Artificial Agents to Help Address the U.S. K–12 Math Gap Between Economically Disadvantaged vs. Advantaged Youth

(extended abstract)

Selmer Bringsjord • John Angel • Naveen Sundar Govindarajulu
Michael Giancola • Elizabeth Bringsjord
Rensselaer AI & Reasoning (RAIR) Laboratory, Rensselaer Polytechnic Institute (RPI), Troy NY 12180 USA
Contact: `Selmer.Bringsjord@gmail.com`

Abstract

Proficiency in math among U.S. pre-college students is overall undeniably low, as shown by reliable and long-established empirical data (e.g., PISA 2015); this is especially true for students in lower-socioeconomic levels. We herein present some of the data in question at the U.S. State level; explain the math tests that generate said data; encapsulate the three parts of the particular paradigm we bring to bear to address the crisis; describe in a bit more detail how the artificial agents in this paradigm operate; make a few remarks regarding related work; anticipate and rebut two inevitable objections; and wrap up with comments regarding next steps.

The Problem, Encapsulated

Proficiency in math among U.S. pre-college students is overall undeniably low, as shown by reliable, longstanding empirical data collected on an ongoing basis; this is especially true, and distressingly so, for students in lower-socioeconomic levels. For example, the standing of U.S. youth in math, as measured against other nations and economies, is given in the PISA series; see e.g. PISA 2015. In the case of data internal to the U.S., i.e. data at the State level, we note and exploit the convenient fact that New York State (NYS), specifically its Department of Education (NYSED), annually tests the mathematical capability of Grade 3–12 students against "common-core" content in mathematics, and provides an informative, web-based interactive system for viewing assessment data. NYS has five (known as the "Big Five") urban school districts, and in the case for instance of Big-Five-member Rochester Central School District (RCSD) in 2019, only 13% of Grades 3–8 students were proficient in math, and in fact 68% of these students scored at Level 1 on their math tests — this being a level that no test-taker, however scant their progress, can fail to achieve. In stark contrast, in districts where residential real-estate is very expensive, and correspondingly the vast majority of K–12 students hail from households of a high socio-economic level, math proficiency is markedly higher. For example, in 2019, overall, 88% of students in Scarsdale's school district were proficient in math. The racial composition of public school districts is shown in NYSED data, and for economy we simply mention here that the lower socio-economic levels seen in the Big Five do unfortunately correlate highly with higher percentages of people of color.

Such inequity is profoundly disturbing. After all, it is extremely difficult to see how youth from lower socio-economic levels will have a decent chance to avoid poverty in an economy increasingly driven by mathematics, computation, and specifically by the development of and collaboration with AI/ML itself. To solve the problem, we propose bringing to bear a distinctive class of artificial agents that can augment the conventional math education of U.S. students, in particular math education tied tightly to common-core content.

Plan of the Extended Abstract

The sequel will unfold as follows. We begin in earnest by presenting some of the depressing data in question; then explain the math tests that generate said data; next, explain the three parts of the particular paradigm we bring to bear to address the crisis; describe in a bit more detail how the artificial agents in this paradigm that we envisage will operate; anticipate and rebut two inevitable objections; make a some quick comments about related work; and wrap up with remarks regarding next steps.

The Math Tests in Question

Content to be Taught: Common-Core Math

New York State's K–12 public-education system conforms to common-core standards, and Grade 4 mathematics is no exception; these standards are set out in natural-language in (New York State Education Department 2012). Space doesn't permit us to recapitulate the standards in question, obviously; we shall need to rest content with conveying to the reader a decent sense of the standards, in order to inform our key claim (below) that these standards are all capturable[1] in computational formal logic. Specifically, common-core standards are captured by a set $\mathcal{CC}_n$ of formulae expressed in the formal language $\mathcal{L}$ of some formal logic $\mathscr{L}$. The

[1]The concept of capture is a technical one in formal logic, the explanation of which is out of scope here. A nice exposition is provided in (Smith 2013).

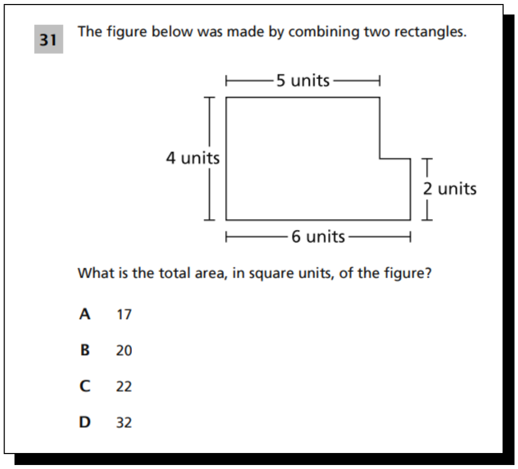

Figure 1: Sample Grade-3 Math Question *Note that there is more than one way to solve this problem. Specifically, there are two different pairs of triangles that one can use/imagine in order to select the answer α. Accordingly, G3Solver (see below), implemented, will find two different minimal proofs automatically that establish 22 = α.*

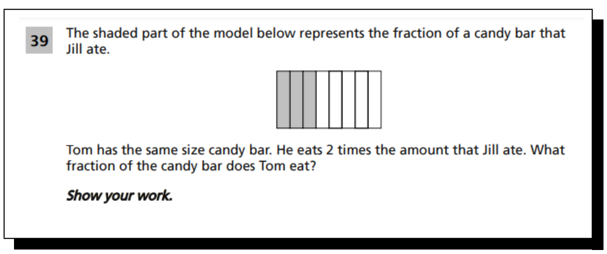

Figure 2: Grade-4 NYS Math-Test Question Calling for Work to be Shown *What the student provides as a justification here is understood to be a proof (or at least an argument) for the selected option, and is assessed by our AI technology.*

subscript n here refers to grade level; so e.g. $\mathcal{CC}_3$ are the common-core standards for Grade 3, captured. Obviously the collection of all common-core standard, which the Master Method has as a static resource, is

$$\bigcup_{i=1}^{12} \mathcal{CC}_i.$$

Sample NYS Math Questions

The specific math tests with which we are concerned are, with minor variation, ubiquitous across public K–12 education in the United States, but we key in on these tests as they are developed, administered, and statistically analyzed in New York State (NYS). As a first example, a Grade-3 question is shown in Figure 1. A harder, more interesting problem is this Grade-4 math question in NYS shown in Figure 2.

Extending the Questions: Justification

While such tests as created and provided by NYSED do not include requests to test-takers that they provide justifications (some questions do request that "work" be shown), in our approach to AI-infused math education in the early grades we believe it's important that students attempt to provide justifications that are precursors to the proofs that will eventually sometimes be requested in high-school (and if not there, then certainly in college-level mathematics).

Our Three-Part Paradigm

In general, there are three main parts to our paradigm. First, we view tests as being at the heart of what AI is, and how it should be pursued, concretely. Second, we specifically pursue logic-based (or as we prefer to say, *logicist* AI (and here, our brand of logicist AI has six important elements). And third, we specifically have invented a class of artificial agents that have (at least least we see it) some very promising attributes when, as is the case herein, the domain of application of pedagogy. We now briefly describe each of these three parts of our paradigm, in sequence.

Psychometric AI

The first part of the paradigm we bring to bear upon the math gap in question is *Psychometric AI* (PAI), first introduced in (Bringsjord and Schimanski 2003). The basic idea underlying PAI is that attempts to build intelligent artifacts should consist in concrete attempts to artifacts that can excel on tests of cognitive ability. The fit between PAI and an attempt to engineer artificial agents able to issue tests and process performance on them is doubtless clear to the reader.

Logicist AI: Our Brand Thereof, & Its Six Key Elements

What is AI? And Specifically Logic-based/Logicist AI? We affirm the uncontroversial, orthodox conception that the discipline of AI consists in the attempt to understand, design, and implement *artificial agents*, i.e. entities that receive *percepts* from the environment in which they are in, and produce as outputs *actions* performed in this environment. This conception is firmly at the heart of all comprehensive and authoritative accounts of AI we are aware of; see e.g. (Russell and Norvig 2020, Bringsjord and Govindarajulu 2018, Luger 2008). But what mediates between percepts and actions in a given artificial agent $\mathfrak{a}$? From a high-level perspective, there is no loss of generality here in stipulating that the percept-to-action mapping corresponds to some Turing-computable function, and that what mediates the inputs and outputs is then naturally said to be a Turing machine m (or an equivalent machine, e.g. a Register machine) that computes function. Typically, the construction of a given artificial agent will then consist in no small part in the engineering of a computer program p corresponding to such a Turing machine m. The picture so far can be encapsulated if we write:

$$\mathfrak{a}_{m/p} : \mathcal{P} \mapsto \mathcal{A},$$

where of course $\mathcal{P}$ is a set of percepts and $\mathcal{A}$ a set of actions.

Now how do we specifically define *logicist* (or *logic-based*) AI, which is the brand of AI we pursue (at least in our approach to AI-augmented education in the formal sciences)? The answer is straightforward and efficient: We simply specify that instead of any Turing-level machine m in the foregoing, the machine in question is an automated reasoner r; this yields:

$$\mathfrak{a}_r : \mathcal{P} \mapsto \mathcal{A},$$

which directly reflects our work, since it is automated reasoning that is the backbone of our artificial pedagogical agents. Below, when we describe the algorithm G3Solver, and demonstrate it in action, it will be seen that we are dealing with an agent of the type $\mathfrak{a}_r$.

We turn now to a rapid enumeration of the six key elements in our brand of logicist AI for seeking to meet the education challenge at hand.

Element #1: Cognitive Calculi Essentially, a (deductive) cognitive calculus is a quantified multi-operator modal logic such that its: proof/argument theory is specified in "natural deduction" form (traceable back to (Gentzen 1935, Fitch 1952)), operators cover all or most of human-level cognition (e.g., *believing*, *knowing*, *perceiving*, *communicating*, and also *obligations*, etc.), and semantics is exclusively proof-theoretic in nature.[2] Proof-theoretic semantics eschews model-theoretic and possible-worlds semantics in favor of the basic idea that meaning is provided to formulae and their constituents solely by virtue of the nature of proofs in which these things appear.[3]

In the present work, we specifically utilize elements of the Deontic Cognitive Event Calculus ($\mathcal{DCEC}$) to model the perceptions and beliefs of students, denoted **B** and **P** respectively. A dialect of $\mathcal{DCEC}$ is specified and used in (Govindarajulu and Bringsjord 2017). can be thought of roughly as a quantified multi-operator modal logic with all of the introduction and elimination rules for first-order logic, plus a host of inference schemata to cover its many modal operators.

Element #2: Automated Reasoning for Cognitive Calculi: ShadowProver An automated reasoner for $\mathcal{DCEC}$ — ShadowProver (Govindarajulu, Bringsjord, and Peveler 2019) — has been created and is under active development. Soundness proofs for cognitive calculi have been obtained but are out of scope.

Element #3: Visual Logics and Associated Automated Reasoning This third element of the kind of logicist AI we are using in order to address the K–12 "math gap" between the economically advantaged and disadvantaged is an ability for artificial agents to represent and reason over not just symbolic content, but *pictorial* content as well. More specifically, we use the Vivid framework introduced in (Arkoudas and Bringsjord 2009), and descendants thereof. This formal science and corresponding automated-reasoning technology is beyond the scope of the present short abstract, and must be left for another time, but all alert readers will have noticed from even a cursory look at the questions shown in Figures 1 and 2 that in the tests in question, diagrams play an important role.

Element #4: Background Axiom Systems for Mathematics Itself One of the remarkable turns in professional mathematics and logic, seen at the end of the 20th century, and accelerating greatly in the 21st century, has been to what is known as *{*reverse mathematics}. In a word, reverse mathematics is devoted to proving what suffices to obtain results in the familiar branches of mathematics; here, the "what" specifically refers to what axioms expressed in what formal logic suffice to formally prove the results in question. The current authoritative introduction to reverse mathematics, which shows that number theory expressed in second-order logic is remarkably powerful, is (Simpson 2010). In our approach, we use automated reasoners able to process such fundamental information, in order to generate tests. In particular, for Grade 3 and Grade 4 mathematics and the NYS tests in question, axiom systems of limited scope for arithmetic suffice. For instance, mathematical induction is not needed. So-called Robinson Arithmetic, sometimes just denoted by **Q**, is sufficient; **Q** is nicely introduced and explained in (Boolos, Burgess, and Jeffrey 2003).

Element #5: Capturing What is to Be Learned in an Axiom System In the case at hand, what is to be learned, as we have said above, are the things listed in the common-core learning standards for public K–12 mathematics education, as described by New York State. In our approach, these standards, which are presented informally in English by New York State, are captured in formal logic. We leave aside here for economy how this capture is accomplished, and the formulae that result from doing so.

Element #6: NLP (NLG & NLU) (and formal semantics) of Our Logicist Sort We do not have space to explain how both natural language generation (NLG) and natural language understanding (NLU) is handled. In general, we use direct passing from formal logic to English constrained by a formal grammar for NLU, and NLG is direct semantic parsing from English to formulae in our logics. Our approach to the semantics of natural language is entirely proof-theoretic (not model-based/Montagovian). A nice introduction to proof-theoretic semantics, including its connection to and use for natural language, see (Francez 2015).

The "TIPPAE" Class of Artificial Agents

TIPPAE, as a paradigm, is a class of artificial intelligences that are enabled by cognitive logics to offer a continuous assistive experience though multiple heterogeneous hardware environments for the user whose education they are enhancing. They are designed to accomplish this by fulfilling the properties of their namesake, TIPPAE (Teleportative Intelligent Personalized Persistent Agents for Education). These properties, formalized and verified through the use of logic based artificial intelligence, aim to create agents that are not

[2]For specification of the formal language & proof theory of a the particular deductive cognitive calculus $\mathcal{DCEC}$ we direct readers to e.g. (Govindarajulu and Bringsjord 2017).

[3]For more on proof-theoretic semantics see (Dummett 1981, 1991, Gentzen 1935, Prawitz 1972).

only helpful to their user's education, but open avenues for the potential of human-computer friendship. (Angel, Govindarajulu, and Bringsjord 2019) In brief, the properties of this paradigm can be described as thus: "Teleportative" represents that need of the agent to be able to teleport, also known as migrate, through many potential hardware environments while giving a believable impression of being one, uninterrupted AI identity. Consider, as a hypothetical example, an agent that can start a conversation about a multiplication as a graphical interface on a desktop computer, then continue that conversation via text messages while the user is in transit, then finish the conversation via a toy robot at the user's home. "Intelligent" represents the need of the agent to make reasonable and logical decisions about how it will assist its user. The agent needs a logical understanding of the pedagogical domain of knowledge that it is working to tutor its user toward better mastery. "Personalized" indicates that the agent can provide individualized assistance that is tailored to that particular user. For example, if the agent detects that the user has a specific misunderstanding in elementary mathematics, carrying the one during addition for example, then the agent can then dispense assistance that relates to that specific misunderstanding. "Persistent" indicates the agent's ability to continuously assist a user over longer periods of time and across incremental pedagogical subject matters. After all, an agent that can continuously offer personalized assistance throughout multiple school years will naturally have a more thorough understanding of a student's mastery over the subject matter than an agent that resets its knowledge every year. "Agent for Education" denotes the continuing commitment of those engineering said assistants to keep in mind the myriad real world and practical concerns that come with working in education such as privacy, security, opportunity distribution, and many, many more.

The Operation of our TIPPAE Agent Sketched

The Overarching Architecture

The overarching architecture of the TIPPAE agent we have invented, and are gradually implementing, part by part, is shown in Figure 3; this architecture we refer to as 'The Master Method.' The caption for this figure, given space constraints, must suffice here. Note that the G3Solver algorithm, which we present and demonstrate below, is part of The Master Method, and indeed a key part thereof, because generation of tests $\mathcal{T}$ (as a series of questions Q_i) are verified as correct, and measured for difficulty, by ensuring that (as) G3Solver can find a proof π of the answer/key α, and that (b) a *minimal* proof is of the appropriate length (where minimal proof length gives us a way to gauge difficulty of a given question Q_i; details of this way are out of scope here).

Demonstration

Our demonstration is of a key sub-algorithm for the "Master Method" shown in Figure 3. This sub-algorithm, G3Solver, is shown in Figure 4, and the pseudo-code provided therein, should be for the most part self-explanatory, but we now give a quick summary in prose:

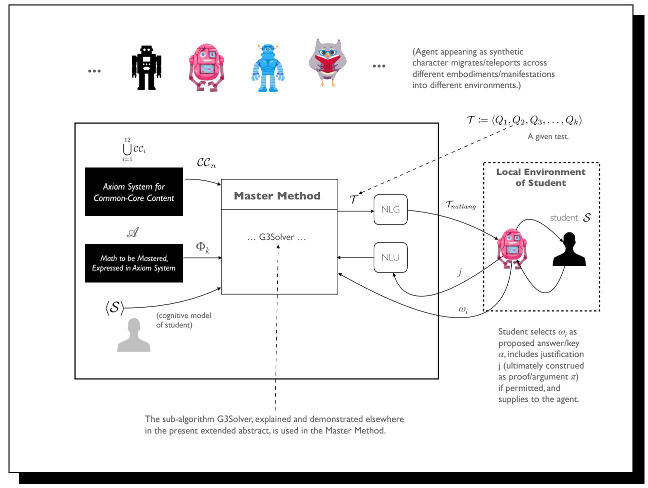

Figure 3: Overarching Architecture of TK12MTA (TIPPAE K–12 Math Testing Agent) *The Master Method, not specified here, begins with the generation of a given test $\mathcal{T}$ using/consistent with common-core standards at the relevant grade level, and also using the particular grade-level axiom system Φ_k selected from reverse-mathematics content (i.e. from $\mathscr{A}$).*

Input: Math Axiomatic System Φ_k, Story Formulae S, Query Formula q, Set of Options Formula $\Omega = \{\omega_1, \omega_2, ..., \omega_j\}$, Student Justification π if given
Output: Key formula ω_i where $\Phi_k \cup S \cup q \vdash \omega_i$ $(1 < i \leq j)$, Student Justification π if given and valid

```
answer = fail;
    for ω ∈ Ω do
        if Φ_k ∪ S ∪ q ⊢ ω then
            answer ⟵ ω;
        end
    end
    if π ≠ null then
        if π is valid then
            justification ⟵ π
        else
            justfication ⟵ fail
        end
    end
    return (answer, justification)
```

Figure 4: G3Solver Algorithm *It should be noted that the algorithm is able to handle selections of an option ω_i in multiple-choice questions,* and *justifications supplied by students to justify their selection.*

```
; Running SNARK from C:/Users/angel/IdeaProjects/prover/snark-20120808r02/snark-system.lisp in Armed Bear Common Lisp 1.0.1-svn-13751 ((
    Trying Agent Closure: (Believes! human t5 (IsAreaOf 22 compRect1))
    Trying Agent Closure: (IsAreaOf 22 compRect1)
Backward Reasoning: Reductio on(not (IsAreaOf 22 compRect1))
        Trying Agent Closure: $Gen_1$
Forward Reasoning: Perceives(P) ==> P[(Knows! human t3 (IsRectangle rect2)), (Knows! human t4 (IsPartRect1Of rect compRect1)), (Knows! h
Backward Reasoning: Reductio on(not (Believes! human t5 (IsAreaOf 22 compRect1)))
        Trying Agent Closure: $Gen_2$
    Forward Reasoning: Knows(P) ==> P[(IsRectWidthOf 2 rect2), (IsPartRect1Of rect compRect1), (IsRectLengthOf 1 rect2), (IsRectangle re
    Forward Reasoning: Perceives(P) ==> P[(Knows! human t3 (IsRectangle rect2)), (Knows! human t4 (IsPartRect1Of rect compRect1)), (Know
    Forward Reasoning: Perceives(P) ==> P[(Knows! human t3 (IsRectangle rect2)), (Knows! human t4 (IsPartRect1Of rect compRect1)), (Know


(Refutation
(Row 178 (SNARK::ISRECTWIDTHOF 2 SNARK::RECT2)                             ASSERTION)
(Row 179 (OR (NOT (SNARK::ISRECTANGLE ?X)) (NOT (SNARK::ISRECTLENGTHOF ?Y ?X)) (NOT (SNARK::ISRECTWIDTHOF ?Z ?X)) (SNARK::ISAREAOF ($$PF
(Row 180 (SNARK::ISPARTRECT1OF SNARK::RECT SNARK::COMPRECT1)               ASSERTION)
(Row 181 (SNARK::ISRECTLENGTHOF 1 SNARK::RECT2)                            ASSERTION)
(Row 182 (SNARK::ISRECTANGLE SNARK::RECT2)                                 ASSERTION)
(Row 183 (SNARK::ISCOMPOSITERECTANGLE SNARK::COMPRECT1)                    ASSERTION)
(Row 184 (OR (NOT (SNARK::ISCOMPOSITERECTANGLE ?X)) (NOT (SNARK::ISPARTRECT1OF ?Y ?X)) (NOT (SNARK::ISPARTRECT2OF ?Z ?X)) (NOT (SNARK::I
(Row 185 (SNARK::ISRECTWIDTHOF 4 SNARK::RECT)                              ASSERTION)
(Row 186 (SNARK::ISRECTLENGTHOF 5 SNARK::RECT)                             ASSERTION)
(Row 187 (SNARK::ISRECTANGLE SNARK::RECT)                                  ASSERTION)
(Row 188 (SNARK::ISPARTRECT2OF SNARK::RECT2 SNARK::COMPRECT1)              ASSERTION)
(Row 189 (NOT (SNARK::ISAREAOF 22 SNARK::COMPRECT1))                       NEGATED_CONJECTURE)
(Row 190 (SNARK::ISAREAOF 20 SNARK::RECT)                                  (REWRITE (HYPERRESOLVE 179 187 186 185) :CODE-FOR-$$PRODUCT))
(Row 191 (SNARK::ISAREAOF 2 SNARK::RECT2)                                  (REWRITE (HYPERRESOLVE 179 182 181 178) :CODE-FOR-$$PRODUCT))
(Row 192 (SNARK::ISAREAOF 22 SNARK::COMPRECT1)                             (REWRITE (HYPERRESOLVE 184 191 183 180 188 190) :CODE-FOR-$$SUM))
(Row 193 FALSE                                                             (REWRITE 189 192))
    Trying Agent Closure: (IsAreaOf 22 compRect1)
CognitiveCompositeRectangleLogic5By4And1By2
Proved!
```

Figure 5: Proof of Answer/Key Returned by G3Solver Algorithm *Readers with superlative eyesight will notice that the proof makes explicit use of aspects of a cognitive model of a given student, by using the available cognitive/modal operators in a cognitive calculus..*

We report that a run of an implementation of G3Solver, in which calls are made to the aforementioned ShadowProver system, returns a proof for the answer/key α in 1.1 seconds. The proof is shown in Figure 5.

Objections

"Your approach falls prey to teaching to the test!"

This is an objection that educators are quite familiar with. In our case, we anticipate hearing it expressed quite loudly, since after all we are undeniably seeing to have AI teach by issuing and assessing performance on tests.

However, the objection does not apply to our work, at least in our opinion, for a very straightforward reason. The reason is that in the case of math (at least at this early level) and the questions composing the tests in question, the content on tests and what is to be learned is one and the same! This is especially true because many of the test questions present scenarios that match problems seen in "real life." Even in the case of the simple math questions shown above, this is true. In trying to obtain the square footage of rooms in a domicile, but where information is not complete, it is quite possible that the figure and problem shown in Figure 1 could be a real-life one.

"A major driver of the disparity you rightly bemoan is technology!"

The objection, fleshed out to a degree: "The economically disadvantaged suffer from a lack of access to computational technology (e.g. decent computers, adequate internet access/speed, etc.). Your solution, since it requires students to be helped by your AI to have access to sophisticated computational technology, is a non-starter without these problems solved first!"

In quick reply: We concede that our AI-based remediation for disparity in mathematical prowess among American youth requires parallel thrusts of effort to improve access to the technology necessary to enable the artificial agents we seek. However, the fact of the matter is that even before the advent of calculators, paper, pens and pencils, and slide rules were things beyond the reach of many disadvantaged students — and yet certainly no one at the time could rationally maintain that pedagogical innovation presupposing the availability of such things should not be pursued. Things should be no different for us. Eventually the penetration of computational technology will advance, and getting busy now on the AI enabled by that technology is simply prudent and planful. At the end of the day, AI-based education will always require computational, and computation will always require hardware and software of some sort. Fortunately there are some countervailing factors in our favor. For instance, given the deployment of artificial agents of the sort we are aiming at, disadvantaged students would be able to sharpen their math skills even if other shortcomings aren't fully addressed (e.g. less competent teachers, unavailability of human tutors, etc.).

Related Work

At present, we are unaware of any related work by others aimed at building artificial pedagogical agents that are based on an axiomatization of mathematics in order to enable such agents to teach mathematics. Indeed, surprisingly, the use of results produced by reverse mathematics seems to be entirely absent in the area of AI-based/AI-infused math education.[4] Of course, the previous sentence refers only to Element #4 in our logicist-AI approach as described above. This element is certainly crucial and prominent (as the overall architecture of our system shows; see again Figure 3. Nonetheless, this element, as the reader has seen, is but one aspect of a number of others that distinguish our attempt to address the math-proficiency disparity between the wealthy and the poor in the U.S. This naturally prompts the question: What about work by others that relates to other aspects of our particular R&D methodology? Quite a bit could be said about related work in the part of AI that is logicist/logic-based, and as most readers will know, intelligent tutoring systems in AI are related to our work. However, discussion of specifics must wait for a full paper that expands the present abstract.

Conclusion and Next Steps

We have introduced the reader to a grave problem in U.S. K–12 math education, and outlined our proposed AI-based solution to it. There is of course much work that remains to be carried out. Concretely, for Grade-3 math questions of the New-York-State variety, we look forward to conducting between-group experiments with subjects in order to ascertain whether treatment (i.e., use of an embodiment of our overall architecture, with which subjects interact in a supplement to their standard in-school education) results in improvement in performance on the annual math tests administered by NYSE.[5] At the theoretical level, next steps in-

[4]This may in part be because reverse mathematics is by any metric formally intricate, and completely different in nature from statistical/neural AI today often called 'machine learning' or just 'ML' for short.

[5]Again: New York State Education Department.

clude pushing forward with prototypes based on more sophisticated instantiations of The Master Method, for example instantiations at the level of elementary algebra, in which even the first simple equations given to students involve both universal and existential quantification in formal logic, and hence call for an understanding of such quantification in order to be proficient.[6]

References

Angel, J.; Govindarajulu, N.; and Bringsjord, S. 2019. Toward Formalizing Teleportation of Pedagogical Artificial Agents. *The Emerging Learning Design Journal* 6(1): 2. ISSN 2474-8218.

Arkoudas, K.; and 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.

Boolos, G. S.; Burgess, J. P.; and Jeffrey, R. C. 2003. *Computability and Logic (Fourth Edition)*. Cambridge, UK: Cambridge University Press.

Bringsjord, S.; and Govindarajulu, N. S. 2018. Artificial Intelligence. In Zalta, E., ed., *The Stanford Encyclopedia of Philosophy*. URL https://plato.stanford.edu/entries/artificial-intelligence.

Bringsjord, S.; and Schimanski, B. 2003. What is Artificial Intelligence? Psychometric AI as an Answer. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI–03)*, 887–893. San Francisco, CA: Morgan Kaufmann. URL http://kryten.mm.rpi.edu/scb.bs.pai.ijcai03.pdf.

Dummett, M. 1981. *Frege. Philosophy of Language (2nd ed)*. London, UK: Duckworth.

Dummett, M. 1991. *The Logical Basis of Metaphysics*. London, UK: Duckworth.

Fitch, F. 1952. *Symbolic Logic: An Introduction*. New York, NY: Ronald Press.

Francez, N. 2015. *Proof-theoretic Semantics*. London, UK: College Publications.

Gentzen, G. 1935. Untersuchungen über das logische Schlieben I. *Mathematische Zeitschrift* 39: 176–210.

Govindarajulu, N.; and Bringsjord, S. 2017. On Automating the Doctrine of Double Effect. In Sierra, C., ed., *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, 4722–4730. International Joint Conferences on Artificial Intelligence. ISBN 978-0-9992411-0-3. doi:10.24963/ijcai.2017/658. URL https://doi.org/10.24963/ijcai.2017/658.

Govindarajulu, N.; Bringsjord, S.; and Peveler, M. 2019. On Quantified Modal Theorem Proving for Modeling Ethics. In Suda, M.; and Winkler, S., eds., *Proceedings of the Second International Workshop on Automated Reasoning: Challenges, Applications, Directions, Exemplary Achievements (ARCADE 2019)*, volume 311 of *Electronic Proceedings in Theoretical Computer Science*, 43–49. Waterloo, Australia: Open Publishing Association. URL http://eptcs.web.cse.unsw.edu.au/paper.cgi?ARCADE2019.7.pdf. The ShadowProver system can be obtained here: https://naveensundarg.github.io/prover/.

Luger, G. 2008. *Artificial Intelligence: Structures and Strategies for Complex Problem Solving (6th Edition)*. London, UK: Pearson. ISBN 0321545893.

New York State Education Department. 2012. New York State P–12 Common Core Learning Standards for Mathematics. The standards in question can be obtained on the Web, at https://www.engageny.org/resource/new-york-state-p-12-common-core-learning-standards-for-mathematics.

Prawitz, D. 1972. The Philosophical Position of Proof Theory. In Olson, R. E.; and Paul, A. M., eds., *Contemporary Philosophy in Scandinavia*, 123–134. Baltimore, MD: Johns Hopkins Press.

Russell, S.; and Norvig, P. 2020. *Artificial Intelligence: A Modern Approach*. New York, NY: Pearson. Fourth edition.

Simpson, S. 2010. *Subsystems of Second Order Arithmetic*. Cambridge, UK: Cambridge University Press. This is the 2nd edition.

Smith, P. 2013. *An Introduction to Gödel's Theorems*. Cambridge, UK: Cambridge University Press. This is the second edition of the book.

[6]E.g., solving membership in simple Diophantine sets requires of (minimally, NYS Grade-6) students that they can e.g. answer correctly whether $\forall x \exists y (2x - y = 0)$, where $x, y \in \mathbb{Z}^+$.