LOGIC

A Modern Approach

Beginning Deductive Logic, Advanced

via HyperSlate™ and HyperGrader™

Larry likes Lucy. Everyone likes anyone who likes at least one entity. Does everyone like Lucy? Prove it!

Selmer Bringsjord

Naveen Sundar Govindarajulu

Joshua Taylor

and camelid art by … KB Foushée

 Published in the United States by Motalen LLC.
www.motalen.com

ISBN 978-0-692-05943-2 (e-book, edition 011519)
ISBN 978-0-692-05944-9 (Slate™ software system, edition 012516)

Preface

0.1 Student/Instructor: This is the Advanced Version!

Please note this immediately: The book you're currently reading, as indicated by its title, is an *advanced* version of beginning deductive logic (BDL) in the LOGIC: A MODERN APPROACH (LAMA; rhymes with "llama") paradigm.[1] College or university students lacking a demonstrable understanding of high-school mathematics in its full breadth (including therefore *discrete* mathematics) are advised to pursue their initial learning of deductive logic using the standard, non-advanced cousin of the present textbook. (Of course, alternatively, those without such understanding can secure it by suitable remedial coverage of high-school mathematics, and can then return here after remediation is complete and understanding is in place.) Needless to say, college-level logic instructors of classes with undergraduate students whose high-school training in math is insufficient are urged to use the standard, non-advanced relative of the present textbook. The standard textbook[2] presupposes exceedingly little on the part of learners who take it up, unfolds at a slower pace, and covers less content. The Slate software system, and Hypergrader (an online AI system that among other things automatically grades files created in Slate), can both be used for the teaching of beginning deductive logic in *both* advanced and non-advanced modes.

The phrase 'demonstrable understanding of high-school mathematics in its full breadth' is rather a mouthful, and stands in need of clarification. What do we mean? Let's first divide the phrase into two parts: *demonstrable understanding* and *high-school mathematics in its full breadth.* Now we can explain each part separately.

As to what's meant by the second part of the phrase, 'high-school mathematics in its full breadth,' that's easy: we simply refer to the branches of mathematics traditionally introduced in high school: viz., algebra, geometry, and trigonometry. Hopefully high school afforded coverage of elementary differential and integral calculus, but this isn't stricly required.[3] In addition, it would be a good thing for

[1] Abbreviating, the present book is affectionately known as LAMA–BDL**A**.

[2] LAMA-BDL

[3] It would nonetheless be a very good thing if the student embarking upon stufy of formal logic via LAMA-BDLA has digested elementary differential and integral calculus. The formal definition of a limit, which 99 times out of 100 is an early foundational concept in this calculus, turns out to involved sophisticated quantification, and thus any student that truly grasps this concept has a leg up.

students to have been taught, as is commendably done in some high schools,[4] some mathematical statistics and computer programming,[5] but neither of these is required either. In fact, the present book includes a brief introduction to logic programming.[6]

Now we come to the first part of our phrase: What is it to have, as we say, *demonstrable understanding* of high-school mathematics? By this is meant that the student is able to *prove* relevant theorems, for while the LAMA™ paradigm is based upon a number of pedagogical principles, first and foremost among them is what can be labelled the Driving Dictum:

$$\boxed{\text{If you can't prove it, you don't get it.}}$$

Using this dictum, it's easy to ascertain whether or not a student (or, for that matter, a machine/AI does[7]) really understands core concepts in, for instance, algebra. Here's a simple example, a problem early in the textbook *Algebra 2* by Charles et al. (2012) that calls for a proof: Prove that the sum of two consecutive positive integers is invariably an odd integer. We assume that our readers can quickly accomplish this.[8] There are many, many requests for proofs in this Algebra-2 textbook (and for that matter in all of its counterparts from other publishers), after this initial, easy one. Here are a few examples, each of which you should be able to establish after

---

[4]To its great detriment as a nation, and to the detriment of many young minds who deserve better, the notion that high-school graduates needn't have more than a smidgeon of mathematics (proposed e.g. in Baker 2013) is now a cancer that has taken firm root in many U.S. schools and households. Perhaps equally pedagogically pernicious is the idea that some college-bound students in the States should be encouraged to follow a route to earning college/university mathematics credit that allows them to dodge algebra, geometry, trigonometry, and calculus. The most prominent such route is one in which college-bound students simply get some elementary coverage of statistics. Unfortunately, one noteworthy consequence of this route is that it produces college graduates constitutionally unable to explain why it's not impossible that an arrow shot by an archer moves (!); see (Bringsjord & Bringsjord 2014).

[5]Notice that we refer here to *mathematical* statistics. Plenty of classes and books devoted to teaching statistics are unfortunately devoid of any mathematical rigor to speak of. As to computer science, we say only that having had some of it can't hurt. Unfortunately, K–12 coverage of computer science in the U.S. is almost invariably informal, and as such doesn't request of students that they prove anything. Given the Driving Dictum presented just below, this means that computer science in the U.S. is taught in a way that fails to guarantee genuine understanding on the part of students. This situation is made all the worse by the fact that the Advanced Placement computer science curriculum and exam is not based on either of the programming paradigms most naturally associated with rigorous proof (functional programming and logic programming), but rather on the object-oriented paradigm.

[6]The first part of this introduction is provided in §2.8.

[7]These days there's much veneration of so-called "machine learning," or just "ML," on the strength of the claim that computing machines who engage in ML learn sophisticated things (e.g. games of perfect information) with amazing speed. From the perspective of the Driving Dictum, such exuberance about ML is suspicious. For an account of what is called 'real learning,' and an argument for the proposition that ML-powered machines don't really learn anything at all, see (Bringsjord, Govindarajulu, Banerjee & Hummel 2018).

[8]Pedantic, but in the clear, crisp style that will become our bread and butter as we progress:

> **Proof**: Let n and $n+1 = m$ be two arbitrary consecutive positive integers (i.e., $n, n+1 \in \mathbb{Z}^+$). We know from basic arithmetic that $n+(n+1) = 2n+1$. Obviously $2n$ is even, since it can be divided without remainder by 2 (to leave n). But then $2n+1$ must be odd. **QED**

For a somewhat more meaningful diagnostic challenge along the same lines: Prove that the sum of the cubes of any two consecutive positive integers is odd.

some work:

- Prove that a triangle with vertices (3, 5), (-2, 6), and (1, 3) is a right triangle. (page 198)
- Prove that $n^3 - n$ is divisible by 3 for all positive integer values of n. (page 516)
- Prove this factoring formula for the sum of cubes: $a^3 + b^3 = (a + b)(a^2 - ab + b^2)$. (page 704)

The truth of the matter is that high-school mathematics, if one takes it seriously and excels, places the student in a position from which they can in turn excel in beginning deductive logic at the college level, even when introduced in the advanced and comprehensive manner of the present book. The fundamental reason, put starkly, is that this level of mathematics is filled with, indeed is directly based upon, deductive logic, and includes plenty of problems that are solved only by providing valid proofs. This can be seen not only by visiting high-school textbooks for Algebra 2, but also by perusing first-rate textbooks for high-school geometry, a very nice example being *Geometry: Common Core* by Bass & Johnson (2012), whose second chapter is entitled "Reasoning and Proof," with sub-sections mastery of which are ideal for learning deductive logic via the present textbook and its associated software.[9]

We end this section by using a convenient deductive inference schema almost always covered in high-school mathematics. The schema is known as **transposition** or **contraposition**, and allows one to deduce from a conditional statement `if` ϕ `then` ψ, where ϕ and ψ are themselves statements, this new conditional statement, the **contrapositive** of the original one: `if not` ψ `then not` ϕ. With this schema, then, we deduce from the Driving Dictum this statement:

If you do get it, you can prove it.

0.2 So, What's Our "Modern" Approach?

It's likely you're able to read this sentence because you (or a benefactor who operated on your behalf) parted with some hard-earned money. Courtesy of this transaction, you now have in your possession a book-software combination (LAMA-BDLA and HyperSlate™) that, used in conjunction with the AI grading system HyperGrader™, will introduce you to modern, formal, deductive logic in what seems to us to be a particularly, if not a singularly, effective way. This trio has been well over a decade in the making, and the construction — since new logics and applications thereof, in both cases from us and our collaborators, continue to arrive on the scene — is still steadfastly underway. Since we stand on the shoulders of many who came before us, there is a real sense in which the trio in question has been in the making for well over

[9]The sub-sections are, and we reproduce verbatim: 2.1 Patterns and Inductive Reasoning; 2.2 Conditional Statements, 2.3 Biconditionals and Definitions, 2.4 Deductive Definitions, 2.5 Reasoning in Algebra and Geometry, and finally 2.6 Proving Angles Congruent. Any high-school student who has mastered this chapter already knows a fair amount of what is covered in Chapter 2 in the present book.

two millennia. Ancient thinkers may have been greatly mistaken about the *physical* world,[10] but many of them were spot-on right about many aspects of deductive logic. Aristotle, a "mathematical spirit" who worked his wonders over 20 centuries ago, is very much alive in the pages of the present book, and its corresponding software.[11] Of course, as we soon enough explain, the *specifics* of Aristotle's logic are certainly outmoded, and have been superseded by logics that do everything Aristotle's could, and a whole lot more.[12] The "whole lot more" first arrived on the scene, as far as we can tell, in the mind of the last great polymathic genius: Leibniz. We view Leibniz as the inventor of modern formal logic, and as the primogenitor of our approach to reasoning, which is to see this reasoning as graphical and diagrammatic in nature.[13] But while our modern approach to logic is indebted to the past, it in no way asks you, a student of logic in the new millennium, to return to the past. Like you, we find it empowering to use a modern mobile smartphone, and though such a device owes a debt to 19^{th}-century telegraph technology, it makes little sense to tap out a message to your friend in Morse code. Those days are gone. And likewise gone are the days of memorizing Aristotle's syllogisms, since they are swallowed up and exceeded by even a tiny part of modern (formal) deductive logic.

Why didn't we just take the rather less strenuous route of circulating and posting a recommendation on the Internet to the effect that all beginning students of formal deductive logic (you included) ought to simply purchase one of the many *other* logic textbook/courseware pairs on the market? We cheerfully own up to the answer, one that you you will no doubt have anticipated: Because we believe you are better off learning beginning deductive **l**ogic under **a** **m**odern **a**pproach (i.e., abbreviating, under LAMA™; pronounced to rhyme with the name of the fascinating sure-footed animal), with the software system known as 'HyperSlate™,' which you now (or soon will) have, and with the third member of the trio: the online AI system able to diagnose and guide student work: HyperGrader™. We

---

[10] Case in point, Aristotle's attempt to solve Zeno's paradoxes of physical motion is entirely deficient. If interested, see (Bringsjord & Bringsjord 2014).

[11] The phrase 'mathematical spirits' is due to mathematician David Hilbert, who used it in the dramatic opening of his famous 1900 Paris lecture, delivered to direct future mathematical minds to such problems as the 23 grand ones he then proceeded to list. Hilbert was addressing the International Congress of Mathematicians. An online English version of Hilbert's address is available at http://aleph0.clarku.edu/~djoyce/hilbert/problems.html. The logician Kurt Gödel, some of whose seminal work we shall have occasion to study later, was soon able to make stunning progress on some of Hilbert's challenges.

[12] We thus for good reason regard study of Aristotle's logic to be otiose for non-historians, and hence refrain from covering Aristotle's logic in the present textbook. And excellent, short overview is provided in (Smith 2017). This overview includes presentation of a roadmap that will enable interested readers to study all the parts of Aristotle's writings that deal with logic. More specifically, all that Smith (2017) deals with can be found in (McKeon 1941).

[13] The classic starting point for the study of Leibniz and his contribution to modern formal logic is (Lewis 1960). However, this starting point, we now know, massively understates the degree to which Leibniz invented key aspects of modern formal logic. As detailed more recently by Lenzen (2004), Leibniz anticipated by approximately one-and-a-half centuries the propositional calculus (!), the inventor of which had been been presumed to be Boole (a presumption made courtesy of Boole 2010). (Boole's classic book is available free from http://www.gutenberg.org.) As Lenzen also points out, Leibniz also anticipated another part of modern formal deductive logic: *modal* logic. Modal logic is introduced in the present book (Chapter 5).

firmly believe that this approach is pedagogically superior to all the alternatives. In the book you're viewing, our approach, and HyperSlate™ and HyperGrader™, are specifically harnessed to teach you beginning (formal) deductive logic. There is plenty of important logic that isn't deductive, and of course plenty of deductive logic that's well beyond the introductory level (and plenty of informal stuff that's called "logic"; in fact the phrase 'informal logic' is sometimes used), but basic (formal) deductive logic is the absolute cornerstone of all rational thought on our planet (and for that matter on any others that might have intelligent life!) — yesterday, today, and tomorrow. Take it away, and in one fell swoop you sweep away the formal (= exact) sciences, computer science (composed, so the not-inaccurate slogan goes, of logic + electrical engineering), the physical sciences (which to the extent that they are sound, are based in the exact sciences), rigorous philosophy, and more.[14] Some of the additional things that would be destroyed by the absence of deductive formal logic are exceedingly practical. For instance, since the only way we have, and will ever have, to establish that software is correct and will behave in the future as we wish it to, is to deduce that the software in question is correct using formal logic, the ubiquity and centrality of software in human life confirms the importance of what you are about to learn. While at present our daily lives are only sprinkled with the appearance of autonomous AIs/robots, the time is fast coming when our existence will be intertwined with, and increasingly dependent upon, such artificial agents, and we will want to be quite certain that these agents will behave correctly. But since the behavior of these agents will hinge in turn on the correctness of the computer programs that regulate them, and since, again, such correctness can be established only via the use of formal, deductive logic, our lives will in a very real sense rest in the hands of logic. If no one has told you this before, that is most unfortunate, but now you know. Good thing you're reading what you're reading.[15]

We do gladly concede that there is much value in *in*formal logic (and indeed Slate has been invented, refined, and implemented to support humans in doing informal reasoning), but the power of informal logic is in our view enjoyed in earnest only by one who *first* understands the formal, rigorous machinery that gives the informal level the matchless power it has. Learning informal logic before learning the basics of formal logic is by our lights like learning to count before learning such arithmetical facts as that (a) 3 is greater than 2 and that (b) 2 is greater than 1: Perhaps a child can *in some sense* count by mouthing "1, 2, 3" without a prior understanding of and implicit assent to (a) and (b), but on the face of it such a pedagogical progression is backwards. The child here is just making sounds without real understanding. Likewise, in the realm of logic, the student of informal logic shouldn't be taught to mouth assent to such things as that from "All men are mortal" and "Matt is man" it follows that "Matt is mortal" without understanding that this is true because the *underlying formal pattern*

[14] You would also sweep away rigorous economics, since without formal logic we lose real analysis, game theory, decision theory, and so on, but we have found that our readers are profoundly mixed on whether or not it would be bad to lose economics.

[15] For further reading, see (Bringsjord 2015, Govindarajulu & Bringsjord 2015, Arkoudas & Bringsjord 2007).

$$\begin{array}{c} \text{All } As \text{ are } Bs. \\ e \text{ is an } A. \\ \hline e \text{ is a } B. \end{array}$$

is valid. This particular formal structure was seen to be at the heart of human reasoning and rationality by Aristotle over (to say it again) two millennia ago. What makes the pattern a *formal* one, and therefore serves to classify the pattern as part of formal logic? You doubtless know the answer: because no matter what you substitute in for the collections A and B and the entity e, the conclusion will indeed follow from the two premises. Unless the student comes to understand the domain-independence of formal logic, no real learning of formal logic has happened. This implies that those educational games and activities that seek to foster learning of logic and mathematics by "concretizing" problems and problem-solving, in the absence of cultivating abstract, domain-independent capability, are at best profoundly incomplete and misleading, and profoundly destructive at worst. A five-year-old who is told such things as that the small yellow wooden circle can go either on this square or that square (where e.g. the squares are cells in a grid that can either be filled or vacant), and then dutifully places the circle in one of the two allowed locations, doesn't need to have a genuine and general understanding of Boolean disjunction, which applies as nicely to simple games as it does to advanced mathematical physics. Without such understanding, our five-year-old really doesn't know what he's doing. And worse, if the human avoids sustained training in formal logic, he or she can fail to do any better at 55 than at five.[16]

Summing up: Both because of the nature of our approach and courseware, and the importance of what you will be studying, you (and perhaps your benefactor) have made a wise investment.

But why *specifically* is the LAMA™ paradigm, and specifically HyperSlate™ and HyperGrader™, better? We have already hinted at our answer above, and in the title of the present book: because this pair, falling squarely in the pardigm in question, takes a *modern* approach to teaching logic. But what does *that* mean? The hallmarks of what we call a "modern" approach, in addition to the Driving Dictum introduced above (again: "If you can't prove it, you don't get it"), include the following ten properties:

1. *Use and Learning of Automated Reasoning Technology.*
2. *Coverage of Large Space of Modern Logics, and Three Families Thereof.*
3. *Use of Flexible, Graphical Workspaces for Human-Computer Collaboration.*
4. *Learning of Logic via Playing* ***Truly*** *Challenging Games on Planet Earth — Rather Than The Sort of Simple Games Traditionally Targeted by AI.*
5. *A Grounding in the Cognitive Science of Human Reasoning and Decision-Making.*
6. *LAMA™, HyperSlate™, and HyperGrader™ Are Based on Deep Mathematical Theories of Cognition and Computation.*

[16] Beware, then, of educational games to *purport* to cultivate, in those that play them, skill at formal reasoning.

7. *Extending Formal Deductive Logic Beyond the Merely Symbolic/Linguistic to the Pictorial/Diagrammatic.*
8. *Extending Formal Deductive Logic to Bridge to Uncertainty/Probability in* ***Inductive*** *Logic.*
9. *Seamless Innteroperability and Integration with a Corresponding Type of Logic-based Computer Programming.*[17]
10. *Logic as the Emerging Foundation for Rigorous Science.*

We now unpack each of these ten properties in turn, albeit very briefly in each case.

0.2.1 The Hallmarks of Our Modern Approach, Unpacked

1. *Use and Learning of Automated Reasoning Technology.* Today, humanity has computer programs able to do things which in the past were the sole province of human beings. It has probably not escaped your notice, for example, that the best checkers player, chess players, and Go players on our planet are no longer human but rather are computers; and perhaps you know that recently even in a game like *Jeopardy!*, which puts a premium on the use of human language, computers can perform at an impressive level.[18] One of the things that computing machines can now do that not all that long ago used to require rarefied human cognition is proving theorems. Machines can't yet match humans in the proving of *robust* theorems (in fact, they aren't even close; we'll get to this issue later, when you know a bit more formal logic), but nonetheless it's undeniable that machines can now routinely prove theorems that not all that long ago required significant human ingenuity.[19] HyperSlate™ and HyperGrader™ have been engineered, and this textbook written, to leverage this progress in support of student learning.[20] In short, and put bluntly: No student these days should learn logic without specifically learning how to collaborate directly and deeply with intelligent machines able *themselves* to reason deductively. Put in analogical terms: the point is little different than the fact that no pilot should learn how to fly a plane without learning specifically how to collaborate with smart software in doing so (e.g. the software that controls autopilot functionality), or that no engineer should be trained without learning how to use a calculator and beyond. One of the most amazing (and by our lights, distressing) facts about American post-secondary education is that, even today, as the field of artificial intelligence (AI) is thriving and marching steadily upwards, many students at high-schools, colleges, and universities in the U.S. are learning logic using only paper and pencil.
2. *Focus on Modern Logics, and Three Families Thereof.* Lots of logic textbooks spend considerable time on logics like Aristotle's **Theory of the Syllogism**, which we alluded

[17]The type of computer programming is a generaliztion of what is called *logic programming*: viz. **pure general logic programming**, or just 'PGLP' for short.

[18]IBM's Watson system beat the best human *Jeopardy!* players; the system is explained at the "Scientific-American" level in (Ferrucci et al. 2010). For commentary on Watson in connection with the AI HAL 9000 in Kubrick's *2001*, commentary offered *before* the actual competition that Watson won, see (Bringsjord, Clark & Taylor 2010).

[19]Luger (2008) makes the point that automated theorem proving is a crucial component of the sort of human-level AI that many are after, and provides a new overview of the earliest days of automated theorem proving.

[20]LAMA™ is also attuned to progress in what is known as *model finding*, which we see as included within automated reasoning, at least for the family of extensional logics.

to earlier. This simple logic debuted, as has been noted, rather long ago. It's always good to know one's history, but the fact is that logic has advanced rather far in the 2000-plus years since Aristotle's small (but seminal) system was presented. In fact, as Glymour (1992) has elegantly explained, Aristotle's primitive logic was insufficiently expressive to rigorously explain why even ancient Euclid's reasoning (in showing such propositions as that the sum of the interior angles of a standard triangle is 180 degrees) is so compelling. To hammer home the point again: The fact is, when students start learning physics at the college level today, they don't go back and learn what the ancients were doing. The reason is simple: the work of the ancients has been forever superseded by a host of innovations. This book includes coverage of modern logical systems that have superseded old logics like those presented by Aristotle and others. More specifically, we cover modern **classical extensional logic**, invented in the 20th century to represent and rigorize mathematical reasoning (and now used to do lots of real-world things; we'll get to that); and we also cover **basic intensional logic**, which also entered the picture in the 20th century, but for different purposes: to systematize the concepts of possibility, necessity, knowledge, belief, obligation, and so on, and to systematically represent what is meant by sentences in **natural languages** like English.[21] These two families of logical systems are unified by the use of HyperSlate™, which allows the user to select a particular logical system when starting a new workspace. As far as we know, HyperSlate™ is the only software in existence that allows a human to pilot an intelligent computer which provides, at once, access to many different logics at the same time.[22]

We said there are *three* families of logics introduced herein. What's the third? It's **inductive logic**, and while the the present version of this textbook goes very light on this third family, the some key information is provided to the student. The reason for the light coverage is that inductive logic is non-deductive, and this textbook, as its title suggests, is intended to serve mostly as an introduction to beginning deductive logic. Nonetheless, it's important for the student of elementary deductive logic to understand some inductive logic, if for no other reason than to see that beginning inductive logic is quite a different beast.[23]

3. *Use of Flexible, Graphical Workspaces for Human-Computer Collaboration.* Real humans doing real deductive reasoning in the real world don't simply list out, line-by-line,

[21] Natural languages stand in contrast to **formal languages**. The latter, examples of which you will soon get to know quite well, have a fully specified alphabet and grammar. A programming language is for instance a formal language.

[22] The deductive logics you will be learning each fall into one of two main families just mentioned: **extensional** or **intensional**. Within the first family, coverage is provided first of logics that have no quantification (viz., **propositional calculus** and **pure predicate calculus**), and then of **full first-order logic** (with subsequent coverage of **second-order logic** and beyond: see Chap. 4). Within the family of intensional logics, this book introduces **propositional modal logic**, **propositional deontic logic**, **propositional epistemic logic**, and finally **quantified modal logic**. All of these logics are finitary in nature: they don't permit expressions that are infinitely long. Yet some of the most interesting and powerful logics are infinitary ones. E.g., the logic $\mathscr{L}_{\omega_1\omega}$ allows conjunctions (such as '1 is greater than zero and 2 is greater than zero and ...'), and also allows proofs that are infinitely long. Summing up, then, we can picture the Universe of Logics as shown in Figure 2.

[23] Very briefly, deductive reasoning, when valid, can't possibly be called into question, because it partakes of inference schemata, such as *modus ponens*, that are indubitable and invulnerable.[24] This schema says that if we know that ϕ implies ψ, and we have ϕ, it absolutely, positively *must* be the case that ψ. Inductive reasoning, in stark contrast, includes such inference schemata as that if it's overwhelmingly likely that if ϕ holds so does ψ, and ϕ does in fact hold, it follows (inductively!) that ψ holds. Chapter 9 is where a glimpse of inductive logic is given.

sequences of formulas when doing so. Yet if you look at standard approaches to teaching deductive logic, and the software touted in these approaches, you would think that real-world deductive reasoning is all and always nothing more than a one-dimensional list of formulas, with a bit of indentation. No; the real world is richer than this, beyond belief. If you look at the notebooks, scratchpads, blackboards, whiteboards, and even computer screens of highly effective human reasoners, you rarely find the kind of rigid text-only linear lists of formulas that have been, and unaccountably continue to be, the hallmark of most if not all other logic books and courseware. Slate is a step in the direction of our vision for a cognitively plausible workspace that is tremendously open-ended and genuinely intelligent and collaborative. In Slate, formal proofs are directed hypergraphs whose vertices contain sentences and, often, other relevant information. (Don't worry in the least if you don't presently know what a hypergraph is. All will be gradually explained in due course!) At the same time, Slate can be used to constuct trees. This means for instance that instead of the traditional harsh division seen in standard approaches between a linear step-by-step proof on the one hand, and entirely different formats for semantics (e.g., truth tables) on the other, Slate's graphs can be used to express proofs *and* express semantics in the form of trees. In addition, in Slate, formal interpretations or models can also be constructed in graphical form.

4. *Learning of Logic via Playing **Truly** Challenging Games — Rather Than the Sort of Simple Games Traditionally Targeted by AI.* The LAMA™ paradigm includes a commitment to learning formal logic in part by playing games; but not easy (or, for that matter, so-called "serious") games, and not games on which mindless trial-and-error and plug-and-chug can secure success. Educational games under the LAMA™ approach must be sufficiently hard and, in keeping with the Driving Dictum, success in LAMA™ games must require the gamer to demonstrate mastery by justifying the correct answers given, or moves made.[25] LAMA™ games are provided separately from the present LAMA-BDLA textbook, and separate as well from HyperSlate™ and online intelligent software that complements HyperSlate™ (e.g., Hypergrader™). It is recommended that the first game students play in the LAMA™ series of games is Catabot Rescue 1, Beginner.[26] For more on genuinely challenging reasoning games, see (Govindarajulu 2013).

5. *A Grounding in the Cognitive Science of Human Reasoning and Decision-Making.* LAMA™ is based on what decades of research in cognitive science and artificial intelligence has informed us about human reasoning and decision-making. Piaget believed that neurobiologically normal human beings provided with a general and standard K-through-12 education (and with sufficient nutrition, shelter, and so on) would effortlessly become competent deductive reasoners and problem-solvers at a level beyond

[25] Some readers may be familiar with logic puzzles in which a story that frames the problem is first given, then a set of clues, and the challenge is to e.g. pin down who did what when with what. Before the advent of apps on smartphones and tablets, these puzzles were seen almost exclusively in magazine-style hard-copy form. These sorts of puzzles have also been a mainstay in certain standardized exams, e.g. the GRE and LSAT. Those who learn sufficient formal logic thereby cultivate a capacity to excel on these problems (Bringsjord & Bringsjord n.d., Bringsjord 2001), but since the customary presentation of these problems lacks a demand that the supplied answer be justified by a proof, the problems in question, in the LAMA™ paradigm, are far from a pedagogically perfect fit. Readers/students who enjoy such puzzles are therefore encouraged to try to prove that their moves and answers are correct, using HyperSlate™.

[26] For more information, visit http://www.catabotrescue.com. Human prowess at Catabot Rescue, which can become quite remarkable, has a number of interesting implications for education and cognitive science; for more on this, see (Govindarajulu & Bringsjord 2016).

full first-order logic (e.g., see Inhelder & Piaget 1958).[27] Unfortunately, Piaget was a little too sanguine. While he carried out numerous experiments showing that many young people do indeed develop an impressive capacity for *some* deductive reasoning, even Piaget himself conceded toward the end of his career that his subjects were not representative of the general population. LAMA™ is based on the *neo*-Piagetian view that *if* students are appropriately trained (courtesy of books and software of the sort that you now own and/or have access to!), they can indeed become first-rate deductive reasoners. There is a lot of evidence supporting this view, including recent neuroscientific evidence.[28] As to theories of reasoning directly at odds with the LAMA™ paradigm, there are some. The best-known of these is so-called *mental models theory* (MMT), originally introduced by Johnson-Laird (1983). A core tenet of MMT is that the human mind isn't fundamentally equipped with the capacity to reason deductively in a manner aligned with classical deduction as seen in formal logic and mathematics. Johnson-Laird and colleagues are quite aggressive and open about their disdain for formal logic (e.g. see Khemlani, Byrne & Johnson-Laird 2018). For work in the cognitive science of reasoning that accords with the LAMA™ paradigm, and hence is at odds with MMT, the reader can consult the truly excellent work of Lance Rips (**?**, 1994).[29]

The LAMA™ paradigm in addition reflects the theoretical position, in the cognitive science of comparative reasoning and decision-making, that reasoning and decision-making at the human level is qualitatively distinct from such activity that is merely at the animal level. The essence of the mastery of formal logic is command of abstract structures (e.g., the inference pattern alluded to in the previous paragraph) that, as the empirical research in animal cognition shows, no animals can grasp. A clever canine can come to know that if the ice on a pond looks a certain way, walking on that ice is a bad idea, and so if Rover comes upon such a pond, it clicks for him that walking on it is a bad idea, and he remains on the firm ground of shore.[30] But the dog will never grasp the abstract inference schema operative here, which is none other than *modus ponens*. The LAMA™ paradigm, put simply, is one suitable for humans, not nonhuman animals. For more on the cognitive science of human versus animal reasoning and decision-making, see (Penn, Holyoak & Povinelli 2008), which is written in opposition to Darwin's (1997) exuberant but misplaced praise for the reasoning

[27] Piaget held that humans develop cognitively through a serious of stages. In the fourth stage, human cognition is distinguished by a use of formal deductive logic in the third logical system we will be studying herein: namely, full first-order logic. See note 22.

[28] For an account of, and evidence in favor of, neo-Piagetianism, see (Bringsjord, Bringsjord & Noel 1998, Rinella, Bringsjord & Yang 2001, Bringsjord & Yang 2003). For neuroscientific evidence, see (Mukherjee 2009). On the other hand, there is *also* a mountain of evidence that if students *don't* receive suitable training in formal logic, they often rely on forms of inference that are invalid. This is of course entirely unsurprising. Without effective mathematical education of the right sort, ostensibly bright humans naturally believe such things as that when an arrow is in motion after being shot, it occupies a series of discrete places at each of the discrete moments composing the interval of time consumed by its trajectory. But such a belief implies that motion is impossible. Given the differential and integral calculus, we know this view is silly, and indeed self-contradictory (for a full discussion, see Bringsjord & Bringsjord 2014).

[29] The inquisitive reader may wonder why, specifically, Rips's work is in line with LAMA™, whereas MMT is inconsistent with this paradigm. The answer can be barbarically (but hopefully informatively!) encapsulated by way of reference to the proof pattern known as **indirect proof** (also known as **proof by contradition** and **reductio ad absurdum**), as follows: Rips's position is that human beings are capable of for instance solving reasoning problems that call for inference patterns such as this, which is at the very heart of the LAMA™ paradigm, while MMT maintains that the natural or "logically untrained" human mind can make no sense of this pattern whatsoever.

[30] This is an example given by Darwin himself, in an attempt, a failed one, alas, to show that the canine mind isn't qualitatively different that yours and ours; see (Darwin 1997).

power of dogs.

6. *LAMA™, HyperSlate™, and HyperGrader™ Are Based on Deep Mathematical Theories of Cognition and Computation.* This isn't the place to explain the mathematics in question; we convey only the core idea here in rapid, intuitive fashion; sedulous readers are invited to carry out further investigation.[31] Our quick, intuitive explanation is based on Figure 3.

   This figure takes the agent and HyperSlate™ workspace, together, as capable of solving problems supplied from some outside source; we can refer to this source as the **environment**. The environment supplies problems to be solved by the duo of human and HyperSlate™; this collaboration, note is indicated in Figure 3. In the abstract, challenges are in the form of a question; this is also shown in the figure. The question is usually whether some given statement p, expressed as a formula ϕ_p, can be proved from some given set of statements, expressed as a set Φ of formulae. When the given information is absent, we have a case where Φ is the empty (= null) set $\emptyset$. In that case the challenge is to prove ϕ_p without the benefit of some starting premises or axioms.

   Sometimes the human agent in this scheme can operate mechanically, with minimum ingenuity. This is the case when the human operating with HyperSlate™ can follow an algorithm. In the course of this book, students will acquire a number of algorithms which, when followed (or executed), allow significant progress to be made toward a solution to the problem confronting the student. This said, even beginning deductive logic includes challenges that cannot be met by *any* algorithm. (The discovery of this fact was momentous in the history of formal logic and computation, and is discussed later.) We'll get to such challenges soon enough; they are part of the fun of studying logic!

7. *Extending Formal Deductive Logic Beyond the Merely Symbolic/Linguistic to the Pictorial/Diagrammatic.* One of the most astounding things about the field of logic is that while it has been going strong for well over two millennia, only the tiniest speck of the discipline during that rather long stretch has been devoted to the study of reasoning over *pictorial* information. Everything other than this speck has been taken up in reasoning over representations that are *symbolic* (or, as it's sometimes said, *sentential* or *linguistic*). In LAMA™, logic is based on representations that are symbolic *and* pictorial, and in Chapter 8, such **heterogeneous logic** is introduced, and students are taught how to construct proofs that employ both diagrams and formulae.

8. *Extending Formal Deductive Logic to Intersect with Uncertainty/Probability.* Standard practice in the coverage of high-school mathematics is to try to distinguish between deductive and inductive reasoning.[32] Deduction is said to be a process of inferring "facts" from "facts" in such a way that the inferences are certain. In constrast, to inductively reason to some proposition from some information means that that proposition

[31] Where should they start? They will need to learn formal deductive logic one "notch" beyond a standard introductory logic course, but full assimilation of the present book, which as we have noted is advanced, will suffice. Assimilation of only the *non*-advanced introductory content under the LAMA™ paradigm would require a second course in formal logic, often called "intermediate" formal logic. Autodidacts can e.g. consult either or (better) both of (Boolos, Burgess & Jeffrey 2003) and (Ebbinghaus, Flum & Thomas 1994). Second, they will need a solid upper-undergraduate course in the theory of computation; a candidate book is (Lewis & Papadimitriou 1981). Third, they will need to become familiar with what are today called **Kolmogorov-Uspenskii machines**; one should begin with the *locus classicus*: (Kolmogorov & Uspenskii 1958).

[32] Note e.g. the first section heading '2.1 Patterns and Inductive Reasoning' given in the progression reproduced in footnote 9. In addition, online mini-lessons available from The Khan Academy attempt to distinguish between inductive versus deductive reasoning (e.g., see this tutorial).

is only for instance *likely* given the information, even if the information is 100% true. This distinction is invariably made in such a way that an insuperable divide is formed between the two modes of reasoning, with logic pushed completely over to the deductive side of the house, and things like statistics and probability are left to tend to the other side. It's amazing but true that this unwise and inaccurate division is made as well by seemingly competent scientists and engineers. The problem with the division is simply that logic covers *both* deductive and inductive reasoning; it's just that the latter type of reasoning is covered not by standard deductive logic, but rather by *inductive* logic. LAMA™ is based on a treatment of logic that unifies deductive and inductive reasoning, and the kernel of the unification is explained in Chapter 9, in which some simple inductive logics in the LAMA™ style are presented. (Naturally enough, since the present book is an introduction to deductive logic, the treatment of inductive logic is minimal.) It is these inductive logics that allow the deductive-inductive unification to be accomplished.

9. *Seamless Innteroperability and Integration with a Corresponding Type of Logic-based Computer Programming.* Some of our readers will have some familiarity with computer programming; and some of these readers will specifically be acquainted with the fact that computer programming can be pursued within fundamentally different (indeed, in many ways radically different) paradigms. For example, in the **imperative** or **procedural** paradigm, which is unfortunately almost without exception what is taught to youth in the pre-secondary and secondary grades in technologized nations, the basic idea, put simply, is that a computer program consists in a series of instructions to do this, then that, then this, with the possibility for branching and iteration included in this step-by-step approach. For instance, a procedural program to compute the exponentiation function $exp(n,m) = n^m$ on positive integers might codify the following basic do-this-then-that algorithm:[33]

   A: Set the value of `result` to n.

   B: If $m = 1$, PRINT `result` and STOP.

   C: Set `counter` to 2.

   D: Multiply `result` $\times n$; and set `result` to this product.

   E: If `counter` $= m$, PRINT `result` and STOP; otherwise, increment `counter` by 1, and go to line D.

   Let's call this algorithm 'A,' and let's assume that some procedural program P_A that regiments the algorithm A is available. Now suppose that Johnny is wondering whether 2^5 equals 32; perhaps he has received this question from his teacher: "Is 2^5 equal to 32?". One possible move in response is for Johnny to follow A by hand and see what he gets, and then give his answer to the teacher. Another possible move is to give as input 2 and 5 to P_A and see what he gets as output. If he doesn't make a mistake in his manual implementation of A on this input, and if the program P_A doesn't malfunction (and he has made no mistake when suppplying the input to the program), Johnny will reply with a Yes to his teacher. However, what if the teacher is just a bit more demanding? Specifically, what if she says to Johnny: "Are you sure? How do you *know* that the correct answer is Yes?" Now Johnny has a problem. He will need to move beyond the procedural paradigm in order to answer the teacher's follow-up question. He will need to somehow *prove* that the affirmative answer is correct, and

[33] This is a very naïve algorithm, but for expository purposes is adequate.

the only way he can do that is to turn to logic. The beauty and power of the approach to programming that LAM™ fits is that when the computer provides an answer on the strength of executing a **pure general logic program**, a proof that that answer is correct is automatically provided as well. This is so because the answer is produced via automated deduction, and the deduction to produced the answer is in fact a proof, one that is retained, checked, and verified. Such programs are at the heart of what Bringsjord calls **pure general logic programming**, or just PGLP for short.

PGLP is an extension of the programming paradigm known simply as 'logic programming.' Logic programming is a programming paradigm is one in which a program that captures the exponentiation function includes two components: namely, one consisting of a set Φ of declarative formulae that formally define the positive integers and the relevant arithmetic (addition, multiplication, exponentiation) on them; and, secondly, a process that searches for a proof of an given individual formula ϕ from Φ. (In the case of Johnny, ϕ would of course express the informal statement that two raised to the fifth power is 32.) In the **pure logic programming** paradigm for computer programming, only the first of these components is something humans need to bother themselves with. Unfortunately, pure logic programming is something that's not achieved by the most popular "logic progrmmaing" language, Prolog. The reason is that in using Prolog, humans must in fact concern themselves with the second component. Beyond pure logic programming is what we have alluded to already, just above: again, **pure general logic programming**, or again, simply the acroynm 'PGLP.'. PGLP includes the purity that inheres in being able to blithely ignore the second component, *and* includes the formalisms that allow the human programming to generalize away from standard first-order logic and fragments thereof to *any* formal logic obeying certain conditions. Given that there are logic-programming languages that are less expressive and hence less powerful even than Prolog (e.g., Proplog, which is has no provision for quantifiers and identify), we have the following basic progression $\mathscr{P}$ as what summarizes our discussion:

Figure 1: The Logicist Programming Progression $\mathscr{P}$

PGLP (pure general logic programmg)

$\uparrow$

pure logic programming

$\uparrow$

Prolog (logic programming)

$\uparrow$

Proplog etc. (restricted logic programming)

Now, what is the relationship between the LAMA™ paradigm, specifically including the present book, and the progression shown in Figure 1? The present book, LAMA-BDLA, will position those readers who assimilate it to move on directly to a deep understanding of the entire progression $\mathscr{P}$, and to corresponding skill at using logic-programming languages like Prolog. However, notice that we say that the reader will be in *position* to acquire deep understanding. We don't say that assimilation of LAMA-BDLA gives the student all that is needed. The reason is that the key theorems the underlie logic programming and Prolog aren't covered in the current version of LAMA-BDLA.[34]

10. *Logic as the Emerging Foundation for Rigorous Science.* Finally, our modern approach reflects the fact that formal logic is the *lingua franca* of all intellectual pursuits that strive for true rigor. As an example, physics, at the start of our new millennium, is being gradually expressed in formal logic; this is true, for instance, for quantum mechanics and special and general relativity (e.g. see Andréka, Madarász & Németi 2007, Andréka, Madarász, Németi & Székely 2011). A second example is one we have already alluded to above: The verification that software and software-based beings (e.g., robots) will behave as we desire them to behave is advancing on the strength of the only effective resource we have to meet such a challenge: formal logic. We believe that students of beginning deductive logic must be given a sense of the connection between formal logic and rigorous science and engineering. This textbook reflects this belief.

Figure 2: The Universe of Logics

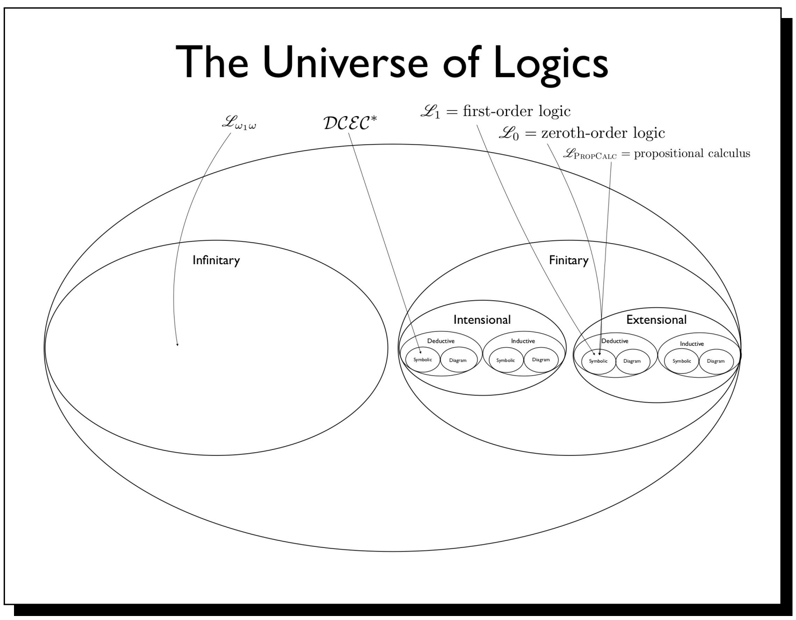

We present now a simple, anticipatory example that briefly shows the first five of the properties just enumerated, in action.

0.3 The Hallmarks in Action in an Example

To begin, take a look at Figure 4. This figure shows an argument, built in a **graphical workspace**, where that argument is expressed as a hypergraph, for a conclusion

[34] E.g., Herbrand's Theorem is fundamental to logic programming, and is not covered in the current version of LAMA-BDLA.

Figure 3: Pictorial Encapsulation of the Underlying Mathematics of LAMA and Slate

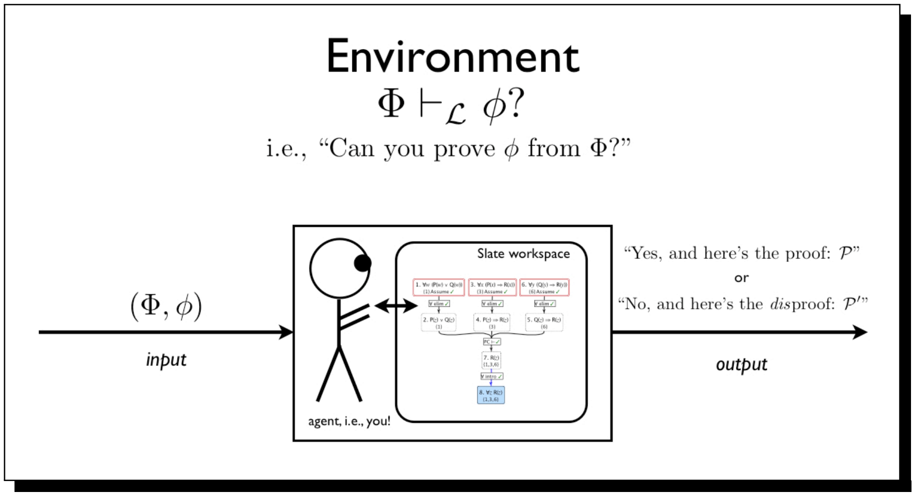

that we claim follows conclusively from the three shown premises. The graph has five nodes; three are labelled `PREMISE1`, `PREMISE2`, and `PREMISE3`; each of these, as you might guess, contains a premise. A fourth node is labelled `CONCLUSION`, and naturally enough contains the conclusion of the argument. A fifth node, rather more mysterious, contains this: `FOL` ⊢ X . We'll return to the meaning of this information momentarily. Study the graph a bit more. Do you agree that the conclusion, which expresses the proposition that Larry and Lucy like each other, can be correctly deduced from these premises?

Perhaps you aren't sure how to answer, because the phrase 'can be correctly deduced' is a bit unclear to you. (It *should* be at least *somewhat* unclear to you, since making it clear is one of the reasons for the remainder of the present book!) We can begin to fix that immediately. What we mean is this: Is it true that if the three premises are assumed to be true, then it *must* be true that Larry and Lucy like each other? This is just another way of asking whether mutual liking between Larry and Lucy can be correctly deduced from the three premises, because any time a statement can be correctly deduced from a set of premises, that statement absolutely, positively *must* be true if the premises are true.

We assume that your answer to this question is in the affirmative. (If you disagree, you haven't done enough thinking about the figure. So you might wish to return to it now before proceeding.[35])

Now, let's see how the first five properties above, all hallmarks of LAMA, are put into action in connection with this example. Giving each of these five properties mnemonic labels, and keeping them in the order in which they were presented

[35] If, after sustained thinking about whether, given the premises here, Larry and Lucy must like each other, you fail to apprehend that indeed they must, you should probably find a fellow human fluent in English, and in conversation with him or her make sure you understand in this case what the English means.

above, we can list them now economically as follows. After the shorthand list, we deploy the label in our discussion, bolding each occurrence of a label.

1. **automated reasoning**
2. **modern logics**
3. **graphical workspaces**
4. **cognitive science and AI**
5. **underlying mathematical theory**

First, it should come as no surprise to you that despite whatever other powers Slate has, it doesn't have the power to understand all of English as easily as a human can. Slate does understand many formal languages, and advanced versions of Slate do understand parts of English, but here in this introductory book we will be working with a version of Slate that doesn't automatically convert English into formal logic. More generally, we will be working with a version of Slate that doesn't automatically convert sentences in — what we have already called — **natural languages** (a class that includes not only English, but German, Norwegian, Japanese, etc.) to formulae in formal logic. Given this, for us in this book, part of the point of studying formal logic is to gain an ability to cast important parts of a natural language (in our case: English) into formal logic. And specific to the case at hand, we have therefore taken it upon ourselves to convert the English you see in Figure 4 to formal logic.

It turns out that though Aristotle was a smart chap, his logic (which you'll recall is The Theory of the Syllogism) is not up to the challenge of representing the argument about Larry and Lucy seen in Figure 4. But one of the **modern logics** that Slate has facility with, aforementioned **full first-order logic** (or, for short, **FOL**), has little problem with our specimen involving Larry and Lucy. If you now look again at Figure 4, you'll see that three arcs flow into that mysterious node that contains a red X immediately after this:

$$\text{FOL} \vdash$$

What on earth does this mean? Well, we just introduced you to the abbreviation 'FOL.' What you know at this point is that this is an abbreviated name for a particular logic (first-order logic). (Of course, you don't know anything significant *about* FOL, but that's fine; this is after all just the Preface.) What Slate as an **automated reasoner** is saying here via its red cross after FOL $\vdash$ is this: "Dear User, the argument that you have given me is not a valid deduction in the logic FOL. Sorry."

However, when we translate the English into FOL for Slate, the system is smart enough to instantly declare that the conclusion can indeed be deduced from the three premises. This declaration is shown in Figure 5. Notice specifically that now Slate has presented an encouraging ✓; this signals that the proposed deduction is indeed valid in FOL. Of course, you don't yet know exactly what the funny symbols that you now see in the figure mean; that understanding will come. You will simply have to trust us at this early point that Slate is indeed right that the deduction is valid, on the strength of its **automated reasoning** power. Later, after you've learned some formal logic, you will yourself be able to prove that Slate is indeed correct. And after gaining a facility to verify that Slate is right when it issues an oracular declaration,

Figure 4: An Argument for Mutual Liking Between Larry and Lucy

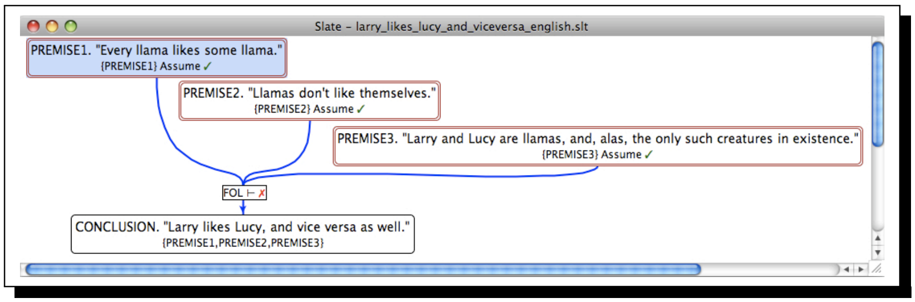

you will come to see that Slate is pretty close to infallible. Slate has considerable intelligence, and as we said above, in our modern approach to learning logic, the student learns by interacting with a machine "being" that itself has sub-human but nonetheless considerable — and helpful — intelligence.

Figure 5: Slate Confirms Mutual Liking Between Larry and Lucy

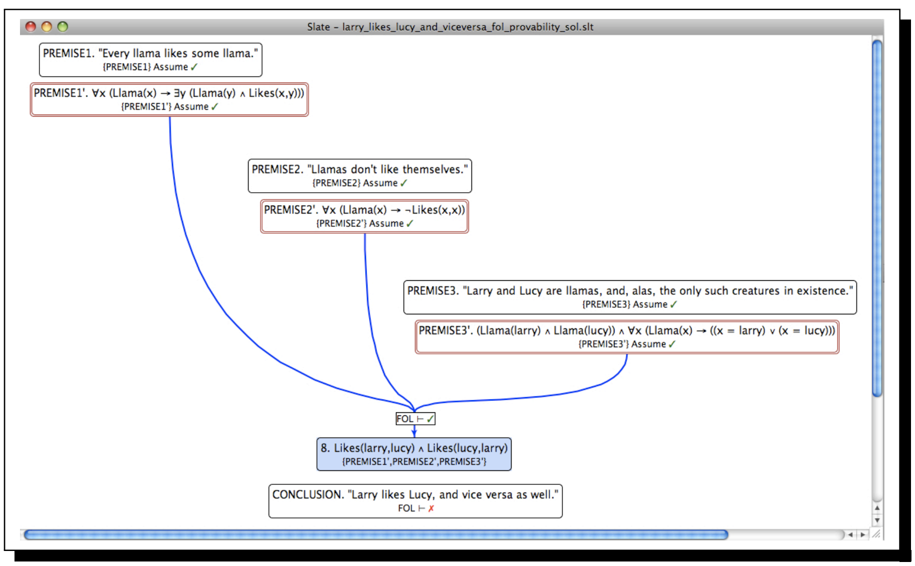

Notice that in the course of our discussion, the third property on the list above, that is **graphical workspaces**, has been used. Well, even if you haven't had any formal logic before, it's a good bet that you have been exposed to arguments and proofs that are both textual and linear in form, rather than what you have seen in Figures 4 and 5. In this old, linear, purely textual style, one sees things like this:

	(1)	Llamas don't like themselves.	
	(2)	Larry is a llama.	
∴	(3)	Larry doesn't like himself.	(from (1) & (2)

We have already said that when you inspect the notebooks of sophisticated human problem solvers, you rarely find anything as stark and one-dimensional as this kind of sterile progression. In general, line-by-line textual proofs reflect a slavish conformity to formats sometimes seen in computer science, where the old and venerable schemes for understanding the essence of computation do include such rigid ones at the **Turing machine**,[36] and where a finished computer program often does take exactly the form of a line-by-line progression of text. In addition, there is a longstanding tradition in the teaching of elementary, formal logic to ask students to proceed in rigid, line-by-line form, using indentation when necessary to indicate sub-proofs. But in learning logic in this day and age we shouldn't be limited by such straight-jackets. And fortunately, with Slate to help us, we aren't. In fact, the great news is that it can be proved that Slate is powerful enough to express any of the old-fashioned line-by-line proofs, should anyone want to work in the antiquated style. (It can also be proved that other graph-based formats for describing deduction can *also* be expressed in Slate.[37]) In the chapters to come, a lot more will be said about, and done in, Slate's workspaces.

What about the remaining attributes in the quintet? Well, in our little example we have in fact seen in action, at least implicitly, the fourth and fifth attributes. The ability to translate the English statements contained within the nodes of the graph shown in Figure 4 into the formulas shown in the nodes of the graph shown in Figure 5 was of course required. You don't yet have this ability, but **cognitive science** tells us that that ability can indeed be cultivated, and of course the adventure you are about to take is designed to achieve exactly that cultivation. And what of the fifth attribute? Here you will have to trust us for the time being that anyone able to move from Figure 4 to Figure 5 in Slate has indeed conformed to our **deep underlying mathematical theory** of problem solving (which, as you'll recall, is at least intuitively encapsulated in Figure 3). That said, we do now briefly describe one algorithm that we executed in order to convert the English version of `PREMISE3` to some of the symbolization in `PREMISE3'`. We conceive of the algorithm as receiving **arguments** as input, and producing **values** as outputs. The particular algorithm in question receives an English statement s of the form "*Name* is an *R*"; and it runs as follows.

1. Convert *Name* to a constant c in the logic in question. (In this case, the logic is FOL.)
2. Convert *R* to a relation symbol R in the logic in question. (In this case, the logic, again, is FOL.)
3. Assemble the formula $\mathsf{R(c)}$ from the results of 1. and 2.; and return this formula.

[36] Introduced by Turing (1937), and thereafter named in his honor. These machines are briefly described in Chapter 1.

[37] E.g., in §2.7.2.1 it's shown that so-called **resolution graphs**, which use resolution inference schemata, can be easily and accurately expressed as Slate hypergraphical proofs.

Of course, this is only one "building-block" algorithm that we ran in order to produce the symbolization shown in Figure 5 from the raw English shown in Figure 4, but this is after all only the Preface, and we thus rest content with your having but an intuitive sense of what's involved in moving from all of the English shown in the former figure to all of the logical notation shown in the latter one.

We anticipate this from the alert reader: "Ok, but what about the *ninth* property of a modern approach to deductive logic?" In what sense and in what way is the LAMA paradigm reflective of your claim that formal logic is the basis of even rigorous physical and natural science? Good question. At this point you will need to take us at our word, and trust us that logic is the steadily emerging universal language of scientific rigor. (Hence, if you perceive scientific rigor where no logic is used, you are misperceiving.) We will provide the evidence in due course. We will also provide in due course the material you must master to speak this universal language.

Finally, we humbly acknowledge in gratitude those devoted and brilliant intelligence analysts who, in the embryonic days of the LAMA paradigm and Slate, provided helpful comments and insights; and also the many RPI students who have likewise provided then, and indeed still provide — in connection with a more mature version of the LAMA paradigm and its technologies — now, valuable feedback.

Selmer Bringsjord
Naveen Sundar G.
Joshua Taylor