

Introduction to Logic-Based AI

(= ILBAI, can rhyme with “eye’ll bay”)

Crosslistings:
COGS-4961-01
COGS-6961-01
64652-CSCI-4972-01
63898-PHIL-4961-01

Fall 24, Mo Th 4:00p–5:50 NY time; DCC 330

Selmer Bringsjord

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

Selmer.Bringsjord@gmail.com

version 1111240940NY

Contents

1	Course Encapsulation	1
2	Four Key Aspects of the Course	1
3	Prerequisites	1
4	Teaching Assistant and Guest Lecturer	2
5	Readings/Films/Textbook/Courseware	2
6	Schedule	3
7	Grading	8
8	Some Learning Outcomes	8
9	Academic Honesty	9
	References	10

1 Course Encapsulation

This course is an introduction to logic-based (= logicist = logical; the three adjectives have all been and continue to be used) artificial intelligence (AI). We learn techniques for designing and engineering AIs with human-level (or higher) cognitive intelligence, enabled by automated reasoning as the cornerstone for: planning, learning, decision-making, communicating, creativity, and perceiving. A special emphasis is placed upon giving AIs intellectual powers that are beyond the reach of large language models (LLMs) like GPT-4 and other so-called “foundation models,” which, based as they are on statistical/numerical machine learning (e.g. deep learning, which is driven by data stripped of logical meaning and structure), are congenitally (and dangerously) prone to poor performance in applications that require high precision and accuracy, and/or require formally verified correct behavior. We thus direct our attention to solving the very problem currently occupying the first-rate minds of many at companies in the AI sector of the economy, e.g. Google. Thus those who succeed in this course will be in position to offer such companies skills that are increasingly sought, but are in very short supply. We explore how to remedy the deficiencies of LLMs¹ with AI based on computational logics, from the propositional calculus, through fragments of first-order logic crucial for the World Wide Web’s productive operation, on up to logics needed to model and simulate very high levels of human and machine intelligence. Our programming paradigm is S Bringsjord’s generalization and “purification” of logic programming, and is introduced and taught from scratch, starting slowly from so-called “Horn Logic.”

Fittingly, AI plays a significant role in advancing learning in the class, for instance through use of logic-based “oracles” in the AI platform we shall be using.

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.

2 Four Key Aspects of the Course

Four key aspects of the course are that:

1. a crucial source of learning in this course will be the cinematic arts, primarily belletristic sci-fi films about AI/AIs (for class meetings that are at least in part cinematic in nature, look for the 🎬 icon);
2. much of the teaching in this course will revolve around playing and analyzing fun games of logic and logical reasoning (look for the ♟️ icon in the description of our class meetings)
3. coverage of AI-relevant quantum computing, analyzed by way of formal logic; and
4. coverage as well of not only “standard” AI, but also so-called artificial /general/ intelligence (AGI).

3 Prerequisites

Standard high-school math progression with Algebra 2 (or equivalent) through some calculus; some prior study of formal logic and proofs; and some prior programming (in at least one or more

¹See e.g., GPT-4 simply can’t reason, period; see e.g. from Arkoudas: “[GPT-4 Can’t Reason](#)”.

procedural or functional languages; no prior experience with logic programming necessary). No particular courses must've been taken in order to qualify.

4 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 already an internationally recognized expert on AI, especially AGI. He will be not just a TA in this course, but operator as a lecturer. His email address is oswalj@rpi.edu, and James' office hours are Thursdays 2–4 pm, Carnegie 3rd floor, in the open meeting area there, which is nicely equipped with projection technology.

5 Readings/Films/Textbook/Courseware

This course is based directly on the in-progress textbook *Introduction to Logic-Based AI* by S. Bringsjord, Naveen Sundar Govindarajulu, & A. Bringsjord. This in-progress textbook corresponds *directly* to the slide decks/slides used in the class. These slide decks and corresponding lectures/tutorials (including in some cases video versions of such) compose crucial content for this course, and will be linked-to from the course web page.

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.

We have a recommended but fully optional textbook, the fourth edition of *Artificial Intelligence: A Modern Approach, 4th Edition* (Russell & Norvig 2020). While this book is on the bookshelf or every serious AI developer/researcher in the world today (and should be located in such a position in the case of managers/executives in and partaking of AI), the key logic-based parts of the book will be covered in concrete fashion in our use HyperSlate[®] 4.0, since the platform now has many built-in problems from AIMA4E. In addition, students should read the overview of AI provided [here](#) (Bringsjord & Govindarajulu 2018), which follows AIMA4E very closely.

Finally, students who opt to stay in the course after all mechanics and policies are shared on Sep 3 will purchase a license giving access to the inseparable and symbiotic triadic combination published and maintained by Motalen:

- access to and use of the HyperSlate[®] 4.0 AI interactive AI environment (for, among other things, engineering proofs and logic programs in collaboration with AI “oracles”). This environment is available on said platform.
- the e-textbook *Logic: A Modern Approach; Beginning Deductive Logic, Advanced via HyperSlate[®] and HyperGrader[®]* (LAMA-BDLAHSBG);
- access to and use of the HyperGrader[®] 4.0 AI platform (for, among other things, assessing student work as the course proceeds).

All three items will be available after purchase in the RPI Bookstore of a barcoded envelope with a personalized starting code/key for registration. Logistics of the purchase, and the contents of the envelope that purchase will secure, will be encapsulated in the class meeting of Sep 3 2024, and then gone over in more detail repeatedly in class, including on Sep 5 & Sept 9 2024. The codes will not be available for purchase until at least Sept 10 2024. The first use in earnest of HyperSlate[®] 4.0

and HyperGrader[®] 3.0 in class will happen on Sep 12 2024, so certainly after that class students should have LAMA-BDLAHSYG, and be able to open both HyperSlate[®] and HyperGrader[®] on their laptops, with the slide deck for that day (and video tutorials) to help get started.

Updates to LAMA-BDLAHSYG, 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. These files will be in the Cloud for you. You will specifically need to assemble your own personalized library of completed and partially completed proofs/arguments/programs/truth-trees etc. in the Cloud provided to you, so that you can use them as building blocks and a reference guide as your progress. Building up your own “ILBAI 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.

6 Schedule

Our schedule now follows.²

²Note that the Rensselaer Academic Calendar is available [here](#).

- **Aug 29:** *General Orientation.* The essence of the course is communicated and discussed, impressionistically. This specifically includes the “cinematic nature” of the course, and our first important film, *Blade Runner* (original), is cited in connection with the Worldcoin initiative afoot in the world today — the central issue being how do humans tell humans apart from AIs.
- **Sep 3** 🎬 ⚙️: *Syllabus and Course Mechanics.* There is first a recap of the general orientation from the prior meeting, and Q&A on that. Next, *Blade Runner* and Worldcoin etc. are discussed. Then, passing to the rather more tedious for the remainder of this meeting (which is obligatory), the syllabus is reviewed. It’s explained that logic-based AI will here be taught in connection with crucial help from Motalen’s

HyperGrader[®]

AI platform (comprising HyperSlate[®]) and an e-textbook), and also ShadowProver (automated deductive reasoner) and Spectra (automated planner).

- **Sep 5:** *What is AI/Logic-Based AI; The History of Logic-Based AI.* We quickly lay out an informal but serviceable version of the “intelligent agent” definition of what AI is, and explain how it is refined and specified to as to produce a definition of logic-based AI. We then use a century-spanning timeline to achieve an overview of the history of logic-based AI. We also look a bit at Martin Gardner’s *Logic Machines*.
- **Sep 9** 🧩: *The Failure of Deep Learning; GPT-4o Bites the Dust.* We challenge GPT-4o with some problems that are cakewalks for logic-based AI systems, and observe well its abject failure. The problems used serve to provide the start of a gentle introduction to the logic known as the *propositional calculus*, or just \mathcal{L}_{pc} — and we also take a look ahead to the problem of layered quantification for (artificial) neural network-based AI, which 45 years ago was laid down as the chief challenge to such AI by the logicians. This challenge is remains completely unmet today. We note, quickly, that the AIs in HyperSlate[®], known as *oracles*, have no trouble at all with the problems that cripple GPT-4o.
- **Sep 12** ⚙️: *The Propositional Logic via Logical Journey of the Zoombinis.* We look more formally at the propositional calculus. Then we see how this calculus can be brought to life by play in /textitZoombinis, and use of HyperSlate[®]. Students by this point should have HyperSlate[®] 4.0 running on their laptops, have their codes registered, have put in their RINs to HyperGrader[®], and have signed and accepted their LA.
- **Sep 16:** *Automated Reasoning in the Propositional Calculus; Our PC Oracle* Having said artificial agents in logic-based AI are powered fundamentally by automated reasoning as computation, we now look at this reasoning in more detail. The game Catabot Rescue helps us understand matters.
- **Sep 19:** *Zero-Order Logic and AI.* We now move up from the propositional calculus = \mathcal{L}_{pc} to zero-order logic = \mathcal{L}_0 . Can we still use *Zoombinis*? How do we handle names and identity?!
- **Sep 23:** *What About AGI?* Before passing on to more powerful logics as additional bases for logic-based AI, we introduce a more ambitious form of AI. So far we have discussed AI, not AGI. What is AGI? How is it different from its two-letter cousin? Is there a cinematic way to portray and bring understanding of AGI versus AI? What AIs in film have AGI? James Oswald lectures, and he will present and discuss The Singularity, that moment in future time when AIs suddenly create AIs with superhuman intelligence and leave us humans in the dust. (S Bringsjord is skeptical. Oswald will include coverage of S Bringsjord’s (Bringsjord 2012) argument that the Singularity is essentially mathematically impossible.)
- **Sep 26:** *Human Disemployment at the The MiniMaxularity.* Will AI disemploy humans, or at least most of them? (By ‘disemployment’ is meant here not job displacement. We understand displacement to be consistent with finding another, different job. Disemploying some human is to make that human unemployable.) The MiniMaxularity, in a word, is that time in the future when AI is able to do well all that it currently is doing to an appreciable degree, but

not well (e.g. drive vehicles on tricky roadways during inclement weather).


- **Sep 30: *First-Order Logic I*.** Now we move on from the propositional calculus for AI to first-order logic (FOL), or — to use the fancy abbreviation — \mathcal{L}_1 . This means that we now have the all-important two quantifiers in play, \exists (for, with a variable x immediately after it, approximately ‘there exists a thing x such that’) and \forall (for, with a variable x immediately after it, approximately ‘for all things x ’). FOL is an absolutely crucial part of logic-based AI; its importance can’t be overstated.
- **Oct 3: *FOL 2; Oracle Power; ShadowProver*.** In our second meeting we make some rather intense use of the FOL oracle in X, and in addition the automated deductive reasoner ShadowProver is introduced, and made available.
- **Oct 7: *FOL 3 and The Failure of ML, Round 2; Intensional Logic*.** We wrap up our coverage of $\text{FOL} = \mathcal{L}_1$ here, and provide an overview of logics that can model cognitive attitudes (knowing, believing, intending, desiring, communicating, perceiving, etc.) and the human emotions. We assess how deep neural nets, LLMs etc. do when confronted with quantifier complexity, the same challenge issued over four decades back to AI systems based on artificial neural networks. We also see how they do on challenges involving cognitive attitudes emotions. Selmer’s Pawn Shop microworld is introduced and used, after the false-belief test (FBT) is introduced. For an introduction to FBT, and the handling of the first level of this test by an AI system, see (Arkoudas & Bringsjord 2009). HyperSlate[®] 4.0 is used to make intensional logic coverage more concrete, in the


DCEC

workspace available therein.

- **Oct 10: *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.

Now we see clearly for the first time that we shall need more than deductive logic for logic-based AI.

- **Oct 14: *No class: Columbus Day*.**
- **Oct 17: *Argument-Based Inductive Logic for Logic-Based AI; ShadowAdjudicator Introduced*.** 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 motivate matters. We also cover here S Bringsjord et al.’s approach to inductive logic in the argument-based mode. In this line, case studies are included, e.g. (Bringsjord, Govindarajulu & Giancola 2021, Bringsjord, Giancola, Govindarajulu, Slowik, Oswald, Bello & Clark 2024). The logicist AI system capable of inductive reasoning, ShadowAdjudicator, is introduced.
- **Oct 21 : *AI, Consciousness, and Universal Cognitive Intelligence*.** 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 the theory of cognitive consciousness and Λ can be used to build the paradigm of *universal cognitive intelligence* in the science of universal intelligence (Bringsjord, Govindarajulu & Oswald 2023).
- **Oct 24: *Logic-Based Automated Planning*** This class introduces one of the central parts of logic-based AI: automated planning. Here, an AI must automatically find a plan, using logical techniques. The automated planner Spectra is introduced, and made available to the class. We investigate how well LLMs can do on planning tasks (spoiler alert ... not well). On this last topic, James Oswald shares his experience and expertise.

- **Oct 28:** *Hutter's (Logic-less!) Theory of Universal Artificial Intelligence, With Harder Environments for AGIs.* James Oswald presents one of the most prominent paradigms for rigorously modeling and systematizing the science of universal intelligence. He includes discussion of how he led the way in recently expanding this paradigm to include provision for a given intelligent agent to find itself in very demanding environments (e.g. environments in which problems at the level of first-order logic are encountered).
- **Oct 31:** *The NARS Approach to AGI.* 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. Because NARS isn't exclusively deductive, the stage has been suitably set via our prior coverage of inductive logic.
- **Nov 4:** *TBD.* James Oswald will give a lecture here on a topic on the frontiers of AGI and AI research. The topic will be inserted in a subsequent version of the present syllabus.
- **Nov 7:** *AI and God.* The bottom line is that plenty of people believe God exists — and should they be a bit concerned about the amazing advance of AI, since presumably AIs don't have souls and such? At the same time, there are many atheists around — and should *they* be concerned about AI as a threat to their beliefs? What about the ethical realm? It seems that the majority of people on Earth base their ethics on supernatural conceptions — but how will this work for an AI-infused world? This class tackles these questions, with a special focus on AI & idolatry (in connection with /TOS/ episode "The Apple"), and includes coverage (and hopefully debate!) about S Bringsjord's *European Journal of Science & Theology* paper, a preprint of which is available [here](#).
- **Nov 11:** *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). From there we move to tackling the great threat to those who adopt logics as way to make sense of human intelligence: Isn't the brain what we are? And look, the brain hasn't got any logic in it whatsoever! It's a neural network!
- **Nov 14:** *Logician Agent-based Economics (LAE)* Economics itself, at least microeconomics, can be based on logic-based AI; this is explained in a single class. The class is anchored by classic decision-theory "puzzlers," such as the St Petersburg Paradox, the Ellsberg Paradox, and the Monty Hall problem. We also consider some applied questions: e.g., 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?
- **Nov 18:** *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 Σ that is part of universal cognitive intelligence.)
- **Nov 21:** *Reflections on Logic-Based AI and CogSci in the Context of Quantum Computing.* Is quantum computing more powerful than standard "Turing-level" computing? Can formal logic help us answer this question? Is there such a thing as "quantum cognition," as some have claimed? What is it?
- **Nov 25:** *Gödel's Great Theorems in a Single Class!* Is this really possible? Yes! Coverage includes the most famous theorem, the first incompleteness theorem, which stands at the heart of such unforgettable books as *Gödel, Escher, Bach*. But all the other great theorems proved by the greatest logician (and likely mathematician as well) who ever lived are covered.
- **Nov 28:** **No Class** (Thanksgiving).
- **Dec 2**  **4:** *Will Killer Robots Kill Us All?* We begin here by stating "The **PAID** Problem," which shows us that the answer to this question is "Yes, unless logic-based AI is suitably brought in to rescue us. In short, only logic can

save us. We need to draw from the cinematic arts to grasp the danger/destruction humanity faces — and we shall do so. (A number of movies are currently under voting and consideration.)

- **Dec 5:** *The Four Steps to Solve the **PAID** Problem — And Two Remaining Challenges to Save the Human Race* We here discuss The Four Steps that offer a saving road to humanity, a road that will avoid The **PAID** Problem. But! — two problems remain unsolved at this point, and must be cracked in the near future: viz., the problem of engineering human-level perception into AIs, and also *moral creativity* into them as well.
- **Dec 9** 🍷: *The Future of AI & Gödel’s “Diophantine Disjunction”; S Bringsjord vs. Rapaport.* We here seek to answer the question (Q) as to whether an AI could ever match or even exceed the human mind. We consider Gödel’s view on this question, which he connected to a certain disjunction involving Diophantine problems. Bringsjord answers Q in the negative, and provides a case for this position. Bringsjord’s position is contrasted with Bill Rapaport’s contrary stance that AI will minimally match the human mind.

7 Grading

All students are expected to earn a grade of A (or of course Pass if taking the course in that mode). However, class attendance is required. Grades are based on six factors:

1. Attendance and participation in class will earn an A for 20% of one's grade. Participation will gradually get easier for everyone, and will be crucially facilitated by reading and watching (movies). It's important for the survival of the human race that sufficiently intellectually talented human beings learn: what AI of the logic-based/logicist/logical variety (and the "ML" numerical variety — but that you'll have to get elsewhere) is, judge/assess it, and openly discuss and debate it with other humans in a search for wise human behavior in the face of AI's advance.
2. All required problems in HyperSlate[®] 4.0, when completed and certified correct by HyperGrader[®] 3.0, earns the student an A/4 for 20% 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 at all onerous; in this regard, help sufficient to ensure success will be provided.
3. Decent performance on live MCQ problems in HyperGrader[®] 3.0, with help and discussion being given on the spot in class. 20% of one's grade.
4. A 2-page paper (not counting references) written as a critique of a position on AI, and the mind advanced by S Bringsjord. (It will be easy to find a position that you vehemently disagree with. The topic must be pre-accepted by S Bringsjord.) 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). 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:

\bar{S} It is logically/mathematically impossible for AI (as defined today in the textbooks and primary literature of the field of AI) to match and then exceed human intelligence in the event known as "The Singularity."
5. Finally, a final project will be submitted. Students will have great flexibility in what their final project, and will be encourage to choose what they do based on their backgrounds and interests. 20% of one's grade.

8 Some Learning Outcomes

There are three desired outcomes:

- *One*: Students will be able to understand in general how to understand and — to a degree — carry out/execute proofs and disproofs and arguments and logic programs in logics relevant to Logic-Based AI, in collaboration with AIs available within the HyperSlate[®] 4.0 system.
- *Two*: Flowing in no small part from outcome one, students will understand the main formalisms, modes of computation, logic-programming languages/environments of Logic-Based AI — and they will thus be in position to bring these things to the attention of their colleagues

in the marketplace outside of and after this course. (In short, they will be able to ride, to an appreciable degree, the two horses of neural/statistical AI (Numerisk, in the textbook that is this course) and logic-based = logicist = logical AI (Logisk, in the textbook that is this course) in their careers upon graduation from this course and from their degree program.)

- *Three*, students will be able to understand, ponder further debate, verbally and in cogent prose, many of the profound, pressing questions raised by AI of today.

9 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 answers to MCQ problems, and in general solutions submitted on the HyperGrader[®] 3.0 AI platform 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), ‘Propositional Attitudes and Causation’, *International Journal of Software and Informatics* **3**(1), 47–65.
URL: http://kryten.mm.rpi.edu/PRICAI_w_sequentialcalc_041709.pdf
- Bringsjord, S. (2012), ‘Belief in The Singularity is Logically Brittle’, *Journal of Consciousness Studies* **19**(7), 14–20.
URL: http://kryten.mm.rpi.edu/SB_singularity_math_final.pdf
- Bringsjord, S., Giancola, M., Govindarajulu, N. S., Slowik, J., Oswald, J., Bello, P. & Clark, M. (2024), ‘Argument-Based Inductive Logics, With Coverage of Compromised Perception’, *Frontiers in Artificial Intelligence* **6**.
URL: <https://www.frontiersin.org/articles/10.3389/frai.2023.1144569/>
- 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. & Govindarajulu, N. S. (2018), Artificial Intelligence, *in* E. Zalta, ed., ‘The Stanford Encyclopedia of Philosophy’.
URL: <https://plato.stanford.edu/entries/artificial-intelligence>
- Bringsjord, S., Govindarajulu, N. S. & Oswald, J. (2023), Universal Cognitive Intelligence, from Cognitive Consciousness, and Lambda (Λ), *in* A. Chella, ed., ‘Computational Approaches to Conscious Artificial Intelligence’, Vol. 5 of *Machine Consciousness*, World Scientific Publishing, Singapore, pp. 127–168. The URL here is to an uncorrected preprint.
URL: <http://kryten.mm.rpi.edu/ch5-main.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.
- Russell, S. & Norvig, P. (2020), *Artificial Intelligence: A Modern Approach*, Pearson, New York, NY. Fourth edition.