Logic-Based Modeling of Cognition*

Selmer Bringsjord • Michael Giancola • Naveen Sundar Govindarajulu

ver 1210211555NY

Abstract

After a brief orientation to logic-based (computational) cognitive modeling (LCCM), we carry out the necessary preliminaries (e.g., we explain what a logic is, and what it is for one to "capture" some human cognition). We introduce three "microworlds" or domains that all readers should be comfortably familiar with (natural numbers and arithmetic; everyday vehicles, and residential schools, e.g. colleges and universities), in order to facilitate exposition in the chapter. We then introduce and briefly characterize the vast and ever-expanding universe $\mathscr{U}$ of formal logics, with an emphasis on three categories therein: deductive logics having no provision for direclty modeling cognitive states, *non*-deductive logics suitable for modeling rational belief through time without machinery to directly model cognitive states such as *believes* and *knows*, and finally non-deductive logics that enable the kind of direct modeling of cognitive states absent from the first two types of logic. We then focus specifically on two important aspects of human-level cognition to be modeled in logic-based fashion: the processing of *quantification*, and *defeasible* (or *nonmonotonic*) reasoning. To wrap things up, we briefly evaluate logic-based cognitive modeling, offer in that connection some comparisons with other approaches to cognitive modeling, and consider the future of LCCM. The chapter presupposes nothing more than high-school mathematics of standard sort on the part of the reader.

*We are deeply grateful for trenchant feedback from four anonymous reviewers, and from Editor R. Sun.

Contents

List of Figures

List of Tables

1 Introduction, our Focus, and Plan of the Chapter

This chapter explains the approach to reaching the overarching scientific goal of capturing the cognition of persons in computational formal logic.[1] The cognition in question must be coherent, and the person must be at least human-level (i.e., at least a human person).[2]) In what can reasonably be regarded to be the prequel to the present chapter, (Bringsjord 2008), a definition of personhood was provided; for economy here, we don't recapitulate that definition. We shall simply take *faute de mieux* a person to be a thing that, through time, in an ongoing cycle, perceives, cognizes, and acts (Sun & Bringsjord 2009).[3] The cognizing, if the overarching goal is to be reached, must be comprised, all and only, of that which can be done in and with computational formal logics. Since it has been proved that Turing-level computation is capturable by elementary reasoning over elementary formulae in an elementary formal logic,[4] any cognition that can be modeled by standard computation is within the reach of the methodology we describe herein.[5] However, it is importat to note a concession that stands at the heart of the logicist research program we explain herein: viz. that even if this program completely succeeds, the challenge to cognitive science of specifying how it is that logic-based cognition emerges from, and interacts with, sub-logic-based processing in such things as neural networks will remain. Theoretically, in the artificial and alien case, where the underlying physical substrate may not be neural in nature, this challenge can be avoided, but certainly in the human case, as explained long ago by Sun (2001), it cannot: we are ultimately brain-based cognizers, and have a "duality of mind" that spans from the sub-symbolic/neural to the symbolic/abstract.

The remainder of the chapter unfolds straightforwardly as follows. We next (§2) quickly deal with some needed preliminaries, including an answer to a question that surely must be dealt with at the outset: What is a logic? @@@

We next (§**??**) use a simple example to anchor the concept of *non*-defeasible/monotonic reasoning — a kind of reasoning that we build upon to explain defeasible/nonmonotonic reasoning.

[1]As some readers may know, there is such a thing as *in*formal logic; but our work ignores this field, entirely. Whatever virtues informal logic may have, because it cannot be used to compute (which is true in turn simply because informal language, the basis for informal logic, cannot be a basis for computing, which by definition is formal), it is of no use to us. A introduction to and overview of informal logic, which confirms its informal linguistic basis, is provided in (Groarke 1996/2017).

[2]It is of course entirely possible that our universe contains persons who aren't humans; this possibility, as the reader will no doubt well know, is a prominent driver of science-fiction and fantasy literature. In addition, many religions of course claim that there are non-human persons. (In the case of Christianity, e.g. The Athanasian Creed asserts that God is a person.) Even if all such religious claims are false, things clearly could have been such that some of them were true, so the concept of personhood outside of *H. sapiens* is perfectly coherent. In fact, the field of AI, which is intimately bound up with at least computational cognitive science and computational psychology, is a testament to this coherence, since, in the view of many, AI is devoted to building artificial persons (a goal e.g. explicitly set by Charniak & McDermott 1985); see (Bringsjord & Govindarajulu 2018) for a fuller discussion. Finally, it is very hard to deny that we will modify our own brains in ways that yield "brains" far outside what physically supports the cognition of *H. sapiens*; see in this regard (Bringsjord 2014).

[3]Cf. the similar cycle given in (Pollock 1995).

[4]There are multiple proofs, in multiple routes. A direct one is a proof that the operation of a Turing machine can be captured by deduction in first-order logic; e.g. see (Boolos, Burgess & Jeffrey 2003). An indirect route is had by way of taking note of the fact that even garden-variety logic-programming languages, e.g. Prolog, are Turing-complete.

[5]One of the advantages of capturing cognition in formal logic is that it is the primary way to understand computation *beyond* the level of standard Turing machines, something that, interestingly enough, is exactly what Turing himself explored in this dissertation under Alonzo Church, a peerless introduction to which, for those not well-versed in formal logic, is provided by Feferman (1995). For a logic-based introduction to computation beyond what a Turing machine can muster, we recommend (Davis, Sigal & Weyuker 1994).

Then we briefly ensure that the reader has a high-altitude view of the vast universe of all formal logics, which forms the necessary context for the present chapter, as well as its precursor (§3). In the section that follows (§6), we give the basic idea behind defeasible logic (specifically, of an argument-centric, and hence more cognitively plausible, sort), in connection with a meteorological example. We then, in §6.1, provide an overview of the argument-adjudication system we use to enable, in precise, computational terms, the modeling and simulation of human defeasible reasoning. (This system is based on the *inductive cognitive event calculus*, $\mathcal{IDCEC}$, which is encapsulated in the appendix.) Next, in 6.2, we return to the tornado example, and delve into it in more detail by showing how our argument-adjudication system can handle it. We then introduce the reader to the suppression task (§6.3), and proceed to show how Stenning & van Lambalgen (2008) model and simulate it in their *purely extensional* nonmonotonic logic (§6.3.1). Next we explain that a proper understanding of the suppression task calls for *intensional logic* (§6.3.3), but rather than spend any time debating whether our interpretation of the original suppression task is correct, we instead simply introduce a new version of the suppression task that is *explicitly* intensional, and model and simulate tackling of this task (§**??**). The chapter then concludes with some final thoughts.

2 Preliminaries

For the goal of capturing the cognition of persons in computational formal logic to be informative to the reader, it is naturally necessary to engage in preliminary exposition to explain what a logic is, what specifically a *computational* logic is, what cognition is herein taken to be, and finally what capturing cognition via formal logic amounts to.

2.1 Anchoring Domains for Exposition: Numbers; Vehicles; Universities

In order to facilitate exposition, it will be convenient to rely upon straightforward reference to three different domains of discourse, each of which will be familiar to the reader: viz., the natural numbers and elementary arithmetic with them, which all readers were presumably learned about when very young; everyday vehicles (cars, trucks, etc.); and residential schools, such as colleges and universities.

The natural numbers, denoted by '$\mathbb{N}$,' is simply the set

$$\{0, 1, 2, 3, \ldots\},$$

and by 'elementary arithmetic' we simply refer to addition, subtraction, multiplication, and so on. Readers are assumed to know for instance that zero $\in \mathbb{N}$ multiplied by $27 \in \mathbb{N}$ is zero.

As to the domain of vehicles, readers are assumed to understand the things represented in Figure 1, which the reader should now view, taking care to read its caption. We invoke three types of familiar vehicles that can be either of two colors (black or grey). Each vehicle is either located at a particular position in the grid shown, or is outside and adjacent to it. The grid is oriented to the four familiar directions, of North, East, South, and West.

What about the domain of residential schools? We assume nothing more than a generic conception, according to which such institutions, for instance colleges and universities, include agents that fall into the categories of student, teacher, and staff; and include as well the standard buildings are in place in accordance with the standard protocols. For example, residential universities have dormitories, classrooms, and libraries. We specifically assume that all readers have common

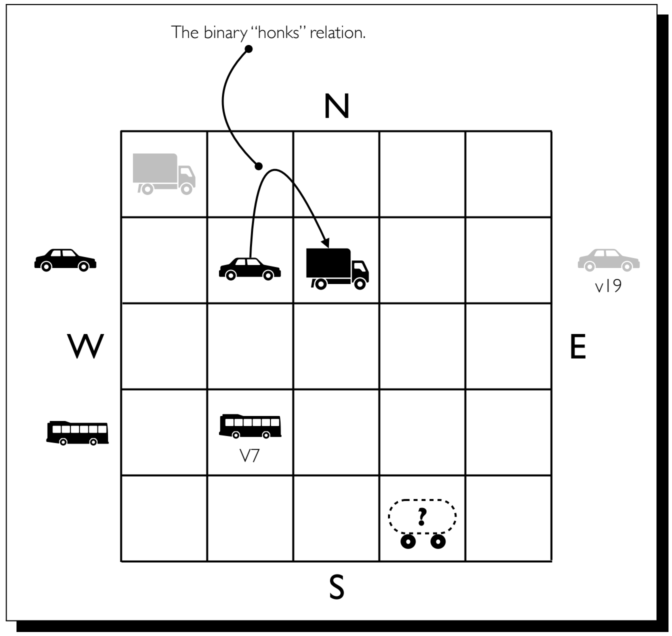

Figure 1: The Vehicular Domain. *The three types of vehicle are shown: cars, box trucks, and buses. The reader will note that there is also a diagram that indicates the existence (and perhaps location) of a "mystery" vehicle. Each vehicle is either colored black or grey (there is one grey vehicle in the grid (a box truck), and one such vehicle outside the grid (a car). Notice that vehicles can be denoted by names (or constants). Finally, we have the standard four directions.*

knowledge of the invariants seen in such schools, for instance that they commonly have classes in session, during which time students in the relevant class perceive the teacher, hold beliefs about this instructor, and so on.

2.2 What is a formal logic?

We rest content with providing two necessary conditions for something's being a formal logic.[6]

The first of these two necessary conditions is that we can't have a formal logic unless we have a formal specification of what counts as a *formula*, and in the vast majority of cases this specification will be achieved by way of the definition of a formal language **L** composed minimally of an alphabet A and a grammar G.[7] Without this, one simply doesn't have a formal logic; with this, one has the ability to determine whether or not a given formal logic is expressive enough to represent some declarative information. Importantly, it is often the case that some natural-language content to be expressed as a formula in some (formal) logic $\mathscr{L}$ cannot be intuitively and quickly expressed

[6] We literally have no idea how to formally define an *in*formal logic, and suspect that at any rate doing so in anything like a scientific manner is likely simply conceptually impossible. On the other hand, please note that everything we say in the present section is perfectly consistent with conceptions of a formal *inductive* logic, which is distinguished by reasosning that is non-deductive. For a nice, non-technical introduction to inductive logic see (Johnson 2016). For a sustained rigorous introduction to formal inductive logic of the model-theoretic variety, which subsumes probability theory, see (Paris & Vencovská 2015).

[7] Please note that this pair $\langle A, G\rangle$ needn't be purely symbolic/linguistic. The pair might e.g. include purely visual or "homomorphic" elements. See the logic Vivid as a robust, specified example (Arkoudas & Bringsjord 2009).

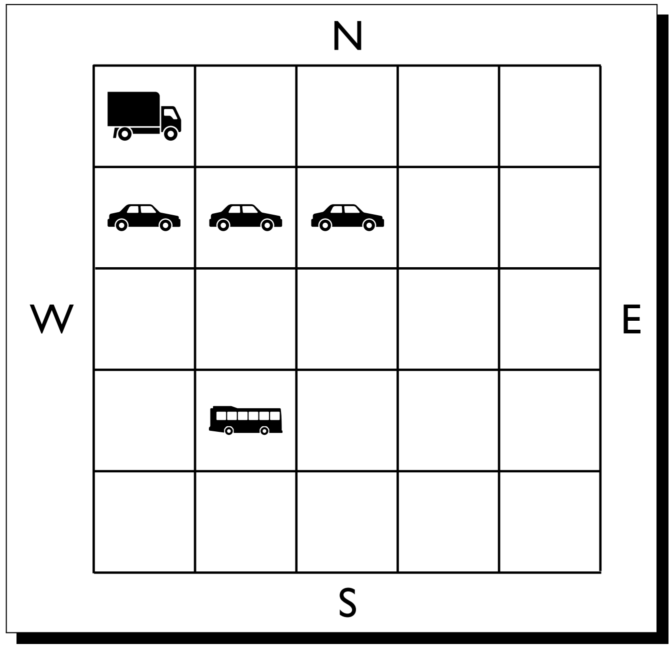

Figure 2: Vehicluar Scenario #1. *Here C is a unary relation symbol used to express being a car, etc.'*

correctly by a simple formula in the formal language for $\mathscr{L}$, so that the formula can then be used (for example by a computer program) instead of natural language. For example, the (declarative) natural-language sentence (1_n)"Every car is north of some bus that's south of every truck," which is true in Vehicular Scenario #1 shown in Figure 2, can't be represented in any dialect of the propososoitional calculus = $\mathscr{L}_{pc}$, since no object variables are permitted in this logic. But this natural-language sentence is easily expressed in first-order logic = $\mathscr{L}_1$ by the following formula in its formal language:

$$(1_l)\ \forall x[C(x) \rightarrow \exists y(By \wedge N(x,y) \wedge \forall z(T(z) \rightarrow S(y,z)))].$$

Here x and y are object variables, C is a unary relation symbol used to express being a car, B denotes the property of being a bus, and N is a binary relation symbol that represents the property of being north-of. In addition, we have in $\mathscr{L}_1$ the two standard and ubiquitous quantifiers: Where φ is any object variable, $exists\varphi$ says that there exists an object φ, and $\forall\varphi$ says that for every φ. We refrain from giving the formal grammar of $\mathscr{L}_1$, since the level of detail required for doing so is incompatble with the fact that the present chapter is first and foremost an overview of cognitive modeling via logic, not a technical overview of logics themselves. The reader should take care to verify for himself, now, that the formula (1) does in fact hold of the scenario shown in Figure 2.

Note that without having on hand a precise definition of the formal language $\mathcal{L}$ that is the basis for a given formal logic $\mathscr{L}$, there is simply no way to rigorously judge the expressive power of $\mathscr{L}$, and hence no way to judge whether $\mathscr{L}$ (or for that matter some theory in cognitive science that purports to subsume $\mathscr{L}$) is up to the task of modeling, say, some proposition that some humans apparently understand and make use of.

Now, what is the second necessary condition for $\mathscr{L}$'s being a formal logic, over and above the one saying that $\mathscr{L}$ must include some formal language? This second condition is disjunctive (*inclusive* disjunction used: i.e. either disjunct, or both, must hold) in nature, and can be stated informally thus:

- Any *bona fide* logic must have a fully specified system for checkable inference (chains of which are expressed as proofs or arguments), and/or[8] a fully specified system for checkable assignments of semantic values (e.g., TRUE, FALSE, PROBABLE, PROBABLE AT VALUE (some number) k, INDETERMINATE, etc.) to formulae and sets thereof.

Note that we have already made some use of truth and falsity in connection with first-order logic $= \mathscr{L}_1$, since we have said that the formula (1) in this formal logic is true on Vehicular Scenario #1. Note as well that the semantic categories for a given logic can often exceed the standard values of TRUE and FALSE. To make this concrete and better understand, take a look back at Figure 1 now, and consider the natural-language statement $(2)_n$ "Car v19 is east of every truck." If we express this declarative sentence in $\mathscr{L}_1$ as a formula, we have

(2_l) $\forall x[T(x) \to (E(vr19, x) \wedge C(v19))]$,

and what is the semantic value of this formula on the scenario shown in Figure 1? There is no way to know, because while we know that vehicle v19 is a car, it's not in the grid. We thus can add the semantic value INDETERMINATE to what we have available for modeling; and this is the value of (2_l) on the scenario in question. For those in favor of couching formal theories of meaning for natural language (and of cognition relating to the use of natural language) in terms of proof, (2_l) is indeterminate specifically because it can't be proved from the information given in Figure 1, nor can the negation of this formula be proved from this information. However, notice something interesting about the scenario in this figure: Suppose that we knew what kind of vehicle the mystery vehicle in Figure 1 is; specifically, suppose that that vehicle is a bus. In addition, assume that vehicle v19 is located in some square in not the eastmost column, but the column one column to the west of the eastmost column. Given this additional information, we can easily prove (2_l) from the information we have under these suppositions. For some, for instance **?** (and such thinkers are aligned with the purely inferential understanding of what a logic is within the disjunction given in our second necessary condition), the meaning of the natural-language sentence (2_n) for an agent consists in its being inferable from what is known by that agent. We spare the reader the formal chain of inference in $\mathscr{L}_1$ that constitutes a formal proof of (2_l). Such a proof is by cases, clearly. The proof starts with noting that v19 will be in one of four different locations in the column in question, and then proceeds to consider each of the only two trucks in the scenaio; both of them are west of each of these four locations.

2.3 What is a *computational* formal logic?

Since the topic at hand is cognitive modeling via logic, and cognitive modeling is by definition a computational affair, we need now to understand what a computational logic is. All readers will have come to this chapter with at least an intuitive conception of what a logic is (and now, given the foregoing, they will have deeper understanding), but no doubt some will be quite puzzled by our reference to a "computational" logic. This is easy to address: a computational logic is just a logic

[8] Again, this is inclusive disjunction. The two disjuncts represent the two major, sometimes-competing schools in logic, namely proof-theory and model-theory. Proponents of the first school avoid semantic notions. The reason why the disjunction is inclusive is that some logicians would desire to see *both* disjuncts satisfied. In particular, model theorists emphasize semantics, but take proofs to be witnesses of validity of formulas.

that can be used to compute, where computing is cast as inference of some sort. Since computing in any form can be conceived of as a process taking inputs to outputs by way of some function that is mechanized in some manner, in the logicist approach to cognition, the mechanization consists in taking inputs to outputs by way of reasoning from these inputs (and perhaps other available content). This is as a matter of fact exactly how logicist programming languages, for instance Prolog, work. Often the inputs are queries, and the outputs are answers, sometimes accompanied by justificatory proofs or arguments. When Newell and Simon presented their system LogicTheorist at the dawn of AI in 1956, at Dartmouth College, this is exactly what the system did. The logic in question was *propositional calculus*, the inputs to LogicTheorist were queries as to whether or not certain strings were theorems in this logic, and the outputs were answers with associated proofs. For more details, see the seminal paper of Simon's (1956), for a recent overview of the history to which we refer, in the context of contemporary AI, see (Russell & Norvig 2020, Bringsjord & Govindarajulu 2018).

2.4 What is cognition?

We come now to the next preliminary to be addressed, which is to answer: What is cognition? And what is it to cognize? Put another way, this pair of questions distill to this question: What is the target for logicist cognitive modeling?

Fortunately we can avail ourselves of an efficient answer under the burden of our space constraints: We shall simply take cognition to consist in instantiation of the familiar cognitive verbs: *communicating*, *deciding*, *reasoning*, *believing*, *knowing*, *fearing*, *perceiving*, and so on, on through all the so-called *propositional attitudes* (Nelson 2015). In other, shorter words, whatever cognitive verb is targeted in human-level cognitive psychology, for instance in any major, longstanding textbook for this sub-field of cognitive science (e.g. see Ashcraft & Radvansky 2013), must, if the overall goal logicist modeling is to be achieved, be captured by what can be done in and with computational formal logics.

@@@

2.5 What is it to capture cogntion in formal logic?

But how do we know when logicist cognitive modeling of human-level cognition succeeds? We can say that such modeling succeeds when selected aspects of human-level cognition are *captured*. But what is it to "capture" part or all of human-level cognition in computational formal logic? After all, isn't 'capture' operating as a metaphor here, and an imprecise one at that? Actually, the concept of formal logic managing to capture some phenomena is *not* a metaphor; it's a technical concept, one easily and crucially conveyed here without going into its ins and outs: Some phenomena P is captured by some formal content C_P, expressed a (formal) logic $\mathscr{L}$ if and only if all the elements p in P are such that from C_P one can provably infer in $\mathscr{L}$ the formal counterpart C_p that expresses p. @@@ To illustrate with a simple example, suppose that the phenomena in question is the appearance of English declarative sentences (in response, say, to some queries) about elementary arithmetic. So an element here could be (s_1) "Twelve is greater than two plus two," or (s_2)"Seven times one is seven," or (s_3) "Any (natural) number times 1 is that number," and so on. It is known that the particular, familiar formal logic *first-order logic*, $\mathscr{L}_1$, can express such sentences rather easily. For instance, if $\dot{n}$ is a constant in this logic's language to denote the number n, and $\times$ is a function symbol in this language for multiplication, the latter two sentences are expressed in $\mathscr{L}_1$ by two formulae ϕ_{s_2} and ϕ_{s_3}, respectively, like this:

- $\phi_{s_2} := \dot{7} \times \dot{1} = \dot{7}$
- $\phi_{s_3} := \forall n(\dot{n} \times \dot{1} = \dot{n})$

And now, what of capturing? There is a rather famous body of content, composed of a set of formulae in first-order logic, known as *Peano Arithmetic*, or just **PA**; it captures all of elementary arithmetic.[9] Given what we said above, this means that every relevant sentence s about elementary arithmetic not only can be expressed by some corresponding formula ϕ_s in $\mathscr{L}_1$, but that every such sentence that's true can be proved from **PA**. This is in fact true of ϕ_{s_2} and ϕ_{s_3}. Elementary arithmetic has been captured,[10] as has content in other fields outside mathematics.[11] For now, this will do in order to provide the reader with some understanding of our ambition to capture the defeasible reasoning of human persons. More specifically and concretely, for this ambition to be reached, we must show that there is some logic such that, whenever such a person defeasibly reasons to some declarative sentence s, there is some content in that logic from which a formula ϕ_s expressing s can be defeasibly inferred. In the present chapter, we endeavor to show this only in connection with a reasoning task that has been much studied in cognitive science: namely, the fascinating *suppression task*, introduced by Byrne (1989). We shall take advantage of the fact that others who worked hard to model and computationally simulate human reasoning and logic have specifically tried their hand at the suppression task, which appears to clearly call specifically for defeasible reasoning, not just purely deductive reasoning. But before we get to this task and its treatment, we shall need to carry out some necessary preparatory work.

3 The Vast Universe of Logics & Our Location Therein

Please at this point take a look at Figure 3. This picture is intended to situate the present chapter within the context of the vast universe of logics that are available for modeling of cognition. We will not concern ourselves with any logics that permit expressions that are themselves infinitely long; therefore we are working outside the "Infinitary" oval on the left of the all-encompassing oval shown in Figure 3. Hence we are working within the "Finitary" oval shown. Notice that within that oval there are shown two sub-categories: "Intensional" versus "Extensional." Roughly speaking, the first of these categories is marked by logics that are tailored to represent such cognitive verbs as we cited above: for example, *believing*, *knowing*, *intending*, and also verbs that are "emotion-laden," such as *hoping*, *desiring*, *fearing'*, and so on. The logics that are up to the task of representing content that is infused with such — to use the phrase that has been popular in philosophy — *propositional attitudes* (Nelson 2015) must be sensitive to a key fact arising from the cognition involved: viz., that when an agent has such an attitude toward a proposition, it's not possible to compute compositionally what the semantic value of the overall attitude is from such values assigned to the target propositions. A simple example illustrates this phenomenon:

[9] Nice coverage is provided in (Ebbinghaus, Flum & Thomas 1994).

[10] For a technical presentation of the concept of capture, including the arithmetic case we have just drawn from, see (Smith 2013).

[11] E.g., formal logic has successfully captured parts of physics; specifically, e.g., classical mechanics (McKinsey, Sugar & Suppes 1953) and — much more recently — special relativity (Andréka, Madarász, Németi & Székely 2011).

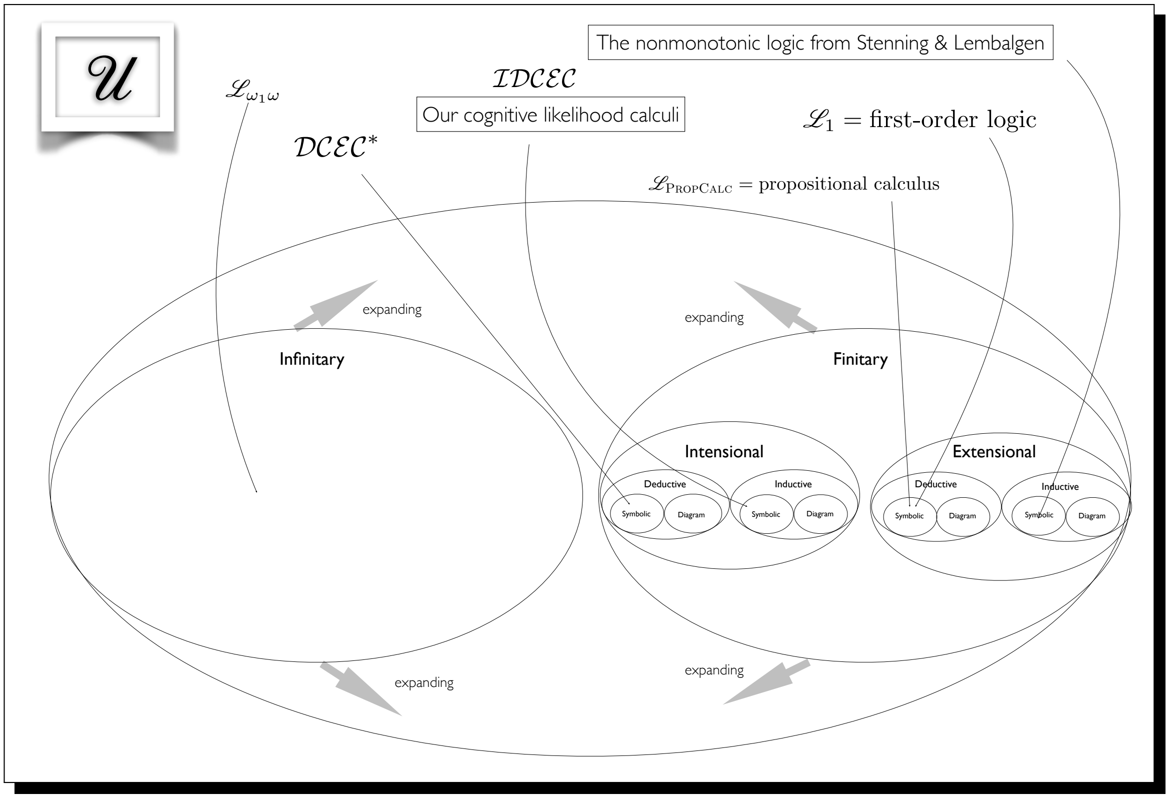

Figure 3: The Vast Universe of Logics. *The universe of formal logics can be divided into those that allow expressions which are infinitely long, and those that don't. Among those that don't, the propositional calculus and first-order logic have been much employed in CogSci and AI. The boxed logics are the ones key to the upcoming analysis and discussion. Note that in the previous section we made crucial use of* $\mathscr{L}_1$.

Consider the proposition p_1 that our Umberto believes that our Terry (recall above) believes that Umberto is brilliant. Now suppose that p_2 it's true that Umberto is brilliant. Does it follow from the fact that p_1 is true that p_2 is as well? Clearly not. Umberto may well believe that Terry thinks that he (Umberto) is quite dim. In stark contrast, every logic in the category "Extensional" is such that the semantic values of molecular propositions built on top of "atomic" propositions are fully determined by the semantic values of the atomic propositions. In the very earliest grades of the study of mathematics, this determination is taught to students; a case in point is shown in Figure **??**, which shows a question on the 2019 Grade 4 math assessment exam used across New York State in the U.S. What the cognizer is hopefully going to do when confronted with this problem is realize, in line specifically with the logic known as the propositional calculus (shown as '$\mathscr{L}_{\text{PropCalc}}$ in Figure 3), that both atomic propositions "the factors of 5×3 is odd" and "both 5 and 3 are odd numbers" are true, which in turn makes the statement **D** true.

We shall be concerned herein with the particular cognition that gave/gives us mathematics. (Without a stage-setting section like this, your readers may proceed under the impression that things like the propositional calculus and first-order logic and old intro-textbook modal logic ARE logic.) We will make sure to mention most of the main categories of logics in U that the referees mentioned. (E.g., there was mention of fuzzy logic. Fuzzy logic as an engineering approach is not logic; but fuzzy logic as logics is a sub-collection of logics within multi-valued logics.) (Without a stage-setting section like this, your readers may proceed under the impression that things like the propositional calculus and first-order logic and old intro-textbook modal logic ARE logic.)

4 Our Focus Herein: Quantification; Defeasible Reasoning

5 Quantification and Cognition

From the perspective of those searching to capture human-level cognition via logic, there can be little doubt that quantification a key, indeed perhaps *the* key, factor upon which to focus. We have of course already seen above some quantification at work, in connection with both our vehicular domain and elementary arithmetic. Hence the reader is now well aware of the fact that 'quantification' in the sense of that word operative here has nothing to do with conventional construals of such phrases as "quantitative reasoning." Such phrases usually refer to quantities or magnitudes in some numerical sense. Instead, in formal logic, and in logicist cognitive modeling, quantification refers specifically to the use of of quantifiers such as 'all,' 'some,' 'many,' 'a few,' 'most,' and so on. In particular, we have placed and will continue to place emphasis in this chapter upon the two quantifiers that are used most in at least deductive formal logics, the two quantifiers that (accompanied by some additional machinery) form the basis for most of the formal sciences, including mathematics and theoretical computer science. These two quantifiers are exactly the ones we have already seen in action above: $\forall$ (read as 'for every' or 'for all') and $\exists$ (read as 'there is at least one' or 'there exists at least one'). When these two quantifiers are employed, almost invariably they are immediately followed by an object variable, so that the key constructions are

$$\forall\varphi\ldots$$

and

$$\exists\varphi\ldots,$$

where, as above φ is some object variable. These constructions, as the reader now knows, are read, respectively, as "For every thing $\varphi\ldots$" and "There exists at least one thing φ such that ....

In our experience, not only students, but also even accomplished researchers outside the formal sciences, are often initially incredulous that something so unassuming as these constructions could be at the very heart of the formal sciences, and at the very heart of cognition. We proceed now to explain why such incredulity is mistaken.

5.1 Quantification in the Study of the Mind

As a matter of empirical fact, a focus on quantification in the study of the mind, at least when such study targets at least human-level cognition, has long been established, and is still being very actively pursued. For example, since Aristotle, there has been a sustained attempt to discover and set out a logic-based theory that could account for the cognition of those who, by the production of theorems and the proofs that confirm them (Glymour 1992), make crucial and deep use of quantification. The first substantive exemplar of such cognition known to us now in the 21st century remains the remarkable Euclid, some of whose core results in geometry are still taught in all technologized societies the world over, and in fact we suspect that most readers will at least vaguely remember that they were asked to learn some of Euclid's axioms, and to prove at least simple theorems from them. If this request met with success, the cognition involved included understanding of quantification (over such things as points and lines, reducible therefore to quantification over real numbers).

What about contemporary study of human-level-or-above cognition by way of quantification?

TODO
Kemp (2009)
(Partee 2013)
The Polish guy

5.2 Quantification in Higher-Order Logic

One of the interesting, apparently undeniable, and powerful aspects of human-level cognition is that it centrally involves not only use of relations such as 'is a bus' or 'is a car' (which are of course represented, respectively, by the relation symbols B and C in our vehicular setup), but also relations that can be applied to relations. A body of cognitive-science work indicates this capacity to be present in, and indeed routinely used by, humans (Hummel 2010, Hummel & Holyoak 2003, Markman & Gentner 2001). Using resources of LCCM, specifically a logic from $\mathscr{U}$ well-known to practicioners of logic-based modeling, this apsect of human-level cognition is quite easy to express in rigorous terms. More specifically, LCCM has available to it higher-order logics. First-order logic $= \mathscr{L}_1$, as we have seen, permits only *object* variables, so named because they refer to objects, not relations (or properties or attributes); this logic doesn't have *relation* variables. To make this concrete, consider Vehicular Scenario #2 for a minute; this scenario is given in Figure 4. Do you agree upon studying the situation that in this scenario the immediately following declarative sentence holds?

($\ddagger_n$) There is at least one relation that holds of every vehicle north of every bus.

We are confident that you, the reader, apprehend the truth of ($\ddagger_n$) in Vehicular Scenario #2. We say this on the strength of the cognitive-science work we have cited above, in the present section. But this natural-language sentence cannot be represented in $\mathscr{L}_))$, since this logic has no provision for expressing "There is a relation that" in this sentence. Second-order logic $= \mathscr{L}_2$ comes to the rescue, because it includes provision for quantification over relation (property) variables. To thus model what the reader apprehends via LCCM, we need to have the formula in second-order logic that expresses ($\ddagger_n$) — and here it is:

($\ddagger_l$) $\forall x[(\forall y(B(y) \rightarrow N(x,y))) \rightarrow \exists X X(x)]$

Notice that, following longstanding tradition in formal logic, we use majuscule Roman letters X, Y, Z etc. for variables that can be instantiated with particular relations. Another look at Figure 4 and the vehicular scenario it holds will reveal to the reader that there are particular relations/properties that can serve as particular instances of X in ($\ddagger_l$). For example, one such relation/property is the color grey, which is indeed the color of every vehicle north of every bus.

The reader may wonder whether there is a level higher than second-order logic $= \mathscr{L}_2$. There is. The next step up, perhaps unsurprisingly, is *third*-order logic $= \mathscr{L}_3$. We strongly suspect

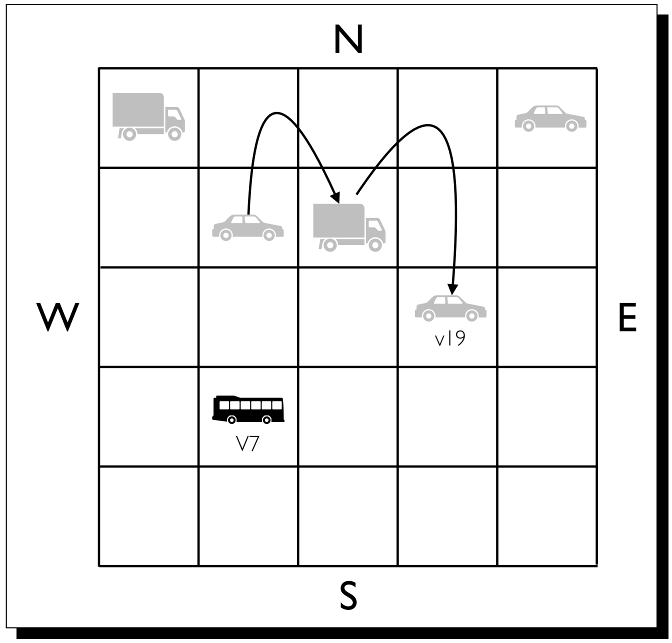

Figure 4: Vehicluar Scenario #2. *Observe that in this scenario there is a relation (property) X which every vehicle north of a bus has. E.g., a witness for such an X could in this scenario be the relation 'Grey.'*

that human-level cognition makes routine use of third-order propositions — though of course we have no idea how such propositions are specifically encoded, in the human case, in the human brains (but see the use made of Clarion for third-order formulae in Bringsjord, Licato & Bringsjord 2016). The distinguishing new feature of $\mathscr{L}_3$ is that it permits, and renders precise, the ascription of relations/properties to relations/properties; this is not permitted in $\mathscr{L}_2$. This feature can be rendered concrete with help from Vehicular Scenario #2, quickly, as follows. First, simply note that grey is a color; hence we can sensisibly write

$$C(G)$$

to represent that fact. Next, to express

($\triangle_n$) There is at least one color property (relation) that holds of every vehicle north of every bus.

($\triangle_l$) $\forall x[(\forall y(B(y) \rightarrow N(x,y))) \rightarrow \exists X(X(x) \land C(X))]$

5.3 Quantification and the Infinite

As is well-known, human-level cognition routinely involves infinite objects, structures, and systems. This is perhaps most clearly seen when such cognition is engaged in the learning and practice of mathematics, and formal logic itself. All readers will for example recall that even basic high-school geometry invokes at its very outset infinite sets and structures. As to such sets, we have $\mathbb{N}$ and $\mathbb{R}$, both introduced above, these being two specimens that every high-school graduate needs to demonstrate considerable understanding of. And as to structures based upon these two infinite sets, all our readers will recall as well that for instance two-dimensional Euclidean geometry is based upon the set of all pairs of real numbers. Within this context, It turns out that cognition associated with even some elementary quantification in $\mathscr{L}_1$ instantly and surprisingly provides an opportunity to zero in on cognition that is compelled to range over infinite scenarios; and an excellent way to acquire deeper understanding of LCCM and its resources is to reflect upon why such scenarios are forced to enter the scene. Notice that so far our vehicular scenarios are been decidedly finite in size.

In order to reveal the quantification in question, consider the following three rather unassuming natural-language sentences our vehicles:[12]

(a_n) No vehicle honks at itself.

(b_n) If x honks at y and y honks at z, then x honks at z.

(c_n) For every vehicle x, there's a vehicle y x honks at.

This trio is quickly represented, respectively, by the following three extremely simple formulae in $\mathscr{L}_1$:

(a_l) $\forall x \neg H(x, x)$

(b_l) $\forall x \forall y \forall z[(H(x, y) \wedge H(y, z)) \rightarrow H(x, z)]$

(c_l) $\forall x \exists y H(x, y)$

Now here is a question: Can a human understand that (a_n)–(b_n), despite their syntactic simplicity, cannot possibly be rendered true by a vehicular scenario that is finite in size? The reader can answer this question, by attempting to build a scenario that does in fact do the trick. A sample try is enlightening. For example, consider the vehicular scenario shown in Figure 5; for the moment, ignore the use made there repeatedly of the ellipsis. The reader should be able to see that the scenario in fact does *not* render (a_l)–(b_l) true, and should be able to see why. In order to construct a vehicular scenario that works, the reader will need to understand that an infinite progression of vehicles will need to be used, with an infinite number of honks. While we refrain from provding details, it is not difficult to see that the cognition that discovers and writes down such an infinite scenario can itself be modeled using the resources of LCCM.

5.4 Quantification as the Heart of the Formal Sciences: Arithmetic and Reverse Mathematics

Axiom 1 $\forall x(0 \neq s(x))$

Axiom 2 $\forall x \forall y(s(x) = s(y) \rightarrow x = y)$

Axiom 3 $\forall x(+(x, 0) = x)$

Axiom 4 $\forall x \forall y(+(x, s(y)) = s(+(x, y)))$

Axiom 5 $\forall x(\times(x, 0) = 0)$

Axiom 6 $\forall x \forall y(\times(x, s(y)) = +(\times(x, y), x))$

Induction Schema Every sentence that is the universal closure of an instance of this schema:

$$[\phi(0) \wedge \forall x(\phi(x) \rightarrow \phi(s(x))] \rightarrow \forall x \phi(x)$$

TODO

[12] We are here guided and inspired by a clever example given by Kleene (1967) (p. 292).

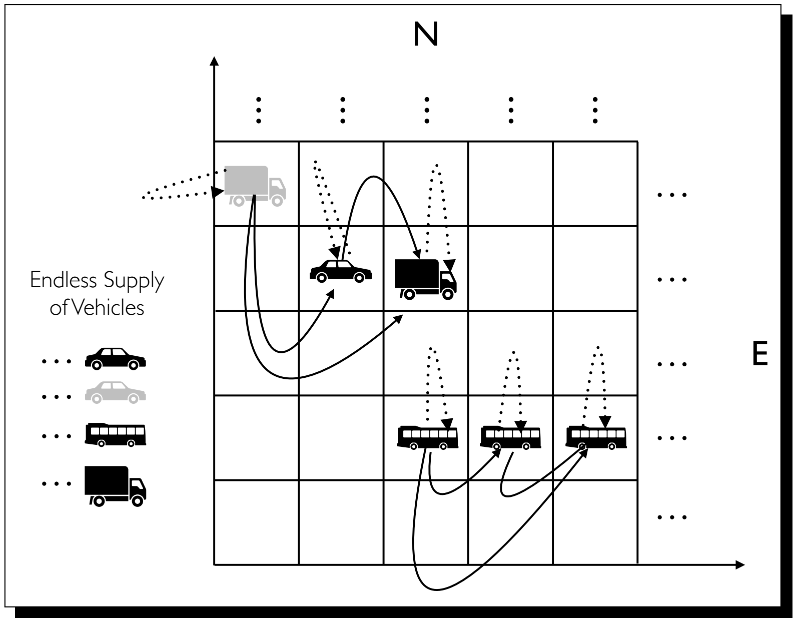

Figure 5: A "Failing" Vehicluar Scenario. *The scenario here fails to model the three rather simple quantified formulas we have specified in the body of the present chapter. The sedulous reader should ascertain why this failure occurs.*

5.5 A Unifying Conception of Quantification and Cognition

Let us now turn back to our vehicular domain; in particular please see Figure 6.

TODO

6 Defeasible/Nonmonotonic Reasoning Logic ~~The Core Idea~~

Deductive reasoning of the sort that we have visited above, in connection with both arithmetic and our vehicular microworld, is *monotonic*. To put this more precisely, to say that if a formula ϕ in some logic can be deduced from some set Φ of formulae (written, recall, $\Phi \vdash_I \phi$, where the subscript I gets assigned to some particular set of inference schemata for precise deductive reasoning), then for any formula $\psi \notin \Phi$, it remains true that $\Phi \cup \{\psi\} \vdash_I \phi$. In other words, when the reasoning in question is deductive in nature, new knowledge never invalidates prior reasoning. More formally, the closure of Φ under standard deduction (i.e., the set of all formulae that can be deduced from Φ via I), denoted by $\Phi_I^{\vdash}$, is guaranteed to be a subset of $(\Phi \cup \Psi)_I^{\vdash}$, for all sets of formulas Ψ. Inductive logics within the universe $\mathscr{U}$ don't work this way, and that's a welcome fact, since much of real life doesn't conform to monotonicity, at least when it comes to the cognition humans; this is easy to see:

Suppose that at present Professor Jones knows that his house is still standing as he sits in it, preparing to teach his class a bit later at his university. If, later in the day, while away from his home and teaching at the university, the Professor learns (along with his students) by notifications pushed to smartphones that a vicious tornado is passing over the town in which his house is located,

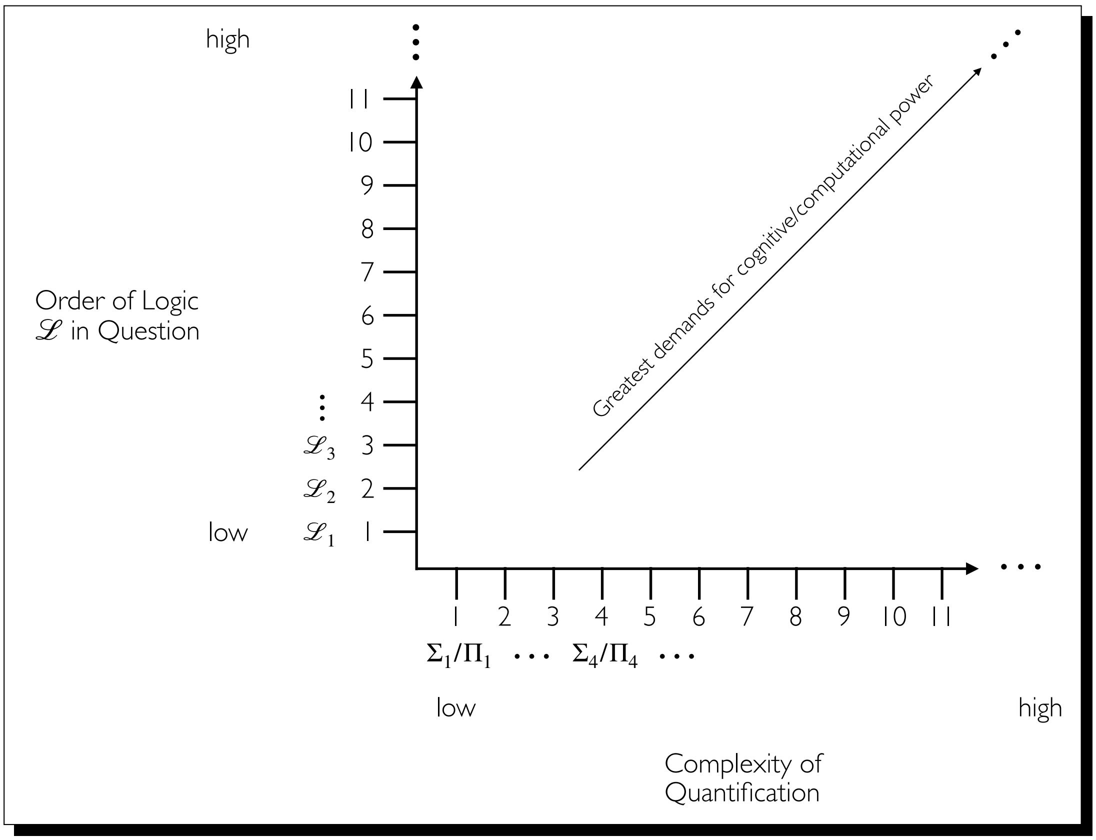

Figure 6: Unifying Conception of Cognition Viewed Through Quantification. *xxxx.*

he has new information that probably leads him to reduce his confidence in the near future as to whether or not his house still stands. Or to take a different example, one much-used in AI (e.g. see the extended treatment in Genesereth & Nilsson 1987), if our Professor Jones knows that Tweety is a bird, he will probably deduce (or at least be tempted to do so) that Tweety can fly, on the strength of a general principle saying that birds can fly. But if Jones learns that Tweety is a penguin, the situation must be revised: that Tweety can fly should now not be among the propositions that Jones believes. Nonmonotonic reasoning is the form of reasoning designed to model, formally, this kind of *defeasible* inference; and some logics within $\mathscr{U}$, all of them non-deductive = inductive in nature, have been devised to specify such reasoning. In the hands of logic-based cognitive modeling, such logics, when computational implemented and run, can then simulate the kind of human/human-level reasoning just seen in the mind of Professor Jones.

There are many different logic-based approaches that have been designed to allow such modeling and simulation, and each approach is associated with a group of logics. Such approaches include: use of default logics (Reiter 1980), circumscription (McCarthy 1980), and the approach we regard to be most cognitively plausible: argument-based defeasible reasoning (e.g. see for an overview, and an examplar of the approach, resp.: Pollock 1992, Prakken & Vreeswijk 2001).[13] An excellent survey, one spanning AI, philosophy, and computational cognitive science, the three fields that work in defeasible/nonmonotonic reasoning spans, is also provided in the Stanford Encyclopedia of Philosophy.[14]) Because argument-based defeasible reasoning seems to us to accord best with what humans actually do as they adjust their knowledge through time (e.g., Professor Jones and his students, if queried on the spot immediately after the notification of the tornado's path as to whether Jones' house still stands, will be able to provide arguments for why their confidence that it does has just declined), this chapter emphasizes the apparent ability of argument-based defeasible reasoning to capture human/human-level defeasible reasoning. It is in fact a rather nice thing about humans and defeasible reasoning that they are often able to explain, and sometimes show, by articulating arguments, why their beliefs have changed through time as new information is known or at least believed, where that new information leads to the defeat of reasoning that they earlier affirmed.

Now, returning to the tornado example, what is the argument that Professor Jones might give to support his belief that his house still stands, while he is in the classroom? There are many possibilities, one respectable one is what can be labeled 'Argument 1,' where the indirect indexical refers of course to Jones:

	(1)	I perceive that my house is still standing.
	(2)	If I perceive ϕ, ϕ holds.
$\therefore$	(3)	My house is still standing.

The second premise is a principle that seems a bit risky, perhaps. No doubt there should be some

[13]From a purely formal perspective, the simplest way to achieve non-monotonicity is to use the so-called *closed world assumption*, according to which, given a set Φ of initially believed declarative statements, what an agent believes after applying the closed world assumption (CWA) to the set is not only what can be deduced from Φ, but also the negation of every formula that *cannot* be deduced. It is easy to verify that it doesn't always hold that $\mathrm{CWA}(\Phi) \subset \mathrm{CWA}(\Phi \cup \Psi)$, for all sets Ψ. I.e., monotonicity doesn't hold. Unfortunately, while this is a rapid route to non-monotonicity, CWA isn't cognitively plausible, at all. To see this, consider our parabular Professor Jones and suppose without loss of generality that he is not a professional logician or mathematician, and hence cannot deduce, say, Gödel's famous first incompleteness theorem (= G1). By CWA, Smith should believe that G1 is false!

[14]At

http://plato.stanford.edu/entries/logic-ai caveats included within it: that when the perception in question occurs, Jones is not under the influence of drugs, not insane, and so on. But to ease exposition, let's leave aside such clauses. So, on the strength of this argument, we assume that Jones' knowledge includes (3), at time t_1.

Later on, as we have said, he finds himself in class at his university, away from home. Let's add that Jones and his students quickly consult smartphone weather apps and learn that the National Weather Service reports this tornado to have touched down somewhere in the town T in which Jones' house is located, and that major damage resulted, in particular some houses were leveled. At this point (t_2, assume), if Jones were pressed to articulate his current position on (3), and his reasoning for that position, and he had sufficient time and patience to comply, he would likely offer something like this (Argument 2):

(4) A tornado has just (i.e., at some time between t_1 and t_2) touched down in T, and destroyed some houses there.
(5) My house is located in T.
(6) I have no particular evidence that my house was *not* struck to smithereens by a tornado that recently passed through the town in which my house is located.
(7) If a tornado has just destroyed some houses in (arbitrary) town T', and house h is located in T, and one has no particular evidence that h is not among the houses destroyed by the tornado, then one ought not to believe that h wasn't destroyed.
$\therefore$ (8) I ought not to believe that my house is still standing. (I.e., I ought not to believe (3).)

Assuming that Jones meets all of his "epistemic obligations" (in other words, assuming that he's rational), he will not believe (3) at t_2. (Actually, and below we will deal with this more plausible modeling, it's more reasonable to imagine that Jones does still believe (3), but that the *strength* of his belief has declined.) Therefore, at this time, (3) will not be among the things he knows. (If a cognitive system s doesn't believe ϕ, it follows immediately that s doesn't know ϕ, in the sense of 'know' with which we are concerned with.) The nonmonotonicity here should be clear.

The challenge is to devise formalisms and mechanisms that model this kind of mental activity through time. The argument-based approach to nonmonotonic reasoning does this. As to how, the main move is to allow one argument to invalidate another (and one argument to invalidate an argument that invalidates an argument, which revives the original, etc.), and to keep a running tab on which propositions should be believed at any particular time. Argument 2 above rather obviously invalidates Argument 1; this is the situation at t_2. Should Jones then learn that only two houses in town T were leveled, and that they are both located on a street other than his own, Argument 2 would be defeated by a third argument, because this third argument would overthrow (6). With Argument 2 defeated, (3) would be reinstated, and back in what Jones knows. Clearly, this ebb and flow in argument-versus-argument activity is provably impossible in straight deductive reasoning.

6.1 An Argument-Adjudication System for Defeasible Reasoning

In order to adjudicate competing arguments, such as those in the tornado example of §6, we first need a system for quantifying the level of subjective uncertainty of declarative statements. To do this, we invoke a system based upon *strength factors* first presented by the first and third author in (Govindarajulu & Bringsjord 2017). These authors in turn were directly inspired by a simpler and

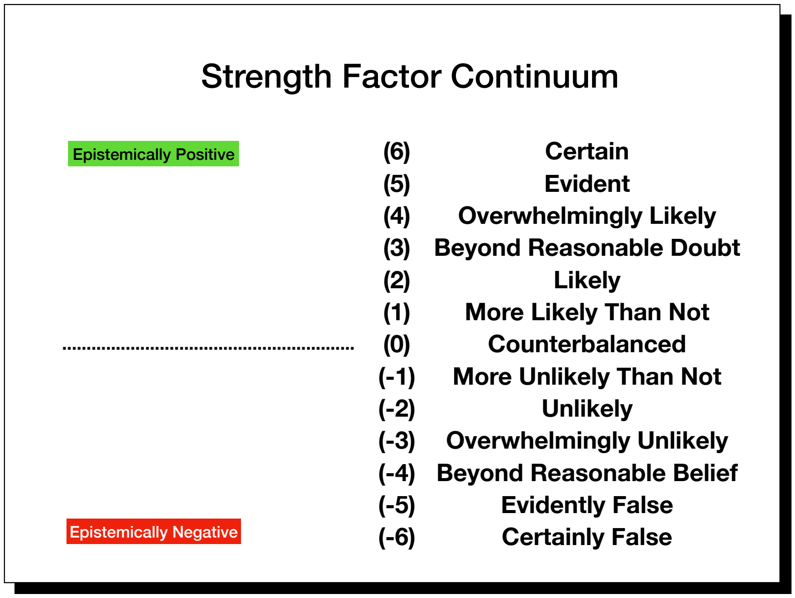

Figure 7: The Current Strength Factor Continuum. *The center value,* counterbalanced, *indicates that there is no evidence for or against belief in the subformula. Increasing positive and negative values indicate increasing and decreasing likelihood of truth in the subformula, respectively.*

smaller system of strength-indexed belief invented over half a century ago by Chisholm (1966).[15] While we have more recently specified a more robust formal inductive logic ($\mathcal{IDCEC}$; note that it is located within $\mathscr{U}$, as Figure 3 indicates) for such processing, implemented it, and demonstrated it (Bringsjord, Govindarajulu & Giancola 2021), the nature of the present chapter means that a "higher altitude" level of detail is prudent, and in what now follows we stay at that altitude. For more details, the reader can consult our detailed work, and the technical survey provided by Prakken & Vreeswijk (2001).

The strength factors we shall now employ consist of 13 values (see Figure 7) that can be used to annotate statements expressing belief or knowledge. For example, we can formalize the sentence "Jones believes it is *more likely than not* at time t_0 that his house is still standing." by the formula $\mathbf{B}^1(jones, t_0, \text{Standing}(home))$.

We note at this point that the introduction of uncertainty measures already forces us to move beyond deductive reasoning into inductive reasoning and logics, as we are no longer producing proofs, but instead, arguments. While a proof guarantees the truth of the formula it proves (as long as the axioms/premises are true), an argument only provides some level of strength that its conclusion is true. Hence, we here find ourselves moving from deductive reasoning to *in*ductive reasoning, in which such arguments are able to be created. The reader may at this point wish to

[15]There are formal logics that subsume probability theory, and theoretically they could be deployed to model the tornado scenario (e.g. there is *uncertain first-order logic*; see Núñez, Murthi, Premaratne, Bueno & Scheutz forthcoming). However, it doesn't seem cognitively plausible that Professsor Jones associates real numbers between 0 and 1 with the proposition that his house is still standing. One could also explore using so-called "fuzzy logic," which emerged out of fuzzy sets first introduced by Zadeh (1965). But here one must be very careful. Most of the things called "fuzzy logics" are not in fact logics at all, and are not in the universe $\mathscr{U}$. The advent of *bona fide* formal fuzzy logics, replete with formal languages, inferential machinery, and so on, came be way of the groundbreaking (Hájek 1998).

note that in Figure 3 inductive logics are denoted. For a recent introduction to inductive logic as an argument-based, as opposed to a proof-based, affair, the reader can consult (Johnson 2016).

We bring to bear here two intensional logics suitable for the type of modeling we need in the tornado scenario. Because the distinguishing purpose of these logics and others like them is the modeling of human-leve cognitive states (such as believing and knowing a proposition at a time), and human-level reasoning, we have long referred to these logics as *cognitive calculi*, and we follow suit here. The first cognitive calculus used here is for purely deductive reasoning; the second supports inductive reasoning. For the encapsulated formal specification of these cognitive calculi, see (Bringsjord et al. 2021). Industrious readers can find these two calculi in the universe $\mathscr{U}$ pictured in Figure **??**; they are named therein as $\mathcal{DCEC}^*$ and $\mathcal{IDCEC}$; the first is a deductive intensional logic, the second an inductive intensional logic.

Note that when we refer to arguments in this chapter, we mean more specifically *formal* arguments. Hence, like in any respectable proof, each step must be sanctioned by the deployment of an inference schema.[16]

When we have multiple such arguments, each of which concludes with the affirmation or rejection of belief in some subformula, the adjudication process is simple: select the argument whose conclusion has the highest strength. This method will be employed in §6.2 to formalize and rigorously model the tornado example first given in §6. More complex adjudication methods for more complex sets of arguments (e.g., where the adjudication process may need to select out subarguments from multiple arguments in order to construct the winning argument and corresponding final conclusion) are the focus of active research outside the scope of the present chapter.

6.2 The Tornado Conquered

Consider again the following scenario, now made a bit more determinate. Jones left his home (at time t_{home}) to go to work, and while there (at time t_{work}) he sees the news and discovers that a tornado has passed through the town (at time $t_{tornado}$) in which his house is located ($town$). What should Jones now believe with regard to whether or not his house is still standing?

We can pose this problem in the argument-adjudication framework we employ for defeasible/nonmonotonic reasoning in order to evaluate the strength of each argument and allow Jones to arrive at a final belief-fixation decision. First, consider an argument Jones might plausibly use to justify his belief that his house is standing at the time that he is about to leave for work, t_{home}, an argument that is now more nuanced and plausible that what we laid about above:

[16] For the relevant lists of such inference schemata, which are outside the scope of this overview chapter, we direct the reader to (Bringsjord et al. 2021).

	(1)	$\mathbf{P}(jones, t_{home}, Standing(home))$	Jones perceived that his home was standing when he left for work.
$\therefore$	(2)	$\mathbf{B}^5(jones, t_{home}, Standing(home))$	Assuming Jones was not dreaming or hallucinating, perception generates *evident* beliefs. Therefore, Jones beleieved it was *evident* that his home was still standing at that time.
$\therefore$	(3)	$\mathbf{O}(jones, t_{home}, \mathbf{B}^5(jones, t_{home}, Standing(home)))$	Hence Jones ought to believe it is *evident* at time t_{home} that his house is still standing.

Table 1: **Argument 1:** Jones determines he ought to believe it is *evident* that his house is still standing at time t_{home}.

We note here that the obligation operator is of an intellectual variety; we are not talking about anything like moral obligations and deontic operators that are at the heart of deontic logic, which is devoted to formalizing human moral reasoning. That one *ought* to believe ϕ here means that there is a rational argument compelling one to believe ϕ as a rational agent. This basic notion of intellectual obligation as part and parcel of an abstract conception of rationality is at the heart of the logic and mathematics of inductive logic (Paris & Vencovská 2015).

Next, consider another sequence of reasoning Jones might go through while driving to work (at time $t_{driving}$). Since he is no longer perceiving his home, his belief cannot be at the level of *evident*. However, his previous belief can persist at the next level down, *overwhelmingly likely*, so long as Jones has not been made aware of any information to the contrary since then.

	(4)	$\neg\mathbf{P}(jones, t_{driving}, Standing(home))$	Jones no longer perceives his home.
$\therefore$	(5)	$\neg\mathbf{B}^5(jones, t_{driving}, Standing(home))$	Hence, Jones no longer beleieves it is *evident* that his home is still standing.
$\therefore$	(6)	$\mathbf{O}(jones, t_{driving}, \mathbf{B}^4(jones, t_{driving}, Standing(home)))$	Assuming Jones' memory is reasonably reliable, and since he has no information to the contrary, he ought to believe it is *overwhelmingly likely* at time $t_{driving}$ that his house is still standing.

Table 2: **Argument 2:** Jones retracts his previous belief that he ought to believe it is *evident* that his house is still standing at time $t_{driving}$, and replaces it with a belief at the level of *overwhelmingly likely*.

Finally, at t_{work}, Jones becomes aware of the tornado which just passed through his town. Therefore he is rationally obligated to retract his previous belief, and replace it with a weaker belief that his house is still standing:

	(7)	$\mathbf{K}(jones, t_{work}, LocatedIn(home, town))$	Jones knows his home is located in his town.
	(8)	$\mathbf{S}(news, jones, t_{work}, TornadoPassedThrough(town, t_{tornado}))$	Jones heard from the news that a tornado passed through the town where his home is located.
	(9)	$\mathbf{K}(jones, t_{work}, \forall h\ a\ t\ (TornadoPassedThrough(a, t) \land LocatedIn(h, a)) \rightarrow \Diamond\neg Standing(h))$	Jones knows that if a tornado passes through an area where a home is located, it is possible that that home is no longer standing.
$\therefore$	(10)	$\mathbf{K}(jones, t_{work}, \Diamond\neg Standing(home)$	Hence Jones knows it is possible that his home is no longer standing.
$\therefore$	(11)	$\neg\mathbf{B}^4(jones, t_{work}, Standing(home))$	Hence Jones no longer believes it is overwhelmingly likely that his home is still standing.
$\therefore$	(12)	$\mathbf{O}(jones, t_{work}, \mathbf{B}^2(jones, t_{work}, Standing(home)))$	However, since Jones has only evidence indicating a possibility that his home has been destroyed, he ought to believe it is *likely* at time t_{work} that his house is still standing.

Table 3: **Argument 3:** Jones determines he ought to believe it is *likely* that his house is still standing at time t_{work}.

6.3 The Suppression Task

The task in question is reported in (Byrne 1989). Three groups of subjects were asked to select which proposition from among a trio of them "follows"[17] from a set of suppositions. Each group of subjects was given a different set of suppositions. Group 1 (= G1) was given this pair of suppositions:

(s1) If she has an essay to finish, then she will study late in the library.

(s2) She has an essay to finish.

This group's options to select from were the following three:

(o1) She will study late in the library.

(o2) She will not study late in the library.

(o3) She may or may not study late in the library.

Among G1, 96% selected (o1). G2 was given suppositions consisting of (s1) and (s2), plus the following supposition:

(s3) If she has a textbook to read, then she will study late in the library.

In G2, again 96% of its members selected option (o1). G3 received (s1) and (s2), plus this supposition:

[17] Unfortunately, 'follows' is a metaphor here — but it's the term Byrne (1989) used. We have no firm conception as to what this term means. From the standpoint of formal logic, what should have been said to subjects is something like 'must necessarily be deducible,' because (i) the hallmark of deduction since first systematically investigated by Aristotle has been appehended as the fact that when deduction from givens/premises/suppositions to (a) conclusion(s) is valid, the former *necessarily* entail the latter, and because (ii) plenty of conclusions are thought by rational agents operating rationally to follow from givens/premises/suppositions that certainly don't necessitate these conclusions (e.g., consider a case in which a conclusion follows from premises by statistical syllogism). However, this being said, for now, the unfortunate use of 'follows' by Byrne (1989) must be left aside.

(s4) If the library stays open, then she will study late in the library.

This time things turned out quite differently: only 38% of G3 selected (o1).

From the perspective of standard zero-order logic = $\mathscr{L}_0 \in \mathscr{U}$,[18] which can accordingly be assumed here to have any standard proof theory, such as is used in early classical mathematics (e.g. high-school mathematics in every technologized society/nation), this result is interesting, since, to begin, in $\mathscr{L}_0$ we might represent the declarative sentences (s1), (s2), (s3), and (s4) as follows, where a represents the female agent in question:

(s1*) $\mathit{ToFinish}(a) \rightarrow \mathit{LateLibrary}(a)$

(s2*) $\mathit{ToFinish}(a)$

(s3*) $\mathit{ToRead}(a) \rightarrow \mathit{LateLibrary}(a)$

(s4*) $\mathit{StaysOpen} \rightarrow \mathit{LateLibrary}(a)$

Next, following suit, the options would be represented thus:

(o1*) $\mathit{LateLibrary}(a)$

(o2*) $\neg \mathit{LateLibrary}(a)$

(o3*) $\neg \mathit{LateLibrary}(a) \vee \mathit{LateLibrary}(a)$

With these representations, easy-to-find proofs in $\mathscr{L}_0$ certify that

$$\{(\text{s1}^*), (\text{s2}^*), (\text{s3}^*)\} \vdash (\text{o1}^*). \qquad (+)$$

However, there is no available proof in this logic of option two from the first three suppositions; that is:

$$\{(\text{s1}^*), (\text{s2}^*), (\text{s3}^*)\} \nvdash (\text{o2}^*). \qquad (-)$$

Option (o3*) is a theorem in this logic, so it's provable from $\{(\text{s1}^*), (\text{s2}^*), (\text{s3}^*)\}$.[19] Because we dealing here with standard deductive reasoning, which as we have noted is non-feasible/monotonic, adding one or both of (s3*), (s4*) to $\{(\text{s1}^*), (\text{s2}^*), (\text{s3}^*)\}$ doesn't change provability/unprovability; that is, neither (+) nor (–) change. This is why the group G3 is interesting from the point of view of $\mathscr{L}_0$, and hence from the point of view of the cognitive science of reasoning. Clearly, the formal modeling just given via $\mathscr{L}_0$ doesn't match what most of the subjects in this group were thinking when they responded.

6.3.1 Stenning & van Lambalgen's Extensional Treatment of the Suppression Task

Byrne, in her presentation of the Suppression Task (Byrne 1989), argues that the findings of her study imply that people don't strictly apply valid methods of logical deduction when reasoning. Therefore, so her story goes, logic is not sufficient for modeling human reasoning. She states that "...in order to explain how people reason, we need to explain how premises of the same apparent logical form can be interpreted in quite different ways" (Byrne 1989).

[18]Obtained by augmenting the formal language of the propositional calculus with provision for relation and function symbols, and the identity symbol =; but no quantifiers are allowed. Like the propositional calculus, $\mathscr{L}_0$ is Turing-decidable; not so any n-order logic $\mathscr{L}_n$ in $\mathscr{U}$, where n is a positive integer.

[19]As a matter of fact it's not appropriate to represent (o3) as having the form $\phi \vee \neg\phi$, but we leave this issue aside here.

Stenning & van Lambalgen (S&V) (2008) formalize this concept of what can be called "premise interpretation."[20] They claim that humans, when presented with a set of premises and possible conclusions, first reason *toward* some rational interpretation of the premises, then *from* that interpretation to some conclusion. They formalize this process in a Horn-style propositional logic, supplemented with a formalization of the Closed World Assumption (CWA). [21] Given this context, when presented with a set of assumptions and a conclusion to prove, S&V follow this three-step algorithm:

1. Reason to an interpretation.
2. Apply nonmonotonic closed-world reasoning to the interpretation.
3. Reason from the result of what step 2. produces.

We now consider the application of these three steps to the first experiment in Byrne's (1989) study, but first we need to have before us the stimuli presented to subjects. In her first experiment, subjects are given the two suppositions:

(s1) If she has an essay to write, she will study late in the library.
(s2) She has an essay to write.

and are then asked to choose from the following set of conclusions which one follows from the premises.[22]

(o1) She will study late in the library.
(o2) She will not study late in the library.
(o3) She may or may not study late in the library.

Now we come to the application of the three-step algorithm.

6.3.1.1 The Algorithm, Applied

Step 1: Reasoning to an Interpretation The first part of this step is appending the antecedent of every conditional with "$\neg ab$" this addition, intuitively, means "no abnormalities." The idea here is that people interpret the conditional $p \to q$ as $(p \wedge \neg ab) \to q$. That is, p implies q, *if* no external factors of which the subject is currently unaware (i.e. the abnormalities represented by ab) subvert the implication.

[20] S&V are not the only LCCMers who have tried their hand at modeling ST: Dietz et al. previously took two distinct logic-based approaches to modeling the Suppression Task. In their first approach, they used a three-valued Łukasiewicz logic which allows the expression of a third truth value beyond *true* and *false*: *unknown.* (Dietz, Hölldobler & Ragni 2012, Dietz, Hölldobler & Wernhard 2014). More recently, they have taken an approach which aims to model the Suppression Task in a more cognitively-plausible way: (Saldanha & Kakas 2020).

Their framework, called *Cognitive Argumentation*, formalizes methods of reasoning used by humans (which may or may not be logically sound) as *cognitive principles*. For example, their "Maxim of Quality" expresses that we (humans) typically assume statements we are told are true if we don't have a reason to believe otherwise (e.g. that the speaker may be lying or incompetent). In the context of the Suppression Task, the Maxim of Quality dictates that the subjects will assume that all of the statements made by the experimenters are true (e.g. "She has an essay to finish.").

[21] Recall that, in a word, CWA is the assumption that everything about a domain is known. Formally, any proposition which is not known to be true (or not provable) is assumed to be false.

[22] Note again that Byrne uses the informal term 'follows' and not one necessitating formal entailment like 'logically deduces.'.

The last part of this step is to collect the assumptions as modified above into a set which S&V refer to as the *logic program* corresponding to the assumptions. Given the foregoing, the output of Step 1 for Experiment 1 would be the set:

$$\{\ EssayToWrite; EssayToWrite \wedge \neg ab \rightarrow StudyLateInLibrary\ \} \tag{1}$$

Step 2: Applying Nonmonotonic Closed-World Reasoning to the Interpretation This step also consists of two sub-parts. First, for all atoms q in the logic program produced in Step 1, if there is no antecedent p such that $p \rightarrow q$, the conditional $\perp \rightarrow q$ is added to the logic program. We note that in S&V's logic, the meaning of an atom p in the assumption base is really $\top \rightarrow p$; but for clarity, they typically just write p; we will do the same. Therefore, in the example above, the only atom for which this step applies is ab; hence we add the conditional $\perp \rightarrow ab$ to the logic program:

$$\{\ EssayToWrite; EssayToWrite \wedge \neg ab \rightarrow StudyLateInLibrary; \perp \rightarrow ab\ \} \tag{2}$$

The second part of Step 2 is what S&V refer to as *constructing the completion* of the logic program. This involves first joining all implications $\phi_i \rightarrow q$ (i.e. those implications whose consequent is q) into a single implication $\vee_i \phi_i \rightarrow q$.[23] Second, all conditionals are converted to biconditionals. Therefore our final logic program (also, our interpretation of the premises) is:

$$\{\ EssayToWrite; EssayToWrite \wedge \neg ab \leftrightarrow StudyLateInLibrary; \perp \leftrightarrow ab\ \} \tag{3}$$

Step 3: Reasoning from the Result of Step 2 The third and final step is fairly straightforward: the subject reasons from the final set of premises using the inference rules of standard propositional logic. Notice that, because $\perp \leftrightarrow ab$, we have $\top \leftrightarrow \neg ab$; hence our logic program above can be simplified to:

$$\{EssayToWrite; EssayToWrite \leftrightarrow StudyLateInLibrary\} \tag{4}$$

Finally, it is obvious that from these premises we can deduce $StudyLateInLibrary$. We note that while the conclusion was obvious in this case, this method of reasoning to and from an interpretation matches the reasoning process of the majority of people in all of Byrne's experiments. We next walk through S&V's algorithm for a slightly more complicated (and more interesting) case, in which an additional premise is introduced.

6.3.2 Applying the Algorithm to the Additional-Premise Case

In the second experiment, recall, Byrne gave her subjects the following set of premises:

> If she has an essay to write, she will study late in the library.
> If the library stays open, she will study late in the library.
> She has an essay to write.

[23] There are no instances of this in this example, but there will be in the next example.

This additional premise is modeled using the same form as the original two premises:

$$LibraryOpen \land \neg ab' \to StudyLateInLibrary \tag{5}$$

However, in this case, S&V also add the following premise:

$$\neg LibraryOpen \to ab \tag{6}$$

This premise is intended to model the belief of those who believed that *modus ponens* applied in Experiment 1, but not in Experiment 2. (In other words, the introduction of the additional premise suppressed their belief.) More specifically, this conditional states that if the library is not open, then it would be abnormal for her to go to study late in the library. We can also add the symmetric condition $\neg EssayToWrite \to ab'$; that is, if she does not have an essay to write, it would be abnormal for her to study late in the library.[24]

Now, performing Step 1 will produce the program:

$$\left\{ \begin{array}{c} EssayToWrite \land \neg ab \to StudyLateInLibrary \\ LibraryOpen \land \neg ab' \to StudyLateInLibrary \\ EssayToWrite \\ \neg LibraryOpen \to ab \\ \neg EssayToWrite \to ab' \end{array} \right\} \tag{7}$$

Next, applying nonmonotonic closed-world reasoning yields:

$$\left\{ \begin{array}{c} (EssayToWrite \land \neg ab) \lor (LibraryOpen \land \neg ab') \leftrightarrow StudyLateInLibrary \\ EssayToWrite \\ (\bot \lor \neg LibraryOpen) \leftrightarrow ab \\ (\bot \lor \neg EssayToWrite) \leftrightarrow ab' \end{array} \right\} \tag{8}$$

And now, using standard logical deduction for the propositional calculus, we can simplify this set to:

$$\{EssayToWrite; (EssayToWrite \land LibraryOpen) \leftrightarrow StudyLateInLibrary\} \tag{9}$$

Finally, the subject reasons from this interpretation of the premises. Note that the second statement says "She will study late in the library if and only if she has an essay to write and the library stays open." Since the premise set doesn't include the proposition *LibraryOpen*, we cannot deduce *StudyLateInLibrary*. This result matches the common human intuition[25] that the additional premise hinders the successful application of *modus ponens* to the original premises.

6.3.3 Modeling the Suppression with *Intensional* Logic

We now quickly demonstrate that the human reasoning in the suppression task can be easily and efficiently modeled in a way simpler than that employed by S&V. In this alternate route, we (1)

[24] This is not necessary but will allow us to simplify the final result.

[25] I.e., the intuition of the majority of the people in Byrne's study.

take timepoints seriously within the narrative implicit in what the subjects are given; and (2), use these timepoints in connection with a simple intensional logic that includes (i) a way to represent and reason with what is *known* and what is *believed*, and (ii) includes an operator for what is *possibly* the case.[26]

This first step in carrying out these two steps is to simply announce a simple set of symbols used to enable the formulae that express what is presented to subjects in the suppression task. We do this by way of the following table, which simply presents the referent in each case intuitively, so that no technical specifications are needed.

Table 4: Symbols for Intensional Mod & Sim of Suppression Task

Symbol	**Referent**
s (variable)	student
e (variable)	essay
b (variable)	book
$t, t', \ldots$ (variables)	timepoints
t_1 (constant)	the particular, initial timepoint
ℓ (constant)	the library
$\mathbf{a}, \mathbf{b}$	two particular agents
$>$ (2-place relation)	later than
$\mathit{ToFinish}(s, t, e)$ (3-place relation)	s at t has e to finish
$\mathit{NearFuture}(t', t)$	t' is in near future of t
$\mathit{LateLibrary}(s, t)$	s works late in the library at t
$\mathit{Open}(\ell, t)$	the library is open at t
$\mathit{ToRead}(s, t, b)$ (3-place relation)	s at t has textbook b to read
$\Diamond$ (alethic operator)	'possibly'
$\mathbf{B}_x$ (epistemic operator)	agent x believes
$\mathrm{ToRead}(s, t, b)$ (3-place relation)	s has at t to read b
$\mathbf{K}_x$ (epistemic operator)	agent x knows

Given this more expressive vocabulary at our disposal, here is how the key propositions from above in the suppression task are expressed in our intensional approach:

$$\exists e\ \mathit{ToFinish}(s, t_1, e) \rightarrow \exists t > t_1\ \big(\mathit{NearFuture}(t, t_1) \wedge \mathit{LateLibrary}(s, t)\big) \tag{s1}$$

$$\exists e\ \mathit{ToFinish}(s, t_1, e) \tag{s2}$$

$$\exists t > t_1(\mathit{NearFuture}((t, t_1) \wedge \mathit{LateLibrary}(s, t)) \tag{o1}$$

$$\neg\big(\exists t > t_1(\mathit{NearFuture}((t, t_1) \wedge \mathit{LateLibrary}(s, t)))\big) \tag{o2}$$

$$\Diamond(\text{o1}) \wedge (\Diamond\neg(\text{o1}) \vee \Diamond O2) \tag{o3}$$

$$\exists b\ \mathit{ToRead}(s, t_1, b) \rightarrow \exists t > t_1\ (\mathit{NearFuture}((t, t_1) \wedge \mathit{LateLibrary}(s, t)) \tag{s3}$$

$$\begin{aligned}&[\mathit{Open}(\ell, t_1) \wedge \forall t > t_1\ (\mathit{NearFuture}((t, t_1) \rightarrow \mathit{Open}(\ell, t)) \wedge \exists e\ \mathit{ToFinish}(s, t_1, e)]\\ &\qquad \rightarrow \exists t > t_1(\mathit{NearFuture}((t, t_1) \wedge \mathit{LateLibrary}(s, t))\end{aligned} \tag{s4}$$

[26] We thus make use of basic constructs from *epistemic* logic (Hendricks & Symons 2006), which formalizes attitudes like *believes* and *knows*; and also basic constructs from *alethic modal logic* (Konyndyk 1986), which formalizes concepts like *possibly* and *necessarily.*

And here is an economical summation of the deductive "facts of the case" under the more expressive rubric afforded by Table 4, where $\Gamma \vdash \phi$ is the ubiquitous way in formal logic, AI, and computer science of saying that ϕ can be deduced from a set Γ of formulae (and $\nvdash$ means 'not deducible'):

- $\{(\mathrm{s1}), (\mathrm{s2})\} \vdash (\mathrm{o1})$
- $\{(\mathrm{s1}), (\mathrm{s2})\} \nvdash (\mathrm{o2})$
- $\{(\mathrm{s1}), (\mathrm{s2})\} \nvdash (\mathrm{o3})$
- $\{(\mathrm{s1}), (\mathrm{s2}), (\mathrm{s3})\} \vdash (\mathrm{o1})$
- $\{(\mathrm{s1}), (\mathrm{s2}), (\mathrm{s3})\} \nvdash (\mathrm{o2})$
- $\{(\mathrm{s1}), (\mathrm{s2}), (\mathrm{s3})\} \nvdash (\mathrm{o3})$

Very well. And now what is the intensional modeling that matches what occurs when subjects are run in the suppression task? Such modeling, as we have said, takes time, possibility, and epistemic attitudes (belief and knowledge) seriously. Specifically, the heart of the matter is a simple inference schema that formalizes the principle that if an agent believes some set Φ of propositions, and knows that from this set it can be deduced specifically that proposition ϕ holds, then the agent will believe ϕ. Here is the inference schema, $\mathcal{S}$, expressed in a manner used in our computational simulations:

$$\frac{\mathbf{B}_a\Phi,\ \mathbf{K}_a\Phi \vdash \phi}{\mathbf{B}_a\phi}\ \mathcal{S}$$

And now, getting down to inferential brass tacks for computational simulation, let '**a**' denote an arbitrary agent in both Group I and Group II in the suppression-task experiment recounted above. We then assume, at the particular timepoint t_1, that

$$\mathbf{B}_a\{(\mathrm{s1}), (\mathrm{s2})\};$$

and in addition that

$$\mathbf{K_a}\{(\mathrm{s1}), (\mathrm{s2})\} \vdash (\mathrm{o1}).$$

Then, by way of crucial use of $\mathcal{S}$, processing automatically locates a proof correpsonding to the responses of agents in Groups I and II: $\mathbf{B}_a(\mathrm{o1})$. In a simulation using an automated theorem prover, this result (and the corresponding proof) was returned in 10^{-4} seconds.[27]

But now, what about the "peculiar" subjects in Group III? That is, what about subjects who clearly reason defeasibly/nonmonotonically, because they go from believing that (o1) "follows" to believing, after receiving new information, that this proposition no longer does? These are of course the subjects that motivated the innovation of S&V. But how is the inferential behavior of these subjects modeled and simulated in our *intensional* approach? The answer is perfectly straightforward. The answer is that, first, Group-III subjects obviously know that when a library is closed (= not open) at some time t, no student can work in that libary at t. This underlying principle is in our modeling expressed thus:

[27]Two automated reasoners were used to generate the simulation results presented in this chapter. The first, ShadowProver (Govindarajulu, Bringsjord & Peveler 2019), uses a novel technique to prove formulae in a modal logic. It alternates between "shadowing" modal formulae down to first-order logic and applying modal inference schemata. The second, ShadowAdjudicator (Giancola, Bringsjord, Govindarajulu & Varela 2020), builds upon ShadowProver, providing the ability to generate *arguments* (as opposed to proofs) which can be justified using *inductive* inference schemata (as opposed to purely deductive inference schemata).

(u) $\forall s \forall t \ [\neg Open(\ell, t) \rightarrow \neg LateLibrary(a, t)]$

In addition, of course, subjects in Group III know from what they have been told that

($*$) $\exists s \exists e \ ToFinish(s, t_1, e)$,

and know as well that at all near-future times relative to t_1 the library is closed;[28] that is:

($\star$) $\forall t(NearFuture(t, t_1) \rightarrow \neg Open(\ell, t))$.

Given the pair of formulae ($*$) and ($\star$) it follows by elementary deduction in $\mathscr{L}_1$ that ¬(s1). Therefore, while we can rationally presume that Group-III subjects — which we denote with **b** — are (like their counterparts in Groups I and II) such that

$$\mathbf{K_b}\{(\mathrm{s1}), (\mathrm{s2})\} \vdash (\mathrm{o1}),$$

they no longer believe (s1), and hence the use of schema $\mathcal{S}$ is blocked. In addition, we can reasonably model that Group-III subjects do believe (s4). But also

$$\{(\mathrm{s4}), (\mathrm{s2})\} \not\vdash (\mathrm{o1}),$$

and these subjects presumably know this. Hence these subjects cannot possibly know that $\{(\mathrm{s4}), (\mathrm{s2})\} \vdash$ (o1), and this too blocks any use of schema $\mathcal{S}$ to arrive at the belief that (s1) holds.[29]

We see little point in asserting that capturing the suppression task via intensional logic as in our case is superior to the extensional-logic approach taken by S&V. However, we do think that it's very important for the student and scholar of computational cognitive science to understand that any such ambition as to capture *all* of human-level-and-above reasoning and decision-making in computational formal logic must early on confront modeling-and-simulation challenges that *necessitate* use of highly expressive intensional logics from $\mathscr{U}$.

7 Evaluating Logic-based Cognitive Modeling, Briefly

Logic-based/logicist cognitive computational cognitive modeling, LCCM, seems in particular to be a rather nice fit when the cognition to be modeled is explicit, rational, and intensely inference-centric. But how accurate and informative is such modeling? And how much reach does does such an approach to cognitive modeling have, in light of the fact that surely plenty of human-level cognition is neither explicit, nor rational, nor inference-centric? This is not the venue for polemical positions to be expressed in response to such questions. We rest content with pointing out that "accuracy" of cognitive models is itself not exactly the clearest concept in science, and that LCCM tantalizingly offers the opportunity to itself provide the machinery to render this concept precise. The relationship of a model M to a targeted phenomenon P to be modeled, in LCCM, should

[28] Actually, as alert readers will apprehend, it's necessary here to use the alethic operator $\Diamond$, since what the subjects in Group III come to know by virtue of the new information given them is that *it might possibly be* that the library is closed in the near future.

[29] Simulations of these lines of reasoning found by our automated-reasoning technology are strikingly fast, but we leave reports aside so as not to have to delve into rather tricky simultaneous use of the alethic operator $\Diamond$ in combination with **K** (knows) and **B** (believes). Please see note 28.

itself be a relation formalized in some logic in the universe $\mathscr{U}$. If we let the relation $\mathcal{A}$ stand for "accurately models" we can then say what is needed is the completion of the biconditional

$$(\star) \quad \mathcal{A}(M, P) \leftrightarrow \boxed{??}.$$

With this accomplished, LCCM would provide the very framework that could be used to assess its own accuracy, because one would be able to prove that ?? holds in the case at hand, and then reason from right to left on the biconditional in order to show $\mathcal{A}(M, P)$. We are not aware of any other approach to cognitive modeling that can hold out the promise of such self-containedness.

As to the reach of LCCM, we report only that some mental phenomena that seem fundamentally ill-suited to this approach, for instance emotions and emotional states, do seem to conform remarkably well at the level of a scientific view of them to collections of formulae from relatively simple modal (i.e. intensional) logics in $\mathscr{U}$ (Adam, Herzig & Longin 2009).

One final point regarding the assessment of LCCM, a point that follows from how we have earlier defined what it is for logicist cognitive modeling to *capture* some aspect or part of human-level cognition. The point is simply this: Whether or not some attempt to cognitively model (in the LCCM approach) some phenomenon succeeds or not can be settled formally, by proof/disproof. The ultimate strong suit of LCCM is formal verifiability of capture. We can know that we have captured some phenomenon, period, because we have proof. Unfortunately, carrying this out in practice in a wide way would require the formalization of $\boxed{??}$ so that $(\star)$ can be employed in the matter we have described.

8 Conclusion and Future Work

It should be clear to the reader that formal computational logic is plausibly up to the challenge of modeling and simulating both quantification-centric reasoning and defeasible (= nonmonotonic) reasoning at the human level and in the human case, even when this challenge is required to be substantively based upon arguments of the sort that human agents routinely form as they adjust their belief and knowledge through time. But for the overarching program of LCCM, is the ambitious long-term goal of capturing *all* rational human reasoning in computational logic reasonable? And if it is, what is next to be done?

While the present chapter extends the rather narrow deduction-focused overview of LCCM given earlier (Bringsjord 2008) into the important realms of quantification and dynamic defeasible reasoning in the human sphere, certinly humans reason and cognize in many additional ways, effectively. These additional ways range from the familiar and everyday, to the rarefied heights of cutting-edge formal science. In the former case, prominently, there is reasoning that makes crucial use of pictorial elements, and hence is reasoning that simply cannot be captured by the kind of symbolic structures we have hitherto brought to bear. As alert readers will have noticed, the universe $\mathscr{U}$ does include logics that offer machinery for representing and reasoning over diagrams and images. For a simple but relevant example, consider the question as to whether

❏

or

□

is more likely to have in front of it and shining upon it a light. Here, the two things centered just above aren't symbols; they are diagrams, and as such denote not as symbols do, but — to use the apt terminology of Sloman (1971) and Barwise (1995), resp. — in a manner that is *analogical* or *homomorphic*. Clearly, humans do routinely reason with diagrams — and yet the logics we have employed above from $\mathscr{U}$ have no diagrams. Therefore further work in LCCM is clearly in order.[30] This work would be a treatment that brings to bear the spaces of pictorial logics indicated in the universe $\mathscr{U}$ of Figure 3.

Now, finally, what about the latter challenge, that of applying LCCM to rarefied reasoning in the formal sciences? Here we must confront the fact that reasoning in logic and mathematics often makes use of expressions and structures that are infinitary in nature. For example, there can be very good reason to make use of formulae that are infinitely long, such as a disjunction like

$$\delta := \exists^{=1} xRx \vee \exists^{=2} xRx \vee \ldots,$$

which says that there is exactly one thing that is an R or exactly two things each of which is an R, or exactly three things each of which is an R, and so on *ad infinitum*. It turns out that however exotic δ may seem, this is about the only way to express that there exist a finite number of Rs. And yet of course we have discussed no logics that allow for infinitely long disjunctions to be constructed; we have left what are classified as "infinitary logics" in the universe $\mathscr{U}$ untouched in the foregoing discussion. Of course, as the reader will rationally suspect, the need for formulae of this nature, given the infinitary expressions presented even in textbooks devoted to bringing humans into serious cognizing about (say) analysis (e.g. see Heil 2019), is undeniable. So again, it would seem that if the general program of logic-based cognitive modeling is to succeed in capturing human reasoning and human-level reasoning across the board, additional effort of a different nature than has so far been carried out will be required of relevant researchers. This effort will need to tap other logics in $\mathscr{U}$ shown in Figure 3, which as the reader can now note by returning to that figure does indeed refer to the space of infinitary logics.[31]

[30] There are very few formal logics that allow, in addition to the standard symbolic/linguistic alphabets and grammars, diagrams/images. For such a logic, see (Arkoudas & Bringsjord 2009), which provided comprehensive references to the relevant literature.

[31] Readers wanting a short, cogent introduction to infinitary logic, should see presentation and explanation of the straightforward infinitary logic $\mathscr{L}_{\omega_1\omega}$ (which can express δ); and those with some logico-mathematical maturity can see (Dickmann 1975).

References

Adam, C., Herzig, A. & Longin, D. (2009), 'A Logical Formalization of the OCC Theory of Emotions', *Synthese* **168**(2), 201–248.

Andréka, H., Madarász, J. X., Németi, I. & Székely, G. (2011), 'A Logic Road from Special Relativity to General Relativity', *Synthese* pp. 1–17.
URL: *http://dx.doi.org/10.1007/s11229-011-9914-8*

Arkoudas, K. & Bringsjord, S. (2009), 'Vivid: An AI Framework for Heterogeneous Problem Solving', *Artificial Intelligence* **173**(15), 1367–1405.
URL: *http://kryten.mm.rpi.edu/KA_SB_Vivid_offprint_AIJ.pdf*

Ashcraft, M. & Radvansky, G. (2013), *Cognition*, Pearson, London, UK. This is the 6th edition.

Barwise, J. & Etchemendy, J. (1995), Heterogeneous Logic, *in* J. Glasgow, N. Narayanan & B. Chandrasekaran, eds, 'Diagrammatic Reasoning: Cognitive and Computational Perspectives', MIT Press, Cambridge, MA, pp. 211–234.

Boolos, G. S., Burgess, J. P. & Jeffrey, R. C. (2003), *Computability and Logic (Fourth Edition)*, Cambridge University Press, Cambridge, UK.

Bringsjord, S. (2008), Declarative/Logic-Based Cognitive Modeling, *in* R. Sun, ed., 'The Handbook of Computational Psychology', Cambridge University Press, Cambridge, UK, pp. 127–169.
URL: *http://kryten.mm.rpi.edu/sb_lccm_ab-toc_031607.pdf*

Bringsjord, S. (2014), 'Review of P. Thagard's *The Brain and the Meaning of Life*', *Religion & Theology* **21**, 421–425.
URL: *http://kryten.mm.rpi.edu/SBringsjord_review_PThagard_TBTMOL.pdf*

Bringsjord, S., Govindarajulu, N. & Giancola, M. (2021), 'Automated Argument Adjudication to Solve Ethical Problems in Multi-Agent Environments', *Paladyn, Journal of Behavioral Robotics* **12**, 310–335. The URL here goes to a *rough*, *uncorrected*, *truncated* preprint as of 071421.
URL: *http://kryten.mm.rpi.edu/AutomatedArgumentAdjudicationPaladyn071421.pdf*

Bringsjord, S. & Govindarajulu, N. S. (2018), Artificial Intelligence, *in* E. Zalta, ed., 'The Stanford Encyclopedia of Philosophy'.
URL: *https://plato.stanford.edu/entries/artificial-intelligence*

Bringsjord, S., Licato, J. & Bringsjord, A. (2016), The Contemporary Craft of Creating Characters Meets Today's Cognitive Architectures: A Case Study in Expressivity, *in* J. Turner, M. Nixon, U. Bernardet & S. DiPaola, eds, 'Integrating Cognitive Architectures into Virtual Character Design', IGI Global, Hershey, PA, pp. 151–180.

Byrne, R. (1989), 'Suppressing Valid Inferences with Conditionals', *Journal of Memory and Language* **31**, 61–83.

Charniak, E. & McDermott, D. (1985), *Introduction to Artificial Intelligence*, Addison-Wesley, Reading, MA.

Chisholm, R. (1966), *Theory of Knowledge*, Prentice-Hall, Englewood Cliffs, NJ.

Davis, M., Sigal, R. & Weyuker, E. (1994), *Computability, Complexity, and Languages: Fundamentals of Theoretical Computer Science*, Academic Press, New York, NY. This is the second edition, which added Sigal as a co-author.

Dickmann, M. A. (1975), *Large Infinitary Languages*, North-Holland, Amsterdam, The Netherlands.

Dietz, E.-A., Hölldobler, S. & Ragni, M. (2012), A Computational Logic Approach to the Suppression Task, *in* 'Proceedings of the Annual Meeting of the Cognitive Science Society', Vol. 34.

Dietz, E.-A., Hölldobler, S. & Wernhard, C. (2014), 'Modeling the Suppression Task Under Weak Completion and Well-Founded Semantics', *Journal of Applied Non-Classical Logics* **24**(1-2), 61–85.

Ebbinghaus, H. D., Flum, J. & Thomas, W. (1994), *Mathematical Logic (second edition)*, Springer-Verlag, New York, NY.

Feferman, S. (1995), Turing in the Land of O(Z), *in* R. Herken, ed., 'The Universal Turing Machine (2nd Ed)', Springer-Verlag, Secaucus, NJ, pp. 103–134.

Genesereth, M. & Nilsson, N. (1987), *Logical Foundations of Artificial Intelligence*, Morgan Kaufmann, Los Altos, CA.

Giancola, M., Bringsjord, S., Govindarajulu, N. S. & Varela, C. (2020), Ethical Reasoning for Autonomous Agents Under Uncertainty, *in* M. Tokhi, M. Ferreira, N. Govindarajulu, M. Silva, E. Kadar, J. Wang & A. Kaur, eds, 'Smart Living and Quality Health with Robots • Proceedings of ICRES 2020', CLAWAR, London, UK, pp. 26–41. Paper vailable at the URL given above. The ShadowAdjudicator system can be obtained here: `https://github.com/RAIRLab/ShadowAdjudicator`.
URL: *http://kryten.mm.rpi.edu/MG_SB_NSG_CV_LogicizationMiracleOnHudson.pdf*

Glymour, C. (1992), *Thinking Things Through*, MIT Press, Cambridge, MA.

Govindarajulu, N., Bringsjord, S. & Peveler, M. (2019), On Quantified Modal Theorem Proving for Modeling Ethics, *in* M. Suda & S. Winkler, eds, 'Proceedings of the Second International Workshop on Automated Reasoning: Challenges, Applications, Directions, Exemplary Achievements (ARCADE 2019)', Vol. 311 of *Electronic Proceedings in Theoretical Computer Science*, Open Publishing Association, Waterloo, Australia, pp. 43–49. The ShadowProver system can be obtained here: https://naveensundarg.github.io/prover/.
URL: *http://eptcs.web.cse.unsw.edu.au/paper.cgi?ARCADE2019.7.pdf*

Govindarajulu, N. S. & Bringsjord, S. (2017), Strength Factors: An Uncertainty System for Quantified Modal Logic, *in* V. Belle, J. Cussens, M. Finger, L. Godo, H. Prade & G. Qi, eds, 'Proceedings of the IJCAI Workshop on "Logical Foundations for Uncertainty and Machine Learning" (LFU-2017)', Melbourne, Australia, pp. 34–40.
URL: *http://homepages.inf.ed.ac.uk/vbelle/workshops/lfu17/proc.pdf*

Groarke, L. (1996/2017), Informal Logic, *in* E. Zalta, ed., 'The Stanford Encyclopedia of Philosophy'.
URL: *https://plato.stanford.edu/entries/logic-informal*

Hájek, P. (1998), *Metamathematics of Fuzzy Logic*, Trends in Logic, Volume 4, Kluwer, Dordrecht, The Netherlands.

Heil, C. (2019), *Introduction to Real Analysis*, Springer, Cham, Switzerland.

Hendricks, V. & Symons, J. (2006), Epistemic Logic, *in* E. Zalta, ed., 'The Stanford Encyclopedia of Philosophy'.
URL: *http://plato.stanford.edu/entries/logic-epistemic*

Hummel, J. (2010), 'Symbolic versus Associative Learning', *Cognitive Science* **34**(6), 958–965.

Hummel, J. E. & Holyoak, K. J. (2003), 'A symbolic-connectionist theory of relational inference and generalization', *Psychological Review* **110**, 220–264.

Johnson, G. (2016), *Argument & Inference: An Introduction to Inductive Logic*, MIT Press, Cambridge, MA.

Kemp, C. (2009), Quantification and The Language of Thought, *in* Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams & A. Culotta, eds, 'Advances in Neural Information Processing Systems', Vol. 22, Curran Associates.
URL: *https://proceedings.neurips.cc/paper/2009/file/82161242827b703e6acf9c726942a1e4-Paper.pdf*

Kleene, S. (1967), *Mathematical Logic*, Wiley & Sons, New York, NY. We recommend a 2002 Dover unabridged republication of the original 1967 book from Wiley, if you cannot obtain the original book.

Konyndyk, K. (1986), *Introductory Modal Logic*, University of Notre Dame Press, Notre Dame, IN.

Markman, A. & Gentner, D. (2001), 'Thinking', *Annual Review of Psychology* **52**, 223–247.

McCarthy, J. (1980), 'Circumscription—A Form of Non-Monotonic Reasoning', *Artificial Intelligence* **13**, 27–39.

McKinsey, J., Sugar, A. & Suppes, P. (1953), 'Axiomatic Foundations of Classical Particle Mechanics', *Journal of Rational Mechanics and Analysis* **2**, 253–272.

Nelson, M. (2015), Propositional Attitude Reports, *in* E. Zalta, ed., 'The Stanford Encyclopedia of Philosophy'.
URL: *https://plato.stanford.edu/entries/prop-attitude-reports*

Newell, A. & Simon, H. (1956), 'The Logic Theory Machine: A Complex Information Processing System', *P-868 The RAND Corporation* pp. 25–63. An almost exactly similar version of this paper can be found in *IRE Transactions on Information Theory*, vol **2**, pages 61–79.
URL: *http://shelf1.library.cmu.edu/IMLS/BACKUP/MindModels.pre_Oct1/logictheorymachine.pdf*

Núñez, R., Murthi, M., Premaratne, K., Bueno, O. & Scheutz, M. (forthcoming), 'Uncertain First-Order Logic: A Classically-Consistent Logic Framework for Modeling, Quantifying, and Tracking Uncertainty Using Dempster-Shafer Theory'.

Paris, J. & Vencovská, A. (2015), *Pure Inductive Logic*, Cambridge University Press, Cambridge, UK.

Partee, B. (2013), The Starring Role of Quantifiers in the History of Formal Semantics, *in* Vít Punčochář and Petr Švarný, ed., 'The Logica Yearbook 2012', College Publications, London, UK.

Pollock, J. (1995), *Cognitive Carpentry: A Blueprint for How to Build a Person*, MIT Press, Cambridge, MA.

Pollock, J. L. (1992), 'How to Reason Defeasibly', *Artificial Intelligence* **57**(1), 1–42.
URL: *citeseer.ist.psu.edu/pollock92how.html*

Prakken, H. & Vreeswijk, G. (2001), Logics for Defeasible Argumentation, *in* D. Gabbay & F. Guenthner, eds, 'Handbook of Philosophical Logic', Springer, Dordrecht, The Netherlands, pp. 219–318.

Reiter, R. (1980), 'A Logic for Default Reasoning', *Artificial Intelligence* **13**, 81–132.

Russell, S. & Norvig, P. (2020), *Artificial Intelligence: A Modern Approach*, Pearson, New York, NY. Fourth edition.

Saldanha, E.-A. D. & Kakas, A. (2020), 'Cognitive Argumentation and the Suppression Task', *arXiv:2002.10149* .

Sloman, A. (1971), 'Interactions Between Philosophy and AI: The role of Intuition and Non-logical Reasoning in Intelligence', *Artificial Intelligence* **2**, 209–225.

Smith, P. (2013), *An Introduction to Gödel's Theorems*, Cambridge University Press, Cambridge, UK. This is the second edition of the book.

Stenning, K. & van Lambalgen, M. (2008), *Human Reasoning and Cognitive Science*, MIT Press, Cambridge, MA.

Sun, R. (2001), *Duality of the Mind*, Lawrence Erlbaum Associates, Mahwah, NJ.

Sun, R. & Bringsjord, S. (2009), Cognitive Systems and Cognitive Architectures, *in* B. W. Wah, ed., 'The Wiley Encyclopedia of Computer Science and Engineering Vol 1', Wiley, New York, NY, pp. 420–428.
URL: *http://kryten.mm.rpi.edu/rs_sb_wileyency_pp.pdf*

Zadeh, L. (1965), 'Fuzzy Sets', *Information and Control* **8**(3), 338–353.