Intermediate Formal Logic and AI

(= IFLAI2, rhymes with "eye fly boo")

Crosslistings:
64651-CSCI-4968-01
64652-MATH-4140-01
63898-PHIL-4140-01
63864-PHIL-4961-01
Fall 23, Mo Th 4:00p–5:50 NY time; DCC 330

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: Wed 10a–12n CA 3rd flr (& sometimes via Zoom during travel); & by appointment

Selmer.Bringsjord@gmail.com

version 1120231249NY

Contents

1 General Orientation

This course is an intermediate class in computational formal logic, with substantive connections made to logicist (= logic-based) AI and "big" issues arising in and because of the AI, within the LAMA® paradigm.[1] (While emphasis is on *deductive* formal logic, there is some coverage of formal *in*ductive and heterogeneous formal logic).[2] AI plays a significant role in advancing learning in the class, and the class includes some coverage of logicist AI and logicist computer programming.[3] While coverage will not be extensive, students are introduced to a new logic-programming language: Hyperlog®. In addition, we shall consider the painful shortcomings of today's "chatbots"/large language models in the area of rigorous reasoning (e.g., GPT-4 simply can't reason, period; see e.g. from Arkoudas "GPT-4 Can't Reason"), and whether these deficiencies can be remedied by logicist AI.

We have referred above to "the LAMA® paradigm." What is that? This question will be answered in more detail later, but we do volunteer 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 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, with which you should already have some familiarity with, and in this course you will be experimenting with **ZFC** in the HyperSlate® system) expressed in the formal logic known as 'first-order logic' (= FOL = $\mathscr{L}_1$, with which you are *also* familiar, and,

[1]'LAMA®' 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 LAMA-BDLAHGHS textbook.

[3]We use 'logicist computer programming' to denote a general approach to computer programming based on formal logic; this general approach covers what is called 'logic programming,' which is connected specifically to such languages as Prolog.

[4]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."

[5]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 (but, we hold, should not) 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-BDLAHSHG 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.

as we shall review 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.)

2 Prerequisites

Students are expected to have taken a serious university-level introductory formal-logic course (which hopefully included second-order logic and some propositional modal logic — but given the rather impoverished reach of the vast majority of such courses, this really and truly is a *hope*), and to have a significant degree of logico-mathematical maturity. Phase I of the course will include a review of some introductory formal logic, via problems in HyperSlate®.

3 Teaching Assistant and Guest Lecturer

The TA for this course is James Oswald, a doctoral student in computer science at RPI, and researcher in the RAIR Lab. James is himself an expert on AI, especially AGI. He will be not just a TA in this course, but a lecturer. His email address is `oswalj@rpi.edu`, and James' office hours are Mondays 2–4 pm.

4 Reading/Videos/Textbook/Courseware

Slide decks and lectures/tutorials (including in some casses video versions of such) are part of the crucial content for this course, and will be linked-to from the course web page; in this regard we have a parallel situation to IFLAI1.

Papers that are required reading will be made available to students as we proceed, usually through hotlinks on the course website, sometimes by direct email.

Students will purchase a license giving access to the inseparable and symbiotic triadic combination published and maintained by Motalen:

- the e-textbook *Logic: A Modern Approach; Beginning Deductive Logic, Advanced via HyperSlate® and HyperGrader®* (LAMA-BDLAHSHG);
- access to and use of the HyperGrader® AI platform (for, among other things, assessing student work); and
- access to and use of the HyperSlate® AI interactive environment (for, among other things, engineering proofs and logic programs in collaboration with AI). This environment is available on said platform.

All three items will be available after purchase in the RPI Bookstore of a barcoded envelope with a personalized starting code/key for registration. Students who previously registered for a version of the online software and ebook will be able to present their email address used in the system and receive a substantial discount. Logistics of the purchase, and the contents of the envelope that purchase will secure, will be encapsulated in the first class meeting, Aug 28 2023, and then gone over in more detail repeatdly in class, including on Aug 31 2023 and Sep 5 2023. The first use in earnest of HyperSlate® and HyperGrade® will happen in class on Sep 5 2023, so by the start of class on that day students should have LAMA-BDLAHSHG, and be able to open both HyperSlate® and HyperGrader® on their laptops in class. Updates to LAMA-BDLAHSHG, 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 your own personalized library of completed and partially completed proofs/arguments/truth-trees etc. in the cloud provided to you, 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® and HyperGrader® are trademarked, copyrighted, and patented 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® and HyperGrader® content, and will be available on the web site for the course for study. Along with slide decks, an appreciable number of video and audio tutorials and mini-lectures will be provided as well.

5 Schedule

5.1 The Four Coverage Areas

This class is divided into four I–IV coverage areas:

I *Review*$^+$. We use HyperSlate® and HyperGrader® to review the logics $\mathscr{L}_{\mathrm{PC}}, \mathscr{L}_1$, $\mathscr{L}_2$, and the propositional modal logics **K**, **T**, **D**. For some students, modal logics may be new; these students will want to pay close attention to, and expend some genuine effort exploring, these logics. Our review will also include brief description of the modal logics **S4** and **S5**; both are still only *propositional* modal logics. Importantly, many students who took "IFLAI1" in the past will not have used a version of HyperSlate® that included **S4** or **S5**. This the reason for the use of '+' after 'Review' in giving a label for Phase I of the class.

II *Metalogic, Including Gödel's Great Theorems.* The standard bulk of intermediate formal logic consists in a series of metatheorems that can be viewed as showing that certain metaproperties hold of certain formal logics and parts thereof. For instance, "completeness" (COMP) is often one of these properties. As we will e.g. see:

- **Theorem**: COMP[$\mathscr{L}_{\mathrm{PC}}$]
- **Gödel's Completeness Theorem/GCT**: COMP[$\mathscr{L}_1$]

One of the distinctive aspects of IFLAI2 is that its coverage of metalogic will include nearly all the great theorems of the greatest logician (Kurt Gödel); this coverage will be from Bringsjord's forthcoming *Gödel's Great Theorems*.

III *Advanced Topics in* HyperSlate®. Recall that as said above Phase I includes some new formal logics to be explored in HyperSlate®. In Phase III of the course we take a jump to *quantified* modal logic, and we also explore, to a degree, a more advanced version of HyperSlate® in which the user if allowed to write Clojure functions that play a role in proofs. (We will only consider relatively simple Clojure functions because of time limitations in our schedule.)

IV *Big Questions: e.g. Will AI Match, or Even Exceed, Human Intelligence?*. Our fourth area of coverage includes "big questions" that AI forces us to consider, if we're thoughtful. One example is given in the heading just above, but there are many others, as the student will see. The first big question we'll ponder is whether The Singularity is going to happen or not. Area IV will also include Gödel's "Either-Or" framework for considering whether standard computing (i.e. Turing-machine-level computing) can ever reach human-level intelligence. Here, Bringsjord shall draw from another forthcoming publication, a debate with Rapaport, who holds that such computing will indeed enable AI to reach human-level heights.

5.2 Fine-Grained Schedule

A more fine-grained schedule now follows.[6]

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

- **Aug 28**: *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-BDLAHSHG textbook,
  2. HyperSlate® AI-infused proof and program construction environment, and
  3. HyperGrader® AI platform (comprising HyperSlate®) for (among other things) automated assessment of proofs,

  are used. Students wishing to learn intermediate formal logic under e.g. the "Stanford" paradigm (or for that matter any other paradigm) and with its associated software and not the aforementioned trio from Motalen are encouraged to drop this LAMA®-based course and take *Intermediate Logic* at Rensselaer from another instructor.

- **Aug 31**: *Tutorials, Mechanics; Historical and Scientific Context re. Formal Logic, AI, and Logic Machines.* A rapid overview of the relevant history and background of formal logic and AI is provided; this content forms the bulk of the context for our coming investigations and learning.

- **Sep 5**: *Review of Extensional Logics.* Note, this is a *Tuesday.* The day before is Labor Day: no classes are held that day. We here first explain the core difference between extensional logics such as $\mathscr{L}_1$ and intensional logics, using Blinky the robot and a cup-switching challenge, and the infinitary False Belief Task. We then proceed to explore

  `SpecialLlamasDisjunction`

  in HyperSlate®, and other poblems.

- **Sep 7**: *Review of Intensional/Modal Logics.* We look at modal-logic **D** in review, and include (for most) new coverage of modal logics **S4** and **S5**. We also look at $\mathcal{DCEC}$ in HyperSlate®.

- **Sep 11**: *Church's Theorem, the Halting Problem, and The Singularity.* Students by this point should have HyperSlate® running on their laptops, have their codes registered, have put in their RINs to HyperGrader®, and have signed and accepted their LA. Church's Theorem tells us that theoremhood for $\mathscr{L}_1$ is Turing-undecidable (= that the *Entscheidungsproblem* is Turing-unsolvable). How do we prove this? And what implications does Church's Theorem have for the future of AI, and The Singularity in particular?

- **Sep 14**: *Completeness Theorems.* We here cover the first of Gödel's great theorems, which says that $\mathscr{L}_1$ is complete (in the sense that every necessary truth has a proof). We consider Henkin's version of this theorem as well, albeit briefly.

- **Sep 18**: *Gödel's First Incompleteness Theorem (G1).* We here cover the second of Gödel's great theorems, the one that has been most discussed and revered among those he obtained: viz., his first incompleteness theorem (G1), taking at its core Peano Arithmetic (**PA**). After first making sure that we have a handle on **PA** via exploration of it in HyperSlate®, we visit a Bringsjordian proof-oriented version of the Liar Paradox, then get clear on Gödel Numbering by simply taking note of how standard dictionaries work, and then we prove the theorem. We end by considering "astrologic" in connection with Goodstein's Theorem, which is an example the kind of mysterious arithmetic sentences that Gödel showed in G1 must exist.

- **Sep 21**: *Gödel's Second Incompleteness Theorem (G2).* We here prove Gödel's second incompleteness theorem (G2), which is a corollary of G1 (Gödel didn't bother to prove it directly in his proof of G1).

- **Sep 25**: *Gödel's Speedup Theorem and Artificial Superintelligence.* We introduce different levels of acceleration, from fast cars (such as electric ones from Tesla and Lucid) to the space shuttle to light-gas guns from NASA to the Ackermann Function to the beyond-recursive speedup of Rado's Σ (= "Busy) function. Then we explain and prove that moving from one logic to a more expressive one can secure speedup that, like Σ, gives non-recursive acceleration!

This is Gödel's Speedup Theorem, the engineering consequences of which (inquisitive and wise) modern minds are still sorting out.

- **Sep 28**: *An Introduction to Non-Axiomatic Logic and Some Quantification.* James Oswald gives an introduction to some levels of Pei Wang's Non-Axiomatic Reasoning System (NARS), a system that is central to numerous 21st-century pursuits of AGI.
- **Oct 2**: *Formal Logic, AI, Computer Science, and the Immaterial.* Formal logic, AI, and computer science all at least appear to entail some non-physical things (e.g. algorithms, infinite cardinal numbers, etc.) exist. Does the entailment go through? And if it does, does this in turn entail that we are immaterial as well? Affirmative answers to both questions are defended by Bringsjord.
- **Oct 5**: *Real Learning ($\mathcal{RL}$).* AI of today has given the world so-called "machine learning," or just 'ML' for short. But do machines doing ML actually learn? A negative answer is given, and defended; and a genuine form of learning (for natural and artificial agents), $\mathcal{RL}$, is introduced and defended. This class is based on (Bringsjord, Govindarajulu, Banerjee & Hummel 2018), among other publications.
- **Oct 9**: *No class: Columbus Day.*
- **Oct 12**: *Can We Ensure That AIs are Ethically Correct?.* This class is to a significant degree based on S Bringsjord's keynote at the 46th annual German national AI conference. The answer given to the title/question, in a word, is: When the AI in question is logicist, yes; and here's how; but in the case of today's so-called "generative" AI, the question is disturbingly open.
- **Oct 16**: *AI, Consciousness, Cognitive Intelligence, and AGI.* This class is based on work by Bringsjord & Govindarajulu in which a new theory of machine consciousness is set out and associated with a scheme (Λ) for measuring this consciousness. B&G also here articulate and analyze purported refutations of the Integrated Information Theory of consciousness advanced by Tononi & Koch, and its associated scheme (Φ) for measuring consciousness. In addition, it is explained how the concept of *cognitive intelligence* can be based upon Λ, and how this has substantive bearing on artificial *general* intelligence = AGI.
- **Supplement**: *What is Formal Inductive Logic?* This class includes compressed coverage of so-called "pure inductive logic" (PIL), which has become nearly the sole province of mathematicians and logicians, with AI activity nearly zero. Why? One reason, which we find compelling, is that PIL is devoid of proofs and arguments build on the basis of the formal structures involved. We use the "Grue Paradox" to help explain matters.
- **Supplement**: *From the Lottery Paradox to Defeasible/Nonmonotonic Logic and AI.* We present and examine The Lottery Paradox as a portal to argument-based defeasible (= nonmonotonic) logic. We also consider the suppression task as a caset study in the applicability of argument-based defeasible logic. Such logic is way to do inductive logic, including automated inductive logic, that is superior to PIL and anything based upon it.
- **Oct 19**: *AI to Surmount Arrow's Impossibility Theorem.* This class is based on (Bringsjord, Govindarajulu & Giancola 2021).
- **Supplement**: *The Argument for God's Existence from AI.* This class is based on in-*European Journal of Science & Theology* paper, a preprint of which is available here.
- **Oct 23**: *TBA.* The second lecture by James Oswald.
- **Oct 26**: *What is the Brain, Computationally and AI Speaking?* We here begin by considering the claim, defended by Richard Granger, that the human brain is fundamentally *less* than a Turing machine (and of course thus its equivalents, e.g. a register machine).
- **Oct 30**: *Logicist Agent-based Economics; AI and Tax Technology.* Can the U.S. federal tax code be captured in formal logic? If so, wouldn't that allow AI to compute minimal tax payments, and certify such payments as minimal? These and other questions are explored in this class meeting.
- **Nov 2**: *Pure General Logic Programming, Functional Programming, Turing-Completeness, and Beyond.* We review the basic paradigms

of computer programming. For the imperative case, we use the simple imperative language of (Davis, Sigal & Weyuker 1994), and also discuss register machines, Turing machines (again), KU machines. We also discuss whether programming beyond the Turing Limit makes sense and can be pursued. In this connection we explore the hierarchy $\mathfrak{LM}$.)

- **Nov 6**: *Hypergraphical Proof and Programming in HyperLog*®. Selmer goes "off script" here, to explain and show parts of an in-progress paper on the relationship between the level of autonomy in an AI, and the level of rational trust that should be placed in that AI. (An earlier slide remains available online for the original plan, which was: We here introduce the availability of writing Clojure functions in the context of proofs in HyperLog®.)

- **Nov 9**: *Quantified Modal Logic.* Selmer again goes "off script" here, to look in more detail at the opposing view of AI from Bill Rapaport, (We were scheduled to here explore quantified **S5**, including the the infamous Barcan Formula. HyperSlate® is used to make this concrete, in the $\mathcal{DCEC}$ workspace available therein.)

- **Nov 13**: *Killer Robots, **D**, and Beyond in HyperSlate® to $\mathcal{DCEC}$.* In the Fall 2023 version of the course, Selmer goes "off script" and shows students what's working on to write a while paper for funding to extend his prior work with others to ensure that LLMs are provably ethically correct — a rather tall order.

  Original plan: We begin here by stating the "PAID Problem," and then the approach to it from Bringsjord et al. advocates. We review that modal logic **D** is painfully inadequate, but now move to some exploration of a version of $\mathcal{DCEC}$ in HyperSlate®.

- **Nov 16**: *The Four Steps (incluing Logicist AI-ification of the Doctrines of N Effect) to Solve the PAID Problem.* In the Fall 2023 version of the course, Selmer stays "off script" and shows progress on said while paper (see entry for previous class).

- **Nov 20**: *Part I of . . . Gödel's Greatest Theorem: The Continuum Hypothesis.* As preparation for Gödel's Greatest Theorem, we review and expand our understanding of axiomatic set theory, and of the relative size of infinite sets. **ZFC** in HyperSlate® is visited and explored.

- **Nov 23**: **No Class** (Thanksgiving).

- **Nov 27**: *Part II of . . . Gödel's Greatest Theorem: The Continuum Hypothesis.* We turn to Sherlock Holmes for help in understanding Gödel's result that CH cannot be proved false on the basis of **ZFC**.

- **Nov 30**: *Gödel's Time-Travel Theorem.* We here visit the world known as "Flatland," and use it to articulate a visual Gödelian proof that backwards time travel is possible. We also consider the the Paradox of Proust/Looping Painter Paradox, and Bringsjord provides his analysis and solution.

- **Dec 4**: *Gödel's "God Theorem".* Did Gödel prove that God exists? We discuss this question, and look in some detail at his attempt to do so, which has become an active area in AI of today.

- **Dec 7**: *The "Games" of Gödel and His "Diophantine Disjunction".* We here assess the theorems of Gödel by considering them in connection with games measured computationally and logically, so as to answer the question (Q) as to whether an AI could ever match Gödel. We also consider Gödel's view on this question, which he connected to a certain disjunction involving Diophantine problems. Bringsjord answers Q in the negative, and provide supporting argument for this position. His position is contrasted with Bill Rapaport's contrary position.

6 Grading

Grades are based on four factors:

1. All required problems in HyperSlate®, when completed and certified correct by HyperGrader®, earns the student an A/4 for 50% of her/his final grade. Students cannot pass the course unless all these required problems are solved and certified correct. It is not expected that passing all of these problems will be onerous; in this regard, IFLAI2 is perhaps a bit different than IFLAI1. All review problems are required.
2. Answers to questions regarding metalogic/metatheorems covered. These questions may in somen cases go out be email; answers will be provided in HS®, and perhaps in some cases as pdfs submitted by email.. This will constitute 20% of one's grade. These questions may ask for (informal) proofs or proof-sketches, or — a concept to be explained — sub-proofs of steps that are/seem mysterious in metaproofs that you are presented with. At most there will be three of these questions for the course.
3. A 3-page paper written as a critique of a position on formal logic, AI, and the mind advanced by Bringsjord. (It will be easy to find a position that you vehemently disagree with. The topic must be pre-accepted by Bringsjord.) This paper will be submitted in two versions, a first version on which feedback is given, and then a final version submitted after that that takes account of this feedback. Overleaf will be used for this process (for proposing topics, clearing topics, for Selmer to write feedback, and for writing papers (thus they must be written in LaTex); space is courtesy of Motalen, and off campus/separate from any RPI technology/infrastructure. This paper is 20% of one's grade. Bringsjord's positions are expressed as declarative propositions, and will often have a philosophical dimension. As an example, here is a position that will be advanced:

   $\overline{\text{AI}^{=\text{HI}}}$ It is logically/mathematically impossible for AI (as defined today in the textbooks and primary literature of the field of AI) to match (let alone exceed) human intelligence.

4. Finally, the remaining 10% of one's grade is based on participation through discussion and email, etc. Cogent critique from students of Bringsjord's positions on "big questions" re. AI and the mind.

7 Some Learning Outcomes

There are three desired outcomes. *One*: Students will be able to/refresh their ability to carry out/execute formal proofs and disproofs, and simple pure logic programs, in collaboration with AI, 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 understand the main metatheorems of intermediate formal logic, and all of those achieved by Gödel (as enumerated above). *Three*, students will be able to debate, verbally and in cogent prose, some of the profound questions raised by AI (questions enumerated above).

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 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*

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., Govindarajulu, N., Banerjee, S. & Hummel, J. (2018), Do Machine-Learning Machines Learn?, *in* V. Müller, ed., 'Philosophy and Theory of Artificial Intelligence 2017', Springer SAPERE, Berlin, Germany, pp. 136–157. This book is Vol. 44 in the book series. The paper answers the question that is its title with a resounding No. A preprint of the paper can be found via the URL given here.
URL: *http://kryten.mm.rpi.edu/SB_NSG_SB_JH_DoMachine-LearningMachinesLearn_preprint.pdf*

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

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

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.