

# PERI.2 Goes to PreSchool and Beyond, in Search of AGI

Selmer Bringsjord<sup>(✉)</sup>

, Naveen Sundar Govindarajulu, John Slowik James Oswald  
Michael Giancola, John Angel, Shreya Banerjee, and Aidan Flaherty

Rensselaer AI & Reasoning (RAIR) Lab, Rensselaer Polytechnic Institute (RPI),  
Troy NY 12180, USA

{selmer.bringsjord,johnkslowik,jamesoswald111,mike.j.giancola,  
angelson1992,shreyabbanerjee,aidan.n.flaherty}@gmail.com

**Abstract.** After introductory remarks, we share our two-part theoretical position, viz. that: (P1) The best overarching approach to suitably defining GI, and obtaining AGI, is via formal logic, including specifically via logic-based learning that is academic in nature; and (P2) AI/AGI is best pursued by seeking artificial agents that pass determinate cognitive tests. We note that in striking harmony with this position is work on AGI by Goertzel et al. that has inspired us; this is work in which PreSchool for would-be AGIs provides an attractive route toward AGI itself. While Goertzel et al. envisage a virtual academic environment, we have in mind physical classrooms, for physical robots. We describe the robot PERI.2, which we have started to send to school.

## 1 Introduction

However one might prefer to define AGI, it seems likely to be a matter of consensus that we have GI,<sup>1</sup> and that you do too.<sup>2</sup> Why are we so fortunate? Many reasons, often competing ones among them, will be offered. One prominent reason, it seems to us, is this: Because we all went to school, year after year, for many years, and learned a lot in the process; and we went there *physically*. From a high-altitude perspective, the present paper revolves around this reason.

The plan for the paper is in general as follows: We begin in the next section (2) by confessing our two-part theoretical position, namely that

---

<sup>1</sup> We note here one vocal objection to that consensus: Yann LeCun has claimed that humans do not have general intelligence [24]. He discusses a hypothetical scenario wherein a human’s visual field is permuted as an example of our lack of general intelligence, arguing (it seems) that the ability to learn this permutation is required of anything which could be considered “general” intelligent. While an attempted refutation of LeCun’s position is out of scope for this paper, we do volunteer here that this “permutation skill” is clearly not particularly intelligent by any reasonable definition of the word (let alone by any reputable test of intelligence/cognitive ability we are aware of), and hence any definition of general intelligence which requires it as a prerequisite is not one we find at all plausible.

<sup>2</sup> If some of our readers are artificial, and not human persons, then they have AGI.

- P1 The best overarching approach to suitably defining GI, and obtaining AGI, is via formal logic, including specifically via logic-based learning that is academic in nature.
- P2 AI/AGI is best pursued by seeking artificial agents that pass determinate cognitive tests.

We then note that in harmony with this position is inspiring work on AGI by Goertzel et al. [17, 18], in which PreSchool (and, in general, grade levels progressing beyond this into at least K–12) for would-be AGIs presents an attractive engineering route toward AGI itself. While Goertzel et al. envisage a virtual academic environment, we have in mind physical classrooms, for physical robots. After explaining our focus on “logical/mathematical” cognition within the underpinnings of Goertzel’s approach, we describe the robot PERI.2, which we have started to send to school in the hopes of it developing AGI. We give an example of an academic challenge for PERI.2 in the logical/mathematical category at the Kindergarten level. PERI.2 succeeds upon this challenge, but, as we admit, much additional work will be needed.

## 2 Our Two Theoretical Pillars

We fully recognize that there isn’t exactly consensus regarding how best to reach AGI. In some cases, for example, non-declarative learning is believed to provide the most, perhaps even the *only*, route to AGI; an exemplar, by our lights, would be [23], an (impressive) approach described in a manner wholly bereft of formal reasoning over declarative knowledge or belief for the intelligent agents in question.<sup>3</sup> However, for better or worse, as the next section confesses, we (or at least the first author) feel differently.

### 2.1 Pillar 1: Logic-based AI and Cognitive Science

The first author has long maintained that logic-based AI is superior to methodological competitors [see e.g. [7]]. Re. computational cognitive science, the unmatched effectiveness of logic-based effort, at least for cognition, has likewise been asserted [see e.g. [11, 6, 8]]. Overall, we posit an infinite collection  $\mathfrak{L}$  of logics (we call them *cognitive calculi*) reasoning in which can constitute any level of GI whatsoever. (Standard logics still used in AI include first-order logic  $\mathcal{L}_1$ , second-order logic  $\mathcal{L}_2$ , etc.) In particular, it seems indubitable that at least for every aspect of human-level cognition that is reasoning-centric, there exists in some cognitive calculus  $\mathcal{L} \in \mathfrak{L}$  that can be tokened, specified, and implemented for concrete use in AGI; for this in action, see e.g. the novel logics specified and implemented in [11]. For notational convenience in the remainder of the present paper, we assume a particular cognitive calculus  $\mathcal{L}^*$  for the AGI science and engineering devoted to PERI.2 we describe and report herein — but for economy forego providing formal specification of  $\mathcal{L}^*$ . For details regarding cognitive calculi, see the Appendix in [9].

<sup>3</sup> On the other hand, among prominent AGI researchers, we are incidentally not alone in our emphasis on logic-based r&d; see e.g. [30], to which we return below.

**Real Learning is Academic Learning** Under the umbrella of logic-based AI & CogSci, we specifically hold that academic learning of and by formal logic and mathematics is key to AGI [3] — and it’s this part of our orientation that aligns with the work of Goertzel et al. (see below).

**Logicist Cognitive Robotics** As to robotics, the logicist approach to it advocated and pursued by the first author can be quickly summed up by tightening the concept of cognitive robotics as defined in [25], wherein it is said that such robotics produces robots whose actions are a function of what they believe. In line with this, but expanded in keeping with  $\mathcal{L}^*$ , we seek to engineer robots all of whose substantive decisions and actions are the result of automated reasoning to these decisions and actions from formulae in some set  $\Phi$  of formulae in  $\mathcal{L}^*$  known or at least believed by these robots, where such knowledge and belief can vary in strength depending upon the underlying likelihood of the formulae in  $\Phi$ .

## 2.2 Pillar 2: Psychometric AI

Our second theoretical pillar is that AI, and AGI, should be fields devoted exclusively to creating and implementing artificial agents able to excel on established tests of cognitive ability and skill, including those used in the Academy for humans; see e.g. [5, 4]. Most recently, this aspect of AI has been used in [1] to have success in solving Bennett mechanical test problems by artificial agents.

## 3 The Goertzelian (et al.) Academic Road to AGI

In general, we seek to follow the road to AGI paved by a progression through academic grade levels at least akin to the progression that brought the reader to a position in which they can understand the present paper; the progression of which we speak has been seminally described in [17, 18].

## 4 PERI.2 in Kindergarten

We give a snapshot of an example of PERI.2 in Kindergarten, being tested in the area “Logical-Mathematical.” (See Figures 1 & 2.) This area is listed in [17] as a specific kind of intelligence according to Gardner’s [13] theory of “multiple intelligences,” and is obviously — given Pillars 1 & 2 for us — pivotal for us.<sup>4</sup>

### 4.1 Automated Reasoning of a Meta-Forms Problem/Solution Pair

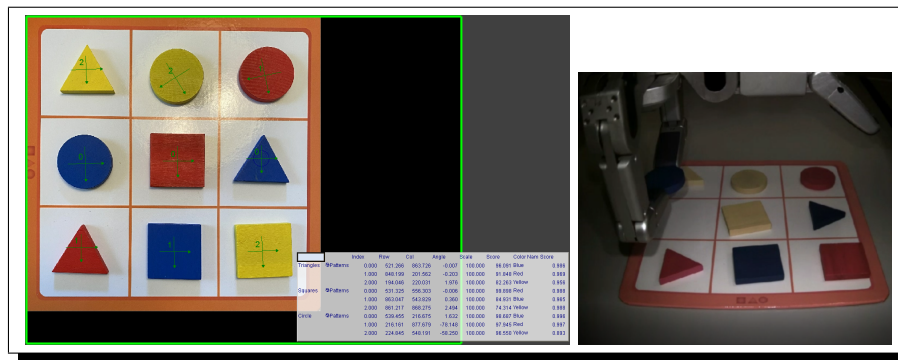
PERI.2 employs the automated deductive reasoner ShadowProver [20] to verify proposed solutions to a given Meta-Forms problem; see Figure 1.<sup>5</sup> Specifically,

<sup>4</sup> [17, 18] also point out that this area finds its way into early education.

<sup>5</sup> ShadowProver has long been used to engineer logic-based intelligent artificial agents in our lab. A robust example can be found e.g. in [19]. While ShadowProver’s rea-



**Fig. 1.** The Meta-Forms Game, from FoxMind. *This game provides a series of “clues” to the would-be puzzle solver, each of which is a visual version of a “logical statement.” The goal is to physically construct a complete configuration of the 3×3 board from these clues. Formally, if  $\Pi$  is a complete configuration of the board, and  $\Gamma$  the collection of formulae that logicize all clues, then necessarily  $\Pi \cup \Gamma$  is provably consistent in  $\mathcal{L}^*$ .*



**Fig. 2.** PERI.2 Sees the Board (left), and Holds a Meta-Forms Piece in One Hand (right). *Machine vision for PERI.2 courtesy of Cognex; hands are from Barrett.*

with a formalization of the clues given to the problem in this figure — as set  $\Gamma$  of formulae in  $\mathcal{L}^*$  — and the proposed configuration of the board shown in this figure — as set  $\Pi$  in  $\mathcal{L}^*$  — ShadowProver was tasked to find a proof of a contradiction (i.e.  $\zeta \wedge \neg\zeta$ ) from  $\Gamma \cup \Pi$ . It failed to find a proof in 3.16 seconds, which entails that the configuration of the board is consistent with the clues, and hence a solution.

soning is deductive, it is the basis for types of reasoning we believe are key to AGI r&d, e.g. nonmonotonic/defeasible reasoning. See [10] for an example of an inductive logic and an inductive automated reasoner (ShadowAdjudicator).

## 5 PERI.2, Concretely: A Glimpse

PERI.2 has a pair of dexterous, tactile-sensing Barrett hands attached to powerful Yaskawa arms, which together provide somatic information to his “mind.” For vision, PERI.2 has three Cognex cameras that compose a system for sight capable of several fundamental operations, ranging from simple object recognition given a training example, to color and blob identification, to edge detection and measurement, and beyond. Given our approach to AGI, all such information ultimately is expressed in formulae of  $\mathcal{L}^*$ . In the case of its tackling a Meta-Forms problem, PERI.2 must ultimately transduce the clues it receives for the problem into the formulae (represented internally as s-expressions) composing  $\Gamma$  (see Fig. 1). For example, a clue might appear as `(and (space 0 0) (blue 0 1) (triangle 0 1))`, suggesting that the blue triangle will have a space beneath it. Presently PERI.2 is rigidly engineered to solve Meta-Forms, but even with our target for AGI restricted to the math/logic category, the fact is that Kindergarten presents challenges (e.g.  $2 + 2 = ?$ ) that are arithmetic in nature. Accordingly, our cognitive calculus  $\mathcal{L}^*$  subsumes Peano Arithmetic, but this dimension must be left aside here.

## 6 Related Work

As is well-known, AGI can generally be classified as the field that explores the creation of computational agents possessing some level of *general intelligence*: the ability to exhibit complex problem-solving capabilities in an arbitrary environment, akin to the ability of humans (but not necessarily at the same level as humans) [15, 14, 33]. As AGI focuses on a broad overarching goal, inevitably there are many camps in AGI, each based upon its own approach to the problem [12, 14]. Obviously, camps that are not overtly logicist bear little connection to our approach to AGI. Nonetheless, a simple triadic breakdown of approaches in AGI helps to contextualize the work discussed herein; this is particularly so for the first element of the trio in question, which is:

- **The Symbolic<sup>6</sup> Approach.** Here logic is in fact the basis for memory and reasoning. Knowledge in these systems consists of statements from which new knowledge can be derived by logical reasoning. New statements may also be added by way of fully logic-based perception (e.g. see [32]). Different approaches use different ontologies and different logics with different properties to optimize for the type of reasoning to be executed [21]. Invariably, at least so far, relative to the calculi  $\mathcal{L}$  upon which our AGI r&d is based, logics in this approach to AGI by others are inexpressive, and reasoning is correspondingly simple. In particular, often representation and reasoning in this AGI approach can be reduced to information and processing in (perhaps with tailor-made inference schemata as needed) at the level of only  $\mathcal{L}_1$ , augmented perhaps with a few intensional operators.

<sup>6</sup> Since all symbolic information and processing in AI/AGI can be carried out in a formal logic, feel free to replace ‘Symbolic’ here with ‘Logicist’ or ‘Logic-based.’

- Some notable members of the symbolic camp are Wang’s NARS [31] system and Shapiro et al.’s SNePS and GLAIR architectures [27]; all three encode symbolic representations of knowledge into a graph representation.
- **The Emergent Approach:** This approach focuses on creating agents whose memory and learning take the form of connectionist systems. The emergent approach assumes, naturally enough, an emergent hypothesis: that symbolic reasoning and learning can emerge from basic connections and interactions between nodes, as they perhaps do (at least in part) in the human brain. *Contra* the logicist approach, “knowledge” in emergent systems is encoded within the weights and connections between nodes of a network, which may evolve over time for “learning.”
  - **Hybrid Approaches:** Hybrid AGI systems aim to combine emergent and symbolic approaches. According to [12], hybrid approaches suffer from the same shortcomings as emergent approaches: they have “difficulty in realizing higher-order cognitive functions” such as reasoning over arbitrarily complex/iterated declarative content, which is the hallmark of our  $\mathcal{L}$ .

AGI stands in stark contrast to today’s mainstream “narrow” AI systems, usually machine-learning models trained on massive datasets to excel in one particular task. For our logicist approach to AGI it is important to contextualize “human-level.” Human-level AI can be thought of as a goal of AGI, but from the standpoint of our approach to AGI it is only a point on a spectrum of general intelligence that AGI agents fall on. AGI researchers of either a thoroughlygoingly or even substantive logicist bent can presumably locate their ambitions for future AGI systems in the standard hierarchies (Arithmetic, based on  $\mathcal{L}_1$ ; and Analytical, based on  $\mathcal{L}_2$ ). In our approach to AGI, because we have a scheme for measuring intelligence (*viz.*  $\mathcal{A}$ ; see [2]), we can quantify very well where the level of given agents fall. One particular point worth noting here is that while we are inspired and guided by Goertzel, his conception of intelligence [16, p.5] stands in contrast to ours, since he writes that “Intelligence in general must be considered as an open-ended phenomenon without any single scalar or vectorial quantification.” This runs completely counter to the spirit and specifics of our approach to AGI. Consider e.g. the fact that we commonly compare the intelligence of human and nonhuman animal agents at least roughly in line with how academic learning and the test-measured success of such learning works. Consider for instance the common view that humans are more intelligent than dogs. It seems more than reasonable that the intuitive concept of intelligence underlying such a view is some sort of a single scalar. If a human is asked why he believes that humans are more intelligent than dogs, takes the question seriously, and tries to justify it, it appears to us likely that the rationale provided will in some way appeal to cognition measured in traditionally academic ways. Canines are smart, but as we all know, they don’t start to learn to read, nor do they learn basic arithmetic, these being things routinely taught in PreSchool.

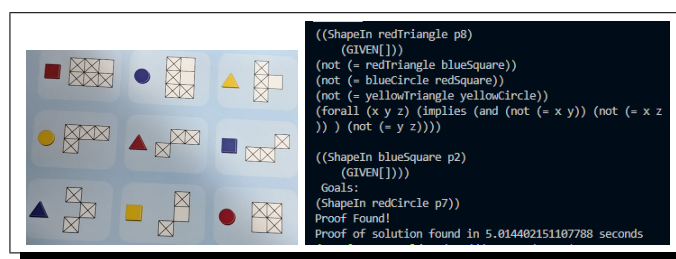
### 6.1 Remarks on NARS w.r.t. Our Theoretical Pillars

NARS [31] is profitable to consider in relation to our two-pillar approach to AGI. NARS is fundamentally logicist in nature and has working implementation that

can solve some preschool-level problems. With regard to our first pillar, Wang has argued for the need for “cognitive logic” rather than “mathematical logic” when capturing human reasoning, and claims that the non-axiomatic logic used in NARS embodies this reasoning [29]. We agree for the most part that cognitive logics are necessary, but hold e.g. that “real learning” in a cognitive logic still needs tools from logics used in mathematics [e.g.  $\mathcal{L}_1$ ,  $\mathcal{L}_2$ , and  $\mathcal{L}_3$ ] to give any serious treatment of what humans do when they go to K–12 school (since a large part of that schooling is in none other than mathematics, and reverse mathematics [28] has disclosed that mathematics ultimately consists of proofs and other structures built from formulae in first-to-third order logic). In taking a cognitive approach, NARS has demonstrated a level of competence in the areas of simple “spatial-visual” and “logical-mathematical” tasks [and has worked toward some basic “linguistic” tasks (see e.g. [22]). Due to the cognitive nature of NARS and its ability to represent knowledge, belief, and self [34], it would in theory be able to realize our working definition of logicist cognitive robotics from [25]. As to our second pillar, Psychometric AI [5], NARS is not necessarily at odds with it, but is focused on achieving intelligence in line with Wang’s working definition of intelligence as “the capacity of an information-processing system to adapt to its environment while operating with insufficient knowledge and resources” [33]. PAI provides by definition a means of meaningfully evaluating incremental progress toward AGI (viz. tests); Wang’s definition doesn’t supply such means.

### 7 Are Harder Problems Computationally Feasible?

For alert readers who may be wondering, Fig. 3 shows that harder Meta-Forms problems are within PERI.2’s intellectual reach, in real time.



**Fig. 3.** A Difficult Meta-Forms Problem. *No positive clues are given (left), yet a proof of the correctness of PERI.2’s proposal found by ShadowProver (right).*

### 8 Future Work: What About Compromised Perception?

Generally intelligent agents are capable of perceiving that they are mis-perceiving, as e.g. when they perceive rather dense smoke, and perceive that their sensors

are therefore compromised; such a situation is shown in Figure 4, for PERI.2’s attempt to perceive Meta-Forms clues. Currently success in this case eludes us (and thus PERI.2), but a new cognitive calculus that formalizes such meta-reasoning is under development, one that reflects the computational science of attention and perception erected by Bello & Bridewell et al. (e.g. see [26]).



**Fig. 4.** Perception of Compromised Perception. *Here the set  $\Gamma$  of clues for the Meta-Form problem are hard to reliably perceive due to ambient smoke.*

### Acknowledgements

We are grateful to: AFOSR for a DURIP award to Bringsjord & Govindarajulu that has brought PERI.2 to physical life, to ONR for sponsorship of r&d devoted to meta-cognitively perceptive artificial agents (award # N00014-22-1-2201), and to both AFOSR (currently award #FA9550-17-1-0191 ) and ONR for longstanding support of r&d in automated reasoning, planning, and logic-based learning.

### References

1. Banerjee, S., Bringsjord, S., Giancola, M., Govindarajulu, N.S.: Qualitative Mechanical Problem-Solving by Artificial Agents: Further Progress, Under Psychometric AI. In: The International FLAIRS Conference Proceedings. vol. 35 (2022)
2. Bringsjord, S., Govindarajulu, N.: The Theory of Cognitive Consciousness, and  $\Lambda$  (Lambda). *Journal of Artificial Intelligence and Consciousness* **7**(1), 155–181 (2020), [http://kryten.mm.rpi.edu/sb\\\_nsg\\\_lambda\\\_jaic\\\_april\\\_6\\\_2020\\\_3\\\_42\\\_pm\\\_NY.pdf](http://kryten.mm.rpi.edu/sb\_nsg\_lambda\_jaic\_april\_6\_2020\_3\_42\_pm\_NY.pdf), The URL here goes to a preprint of the paper.
3. Bringsjord, S., Govindarajulu, N., Banerjee, S., Hummel, J.: Do Machine-Learning Machines Learn? In: Müller, V. (ed.) *Philosophy and Theory of Artificial Intelligence 2017*, pp. 136–157. Springer SAPERE, Berlin, Germany (2018), [\tiny{http://kryten.mm.rpi.edu/SB\\\_NSG\\\_SB\\\_JH\\\_DoMachine-LearningMachinesLearn\\\_preprint.pdf}](http://kryten.mm.rpi.edu/SB\_NSG\_SB\_JH\_DoMachine-LearningMachinesLearn\_preprint.pdf), 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.
4. Bringsjord, S., Licato, J.: Psychometric Artificial General Intelligence: The Piaget-MacGuyver Room. In: Wang, P., Goertzel, B. (eds.) *Foundations of Artificial General Intelligence*, pp. 25–47. Atlantis Press, Amsterdam, The Netherlands (2012),



- [http://kryten.mm.rpi.edu/Bringsjord\Licato\\\_PAGI\\\_071512.pdf](http://kryten.mm.rpi.edu/Bringsjord\Licato\_PAGI\_071512.pdf), This url is to a preprint only.
5. Bringsjord, S., Schimanski, B.: What is Artificial Intelligence? Psychometric AI as an Answer. In: Proceedings of the 18<sup>th</sup> International Joint Conference on Artificial Intelligence (IJCAI-03). pp. 887–893. Morgan Kaufmann, San Francisco, CA (2003), <http://kryten.mm.rpi.edu/scb.bs.pai.ijcai03.pdf>
  6. Bringsjord, S.: Declarative/Logic-Based Cognitive Modeling. In: Sun, R. (ed.) The Handbook of Computational Psychology, pp. 127–169. Cambridge University Press, Cambridge, UK (2008), [http://kryten.mm.rpi.edu/sb\\\_lccm\\\_abstract\\\_031607.pdf](http://kryten.mm.rpi.edu/sb\_lccm\_abstract\_031607.pdf), This URL goes to a preprint only.
  7. Bringsjord, S.: The Logicist Manifesto: At Long Last Let Logic-Based AI Become a Field Unto Itself. *Journal of Applied Logic* **6**(4), 502–525 (2008), [http://kryten.mm.rpi.edu/SB\\\_LAI\\\_Manifesto\\\_091808.pdf](http://kryten.mm.rpi.edu/SB\_LAI\_Manifesto\_091808.pdf)
  8. Bringsjord, S., Giancola, M., Govindarajulu, N.S.: Logic-Based Modeling of Cognition. In: Sun, R. (ed.) The Handbook of Computational Psychology. Cambridge University Press, Cambridge, UK ((forthcoming)), <http://kryten.mm.rpi.edu/Logic-basedComputationalModelingOfCognition.pdf>
  9. Bringsjord, S., Govindarajulu, N.S., Licato, J., Giancola, M.: Learning *Ex Nihilo*. In: GCAI 2020. 6th Global Conference on Artificial Intelligence. EPiC Series in Computing, vol. 72, pp. 1–27. International Conferences on Logic and Artificial Intelligence at Zhejiang University (ZJULogAI), EasyChair Ltd, Manchester, UK (2020). <https://doi.org/10.29007/ggcf>, <https://easychair.org/publications/paper/NzWG>
  10. Bringsjord, S., Govindarajulu, N., Giancola, M.: Automated Argument Adjudication to Solve Ethical Problems in Multi-Agent Environments. *Paladyn, Journal of Behavioral Robotics* **12**, 310–335 (2021), <http://kryten.mm.rpi.edu/AutomatedArgumentAdjudicationPaladyn071421.pdf>, The URL here goes to a *rough, uncorrected, truncated* preprint as of 071421.
  11. Bringsjord, S., Sundar Govindarajulu, N.: Rectifying the Mischaracterization of Logic by Mental Model Theorists. *Cognitive Science* **44**(12), e12898 (2020). <https://doi.org/https://doi.org/10.1111/cogs.12898>, <https://onlinelibrary.wiley.com/doi/abs/10.1111/cogs.12898>
  12. Duch, W., Oentaryo, R.J., Pasquier, M.: Cognitive Architectures: Where do we go from here? In: AGI. pp. 122–136 (2008)
  13. Gardner, H.: *Frames of Mind: The Theory of Multiple Intelligences*. Basic Books, New York, NY (1983)
  14. Goertzel, B.: Artificial General Intelligence. *Scholarpedia* **10**(11), 31847 (2015). <https://doi.org/10.4249/scholarpedia.31847>, revision #154015
  15. Goertzel, B.: Artificial General Intelligence: Concept, State of the Art, and Future Prospects. *Journal of Artificial General Intelligence* **0**, 1–48 (01 2014). <https://doi.org/10.2478/jagi-2014-0001>
  16. Goertzel, B.: The General Theory of General Intelligence: A Pragmatic Patternist Perspective. *ArXiv abs/2103.15100*, 1–64 (2021)
  17. Goertzel, B., Bugaj, V.: AGI Preschool: A Framework for Evaluating Early-Stage Human-like AGIs. In: Proceedings of the 2nd Conference on Artificial General Intelligence (2009). pp. 12–17. Atlantis Press (2009), <https://www.atlantispress.com/proceedings/agi09/1826>
  18. Goertzel, B., Pennachin, C., Geisweiller, N.: AGI Preschool. In: *Engineering Intelligence, Part 1: A Path to Advanced AGI via Embodied Learning and Cognitive Synergy*, pp. 337–354. Atlantis Thinking Machines, Atlantis Press, Amsterdam, The Netherlands (2014), This book is Volume 5 in the series.

19. Govindarajulu, N., Bringsjord, S.: On Automating the Doctrine of Double Effect. In: Sierra, C. (ed.) Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17). pp. 4722–4730. International Joint Conferences on Artificial Intelligence (2017). <https://doi.org/10.24963/ijcai.2017/658>, <https://doi.org/10.24963/ijcai.2017/658>
20. Govindarajulu, N., Bringsjord, S., Peveler, M.: On Quantified Modal Theorem Proving for Modeling Ethics. In: Suda, M., Winkler, S. (eds.) Proceedings of the Second International Workshop on Automated Reasoning: Challenges, Applications, Directions, Exemplary Achievements (ARCADE 2019), Electronic Proceedings in Theoretical Computer Science, vol. 311, pp. 43–49. Open Publishing Association, Waterloo, Australia (2019), <http://eptcs.web.cse.unsw.edu.au/paper.cgi?ARCADE2019.7.pdf>, The ShadowProver system can be obtained here: <https://naveensundarg.github.io/prover/>.
21. Gust, H., Krumnack, U., Schwering, A., Kühnberger, K.U.: The Role of Logic in AGI Systems: Towards a Lingua Franca for General Intelligence pp. 126–131 (06 2009). <https://doi.org/10.2991/agi.2009.28>
22. Hammer, P., Lofthouse, T., Wang, P.: The OpenNARS Implementation of the Non-Axiomatic Reasoning System. In: AGI (2016)
23. Hutter, M.: Universal Artificial Intelligence: Sequential Decisions Based on Algorithmic Probability. Springer, New York, NY (2005)
24. LeCun, Y.: Yann LeCun: Deep Learning, Convolutional Neural Networks, and Self-Supervised Learning. Lex Fridman Podcast (August 2019), transcript of podcast available at: [https://www.happyscribe.com/public/lex-fridman-podcast-artificial-intelligence-ai/yann-lecun-deep-learning-convolutional-neural-networks-and-self-supervised-learning#paragraph\\_2405](https://www.happyscribe.com/public/lex-fridman-podcast-artificial-intelligence-ai/yann-lecun-deep-learning-convolutional-neural-networks-and-self-supervised-learning#paragraph_2405). Relevant portion begins at 00:40:05.
25. Levesque, H., Lakemeyer, G.: Chapter 24: Cognitive Robotics. In: Handbook of Knowledge Representation. pp. 869–882. Elsevier, Amsterdam, The Netherlands (2007), <http://www.cs.toronto.edu/~hector/Papers/cogrob.pdf>
26. Lovett, A., Bridewell, W., Bello, P.: Selection, Engagement, & Enhancement: A Framework for Modeling Visual Attention. In: Proceedings of the 43rd Annual Conference of the Cognitive Science Society. pp. 1893–1899. Cognitive Science Society, Vienna, Austria (2021)
27. Shapiro, S., Bona, J.: The GLAIR Cognitive Architecture. International Journal of Machine Consciousness **02**, 144–152 (12 2010). <https://doi.org/10.1142/S1793843010000515>
28. Simpson, S.: Subsystems of Second Order Arithmetic. Cambridge University Press, Cambridge, UK (2010), this is the 2nd edition.
29. Wang, P.: Cognitive Logic versus Mathematical Logic (2004)
30. Wang, P.: Rigid Flexibility: The Logic of Intelligence. Springer, Dordrecht, The Netherlands (2006), This book is Volume 34 in the *Applied Logic Series*, edited by Dov Gabbay and Jon Barwise.
31. Wang, P.: Non-Axiomatic Logic, A Model of Intelligent Reasoning. Sciendo (2013)
32. Wang, P.: Proceedings of the 6th International Conference on Artificial General Intelligence. In: Artificial General Intelligence. pp. 160–169 (07 2013). [https://doi.org/10.1007/978-3-642-39521-5\\_17](https://doi.org/10.1007/978-3-642-39521-5_17)
33. Wang, P.: On Defining Artificial Intelligence. Journal of Artificial General Intelligence **10**, 1–37 (2019)
34. Wang, P., Li, X., Hammer, P.: Self in NARS, an AGI System. Frontiers in Robotics and AI **5** (2018)