschew talk of propositions in our ontology, *mutatis mutandis*, we take sentence contents to be the proper bearers of truth and falsity. We therefore advise that ArgO may accommodate evaluations of sentence contents by adding a subclass of the BFO class Process, which we here call Act of Evaluation. Each instance of Act of Evaluation is performed by some evaluator, who evaluates a sentence content and determines a corresponding value for it, where this value is a nominal measurement. Accordingly, every act of evaluation has input some sentence content (e.g. a statement) and has output some nominal measurement (e.g. 'True').

This basic strategy has a number of benefits. First, evaluations of statements are treated as information content entities, rather than as qualities of statements carrying unintuitive ontological commitments. Second, our proposal can admit of disagreements over evaluations of the same statements by relativizing evaluations to acts performed by participants. In addition, different classes of evaluation may be introduced that have, as input, different criteria of evaluation, allowing, for instance, different approaches to truth and probability to co-exist within the same ontology. Finally, since evaluations are independent of the extensions of statements, we can state the evaluations of statements independent of questions about whether they correctly refer to reality.

Future Work

It is our wish that ArgO be re-used for many different projects, and for this reason, that it remain a relatively small, mid-level ontology. In our experience, it is simply easier to re-use a small, well-defined, consistent ontology based on an expertly developed analysis of a domain, rather than a large ontology developed for a particular application. Large ontologies require large resources, expertise, and time to be well-developed, and because these are not readily available, large ontologies often take shortcuts, making representation decisions not on the basis of ontological analysis, but rather to get some particular project-based query strategy to work. Such decisions are allowable in particular extensions that facilitate engineering applications, but if ontologists are to maintain a long-term strategy of data integration, such short cuts should be avoided.

In the future, we plan to provide interpreted extensions of ArgO that will offer further classes, properties, and distinctions relative to a system of logic, and we encourage others to do this as well. The formal features of OWL and OWL reasoners such as Hermit and Fact++ will allow for inferences also about the validity and soundness of arguments, relative to a system of logic. This remains an interesting project for future research to pursue.

Acknowledgments

For their valuable comments, we would like to thank Barry Smith, Paul Poenicke, as well as reviewers from the RICH Semantic and Direct Representation for Digital

Collections Work at the Joint Conference on Digital Libraries in Toronto, Canada, where we presented an early draft of this paper on June 22, 2017.

References

Arcangeli, Margherita (2016), “Against Cognitivism about Suppositions”, *Philosophia* 42 (3): 607-624.

Arp, Robert, Smith, Barry, and Spear Andrew D. (2015), *Building Ontologies with Basic Formal Ontology*, MA: MIT Press.

Baronett, Stan (2012), *Logic*, 2nd edition Oxford University Press.

Breuker, Joost, Rinke Hoekstra, and Alexander Boer (eds) (2006), *OWL Ontology of Basic Legal Concepts (LKIF-Core)*, Estrella, University of Amsterdam. Published June 2006. Available at: www.estrellaproject.org/doc/D1.4-OWL-Ontology-of-Basic-Legal-Concepts.pdf.

Ceusters, Werner and Smith, Barry (2015), “Aboutness: Towards Foundations for the Information Artifact Ontology”, in *Proceedings of the Sixth International Conference on Biomedical Ontology (ICBO) CEUR* Vol. 1515: 1-5.

Clark, Tim, Paolo N. Ciccarese, and Carole A. Goble (2014), “Micropublications: a semantic model for claims, evidence, arguments and annotations in biomedical communications” *Journal of Biomedical Semantics*, 5:28. DOI: 10.1186/2041-1480-5-28.

Clark-Younger, H. (2015), “Imperatives and the More Generalized Tarski Thesis”, *Thought: A Journal of Philosophy* 3 (4): 314-320.

Copi, Irving M., Carl Cohen, and Kenneth McMahon (2011), *Introduction to Logic.* Fourteenth Edition. Taylor & Francis.

Descartes, René, *The Philosophical Writings of Descartes Volume 2*. Trans. John Cottingham, Robert Stoothoff, and Dugald Murdoch. Cambridge University Press; reprint edition, 1985.

Dumontier, Michel et al. (2014), “The Semantiscience Integrated Ontology (SIO) for Biomedical Knowledge Discovery”, *Journal of Biomedical Semantics* 6 March 2014, 5:14.

van Eemeren, Frans H. and Grootendorst, Rob (1992), *Argumentation, Communication, and Fallacies: APragma-Dialectical Perspective*. Routledge.

Feng, Vanessa Wei and Hirst, Graeme (2011), “Classifying Arguments by Scheme”, *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pp. 987–996, Portland, Oregon, June 19-24.

Gruninger, Michael, et. al. (2014), “A Sideways Look at Upper-level Ontologies”, *FOIS*. 2014: 9-22.

Hovda, Paul (2009), “What is Classical Mereology?” *Journal of Philosophical Logic* 38: 55-82.

Jørgensen, Jørgen (1938), “Imperatives and logic”, *Erkenntnis* 7: 288-98.

Kennedy, Duncan (1991) “A Semiotics of Legal Arguments”, *Syracuse Law Review, Vol. II*, p. 75.

Mastop, Rosja (2011), “Imperatives as semantic primitives”, *Linguistics and Philosophy,* 34 (4): 305-340.

McCullough, Laurence B., Coverdale, John H., and Chervenak, Frank A. (2004), "Argument-Based Medical Ethics: A Formal Tool for Critically Appraising the Normative Medical Ethics Literature", *American Journal of Obstetrics and Gynecology* 191, no. 4 (2004/10/01/2004): 1097-102.

McGrath, Matthew (2015) “Propositions”, *Stanford Encyclopedia of Philosophy*, Accessed August 27, 2017.

McCune, William, "Prover9 and Mace4", http://www.cs.unm.edu/~mccune/Prover9, 2005-2010.

Menzel, Christopher (2011), “Knowledge Representation, the World Wide Web, and the Evolution of Logic”, *Synthese* 182: 269-295.

Office of the Press Secretary of the United States (2002), Joint Resolution to Authorize the Use of United States Armed Forces Against Iraq. Released October 2, 2002. Accessed on May, 5, 2017 at: https://georgewbush-whitehouse.archives.gov/news/releases/2002/10/print/20021002-2.html.

Parsons, Josh (2013), “Command and Consequence”, *Philosophical Studies* 164 (1): 69-92.

Peroni, Silvio and Vitali, Fabio (2011), *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.

Rahwan, Iyad and Banihashemi, Bita (2008), “Arguments in OWL: A Progress Report” *Computational Models of Argument: Proceedings of COMMA 2008*, Toulouse, France, May 28-30, 2008.

Rapanta, Chrysi, and Douglas Walton (2016), "The Use of Argument Maps as an Assessment Tool in Higher Education." *International Journal of Educational Research* 79 (2016/01/01/ 2016): 211-21.

Simons, Peter (1987), *Parts: A Study in Ontology*. Oxford University Press.

Smith, Barry (2012) “Classifying Processes: An Essay in Applied Ontology.” *Ratio* 25 (4): 463-488.

Stalnaker, Robert (1984), *Inquiry*. MIT Press.

Swift, Jonathan (1729), *A Modest Proposal for Preventing the Children of Poor People from Being a Burthen to their Parents or Country, and for Making Them Beneficial to the Publick*, available at http://www.gutenberg.org/files/1080/1080-h/1080-h.htm.

Tuomela, Raimo (2000), “Belief versus Acceptance.” *Philosophical Explorations* 3 (2): 122-137.

Toulmin, Stephen (2003) *The Uses of Argument*, 2nd edition (1st edition 1958) Cambridge University Press.

Vranas, Peter B. M. (2008). “New foundations for imperative logic I: Logical connectives, consistency, and quantifiers”, *Noûs*, 42 (4): 529-572.

Walton, Douglas (1996), *Argumentation Schemes for Presumptive Reasoning (Studies in Argumentation Theory)*, Routledge.

Walton, Douglas, Reed, Christopher, and Macagno, Fabrizio (2008), *Argumentation Schemes*, Cambridge University Press.