

Introduction to Formal Logic (and AI)

(IFLAI1, pronounced “eye” • “fly” • “one” ’)

Spring 25, M R 12:00–1:50, SAGE 3510

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 3:45–5 in-person CA306 \vee Zoom; Wed by appointment

(518) 276-6472

Selmer.Bringsjord@gmail.com

version 0106241719NY

Contents

1	General Orientation	1
2	Assistance to Bringsjord	2
3	Prerequisites	2
4	AI Platform (w/ textbook)	3
5	Schedule	4
5.1	Why Study Logic?; Its History (I)	5
5.2	Propositional Calculus (\mathcal{L}_{PC}) & “Pure” Predicate Calc. (\mathcal{L}_0) (II)	5
5.3	Extensional Logics in General (e.g. First-Order Logic (FOL = \mathcal{L}_1 , & Second-Order Logic (SOL = \mathcal{L}_2); Some (Extensional) Paradoxes; ZFC (III)	6
5.4	More Theories (= Axiom Systems); Beyond FOL (IV)	7
5.5	Intensional Logic; Deontic Logic and Killer Robots/Vehicles (V)	7
5.6	Beginning Heterogenous Logic & Beginning <i>Inductive</i> Logic (BIL): Glimpses (VI)	7
5.7	Gödel (VII)	7
6	Grading	9
7	Some Learning Outcomes	9
8	Academic Honesty	9
	References	10

1 General Orientation

This course is an accelerated, advanced introduction, within the LAMA[®] paradigm,¹ to deductive formal logic (with at least some brief but informative pointers to both *inductive* and heterogeneous formal logic), and to a substantive degree logic-based AI.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 a gentle introduction to logicist aspects of AI and computer science and logic programming.³ After this class, the student can proceed to the intermediate level in formal deductive (and, for that matter, inductive 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). The system HyperSlate[®], an important system in the class, can be used productively, by the way, in all these disciplines. This system is accessed from the overarching AI platform used in the class: HyperGrader[®].

We have referred above to “the LAMA[®] paradigm.” What is that? This question will be answered in more detail later, but we do say here that 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:⁴

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,⁵ but even this supremely general slogan fails to convey the flexibility and

¹LAMA[®] 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.’

²Sometimes ‘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 the LAMA-BDLAHGHS textbook.

³The formal basis of computer science and computer programming is formal logic (despite claims to the contrary made on the strength of the rise of such machine-learning approaches as “deep learning,” and large language models (e.g. GPT-4) that use deep learning), so this is quite natural.

⁴It’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.”

⁵Warning: 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 *informal* logic, one can remove from the previous sentence the requirements that the links must conform to explicit inference 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 largely ignored in this class and its LAMA-

enormity of the field. For example, a vast part 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[®] system) expressed in the formal logic known as ‘first-order logic’ (= FOL = \mathcal{L}_1 , which you’ll *also* be hearing more about), and, as we shall see and discuss 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.)

The Pedagogical Gift of Illogical AIs Like GPT-4

To complete this section, please note that because of all the attention now being paid world over to AIs like GPT-4 and other “chatbots” (= large language models = LLMs), that logic is the science of reasoning means that a great way to learn logic is by seeing why such AIs are — as Arkoudas (2023) has e.g. pointed out, and as we shall repeatedly confirm — bad reasoners.

2 Assistance to Bringsjord

The TA for this course is PhD student James Oswald (oswalj@rpi.edu), an expert on computational logic and AI, esp. AGI, and in particular therein the formal science of intelligence in artificial agents. James will hold office hours on Thu 10–12 in Carnegie 306; by appointment also, of course. Some guest lectures may be provided by Oswald, and perhaps other researchers working in the RAIR Lab, a logic-based AI lab at RPI directed by Bringsjord.

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.⁶ 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

BDLAHGHS textbook (accessed from the overarching HyperGrader[®] AI platform), 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.

⁶If 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 (course-wise, esp. Algebra 2), 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-BDLHGHS textbook, *not* LAMA-BDLAHGHS. The ‘A’ in ‘LAMA-BDLAHGHS’ is for ‘Advanced.’ Check which textbook you have!

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). (Note in particular that this Algebra 2 textbook has extensive coverage of proof by contradiction.) It's recommended that during the first three weeks of the class, students review their high-school coverage of formal logic, which includes at minimum the rudiments of the propositional calculus = \mathcal{L}_{PC} .⁷

4 AI Platform (w/ textbook)

Students electing to stay in the course will purchase access to the inseparable and symbiotic triadic combination published by Motalen:

- the overarching AI platform HyperGrader[®], AI for, among other things, presenting and tracking student work, and from which authorized users can ...
- access and use of HyperSlate[®] AI system (for, among other things, engineering proofs (and computer programs in the Hyperlog logic-programming language) in collaboration with AI) and;
- access and read the e-textbook *Logic: A Modern Approach; Beginning Deductive Logic, Advanced via HyperGrader[®] & HyperSlate[®]* (LAMA-BDLAHGHS).

All three items will be accessible 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 6 2025, and then gone over in more detail repeatedly after that, including during the key class meeting of Jan 23 2025, by which time students should be registered and ready during class to use their HyperGrader[®] accounts/libraries. The first intense use of HyperSlate[®] and HyperGrader[®] will happen, at the earliest, during class on Jan 23 2025, so by the start of class on that day students should certainly have LAMA-BDLAHGHS, and be able to open both HyperSlate[®] and HyperGrader[®] in a browser on their laptops in class. Updates to LAMA-BDLAHGHS, and additional exercises, will be provided by listings on HyperGrader[®] (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. in HyperSlate[®] so that you can use them as building blocks in harder proofs; in other words, building up your own “logical library” in the cloud will be crucial.

Please note that HyperSlate[®] and HyperGrader[®] are copyrighted, trademarked software based on patented and Pat. Pend. methods: 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.

⁷Sometimes referred to as ‘sentential logic’ or ‘zeroth-order logic.’ (For us, zeroth-order logic, \mathcal{L}_0 , includes relation symbols and function symbols, as well as identity; these things are not part of the propositional calculus.) 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.

In addition, an important part of the course will be coverage of the great theorems of the greatest logician ever: Kurt Gödel. A pre-publication version of Bringsjord’s *Gödel’s Great Theorems*, forthcoming from Oxford University Press, will be made available; this copying or sharing this content in any way is prohibited.

Also, occasionally papers may be assigned as reading. Three background ones, indeed, are hereby assigned: (Bringsjord, Giancola & Govindarajulu 2023, Bringsjord, Taylor, Shilliday, Clark & Arkoudas 2008, Bringsjord 2008).

Finally, slide decks used in class will contain crucial additional content above and beyond LAMA-BDLAHGHS and HyperSlate[®] and HyperGrader[®] content, and will be available on the web site for our 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 (and hence living) well (you can be cured by this course!), and set an historical context for modern formal logic and AI, and then six additional parts II–VII. In the first of these remaining parts, II, we’ll focus on the **propositional calculus** ($= \mathcal{L}_{PC}$) and **zero-order logic** ($ZOL = \mathcal{L}_0$). We will also introduce *Pure General Logic Programming* (PGLP), and the new programming language, Hyperlog (a fragment of which will be available to you in HyperSlate[®]), that makes PGLP concrete.

In Part III we shall focus on **first-order logic** ($= FOL = \mathcal{L}_1$), with substantive study of **second-order logic** ($= SOL = \mathcal{L}_2$) and beyond. Proofs will be constructed in the AI-infused HyperSlate[®] system; and in IV we’ll cover **modal logic**, in the form, specifically, of four closely related modal logics: **T**, **S4**, **D** ($= \mathbf{SDL}$), and **S5**, with the emphasis on **SDL** as a candidate formalism for AI/machine ethics — a candidate that fails. Once we understand the reasons for this failure, we will look at a very expressive quantified modal logic that has been used with considerable success in AI ethics: \mathcal{DCEC}^* . Emphasis will be on learning how to construct hypergraphical 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 *inductive* 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 the greatest logician: Kurt Gödel. Part VII will include private, non-copyable distribution of a pre-publication version of *Gödel’s Great Theorems*, forthcoming from Oxford University Press. Distribution of this content outside of students in the class is prohibited (by the policies of the Press itself).

A more fine-grained schedule now follows.⁸

⁸Note that the Rensselaer Academic Calendar is available [here](#).

5.1 Why Study Logic?; Its History (I)

Disclaimer: While we shall stick very closely to the sequence and topics that follow, variations and even new sub-topics will emerge in connection with up-to-the-moment research being carried out by Bringsjord & RAIR-Lab researchers.

- **Jan 6:** *General Orientation to the LAMA[®] Paradigm, Logistics, Mechanics.* The syllabus is reviewed in detail, and discussed. It's made clear to students that, in this class, there is a very definite, comprehensive, theoretical position on computational formal logic and the teaching thereof; this position corresponds to the affirmation of the LAMA[®] (= Logic: A Modern Approach) paradigm, and that in lock-step with this position the tightly integrated trio listed above, i.e.
 1. LAMA-BDLAHGHS textbook,
 2. HyperSlate[®] AI-infused proof-construction system, and
 3. HyperGrader[®] AI platform for (among other things) accessing the prior two things, and automated assessment of proofs and management of points earned on the leaderboards,

are used. Students wishing to learn under the venerable “Stanford” paradigm (and cognate ones) are strongly encouraged to immediately drop this LAMA[®]-based course and take PHIL 2140 in its alternating spot (i.e., Fall semester, annually).

- **Jan 9:** *Motivating Puzzles, Problems, Paradoxes, \mathcal{R} , \mathcal{H} , Part I.* Here we among other things tackle problems such that: (i) if solvable before further learning, obviate taking the course; and (ii) if solvable after taking the course will empower you. 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 13:** *Motivating Puzzles, Problems, Paradoxes, \mathcal{R} , \mathcal{H} , Part II.* A continuation of Part I; the problems in question get harder!
- **Jan 16:** *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 AI and Automated Reasoning. 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), presented in his *Organon* (for an overview, see Smith 2017) over two millennia back. 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 another round of instructions for purchase and initial use of a personalized code from the Follett Bookstore, which will enable students to obtain access to the on-line AI systems HyperSlate[®] and HyperGrader[®], and via them to LAMA-BDLAHGHS. Codes, in laser-tagged, sealed envelopes, should be on sale at this point in the Follett Bookstore.

- **Jan 20:** *No Class: Martin Luther King Day*

5.2 Propositional Calculus (\mathcal{L}_{pc}) & “Pure” Predicate Calc. (\mathcal{L}_0) (II)

- **Jan 23:** *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 be registered and have HyperSlate[®] running in a browser laptops, and have signed and accepted their LA. This is the start of coverage of the propositional calculus, \mathcal{L}_{pc} . We see AI in action in HyperSlate[®], in the form of the provability oracle for \mathcal{L}_{pc} .
- **Jan 27:** *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 9 & Jan 13). Simple proofs settle these problems. The view that formal logic, in particular some of the rudiments of the propositional calculus, exists in an immaterial world, a view that can be defended with help from the late James Ross (1992), is presented. This view is extended to a conception of all of computer science based on formal logic/logic machines.

- **Jan 30:** *Propositional Calculus III: Remaining Rules of Inference/Inference Schemata; Propositional Trees.* Here we discuss the “harder” inference rules/schemata; e.g. proof-by-cases = disjunction elimination. More substantive proofs achieved. In addition, hypergraphical indirect proof (= proof by contradiction = *reductio ad absurdum*) is introduced in earnest. We also introduce propositional truth trees, explain their superiority over (infernal) truth tables, and show how these trees can be easily constructed in HyperSlate[®].
- **Feb 3:** *PGLP (Pure General Logic Programming) at the Level of \mathcal{L}_{PC} , and Hyperlog, Part I.* Some harder proofs obtained. By this class meeting students will be comfortable using HyperSlate[®] in conjunction with HyperGrader[®]. Demonstrations will be given. Coverage here of resolution, and PGLP/Hyperlog at the level of the propositional calculus. The Motalen game Catabot Rescue will be shown and used.
- **Feb 6:** *Darwin’s Mistake; The Pure Predicate Calculus; Metalogic: Soundness and Completeness of \mathcal{L}_{PC} and \mathcal{L}_0 .* This is zeroth-order logic, or \mathcal{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 \mathcal{L}_0 , and we explore some problems and proofs in this logic. We end with proof-sketches of both soundness and completeness of both the propositional calculus and zero-order logic.
- **Feb 10:** **Discussion & Publication (online, in HyperGrader[®]) of Test #1.** All problems due **11:59pm Feb 18**, so students have a solid week to work on Test 1.



5.3 Extensional Logics in General (e.g. First-Order Logic ($FOL = \mathcal{L}_1$, & Second-Order Logic ($SOL = \mathcal{L}_2$); Some (Extensional) Paradoxes; ZFC (III)

- **Feb 13:** *Extensional Logic; The Need for Quantification, and the Centrality Thereof in Human Thought and Communication; “Proving”*

God’s Existence. We use our picture of the entire, vast universe of logics to establish a context in which to then zero in on *extensional* logics, and then discuss the crucial need for having quantifiers like \forall and \exists , and the centrality of quantifiers to human cognition, which in this regard is discontinuous with the cognition of nonhuman animals.

- **Feb 17:** No class (President’s Day Holiday)
- **Feb 18: THIS WILL BE AN EXCLUSIVELY ONLINE (ZOOM) CLASS MTG.** *New Inference Schemata in $\mathcal{L}_1 = FOL$, I.* We here introduce, discuss, and employ **existential intro** and **universal elim** in their hypergraphical form; these are the two “easy” new inference rules/schemata of \mathcal{L}_1 . But easy as they might be, do they suffice to enable us to prove that God exists? ...
- **Feb 20:** *The Two New (Harder) Inference Schemata in $\mathcal{L}_1 = FOL$, II.* We introduce hypergraphical **universal intro**, and also introduce **existential elim**, and related matters. Some further discussion of games that require and supposedly cultivate logical reasoning is also carried out.
- **Feb 24:** *Quantificational Trees; FOL , IV (Numerical Quantification, Intro to CVQ^+A^+JV); $\mathcal{L}_2 =$ second-order logic = SOL , $\mathcal{L}_3 =$ third-order logic = TOL , and Beyond.* We look at trees with quantification, and see how to build and test them with help from AI in HyperSlate[®]. We also introduce an advanced, interesting form of Visual Question Answering (VQA) that contemporary ML can’t handle. Finally, we discuss logic-based hierarchies (the Polynomial and Arithmetic) for measuring the level of intelligence in intelligent agents, whether natural (like us), artificial, or extraterrestrial/alien.
- **Feb 27:** *Some (Extensional) Paradoxes, and the Rebuilding Thereafter via ZFC* The Liar; Russell’s Paradox/The Barber; Richard’s Paradox; and perhaps some coverage of the more complicated Skolem’s Paradox. Then, the solution: **ZFC**. We see how a computational version of **ZFC** is available for exploration in HyperSlate[®].
- **Mar 3–Mar 7:** Spring Break!
- **Mar 10:** *PA (Peano Arithmetic)/Number Theory.* Theories of Arithmetic I. **PA**, and its simpler parts (e.g. so-called “Baby Arithmetic”).

We here make what I hope are some interesting connections to AI-based math education in the early grades.



5.4 More Theories (= Axiom Systems); Beyond FOL (IV)

- **Mar 13: Discussion & Publication of (online, in HyperGrader[®]) Test #2.** All problems due **11:59pm Mar 20**, so students have a solid week to work on Test 2.
- **Mar 17: Climbing the k-Order Ladder; God and Third-Order Logic; Astrologic.** We here bring a large swathe of our Universe of Logics to life, by moving beyond first-order logic = \mathcal{L}_1 into second- (\mathcal{L}_2) and third-order logic = \mathcal{L}_3 , the latter a logic Gödel used for his famous proof of God's existence, which we cover in the part part of the present course. We end by considering so-called "Astrologic," which is logic for all human-level beings, including those on other planets (if such exist).



5.5 Intensional Logic; Deontic Logic and Killer Robots/Vehicles (V)

- **Mar 20: On to Intensional Logic; Modal Logic: What and Why.** This is a general introduction to the crucial difference between *extensional* logic versus *intensional* logic. The logics \mathcal{L}_{PC} , \mathcal{L}_0 , \mathcal{L}_1 , \mathcal{L}_2 , \mathcal{L}_3 are all extensional. Now we move to the intensional category, which includes modal logics. Five modal logics are introduced, rapidly for now: **K**, **T** = **M**, **D**, **S4**, and **S5**.
- **Mar 24: The Threat of "Killer" Robots; Logic Can Save Us; Here's How.** Here is presented the "PAID" problem: artificial agents (e.g. robots) that are **p**owerful, **a**utonomous, and **i**ntelligent, are **d**angerous (if not capable of **d**estroying us). We also discuss the dangers arising from using logic-less AI to engineer self-driving vehicles. After taking note of the fact that *Star Trek* (original) teaches us that logic can save us, this class introduces an engineered quantified multi-operator modal logic, *DCEC**, invented

by Bringsjord & Govindarajulu (for a technical look, see Govindarajulu & Bringsjord 2017), and explains how use of the computational version thereof, implemented, can be used to enable an AI/robot to adjudicate thorny ethical dilemmas. We explore *DCEC** in HyperSlate[®].



5.6 Beginning Heterogenous Logic & Beginning Inductive Logic (BIL): Glimpses (VI)

- **Mar 27: 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. So-called "Pure Inductive Logic," the modern version of inductive logic stemming from R. Carnap,⁹ is encapsulated, and distinguished from inductive logic in the LAMA[®] paradigm.



5.7 Gödel (VII)

- **Mar 31: Introduction to & Overview of Gödel's Great Theorems; Gödel's Completeness Theorem (GCT).** After an overview of the entire book on Gödel, Selmer presents and teaches Gödel's first theorem, from his doctoral dissertation. This is the completeness theorem (GCT), which early in his dissertation, before delivering the proof, Gödel boldly declares to constitute a demonstration that syntax and semantics are the same; i.e., that the principled organization of symbols and nothing else, if observed and understood, suffices to provide meaning.
- **Apr 3: Gödel's First Incompleteness Theorem (G1).** This is the theorem that gets all the attention, to this day. The main trick here (the taking of which is in line with the dictum that "there's no free lunch") is to take as a given the Fixed-Point Theorem, and to use it

⁹And definitively presented in (Paris & Vencovská 2015).

as the powerful pivot for not only G1 here, but later for G2 and GST. This will be explained. In a nutshell, G1 says that there are perfectly straightforward propositions about arithmetic that are perpetually mysterious, since no matter what they can't be proved, and neither can their negations. Is there any escape from this situation? As a matter of fact, though it is rarely mentioned, there is, and it will be discussed.

- **Apr 7:** *Gödel's Second Incompleteness Theorem (G2)*. Given certain assumptions about the power of our proof methods, along with Selmer's assumption that mathematics must include Peano Arithmetic (PA), we can prove, following Gödel, that we can't prove mathematics to be free of contradiction! Yet again, there is an escape.
- **Apr 10:** *Gödel's Greatest Theorem*. Pretty much everyone outside cognoscenti on Gödel's *oeuvre* believes that Gödel's greatest achievement is his incompleteness result (it's actually 'results' in the plural, of course). Nothing could be further from the truth, as we see. His greatest achievement is proving that how fundamentally mysterious is the claim that, while we know definitively that the set of natural numbers $\mathbb{N} := \{0, 1, 2, 3, \dots\}$, while infinite, is actually "smaller" than the set of all the real numbers \mathbb{R} (the natural numbers plus not only rational numbers/fractions, but irrational numbers such as $\sqrt{2}$), there is no set "in between" these two sets in size. To explain Gödel's proof, we turn to the reasoning carried out by Sherlock Holmes in his most brilliantly solved case, that of the missing racehorse, Silver Blaze.
- **Apr 14:** *Gödel's God Theorem (GGT)*. Thinking that he was upon his death-bed, Gödel gave a notebook of his for a bedside Dana Scott to see and analyze a proof of God's existence that he — Gödel — had been working on for rather a long time. Here began an amazing sequence of work, some involving state-of-the-art AI to try to analyze and check the validity of Gödel's reasoning. We examine GGT for ourselves, and Selmer offers a refined version for consideration.
- **Apr 17:** *Gödel's Time-Travel Theorem, and Gödel's God Theorem*. Gödel proved that travel backwards in time is mathematically possible;

the proof was a present he gave to Einstein. We give an intuitive version of the proof based on visualizations stemming from the famous, must-read fantasy book *Flatland* (Abbott 2006).

- **Apr 21:** *In the Light of Gödel, Will AI Ever Match (or Exceed) Human Intelligence?*. This class includes presentation and discussion of the last chapter in Bringsjord's *Gödel's Great Theorems* book. The chapter does two things. One, it includes an explanation and discussion of Gödel's Either-Or, which he delivered to an audience in 1954. It says that either humanity will be forever unable to solve a certain Diophantine equation, or the human mind is superior to all forms of computational intelligence/AI. We end by wrapping up some loose-ends, and discussing the coming onlinem in HyperGrader[®]) Test #3.

6 Grading

Grades are based in part on three tests. Each of these tests will call for timed use of HyperSlate[®] in conjunction with HyperGrader[®]; they will be sometimes discussed in class on the relevant day in the Schedule (§5); and then released online in HyperGrader[®] later that same day (\approx 12 noon Pacific Time) with a countdown timer (no work is allowed beyond zero time left). The three tests are weighted 10%, 15%, and 25%, respectively.

In addition, grades are based on a series of *Required*, self-paced homework problems to be done in the HyperSlate[®] system, and verified by HyperGrader[®]. These are called ‘Required’ problems in HyperGrader[®]. Every problem in the collection must be certified 100% correct by HyperGrader[®] in order to pass the course, and a grade of ‘A’ is earned for the series when it’s completed, which is 40% of the final grade. All required homework assignments on HyperGrader[®] must be completed and submitted in order to receive a final grade.

The remaining 10% of one’s grade is based on performance on “live logic” problems released online in HyperGrader[®], to be solved before the countdown timer for the problem hits zero.

7 Some Learning Outcomes

There are four desired outcomes. *One*: Students will be able to carry out formal proofs and disproofs, within the HyperSlate[®] 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 (e.g. the quantified multi-modal logic $\mathcal{DC}\mathcal{EC}$, available in a type of workspace in HyperSlate[®]).

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[®] 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 directly upon the work of anyone else.

References

- Abbott, E. A. (2006), *Flatland: A Romance of Many Dimensions*, Oxford University Press, Oxford, UK. *Flatland* was originally published in 1884 by Seeley & Co. of London.
- Arkoudas, K. (2023), ‘GPT-4 Can’t Reason’.
URL: <https://arxiv.org/abs/2308.03762>
- 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. This URL goes to a preprint only.
URL: http://kryten.mm.rpi.edu/sb-lccm-ab-toc_031607.pdf
- Bringsjord, S., Giancola, M. & Govindarajulu, N. S. (2023), Logic-Based Modeling of Cognition, in R. Sun, ed., ‘The Cambridge Handbook on Computational Cognitive Sciences’, Cambridge University Press, Cambridge, UK, pp. 173–209. The URL here goes to an uncorrected preprint.
URL: http://kryten.mm.rpi.edu/SBringsjord-etal-L-BMC_121521.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.
- Govindarajulu, N. & Bringsjord, S. (2017), On Automating the Doctrine of Double Effect, in C. Sierra, ed., ‘Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)’, International Joint Conferences on Artificial Intelligence, pp. 4722–4730.
URL: <https://doi.org/10.24963/ijcai.2017/658>
- 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.
- Paris, J. & Vencovská, A. (2015), *Pure Inductive Logic*, Cambridge University Press, Cambridge, UK.
- Ross, J. (1992), ‘Immaterial Aspects of Thought’, *The Journal of Philosophy* **89**(3), 136–150.
- Smith, R. (2017), Aristotle’s Logic, in E. Zalta, ed., ‘The Stanford Encyclopedia of Philosophy’.
URL: <https://plato.stanford.edu/entries/aristotle-logic>