Introduction to (Formal) Logic (via, and also to, AI)

Spring 19, M R 12–1:50 DCC308

Selmer Bringsjord
Rensselaer AI & Reasoning (RAIR) Lab
Department of Cognitive Science
Department of Computer Science
Lally School of Management
Rensselaer Polytechnic Institute (RPI)
Troy NY 12180 USA
Office Hours: M R 345–5 CA 3rd flr; & by appointment

(518) 276-6472

Selmer.Bringsjord@gmail.com

version 0310191840NY

Contents

1 General Orientation

This course is an accelerated, advanced introduction, within the LAMA$^{\text{TM}}$ paradigm,[1] to deductive formal logic (with at least some brief but informative pointers to both *in*ductive and heterogeneous formal logic).[2] The phrase we use to describe what the student is principally introduced to in this class is: *beginning deductive logic, advanced* (BDLA). AI plays a significant role in advancing learning in the class; and the class includes an introduction to logicist aspects of AI and computer programming. After this class, the student can proceed to the intermediate level in formal deductive logic, and — with a deeper understanding and better prepared to flourish — to various areas within the *formal sciences*, which are all based on formal logic. The formal sciences include e.g. theoretical computer science (e.g., computability theory, complexity theory, rigorous coverage of programming and programming languages), mathematics in its traditional branches (analysis, topology, algebra, etc), decision theory, game theory, set theory, probability theory, mathematical statistics, etc. (and of course formal logic itself).

We have referred above to "the LAMA$^{\text{TM}}$ paradigm." What is that? This question will be answered in more detail later, but we do say here that while the LAMA$^{\text{TM}}$ paradigm is based upon a number of pedagogical principles, first and foremost among them is what can be labelled the Driving Dictum:[3]

If you can't prove it, you don't get it.

Turning back to the nature of formal logic, it can accurately be said that it's the science and engineering of reasoning,[4] but even this supremely general slogan fails to convey the flexibility and enormity of the field. For example, all of classical mathematics can be deductively derived from a small set of formulae (e.g., **ZFC** set theory, which you'll be hearing more about, and indeed experimenting with in the HyperSlate$^{\text{TM}}$ system) expressed in the formal logic known as 'first-order logic' (= FOL = $\mathscr{L}_1$, which you'll *also* be hearing more about), and, as we shall see and discuss

[1]'LAMA$^{\text{TM}}$' is an acronym for 'Logic: A Modern Approach," and is pronounced to rhyme with 'llama' in contemporary English, the name of the exotic and sure-footed camelid whose binomial name is *Lama glama*, and has in fact been referred to in the past by the single-l 'lama.'

[2]Sometimes 'symbolic' is used in place of 'formal,' but that's a bad practice, since — as students in this class will soon see — formal logic includes the representation of and systematic reasoning over *pictorial* information, and such information is decidedly *not* symbolic. For a discussion of the stark difference between the pictorial vs. the symbolic, and presentation of a formal logic that enables representation of and reasoning over both, see (Arkoudas & Bringsjord 2009), which directly informs Chapter 8 of our textbook.

[3]It's profitable to ponder a variant of this dictum, applicable in venues [e.g. legal hearings, courtrooms, reports by analysts in various domains that are not exclusively formal (e.g. fundamental investing, intelligence, etc.)] in which reasoning is not only deductive, but inductive, viz. "If you can't show by explicit argument that it's likelihood reaches a sufficient level, you don't get it."

[4]Warning: Increasingly, the term 'reasoning' is used by some who don't *really* do anything related to reasoning, as traditionally understood, to nonetheless label what they do. Fortunately, it's easy to verify that some reasoning is that which is covered by formal logic: *If the reasoning is explicit, links declarative statements or declarative formulae together via explicit, abstract reasoning schemata or rules of inference (giving rise to at least explicit arguments, and often proofs), is surveyable and inspectable, and ultimately machine-checkable, then the reasoning in question is what formal logic is the science and engineering of.* In order to characterize *in*formal logic, one can remove from the previous sentence the requirements that the links must conform to explicit reasoning schemata or rules of inference, and machine-checkability. It follows that so-called informal logic would revolve around arguments, but not proofs. An excellent overview of informal logic, which will be completely ignored in this class and its LAMA-BDLA textbook, is provided in "Informal Logic" in the Stanford Encyclopedia of Philosophy. In this article, it's made clear that, yes, informal logic concentrates on the nature and uses of argument.

in class, computer science emerged from and is in large part based upon logic (for cogent coverage of this emergence, see Glymour 1992, Halpern, Harper, Immerman, Kolaitis, Vardi & Vianu 2001). Logic is indeed the foundation for *all* at once rational-and-rigorous intellectual pursuits. (If you can find a counter-example, i.e. such a pursuit that doesn't directly and crucially partake of logic, S Bringsjord would be very interested to see it.)

2 Assistance to Bringsjord

The TAs for this course are: Swapnil Khandekar; email address: `khands2@rpi.edu`; and Shreyansh Nawlakha; email address `shreyanshnawlakha@gmail.com`. Swapnil will hold office hours on Tues 10a–12p on CA 3rd flr (and by appointment). Shreyansh will hold office hours on Fri 2p–4p also on CA 3rd flr (and by appointment). Some guest lectures may be provided by researchers working in the RAIR Lab, a logic-based AI lab.

3 Prerequisites

There are no formal prerequisites. However, as said above, this course introduces *formal* logic, and does so in an accelerated, advanced way. This implies that — for want of a better phrase — students are expected to have a degree of logico-mathematical maturity. You have this on the assumption that you understood the math you were supposed to learn in order to make it where you are.[5] For example, to get to where you are now, you were supposed to have learned the technique of *indirect proof* (= proof by contradiction = *reductio ad absurdum*). An example of the list of concepts and techniques you are assumed to be familiar with from high-school geometry can be found in the common-core-connected (Bass & Johnson 2012). An example of the list of concepts and techniques you are assumed to be familiar with from high-school Algebra 2 can be found in the common-core-connected (Bellman, Bragg & Handlin 2012). It's recommended that during the first two weeks of the class, students review their high-school coverage of formal logic, which includes at minimum the rudiments of the propositional calculus = $\mathscr{L}_{\text{PC}}$.[6]

4 Textbook/Courseware

Students will purchase the inseparable and symbiotic triadic combination published by Motalen:

- the e-textbook *Logic: A Modern Approach; Beginning Deductive Logic, Advanced via Hyper-Slate*™ (LAMA-BDLA);
- access to and use of the HyperGrader™ AI system (for, among other things, assessing student work); and

[5]If you happen to be a student reading this as one wanting to be introduced to formal logic, from outside RPI, please examine your own case realistically. If you are not in command of the traditional high-school-level content for algebra, geometry, trigonometry, and at least some (differential and integral) calculus, you will need to go with a standard, non-advanced introduction to logic in the LAMA™ paradigm, or in some other paradigm. Specifically, if in the LAMA™ paradigm, you will need the LAMA-BDL textbook, *not* LAMA-BDL**A**. The 'A' in 'LAMA-BDL**A**' is for 'Advanced.' Check which textbook you have!

[6]Sometimes referred to as 'sentential logic' or 'zeroth-order logic.' (For us, zeroth-order logic, $\mathscr{L}_0$, includes relation symbols and function symbols, as well as identity.) If you are at all confused about how these terms were used before reaching the present course, please discuss asap with the instructor or TAs.

- access to and use of the HyperSlate$^{\text{TM}}$ AI system (for, among other things, engineering proofs in collaboration with AI);

All three items will be available after purchase in the RPI Bookstore of an envelope with a personalized starting code for registration. Logistics of the purchase, and the contents of the envelope that purchase will secure, will be encapsulated in the first class meeting, Jan 10 2019, and then gone over in more detail on Jan 14 2019, after which the envelopes in question will soon be on sale in the Rensselaer Bookstore. The first use in earnest of HyperSlate$^{\text{TM}}$ and HyperGrader$^{\text{TM}}$ will happen in class on Jan 28 2019, so by the start of class on that day students should have LAMA-BDLA, and be able to open both HyperSlate$^{\text{TM}}$ and HyperGrader$^{\text{TM}}$ on their laptops in class. Updates to LAMA-BDLA, and additional exercises, will be provided by listing on the course web page (and sometimes by email) through the course of the semester. You will need to manage many electronic files in the course of this course, and e-housekeeping and e-orderliness are of paramount importance. You will specifically need to assemble a library of completed and partially completed proofs/arguments/truth-trees etc. so that you can use them as building blocks in harder proofs; in other words, building up your own "logical library" will be crucial.

Please note that HyperSlate$^{\text{TM}}$ and HyperGrader$^{\text{TM}}$ are copyrighted software: copying and/or reverse-engineering and/or distributing this software to others is strictly prohibited. You will need to submit online a signed version of a License Agreement. This agreement will also reference the textbook, which is copyrighted as well, and since it's an ebook, cannot be copied or distributed or resold in any way.

In addition, occasionally papers may be assigned as reading. Two background ones, indeed, are hereby assigned: (Bringsjord, Taylor, Shilliday, Clark & Arkoudas 2008, Bringsjord 2008).

Finally, slide decks used in class will contain crucial additional content above and beyond LAMA-BDLA and HyperSlate$^{\text{TM}}$ and HyperGrader$^{\text{TM}}$ content, and will be available on the web site for course for study. Along with slide decks, video and audio tutorials and mini-lectures will be provided as well.

5 Schedule

The progression of class meetings is divided into seven parts: first a motivation/history stretch I, during which we show that the logically untrained have great trouble reasoning well, and set an historical context for modern logic and AI, and then six additional parts II–VII. In the first of these remaining parts we'll focus in II on the **propositional calculus** (= $\mathscr{L}_{\text{PC}}$); in III on **first-order logic** (= FOL = $\mathscr{L}_1$), with a brief look at **second-order logic** (= SOL = $\mathscr{L}_2$) and beyond; and in IV we'll cover **modal logic** (in the form, specifically, of four closely related modal logics: **T**, **S4**, **D** (= **SDL**), and **S5**, with the emphasis on **SDL** as a formalism for AI/machine ethics). Emphasis will be on learning how to construct proofs in each system. Part V of the course looks at formal axiom systems, or as they are often called in mathematical logic, **theories**. Part VI of the course looks at formal *in*ductive logic, and to a degree at logics for reasoning over visual content (e.g., diagrams). The seventh (VII) and final part of the course is a synoptic look at some of the astonishing work of perhaps the greatest logician: Kurt Gödel.

A more fine-grained schedule now follows.[7]

[7]Note that the Rensselaer Academic Calendar is available here.

5.1 Why Study Logic?; Its History (I)

- **Jan 10**: *General Orientation to the LAMA™ Paradigm, Logistics, Mechanics.* The syllabus is reviewed in detail. It's made clear to the students that, in this class, there is a very definite, comprehensive, theoretical position on formal logic and the teaching thereof; this position corresponds to the affirmation of the LAMA™ (= Logic: A Modern Approach) paradigm, and that in lockstep with this position the tightly integrated trio of

    1. the LAMA-BDLA textbook,
    2. HyperSlate™ proof-construction system, and
    3. HyperGrader™ system for (among other things) automated assessment of proofs,

  are used. Students wishing to learn under the "Stanford" paradigm are encouraged to drop this LAMA™-based course and take PHIL 2140 in its alternating spot (i.e., Fall semester, annually).

- **Jan 14**: *Motivating Puzzles, Problems, Paradoxes, $\mathcal{R}$, $\mathcal{H}$, Part I.* Here we among other things tackle problems which, if solvable before further learning, obviate taking the course. We also discuss Bringsjord's "elevated" view of the human mind as potentially near-perfectly rational, and specifically capable of systematic and productive reasoning about the infinite.

- **Jan 17**: *Motivating Puzzles, Problems, Paradoxes, $\mathcal{R}$, $\mathcal{H}$, Part II.* A continuation of Part I; the problems in question get harder.

- **Jan 21**: *No Class: Martin Luther King Day*

- **Jan 24**: *Whirlwind History and Overview of Formal Logic (in intimate connection with computer science and AI), From Euclid to today's Cutting-Edge Computational Logic.* In one class meeting we surf the timeline of all of formal logic, from Euclid to the present. A particular emphasis is placed on Leibniz, the inventor of modern formal logic. Aristotle is cast as the inventor of formal logic in its original form (syllogistic deduction). The crucial timepoint of the discovery of the unsolvability of the *Entscheidungsproblem* by Turing-level computers figures prominently, and supports a skeptical position on The Singularity. The class ends with instructions for purchase of a personalized code in the Bookstore, which will enable students to obtain LAMA-BDLA and gain access to both HyperSlate™ and HyperGrader™. Codes, in laser-tagged, sealed envelopes, should be on sale after this class meeting.

5.2 Propositional Calculus ($\mathscr{L}_{\mathbf{PC}}$) & "Pure" Predicate Calculus ($\mathscr{L}_0$) (II)

- **Jan 28**: *Review from High School: Variables & Connectives; Propositional Calculus I.* This meeting will tie up any loose ends on the history side of things. Students by this point should have HyperSlate™ running on their laptops, have their codes registered, and have signed and accepted their LA. This is the start of coverage of the propositional calculus, $\mathscr{L}_{\mathrm{PC}}$.

- **Jan 31**: *Propositional Calculus II: The Formal Language, First Rules of Inference/Inference Schemata, and Immaterialism.* Application to some of the original problems used to motivate the course (meetings Jan 14 & 17). Simple proofs settle these problems. The view that formal logic, in particular some of the rudiments of the propositional calculus, is based on immaterial world, a view defended by the late James Ross (1992), is presented and defended.

- **Feb 4**: *Propositional Calculus III: Remaining Rules of Inference/Inference Schemata.* Here we discuss the "harder" inference schemata; e.g. proof-by-cases/disjunction elimination. More substantive proofs achieved. In addition, indirect proof is introduced in earnest.

- **Feb 7**: *Propositional Calculus IV: Pure General Logic Programming (PGLP) at the Level of $\mathscr{L}_{\mathrm{PC}}$.* Some harder proofs obtained. This class meeting will probably be the first time HyperSlate™ is used in conjunction with HyperGrader™. Demonstrations will be given. By this time students should be set up to use HyperGrader™ to win trophies. Coverage here of resolution, and PGLP at the level of the propositional calculus.

- **Feb 11**: *The Pure Predicate Calculus.* This is zeroth-order logic, or $\mathscr{L}_0$, for us. What kind of logic do we get if we add to the propositional calculus machinery for relation symbols, function symbols, and identity (=)? The result is

$\mathscr{L}_0$, and we explore some problems and proofs in this logic.

- **Feb 14**: Review, taking stock, prep for first test (next class), sustained Q&A.
- **Feb 18**: No class (President's Day Holiday)
- **Feb 19**: **Test #1**. Note that this is a Tuesday!

5.3 First-Order Logic (FOL = $\mathscr{L}_1$); Glimpse of SOL = $\mathscr{L}_2$, TOL = $\mathscr{L}_3$ (III)

- **Feb 21**: *The Need for Quantification, and the Centrality Thereof in Human Thought and Communication.*
- **Feb 25**: *New Inference Schemata in $\mathscr{L}_1$ = FOL, I.* We here introduce, discuss, and employ `existential intro` and `universal elim`; these are the two easy ones. But easy as they might be, do they suffice to enable us to prove that God exists?
- **Feb 28**: *A New (Harder) Inference Schema in $\mathscr{L}_1$ = FOL, II.* We introduce `universal intro`, the first of the two harder new inference schemata for FOL.
- **Spring Break**: Mar 4–Mar 8
- **Mar 11**: *Proofs/Problems in $\mathscr{L}_1$ = FOl, III.* We here cover the inference schema `existential elim`.
- **Mar 14**: *Proofs/Problems in $\mathscr{L}_1$ = FOl, IV.* Now we're ready for some challenges in FOL!
- **Mar 18**: *The Liar; Russell's Paradox; Skolem's Paradox.*
- **Mar 21**: *ZFC; Second-Order Logic (SOL), Third-Order Logic (TOL), and Beyond (e.g. Type Theory).* We return here to the failed attempt to prove God's existence, and with some help from Kurt Gödel, try again.
- **Mar 25**: **Test #2**

5.4 Theories (= Axiom Systems) (IV)

- **Mar 28**: *Theories of Arithmetic I (e.g.,* **EA**).
- **Apr 1**: *Theories of Arithmetic II (e.g.,* **PA**).

5.5 Deontic Logic and Killer Robots (V)

- **Apr 4**: *Modal Logic: What and Why.* This is a general introduction to the crucial difference between *extensional* logic versus *intensional* logic. The logics $\mathscr{L}_{\text{PC}}$, $\mathscr{L}_0$, $\mathscr{L}_1$, $\mathscr{L}_2$, $\mathscr{L}_3$ are all extensional. Now we move to the intensional category, which includes modal logics.
- **Apr 8**: *The System* **D** = **SDL**. This is the basic system of logic intended to capture central categories in ethics (e.g., obligation, permissibiity, etc.). It will turn out that **SDL** is in need of major improvement.
- **Apr 11**: *The Threat of "Killer" Robots.* Here I present the "PAID" problem: artificial agents/robots that are **p**owerful, **a**utonomous, and **i**ntelligent, are **d**angerous (if not capable of **d**estroying us).
- **Apr 15**: *Logic Can Save Us; Here's How.* After taking note of the fact that *Star Trek* (original) teaches us that logic can save us, this class introduces a engineered quantified multi-operator modal logic, $\mathcal{DCEC}^*$, developed at Rensselaer, and explains how use of the computational version thereof, implemented, can be used to enable an AI/robot to adjudicate thorny ethical dilemmas.

5.6 Beginning Heterogenous Logic & Beginning *In*ductive Logic (BIL): Glimpses (VI)

- **Apr 18**: *Heterogeneous Logic; Whirlwind History & Overview Beginning (Formal) Inductive Logic (LAMA-BIL) in the LAMA™ Paradigm.* A solution to the Lottery Paradox is provided, and recent work in the RAIR Lab devoted to solving the St Petersburg Paradox will also be covered.

5.7 Gödel (VII)

- **Apr 22**: *Gödel's Completeness Theorem & First Incompleteness Theorem*. We seek here to understand the brilliant core of Gödel's CT, from his doctoral dissertation. In addition, we provide the class with a glimpse of Gödel's stunning *in*completeness theorems, and briefly take up the question: Could an AI ever match Gödel?
- **Apr 25**: **Test #3**

6 Grading

Grades are based in part on three in-class tests. Each of these tests will call for in-class use of HyperSlate^{TM} in conjunction with HyperGrader^{TM}. The three tests are weighted 10%, 15%, and 25%, respectively. In addition, grades are based on a series of required problems to be done in the HyperSlate^{TM} system, and verified by HyperGrader^{TM}. Every problem in the series must be certified 100% correct by HyperGrader^{TM} in order to pass the course, and a grade of 'A' is earned for the series, which is 40% of the final grade. All required assignments on HyperGrader^{TM} must be completed and submitted in order to receive a final grade. The remaining 10% of one's grade is based on performance on "pop" problems given in class, to be solved in HyperSlate^{TM} and graded by HyperGrader^{TM}. Finally, please note that class attendance is mandatory. Any more than two unexcused absences will result in a failing grade.

7 Some Learning Outcomes

There are four desired outcomes. *One*: Students will be able to carry out formal proofs and disproofs, within the HyperSlate^{TM} system and its workspaces, at the level of the propositional and predicate calculi, and propositional modal logic (the aforementioned systems **T**, **S4**, **D**, and **S5**). *Two*: Students will be able to translate suitable reasoning in English into interconnected formulae in the languages of these four calculi, and assess this reasoning by determining if the desired structures are present in the formulae and relationships between them. *Three*, students will be able to carry out informal proofs. *Four*, students will demonstrate significant understanding of the advanced topics covered.

8 Academic Honesty

Student-teacher relationships are built on mutual respect and trust. Students must be able to trust that their teachers have made responsible decisions about the structure and content of the course, and that they're conscientiously making their best effort to help students learn. Teachers must be able to trust that students do their work conscientiously and honestly, making their best effort to learn. Acts that violate this mutual respect and trust undermine the educational process; they counteract and contradict our very reason for being at Rensselaer and will not be tolerated. Any student who engages in any form of academic dishonesty will receive an F in this course and will be reported to the Dean of Students for further disciplinary action. (The *Rensselaer Handbook* defines various forms of Academic Dishonesty and procedures for responding to them. All of these forms are violations of trust between students and teachers. Please familiarize yourself with this portion of the handbook.) In particular, all solutions submitted to HyperGrader^{TM} for course credit under a student id are to be the work of the student associated with that id alone, and not in any way copied or based on the work of anyone else.

References

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*

Bass, L. & Johnson, A. (2012), *Geometry: Common Core*, Pearson, Upper Saddle River, NJ. Series Authors: Charles, R., Kennedy, D. & Hall, B. Consulting Authors: Murphy, S. & Wiggins, G.

Bellman, A., Bragg, S. & Handlin, W. (2012), *Algebra 2: Common Core*, Pearson, Upper Saddle River, NJ. Series Authors: Charles, R., Kennedy, D. & Hall, B. Consulting Authors: Murphy, S. & Wiggins, G.

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., Taylor, J., Shilliday, A., Clark, M. & Arkoudas, K. (2008), Slate: An Argument-Centered Intelligent Assistant to Human Reasoners, *in* F. Grasso, N. Green, R. Kibble & C. Reed, eds, 'Proceedings of the 8th International Workshop on Computational Models of Natural Argument (CMNA 8)', University of Patras, Patras, Greece, pp. 1–10.
URL: *http://kryten.mm.rpi.edu/Bringsjord_etal_Slate_cmna_crc_061708.pdf*

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

Halpern, J., Harper, R., Immerman, N., Kolaitis, P., Vardi, M. & Vianu, V. (2001), 'On the Unusual Effectiveness of Logic in Computer Science', *The Bulletin of Symbolic Logic* **7**(2), 213–236.

Ross, J. (1992), 'Immaterial Aspects of Thought', *The Journal of Philosophy* **89**(3), 136–150.