Learning Ex Nihilo

Selmer Bringsjord
Naveen Sundar Govindarajulu
John Licato
Michael Giancola

RAIR
Rensselaer AI and Reasoning Lab

6th Global Conference on Artificial Intelligence (GCAI 2020)
October 30th, 2020

Sponsored by

Office of Naval Research Science & Technology ONR USN

AFOSR AIR FORCE OFFICE OF SCIENTIFIC RESEARCH

Rensselaer

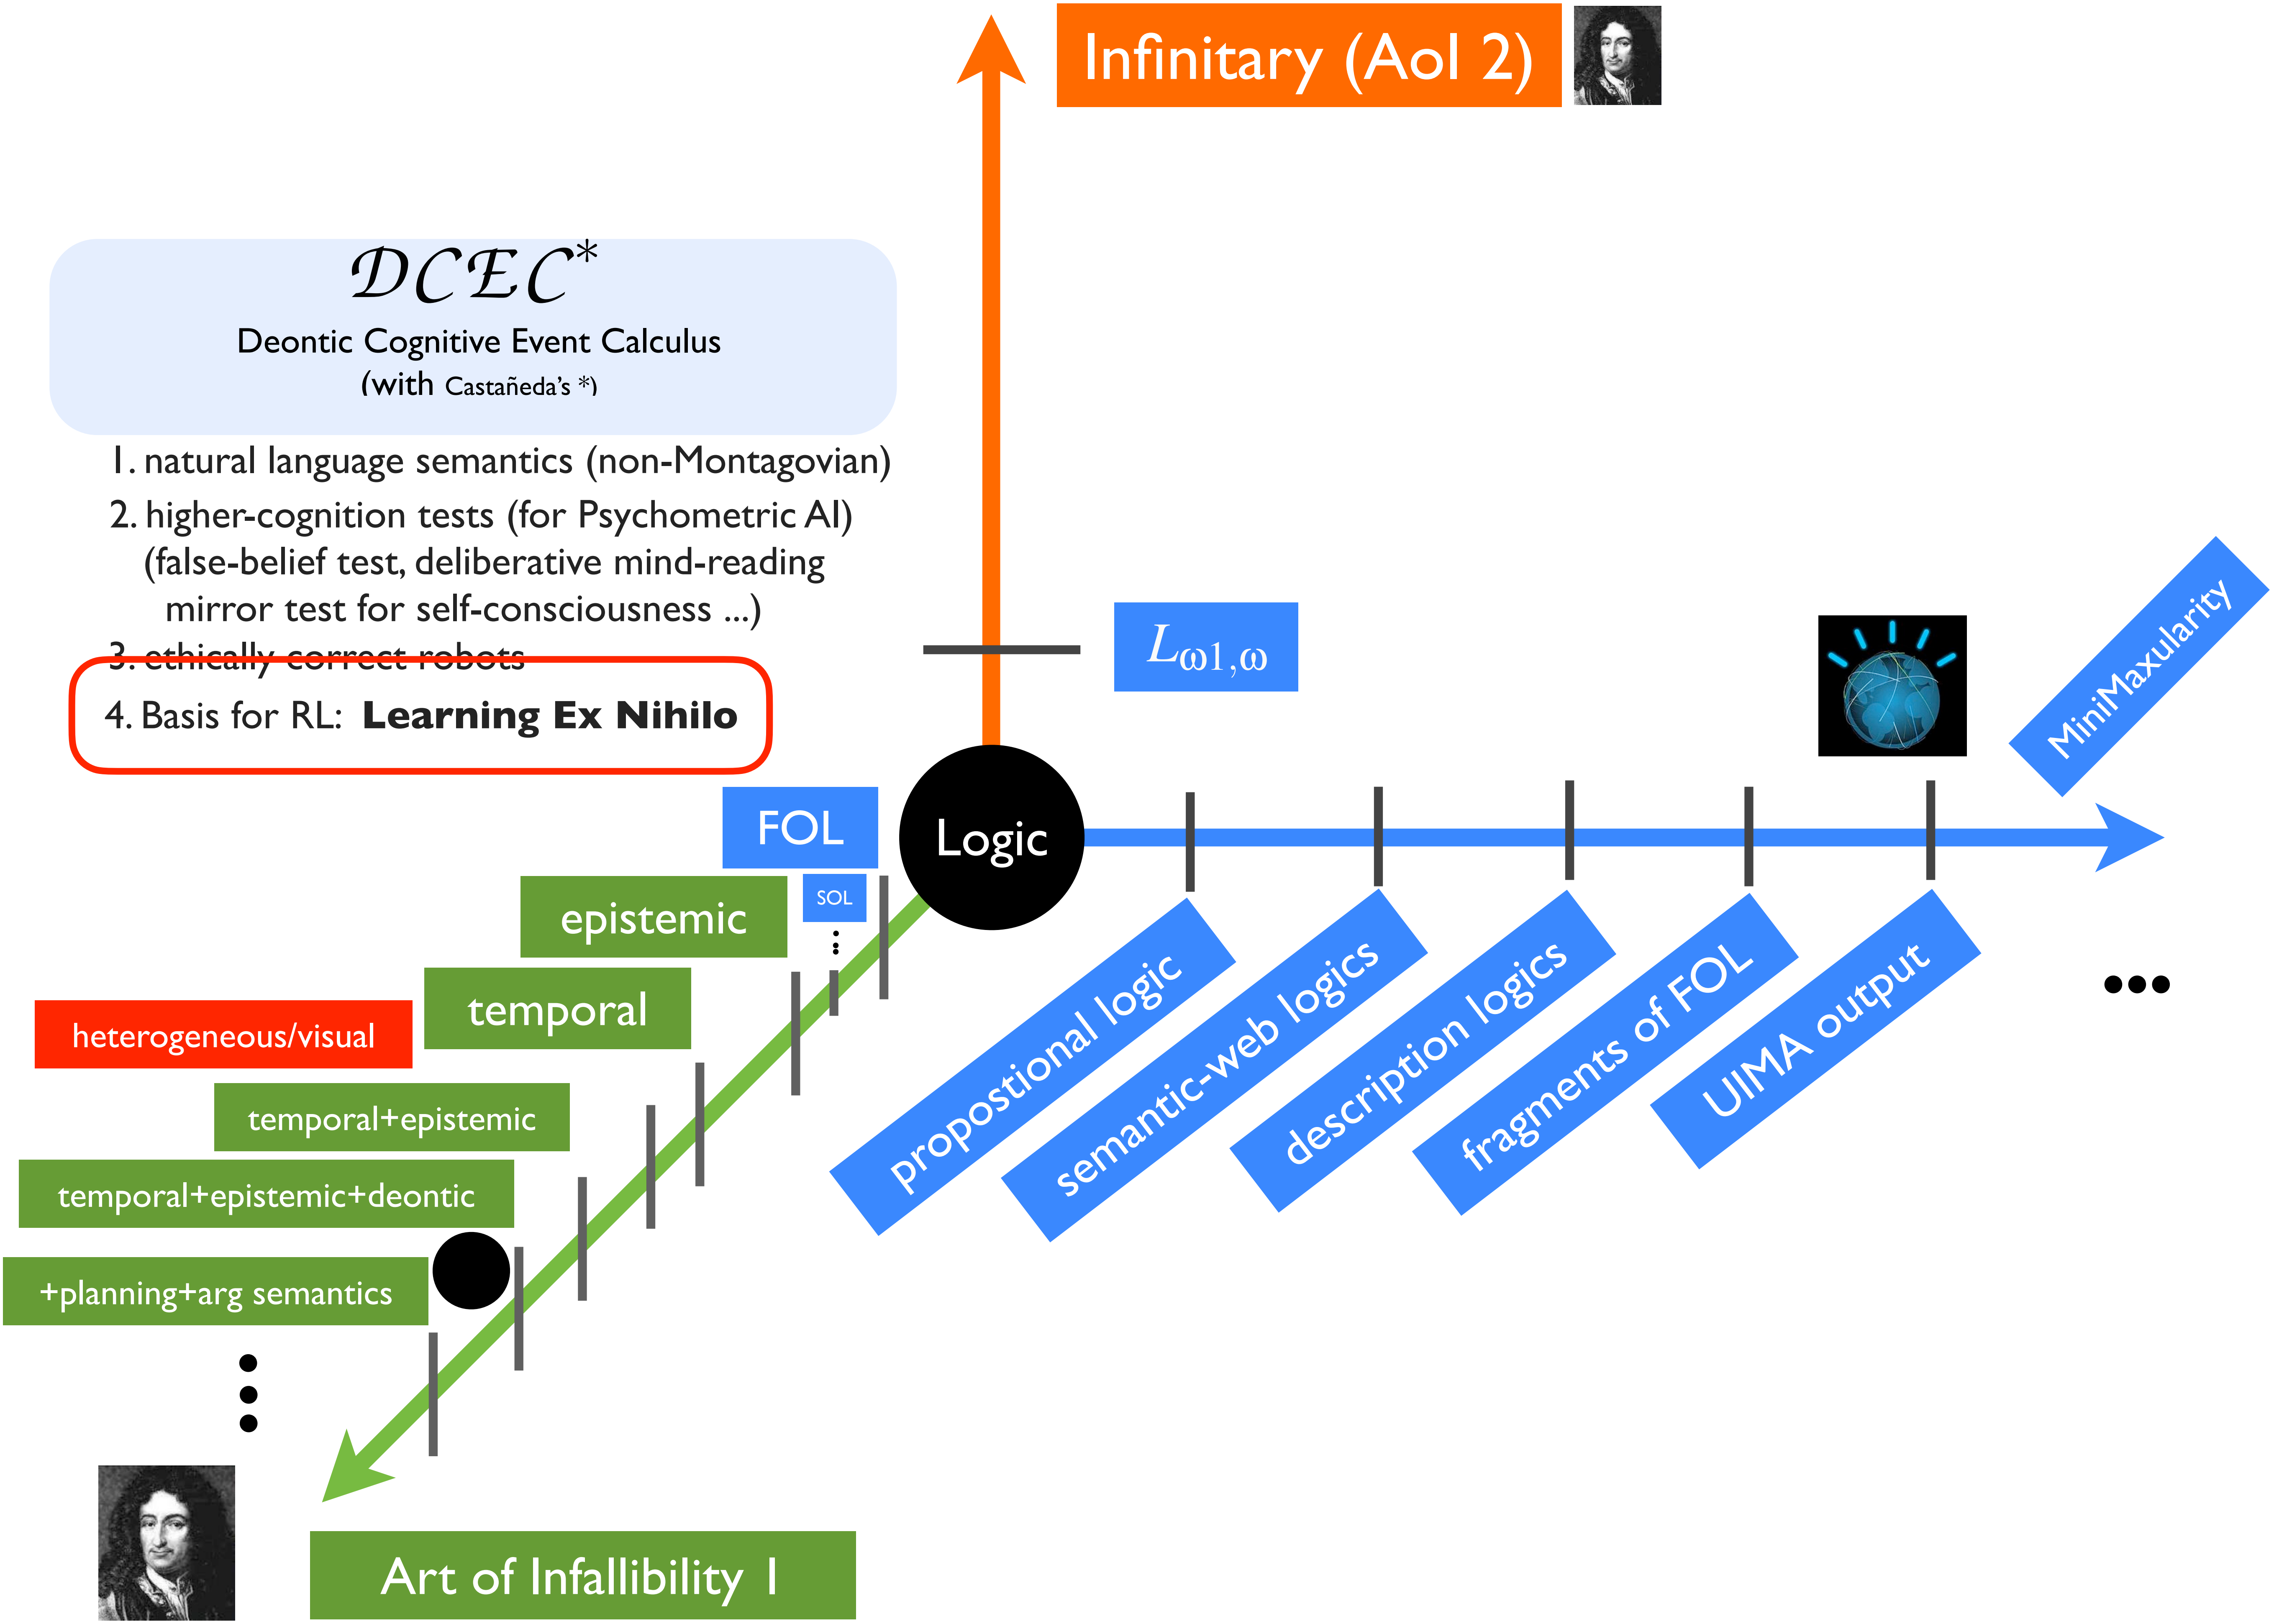

Infinitary (AoI 2)
$\mathcal{DCEC}^*$
Deontic Cognitive Event Calculus
(with Castañeda's *)
1. natural language semantics (non-Montagovian)
2. higher-cognition tests (for Psychometric AI)
(false-belief test, deliberative mind-reading
mirror test for self-consciousness ...)
~~3. ethically correct robots~~
4. Basis for RL: **Learning Ex Nihilo**
$L_{\omega 1,\omega}$
MiniMaxularity
FOL
Logic
SOL
epistemic
temporal
heterogeneous/visual
temporal+epistemic
temporal+epistemic+deontic
+planning+arg semantics
propostional logic
semantic-web logics
description logics
fragments of FOL
UIMA output
Art of Infallibility I

Advanced Logicist (Real) Machine Learning

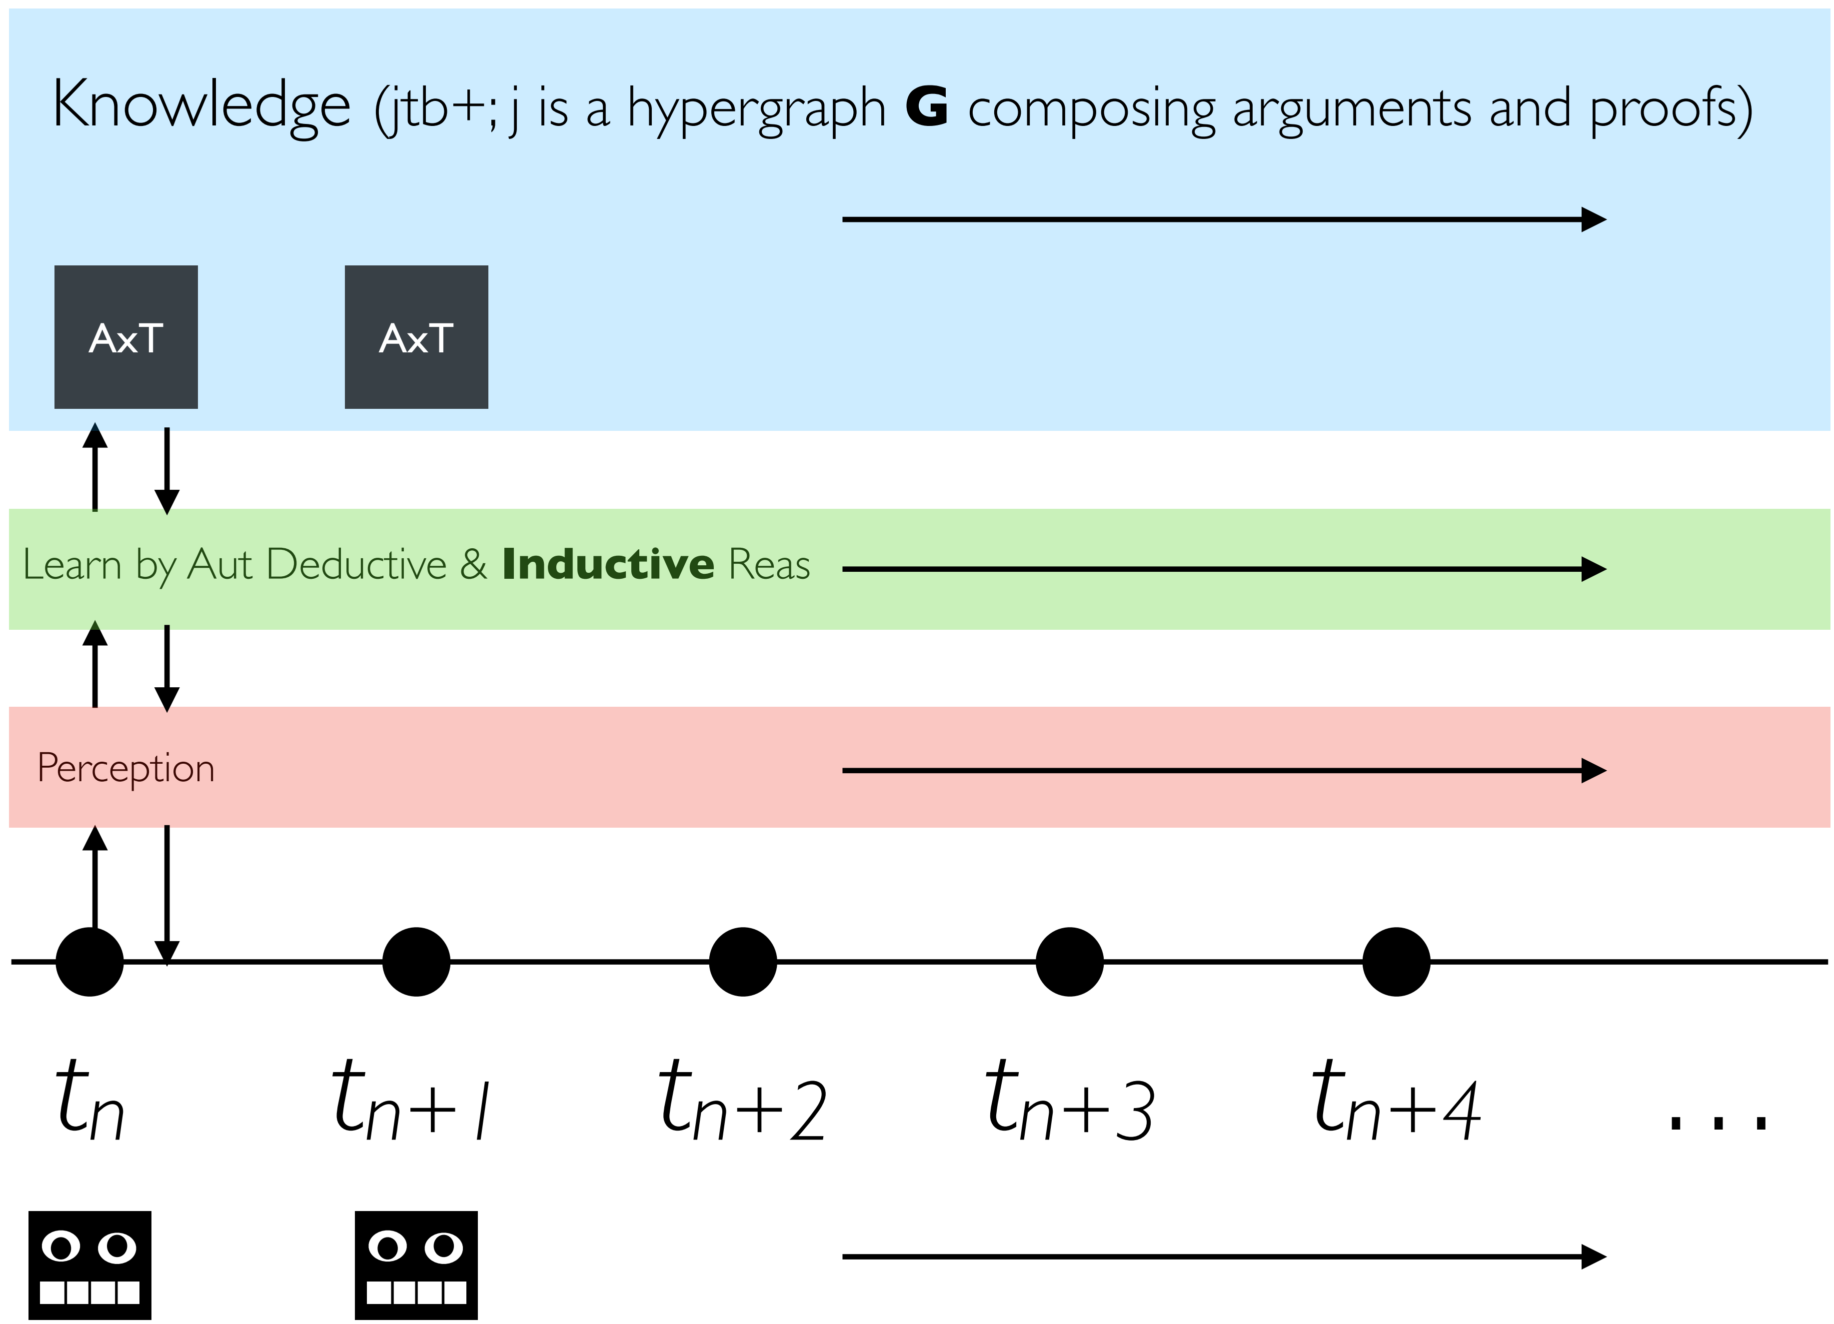

Given This, Do Machine-Learning Machines Learn? No.

Do Machine-Learning Machines Learn?

Selmer Bringsjord and Naveen Sundar Govindarajulu and Shreya Banerjee and John Hummel

Abstract We answer the present paper's title in the negative. We begin by introducing and characterizing "real learning" ($\mathcal{RL}$) in the formal sciences, a phenomenon that has been firmly in place in homes and schools since at least Euclid. The defense of our negative answer pivots on an integration of *reductio* and proof by cases, and constitutes a general method for showing that any contemporary form of machine learning (ML) isn't real learning. Along the way, we canvass the many different conceptions of "learning" in not only AI, but psychology and its allied disciplines; none of these conceptions (with one exception arising from the view of cognitive development espoused by Piaget), aligns with real learning. We explain in this context by four steps how to broadly characterize and arrive at a focus on $\mathcal{RL}$.

Selmer Bringsjord
Rensselaer Polytechnic Institute, 110 8th Street Troy, NY USA 12180, e-mail: selmerbringsjord@gmail.com

Naveen Sundar Govindarajulu
Rensselaer Polytechnic Institute, 110 8th Street Troy, NY USA 12180, e-mail: Naveen.Sundar.G@gmail.com

Shreya Banerjee
Rensselaer Polytechnic Institute, 110 8th Street Troy, NY USA 12180, e-mail: shreyabbanerjee@gmail.com

John Hummel
901 West Illinois Street, Urbana, IL 61801, e-mail: jehummel@illinois.edu

8 Appendix: The Formal Method

The following deduction uses fonts in an obvious and standard way to sort between functions ($\mathfrak{f}$), agents ($\mathfrak{a}$), and computing machines ($\mathfrak{m}$) in the Arithmetical Hierarchy. Ordinary italicized Roman is used for particulars under these sorts (e.g. f is a particular function). In addition, '$\mathcal{C}$' denotes any collection of conditions constituting jointly necessary-and-sufficient conditions for a form of current ML, which can come from relevant textbooks (e.g. Luger, 2008; Russell and Norvig, 2009) or papers; we leave this quite up to the reader, as no effect upon the validity of the deductive inference chain will be produced by the preferred instantiation of '$\mathcal{C}$.' It will perhaps be helpful to the reader to point out that the deduction eventuates in the proposition that no machine in the ML fold that in this style learns a relevant function $\mathfrak{f}$ thereby also real-learns $\mathfrak{f}$. We encode this target as follows:

$$(\star)\ \ \neg\exists\mathfrak{m}\,\exists\mathfrak{f}\,[\phi := \mathit{MLlearns}(\mathfrak{m},\mathfrak{f}) \land \psi := \mathit{RLlearns}(\mathfrak{m},\mathfrak{f}) \land \mathcal{C}_\phi(\mathfrak{m},\mathfrak{f}) \vdash^* (\mathrm{ci}')\text{–}(\mathrm{ciii})_\psi(\mathfrak{m},\mathfrak{f})]$$

Note that $(\star)$ employs meta-logical machinery to refer to particular instantiations of $\mathcal{C}$ for a particular, arbitrary case of ML (ϕ is the atomic sub-formula that can be instantiated to make the particular case), and particular instantiations of the triad (ci$'$)–(ciii) for a particular, arbitrary case of $\mathcal{RL}$ (ψ is the atomic sub-formula that can be instantiated to make the particular case). Meta-logical machinery also allows us to use a provability predicate to formalize the notion that real learning is produced by the relevant instance of ML. If we "pop" ϕ/ψ to yield ϕ'/ψ' we are dealing with the particular instantiation of the atomic sub-formula.

The deduction, as noted in earlier when the informal argument was given, is indirect proof by cases; accordingly, we first assume $\neg(\star)$, and then proceed as follows under this supposition.

	(1)	$\forall\,\mathfrak{f},\mathfrak{a}\,[\mathfrak{f}:\mathbb{N}\mapsto\mathbb{N}\to(\mathit{RLlearns}(\mathfrak{a},\mathfrak{f})\to(\mathrm{i})\text{–}(\mathrm{iii}))]$	Def of Real Learning
	(2)	$\mathit{MLlearns}(m,f)\land\mathit{RLlearns}(m,f)\land f:\mathbb{N}\mapsto\mathbb{N}$	supp (for $\exists$ elim on $(\star)$)
	(3)	$\forall\,\mathfrak{m},\mathfrak{f}\,[\mathfrak{f}:\mathbb{N}\mapsto\mathbb{N}\to(\mathit{MLlearns}(\mathfrak{m},\mathfrak{f})\leftrightarrow\mathcal{C}(\mathfrak{m},\mathfrak{f}))]$	Def of ML
	(4)	$\forall\,\mathfrak{f}\,[\mathfrak{f}:\mathbb{N}\mapsto\mathbb{N}\to(\mathit{TurComp}(\mathfrak{f})\lor\mathit{TurUncomp}(\mathfrak{f}))]$	theorem
	(5)	$\mathit{TurUncomp}(f)$	supp; Case 1
	(6)	$\neg\exists\,\mathfrak{m}\,\exists\,\mathfrak{f}\,[(\mathfrak{f}:\mathbb{N}\mapsto\mathbb{N}\land\mathit{TurUncomp}(\mathfrak{f})\land\mathcal{C}(\mathfrak{m},\mathfrak{f})]$	theorem
$\therefore$	(7)	$\neg\exists\,\mathfrak{m}\,\mathit{MLlearns}(\mathfrak{m},f)$	(6), (3)
$\therefore$	(8)	$\bot$	(7), (2)
	(9)	$\mathit{TurComp}(f)$	supp; Case 2
$\therefore$	(10)	$\mathcal{C}_{\phi'}(m,f)$	(2), (3)
$\therefore$	(11)	$(\mathrm{ci}')\text{–}(\mathrm{ciii})_{\psi'}(m,f)$	from supp for $\exists$ elim on $(\star)$ and provability
$\therefore$	(12)	$\neg(\mathrm{ci}')\text{–}(\mathrm{ciii})_{\psi'}(m,f)$	inspection: proofs wholly absent from $\mathcal{C}$
$\therefore$	(13)	$\bot$	(11), (12)
$\therefore$	(14)	$\bot$	*reductio*; proof by cases

http://kryten.mm.rpi.edu/SB_NSG_SB_JH_DoMachine-LearningMachinesLearn_preprint.pdf

The Four-Step Road to Real Learning

Step 1: Observe the acute discontinuity of human vs. nonhuman cognition. (Only humans understand and employ e.g. abstract reasoning schemas unaffected by the physical; layered quantification; recursion; and infinite structures/infinitary reasoning.)

Step 2: Exclude forms of "learning" made possible via exclusive use of reasoning and communication capacities in nonhuman animals (i.e. exclude forms of "learning" that don't eventuate in bona fide jtb *knowledge*).

Step 3: Within the focus arising from Step 2, further narrow the focus to $HL^{\geq}$ reasoning and communication sufficiently powerful to perceive, and be successfully applied to, both (i) cohesive bodies of declarative content, and (ii) sophisticated natural-language content. Dub this **RC**.

Step 4: Real Learning ($\mathscr{RL}$) is the acquisition of genuine knowledge via **RC**.

But how is this mechanizable?

Well, how about a new form of machine learning?
(by reasoning)

Novel Form of Machine Learning:

$\boxed{\text{Learning } \textit{Ex Nihilo}}$

(or Learning *Ex Minima*)

Example: Learning *Ex Nihilo* at a Dinner Party

Robert arrives at a black-tie dinner party at a massive, manicured stone mansion.

Robert does not know anyone at this party, including the couple hosting the party[1].

Robert is seated at an elegant table; in front of him is a thin, tall, crystal glass.

A white-tuxedo-wearing server pours a bubbly liquid into the tall glass, and says "Your aperitif, sir."

At this point, Robert is in a position to learn an infinite number of propositions *ex nihilo*.

1 Although he does know which couple is hosting.

https://www.collegemagazine.com/wp-content/uploads/2015/02/downton_abbey.jpg

Rensselaer AI and Reasoning Lab

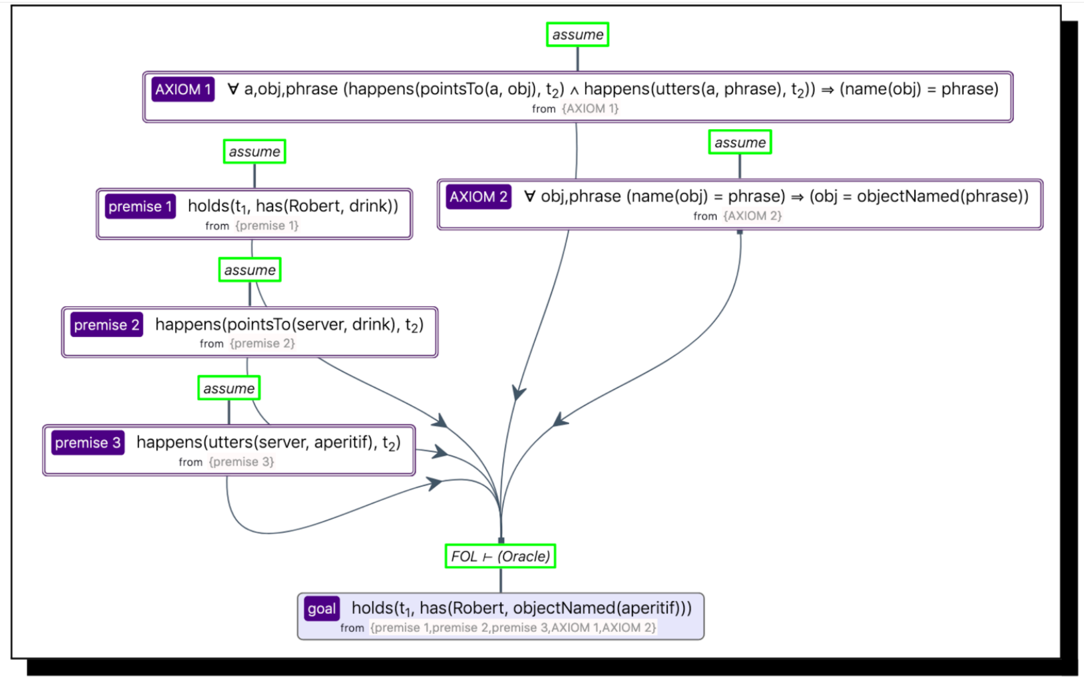

Figure 1: Dinner Party Example Part 1. (The proof here is in the HyperSlate® system (Pat. Pend.) published by Motalen in it's Logic: A Modern Approach (LAMA)® paradigm. See www.logicamodernapproach.com.)

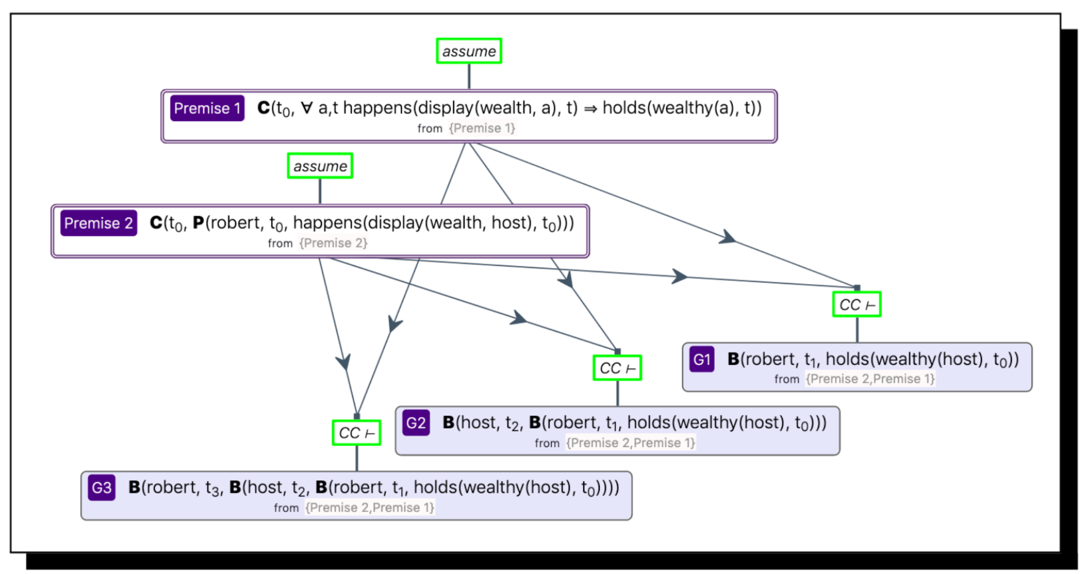

Figure 2: Dinner Party Example Part 2. (The proof here, like its predecessor, is in the HyperSlate® system (Pat. Pend.) published by Motalen in it's Logic: A Modern Approach (LAMA)® paradigm. See www.logicamodernapproach.com.)

Questions?