How'd We Arrive Here?

(Selmer's Leibnizian Whirlwind History of Logic, with Its Impact on Arguments for The Arrival of "The Singularity")

Selmer Bringsjord

Rensselaer AI & Reasoning (RAIR) Lab
Department of Cognitive Science
Department of Computer Science
Lally School of Management & Technology
Rensselaer Polytechnic Institute (RPI)
Troy, New York 12180 USA

Intro to Logic
1/29/2026

Rensselaer AI and Reasoning Lab

Logic (& AI) in the news

…

COMMENTARY

We're Planning for the Wrong AI Job Disruption

If artificial intelligence takes over some of your tasks, that doesn't render you unemployable.

By Stephen Lewarne

Jan 28, 2026 04:18 p.m. ET

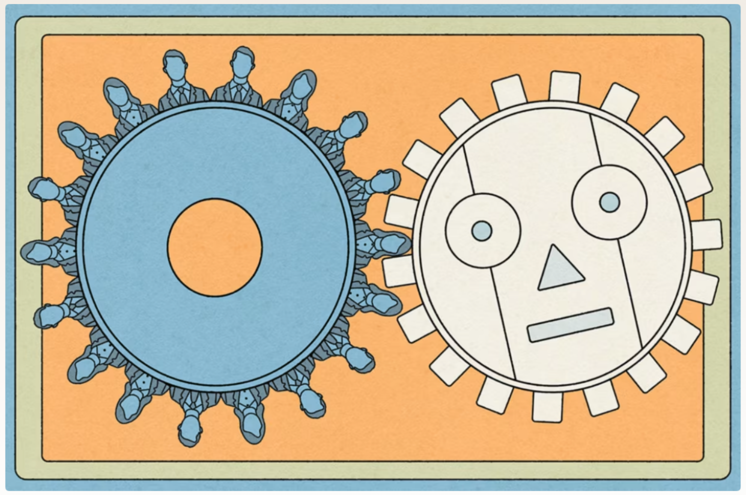

ILLUSTRATION: DAN PAGE

Washington is bracing for an artificial-intelligence employment shock that is unlikely to arrive—and so the government risks spending billions of dollars preparing for the wrong problem. Panicked politicians are making the error of treating task-based occupational

COMMENTARY

We're Planning for the Wrong AI Job Disruption

If artificial intelligence takes over some of your tasks, that doesn't render you unemployable.

By Stephen Lewarne

Jan 28, 2026 04:18 p.m. ET

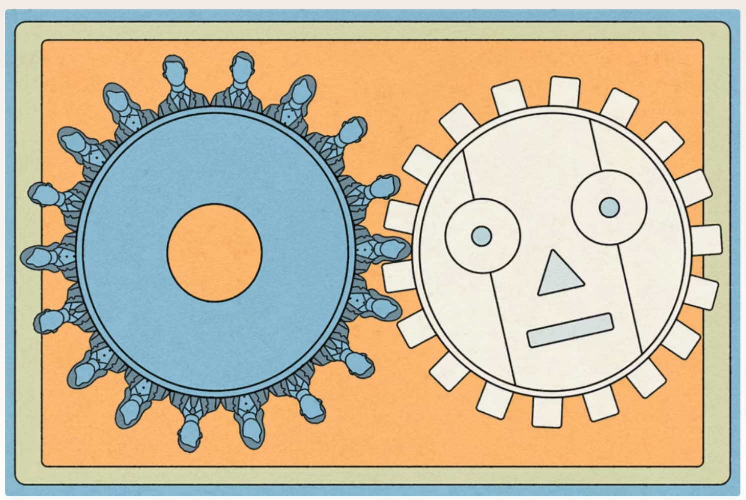

ILLUSTRATION: DAN PAGE

Washington is bracing for an artificial-intelligence employment shock that is unlikely to arrive—and so the government risks spending billions of dollars preparing for the wrong problem. Panicked politicians are making the error of treating task-based occupational

The logic of the Goldman Sachs estimate is similar to that underlying academic studies tracing back to the Frey-Osborne framework, which classifies occupations by susceptibility to computerization. It's the same method behind the Organization for Economic Cooperation and Development's 2023 Employment Outlook, which reports that roughly one-quarter of jobs across advanced economies are highly exposed to AI-driven automation—another report that has raised political alarm. These studies' approaches have a common structure: They map technologies onto tasks rather than onto labor-market outcomes. When repurposed as forecasts of displacement, they are asked to do something for which they were never designed. That is why both the OECD and Goldman Sachs explicitly caution in their reports that measures of AI "exposure" describe task susceptibility, not forecasts of job loss.

COMMENTARY

We’re Planning for the Wrong AI Job Disruption

If artificial intelligence takes over some of your tasks, that doesn’t render you unemployable.

By Stephen Lewarne

Jan 28, 2026 04:18 p.m. ET

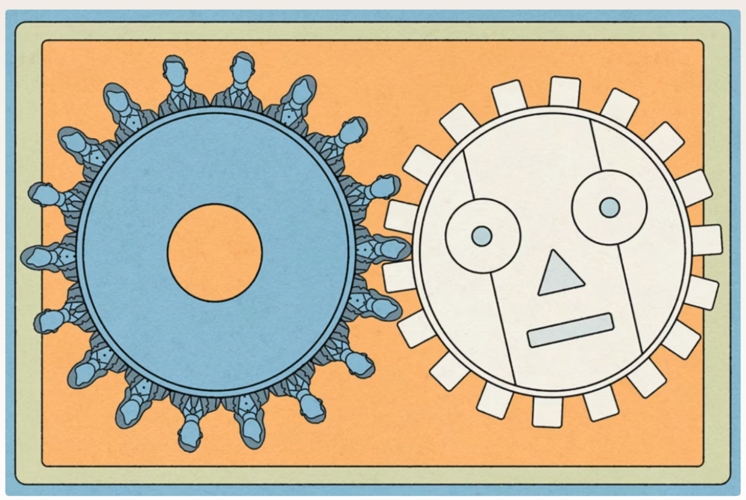

ILLUSTRATION: DAN PAGE

Washington is bracing for an artificial-intelligence employment shock that is unlikely to arrive—and so the government risks spending billions of dollars preparing for the wrong problem. Panicked politicians are making the error of treating task-based occupational

The logic of the Goldman Sachs estimate is similar to that underlying academic studies tracing back to the Frey-Osborne framework, which classifies occupations by susceptibility to computerization. It’s the same method behind the Organization for Economic Cooperation and Development’s 2023 Employment Outlook, which reports that roughly one-quarter of jobs across advanced economies are highly exposed to AI-driven automation—another report that has raised political alarm.

These studies’ approaches have a common structure: They map technologies onto tasks rather than onto labor-market outcomes. When repurposed as forecasts of displacement, they are asked to do something for which they were never designed. That is why both the OECD and Goldman Sachs explicitly caution in their reports that measures of AI “exposure” describe task susceptibility, not forecasts of job loss.

Illogical Reasoning

90% of the tasks that are carried out by a human *h* in doing job *j* will be automated tomorrow by AIs.

Illogical Reasoning

90% of the tasks that are carried out by a human *h* in doing job *j* will be automated tomorrow by AIs.

$\therefore$

Illogical Reasoning

90% of the tasks that are carried out by a human *h* in doing job *j* will be automated tomorrow by AIs.

∴ Human *h* will tomorrow not have job *j*.

Illogical Reasoning

90% of the tasks that are carried out by a human *h* in doing job *j* will be automated tomorrow by AIs.

~~∴~~ Human *h* will tomorrow not have job *j*.

Illogical Reasoning

90% of the tasks that are carried out by a human *h* in doing job *j* will be automated tomorrow by AIs.

X Human *h* will tomorrow not have job *j*.

Later (see syllabus), history w.r.t. Powerful-Autonomous-Intelligent Machines

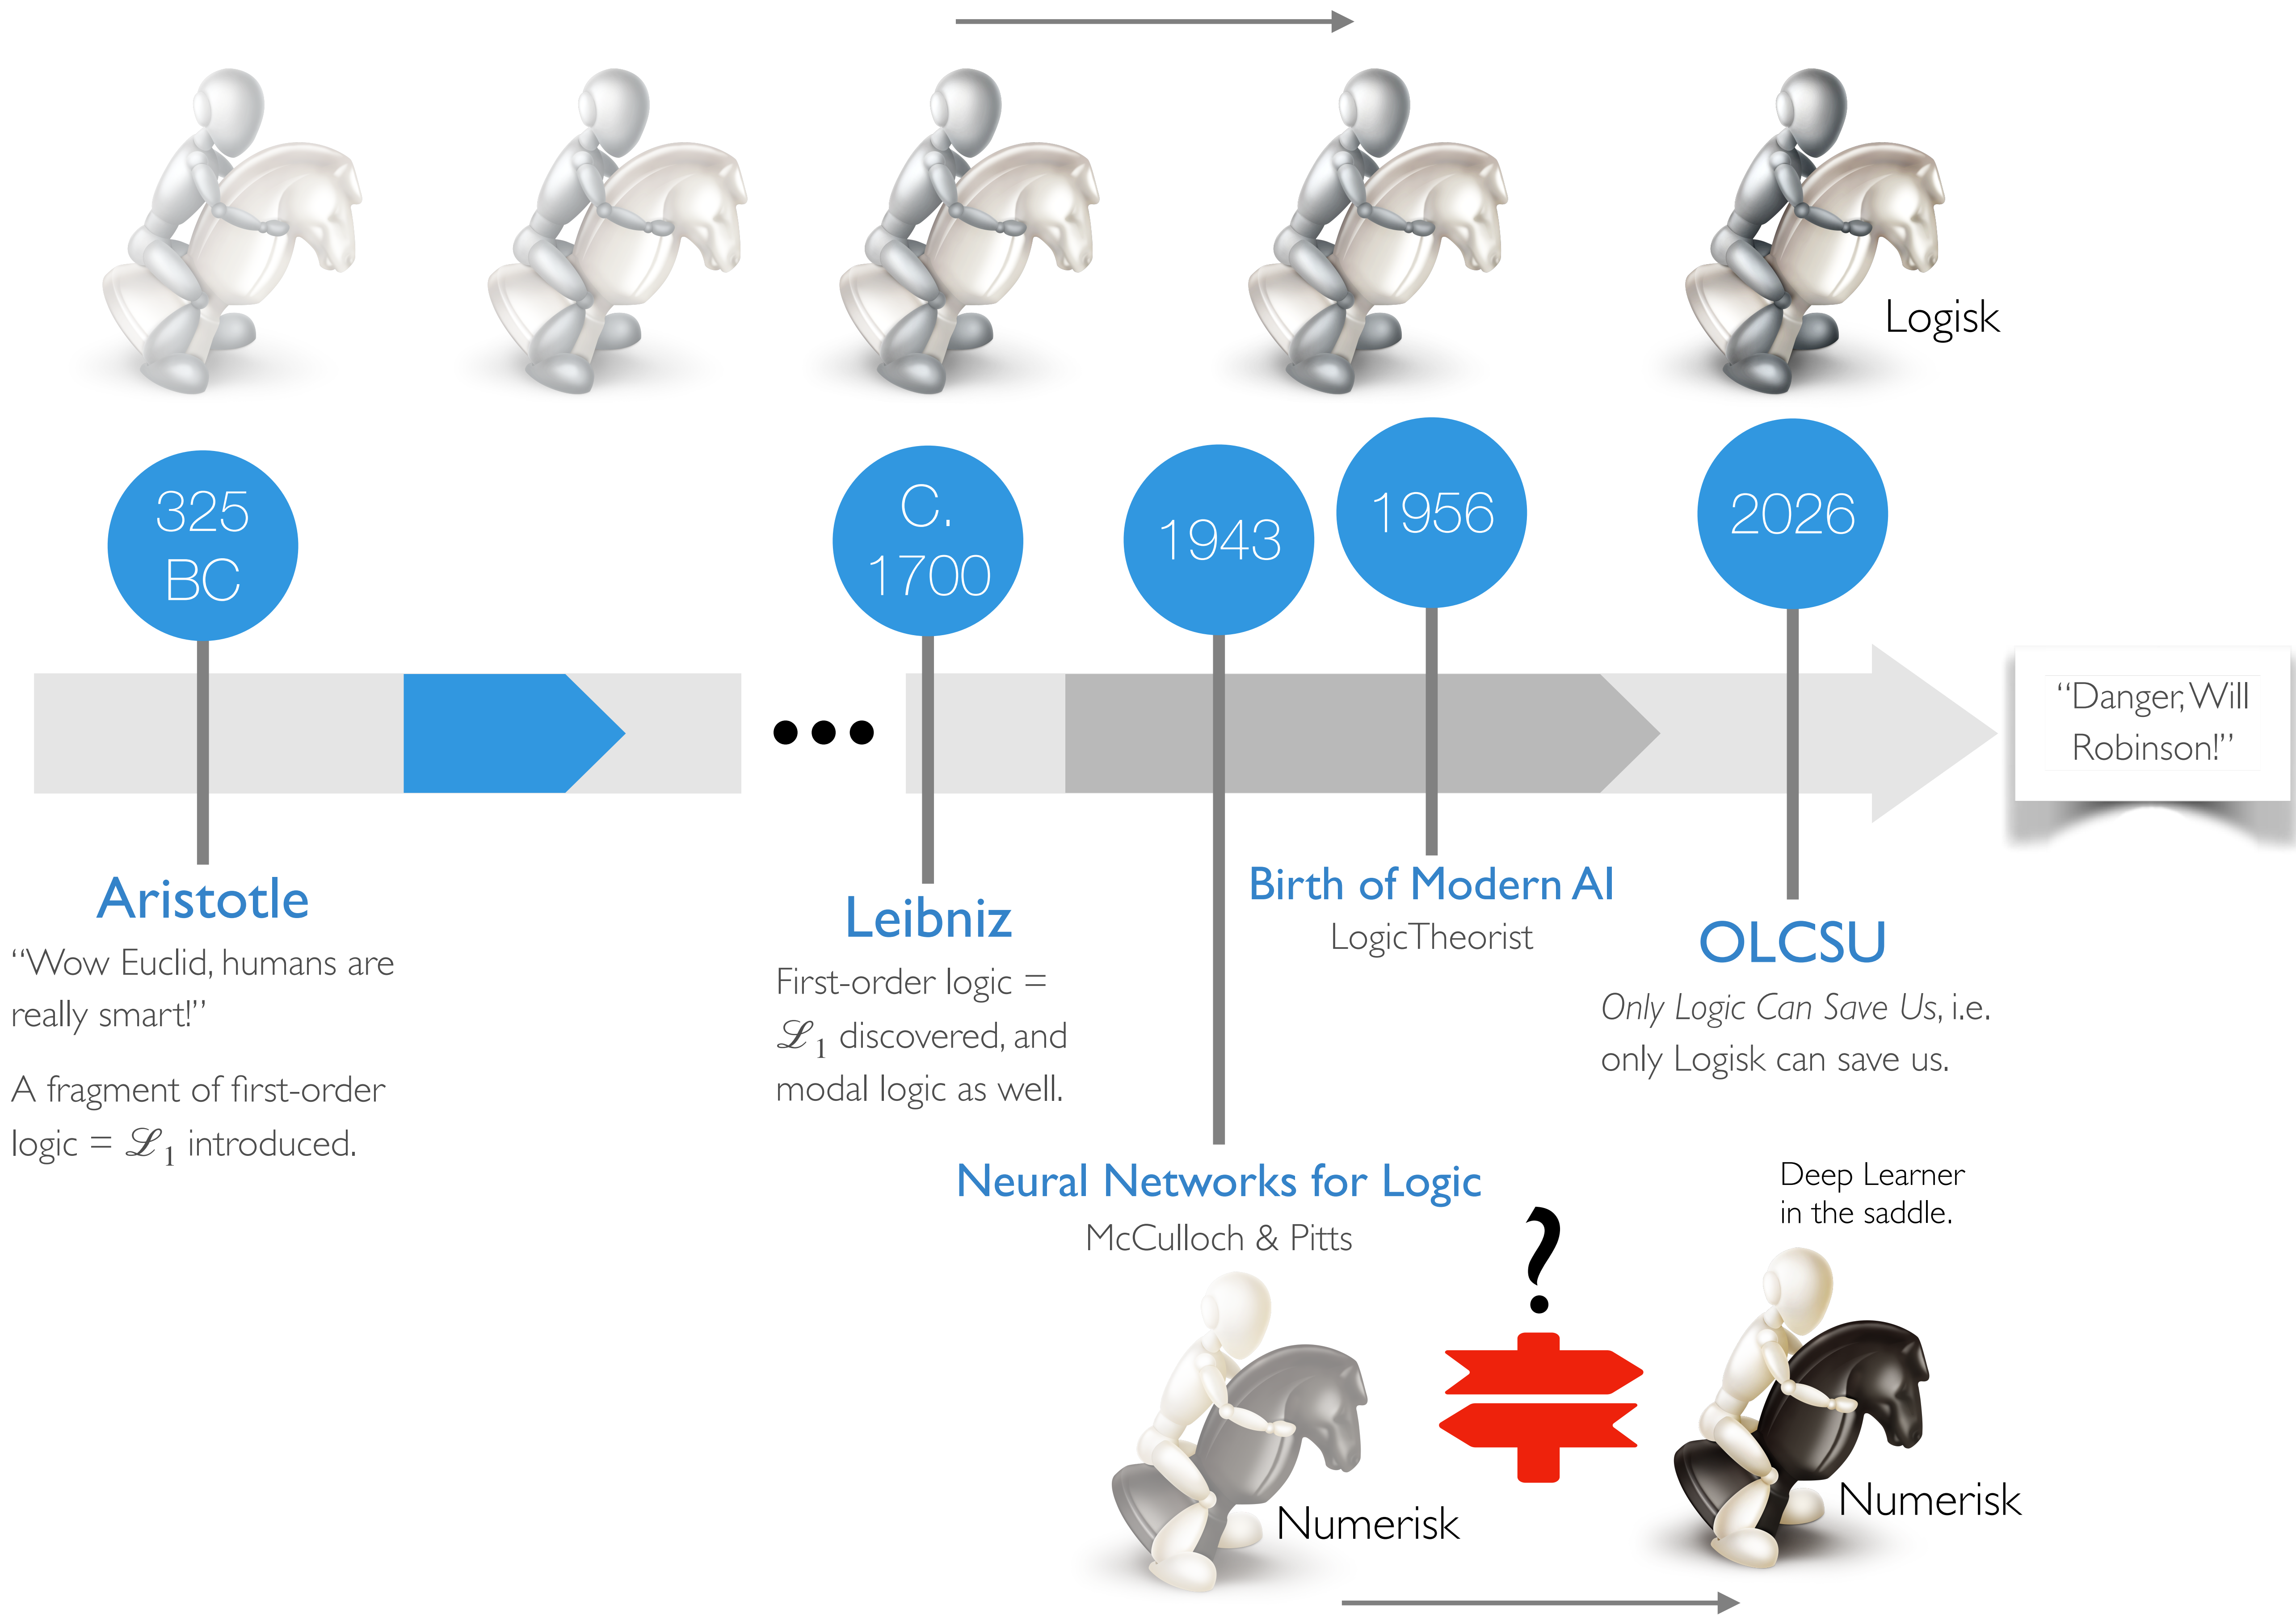
Logisk
325 BC
C. 1700
1943
1956
2026
"Danger, Will Robinson!"
Aristotle
"Wow Euclid, humans are really smart!"
A fragment of first-order logic = $\mathscr{L}_1$ introduced.
Leibniz
First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.
Birth of Modern AI
LogicTheorist
OLCSU
Only Logic Can Save Us, i.e. only Logisk can save us.
Neural Networks for Logic
McCulloch & Pitts
Deep Learner in the saddle.
?
Numerisk
Numerisk

Last time …

A criminal genius nearly a match for Sherlock Holmes (Do you recognize the Dr?) has built a massive hydrogen bomb, and life on Earth is hanging in the balance, hinging on whether you make the logical prediction. Dr M gives you a sporting chance to: make the right prediction, snip or not snip accordingly, and prove that you’re right …

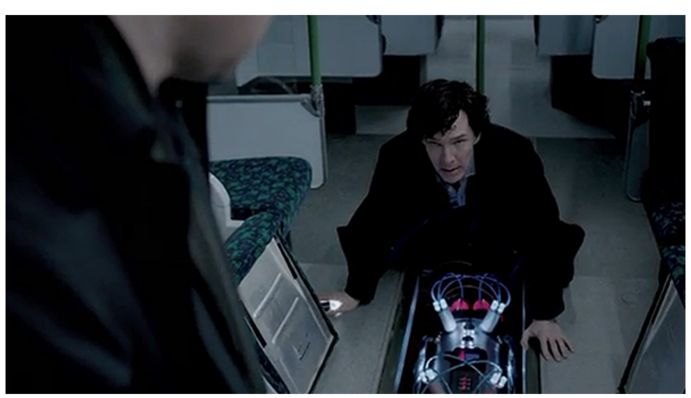

A criminal genius nearly a match for Sherlock Holmes (Do you recognize the Dr?)

A criminal genius nearly a match for Sherlock Holmes (Do you recognize the Dr?) has built a massive hydrogen bomb, and life on Earth is hanging in the balance, hinging on whether you make the logical prediction. Dr M gives you a sporting chance to: make the right prediction, snip or not snip accordingly, and prove that you're right …

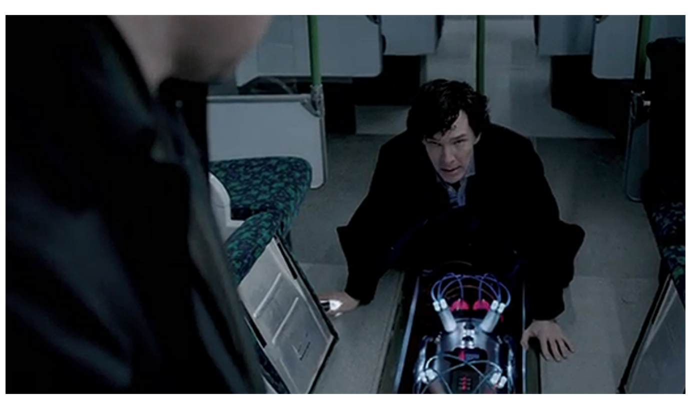

```
If one of the following assertions is true then so is
the other:

(1) If the red wire runs to the bomb, then the blue
wire runs to the bomb; and, if the blue wire runs to
the bomb, then the red wire runs to the bomb.

(2) The red wire runs to the bomb.
```

Given this perfectly reliable clue from Dr Moriarty, if either wire is more likely to run to the bomb, that wire *does* run to the bomb, and the bomb is ticking, with only a minute left! If both are equiprobable, neither runs to the bomb, and you are powerless. Make your prediction as to what will happen when a wire is snipped, and then make your selected snip by clicking on the wire you want to snip! Or leave well enough alone!

Red more likely.

Blue more likely.

Equiprobable.

Snip

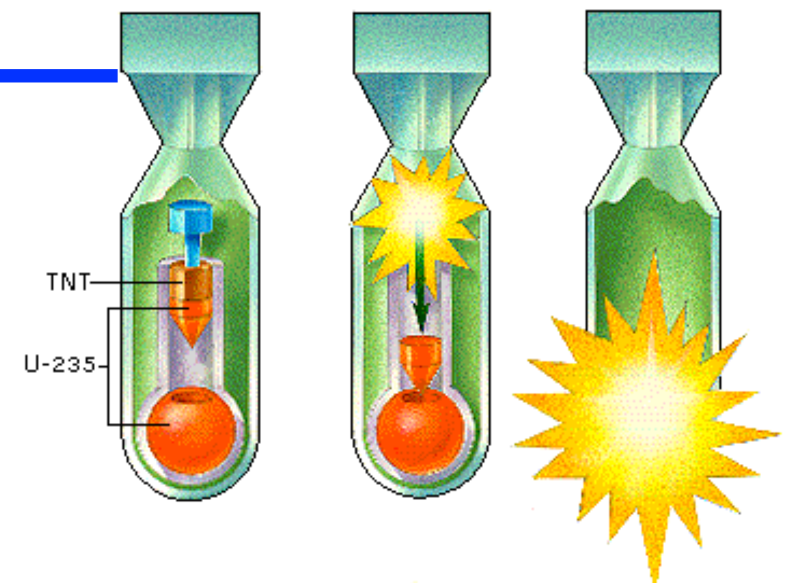

Life on Earth has ended .

advance one more slide to see a proof that you indeed made an irrational decision…

Proposition: The blue wire is more likely!

Proof: (1) can be treated as a biconditional, obviously (R <=> B).

There are two top-level cases to consider: (1) and (2) are both true; or both are false. In the case where they are both true, it's trivial to deduce both R and B. So far, then, R and B are equiprobable. What happens in the case where (1) and (2) are both false? We immediately have ~R from the denial of (2). But a biconditional is true just in case both sides are true, or both sides are false; so we have two sub-cases to consider.

Consider first the case where R is true and B is false. We have an immediate contradiction in this sub-case, so both R and B can both be deduced here, and we have not yet departed from equiprobable. So what about the case where R is false and B is true? The falsity of R is not new information (we already have that from the denial of (2)), but we can still derive B. Hence the blue wire is more likely. **QED**

STOP

Snip

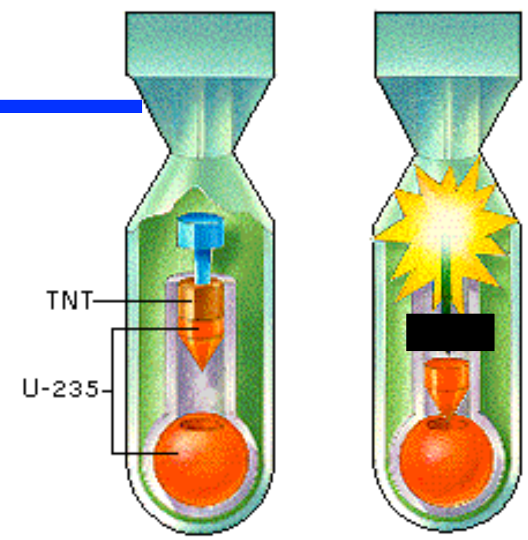

Life on Earth is saved!

if you can now hand Dr M a proof that your decision was the rational one!

Advance one more slide to see a proof from Bringsjord that yours had better match up to …

Proposition: The blue wire is more likely!

Proof: (1) can be treated as a biconditional, obviously (R <=> B).

There are two top-level cases to consider: (1) and (2) are both true; or both are false. In the case where they are both true, it's trivial to deduce both R and B. So far, then, R and B are equiprobable. What happens in the case where (1) and (2) are both false? We immediately have ~R from the denial of (2). But a biconditional is true just in case both sides are true, or both sides are false; so we have two sub-cases to consider.

Consider first the case where R is true and B is false. We have an immediate contradiction in this sub-case, so both R and B can both be deduced here, and we have not yet departed from equiprobable. So what about the case where R is false and B is true? The falsity of R is not new information (we already have that from the denial of (2)), but we can still derive B. Hence the blue wire is more likely. **QED**

STOP

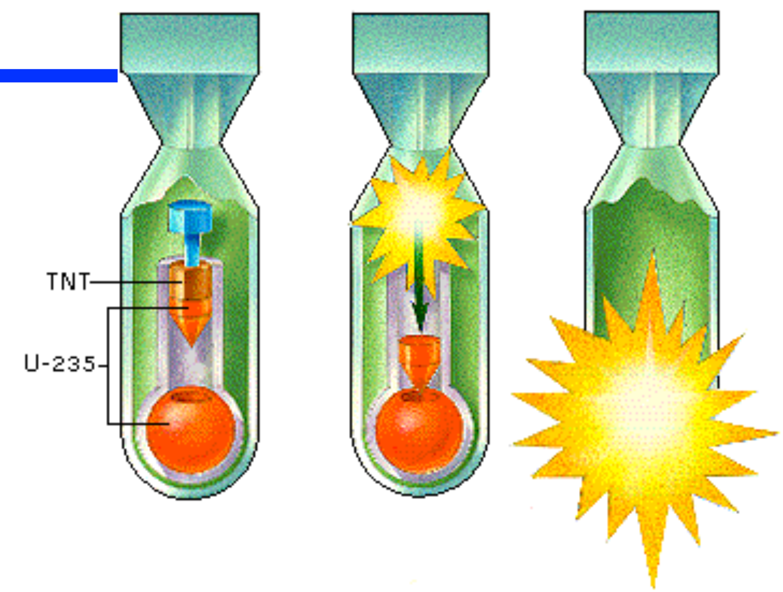

Life
on
Earth
has
ended
.

advance one more slide to see a proof that you indeed made an irrational decision…

Proposition: The blue wire is more likely!

Proof: (1) can be treated as a biconditional, obviously (R <=> B).

There are two top-level cases to consider: (1) and (2) are both true; or both are false. In the case where they are both true, it's trivial to deduce both R and B. So far, then, R and B are equiprobable. What happens in the case where (1) and (2) are both false? We immediately have ~R from the denial of (2). But a biconditional is true just in case both sides are true, or both sides are false; so we have two sub-cases to consider.

Consider first the case where R is true and B is false. We have an immediate contradiction in this sub-case, so both R and B can both be deduced here, and we have not yet departed from equiprobable. So what about the case where R is false and B is true? The falsity of R is not new information (we already have that from the denial of (2)), but we can still derive B. Hence the blue wire is more likely. **QED**

STOP

The Universe of Logics

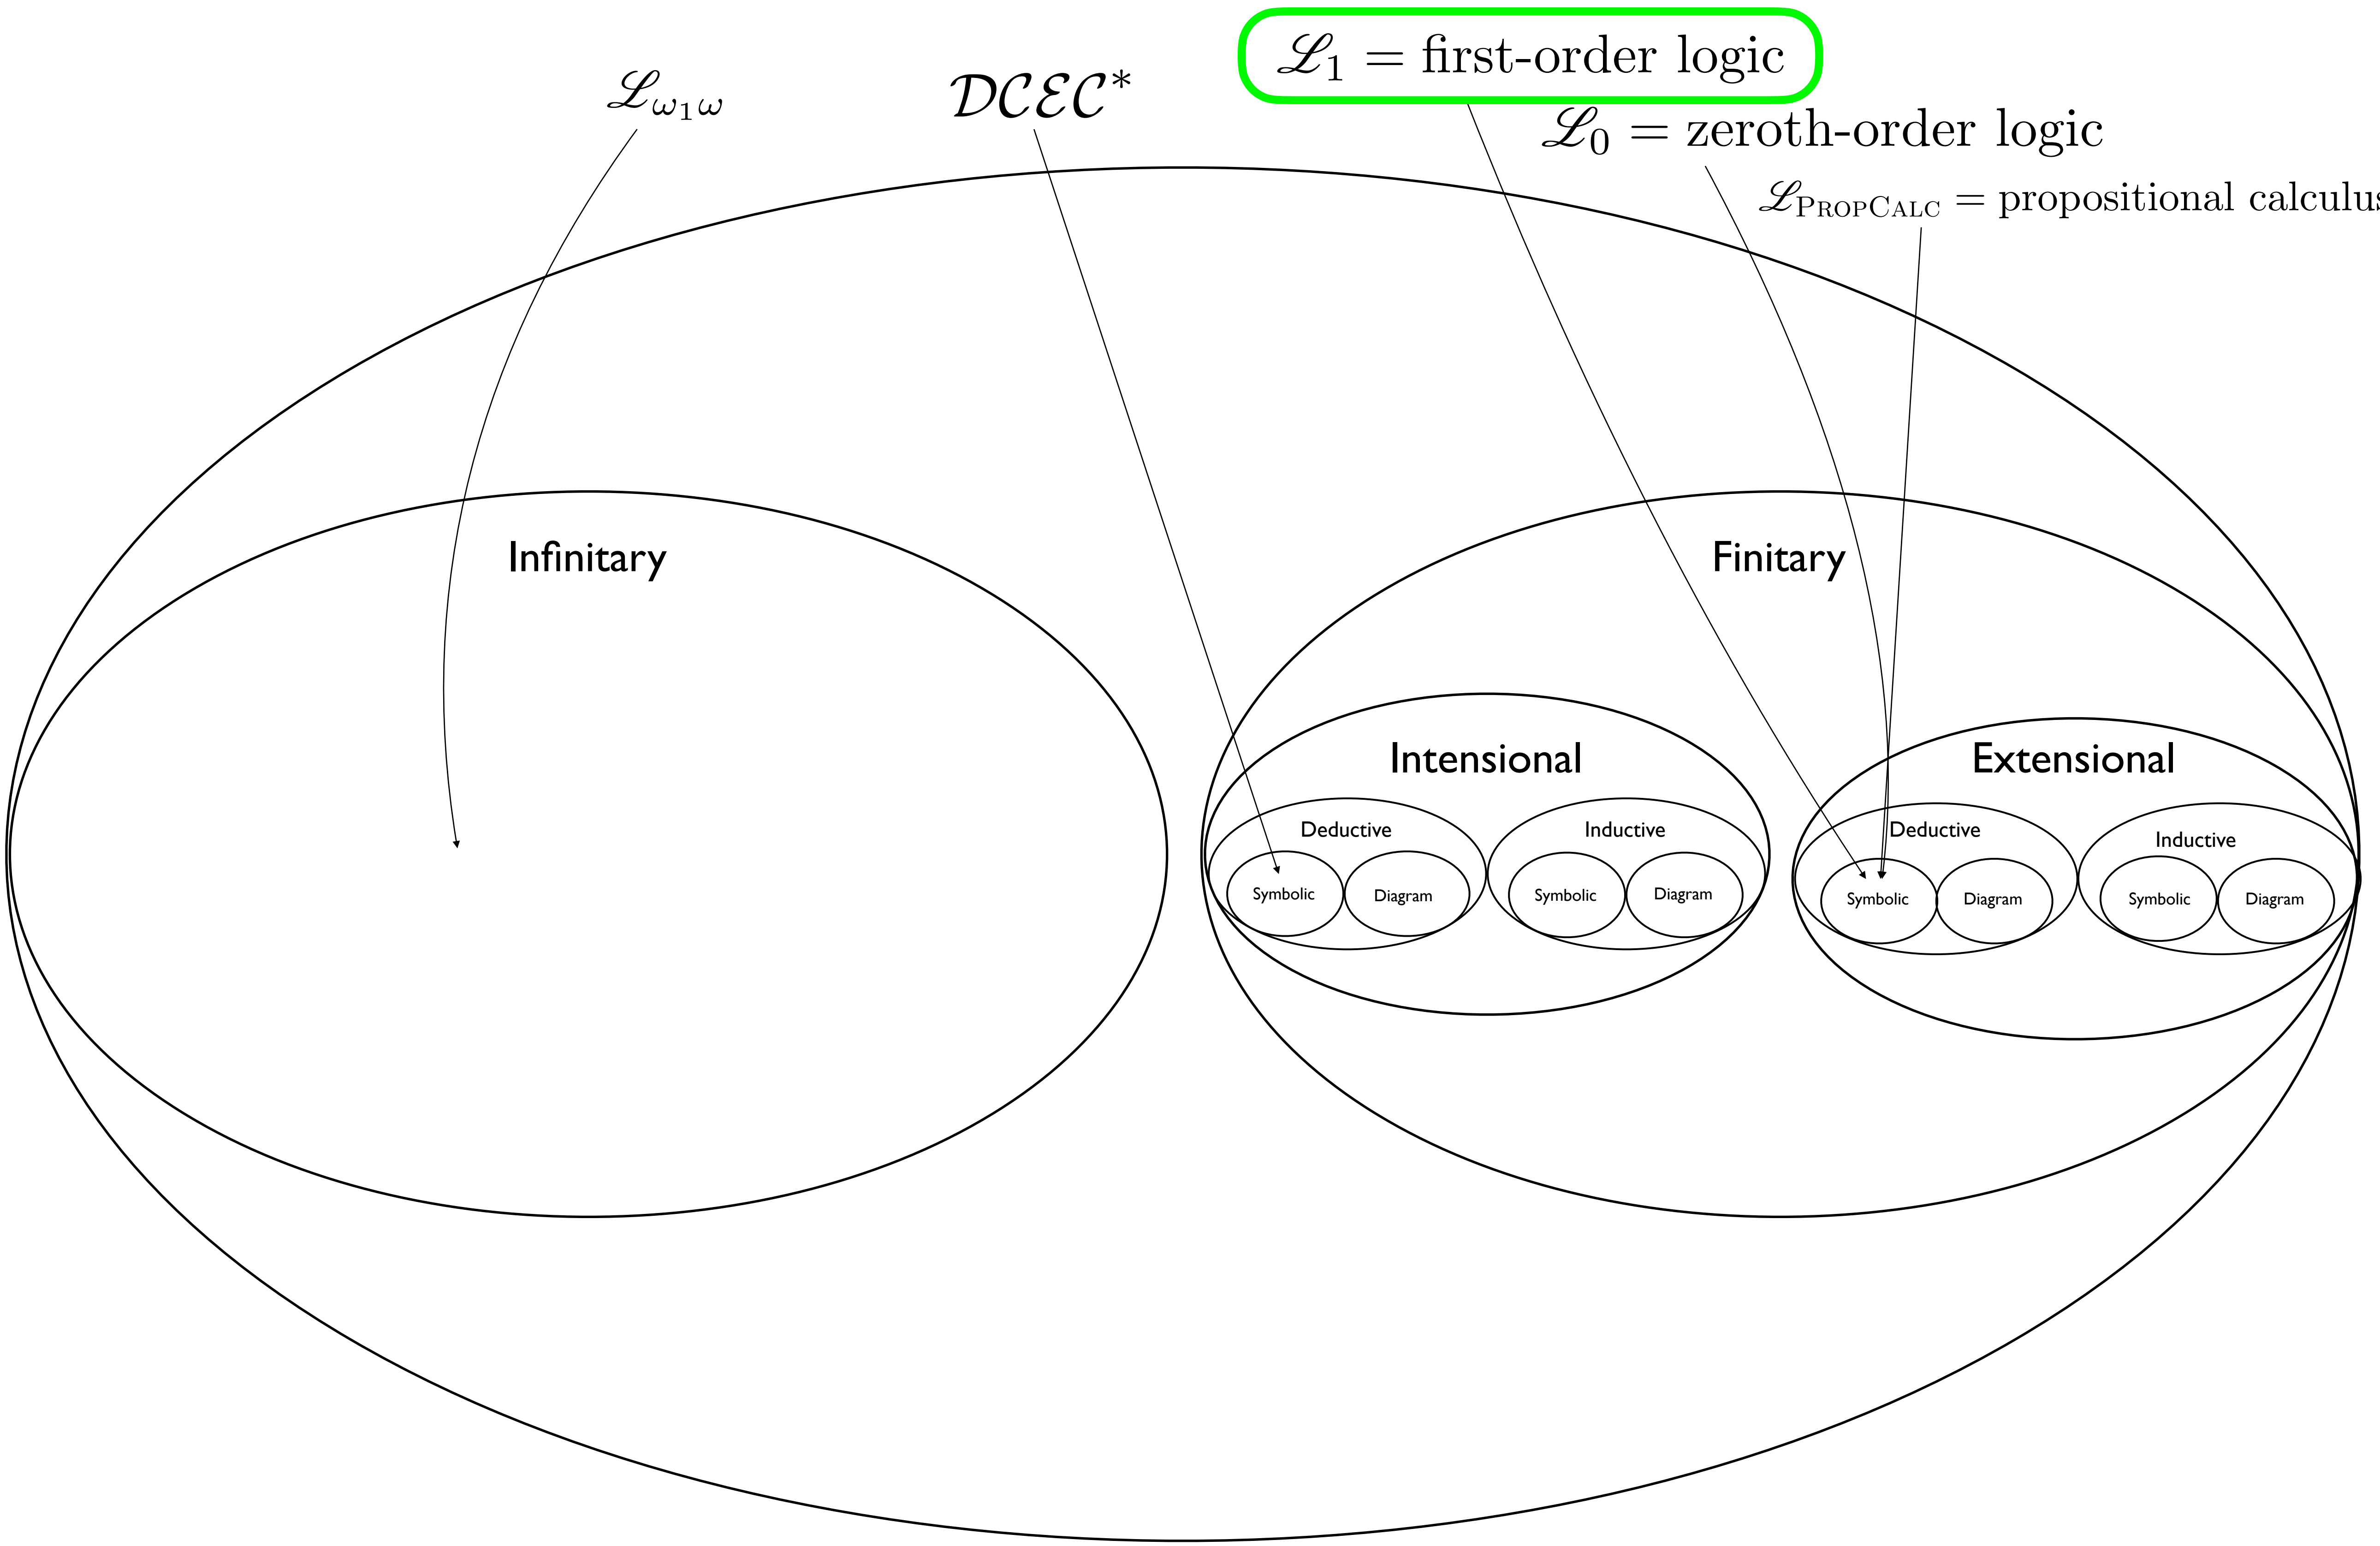

Special Llamas Disjunction

There's a thing such that it's both a llama and a non-llama;
or
there's a thing such that if it's a llama, everything is a llama;
or
there's a thing such that every llama is a non-llama.

Special Llamas Disjunction

There's a thing such that it's both a llama and a non-llama;
or
there's a thing such that if it's a llama, everything is a llama;
or
there's a thing such that every llama is a non-llama.

Is this disjunction TRUE, FALSE, or UNKNOWN?

Special Llamas Disjunction

There's a thing such that it's both a llama and a non-llama;
or
there's a thing such that if it's a llama, everything is a llama;
or
there's a thing such that every llama is a non-llama.

Is this disjunction TRUE, FALSE, or UNKNOWN?

Special Llamas Disjunction

There's a thing such that it's both a llama and a non-llama;
or
there's a thing such that if it's a llama, everything is a llama;
or
there's a thing such that every llama is a non-llama.

Is this disjunction TRUE, FALSE, or UNKNOWN?

Special Llamas Disjunction

There's a thing such that it's both a llama and a non-llama;
or
there's a thing such that if it's a llama, everything is a llama;
or
there's a thing such that every llama is a non-llama.

Is this disjunction TRUE, FALSE, or UNKNOWN?

Later … discover a formal proof!

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

abstract-and-valid inference schemata

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

abstract-and-valid inference schemata

quantification

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

abstract-and-valid inference schemata

quantification

intensional reasoning

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

abstract-and-valid inference schemata

quantification

intensional reasoning

recursion

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

abstract-and-valid inference schemata

quantification

intensional reasoning

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, or create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

recursion

self-reference

abstract-and-valid inference schemata

quantification

intensional reasoning

recursion

self-reference

To infinity and beyond! — routinely

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

abstract-and-valid inference schemata

quantification

intensional reasoning

HS®

recursion

self-reference

To infinity and beyond! — routinely

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, or create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

abstract-and-valid inference schemata

quantification

intensional reasoning

HS®

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

recursion

self-reference

To infinity and beyond! — routinely

abstract-and-valid inference schemata

quantification

intensional reasoning

HS®

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

recursion

self-reference

To infinity and beyond! — routinely

abstract-and-valid inference schemata

quantification

intensional reasoning

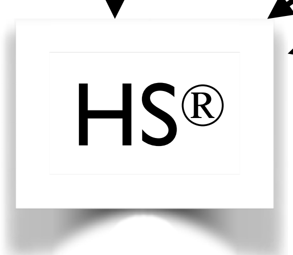

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

recursion

self-reference

To infinity and beyond! — routinely

abstract-and-valid inference schemata

quantification

intensional reasoning

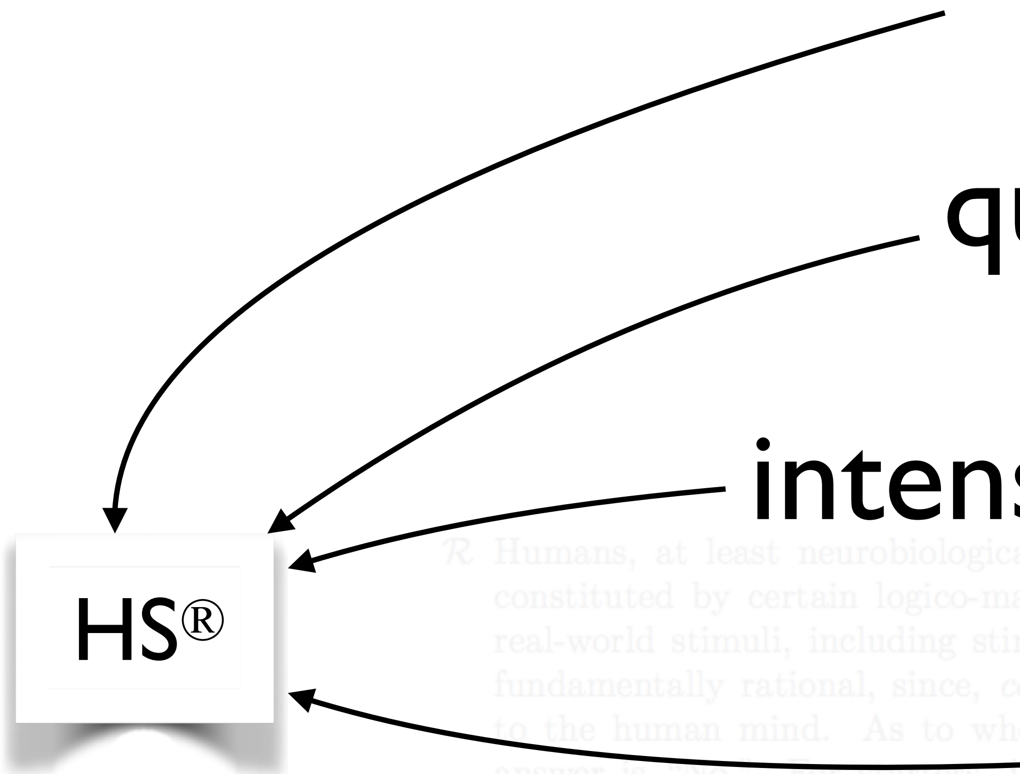

$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

recursion

self-reference

To infinity and beyond! — routinely

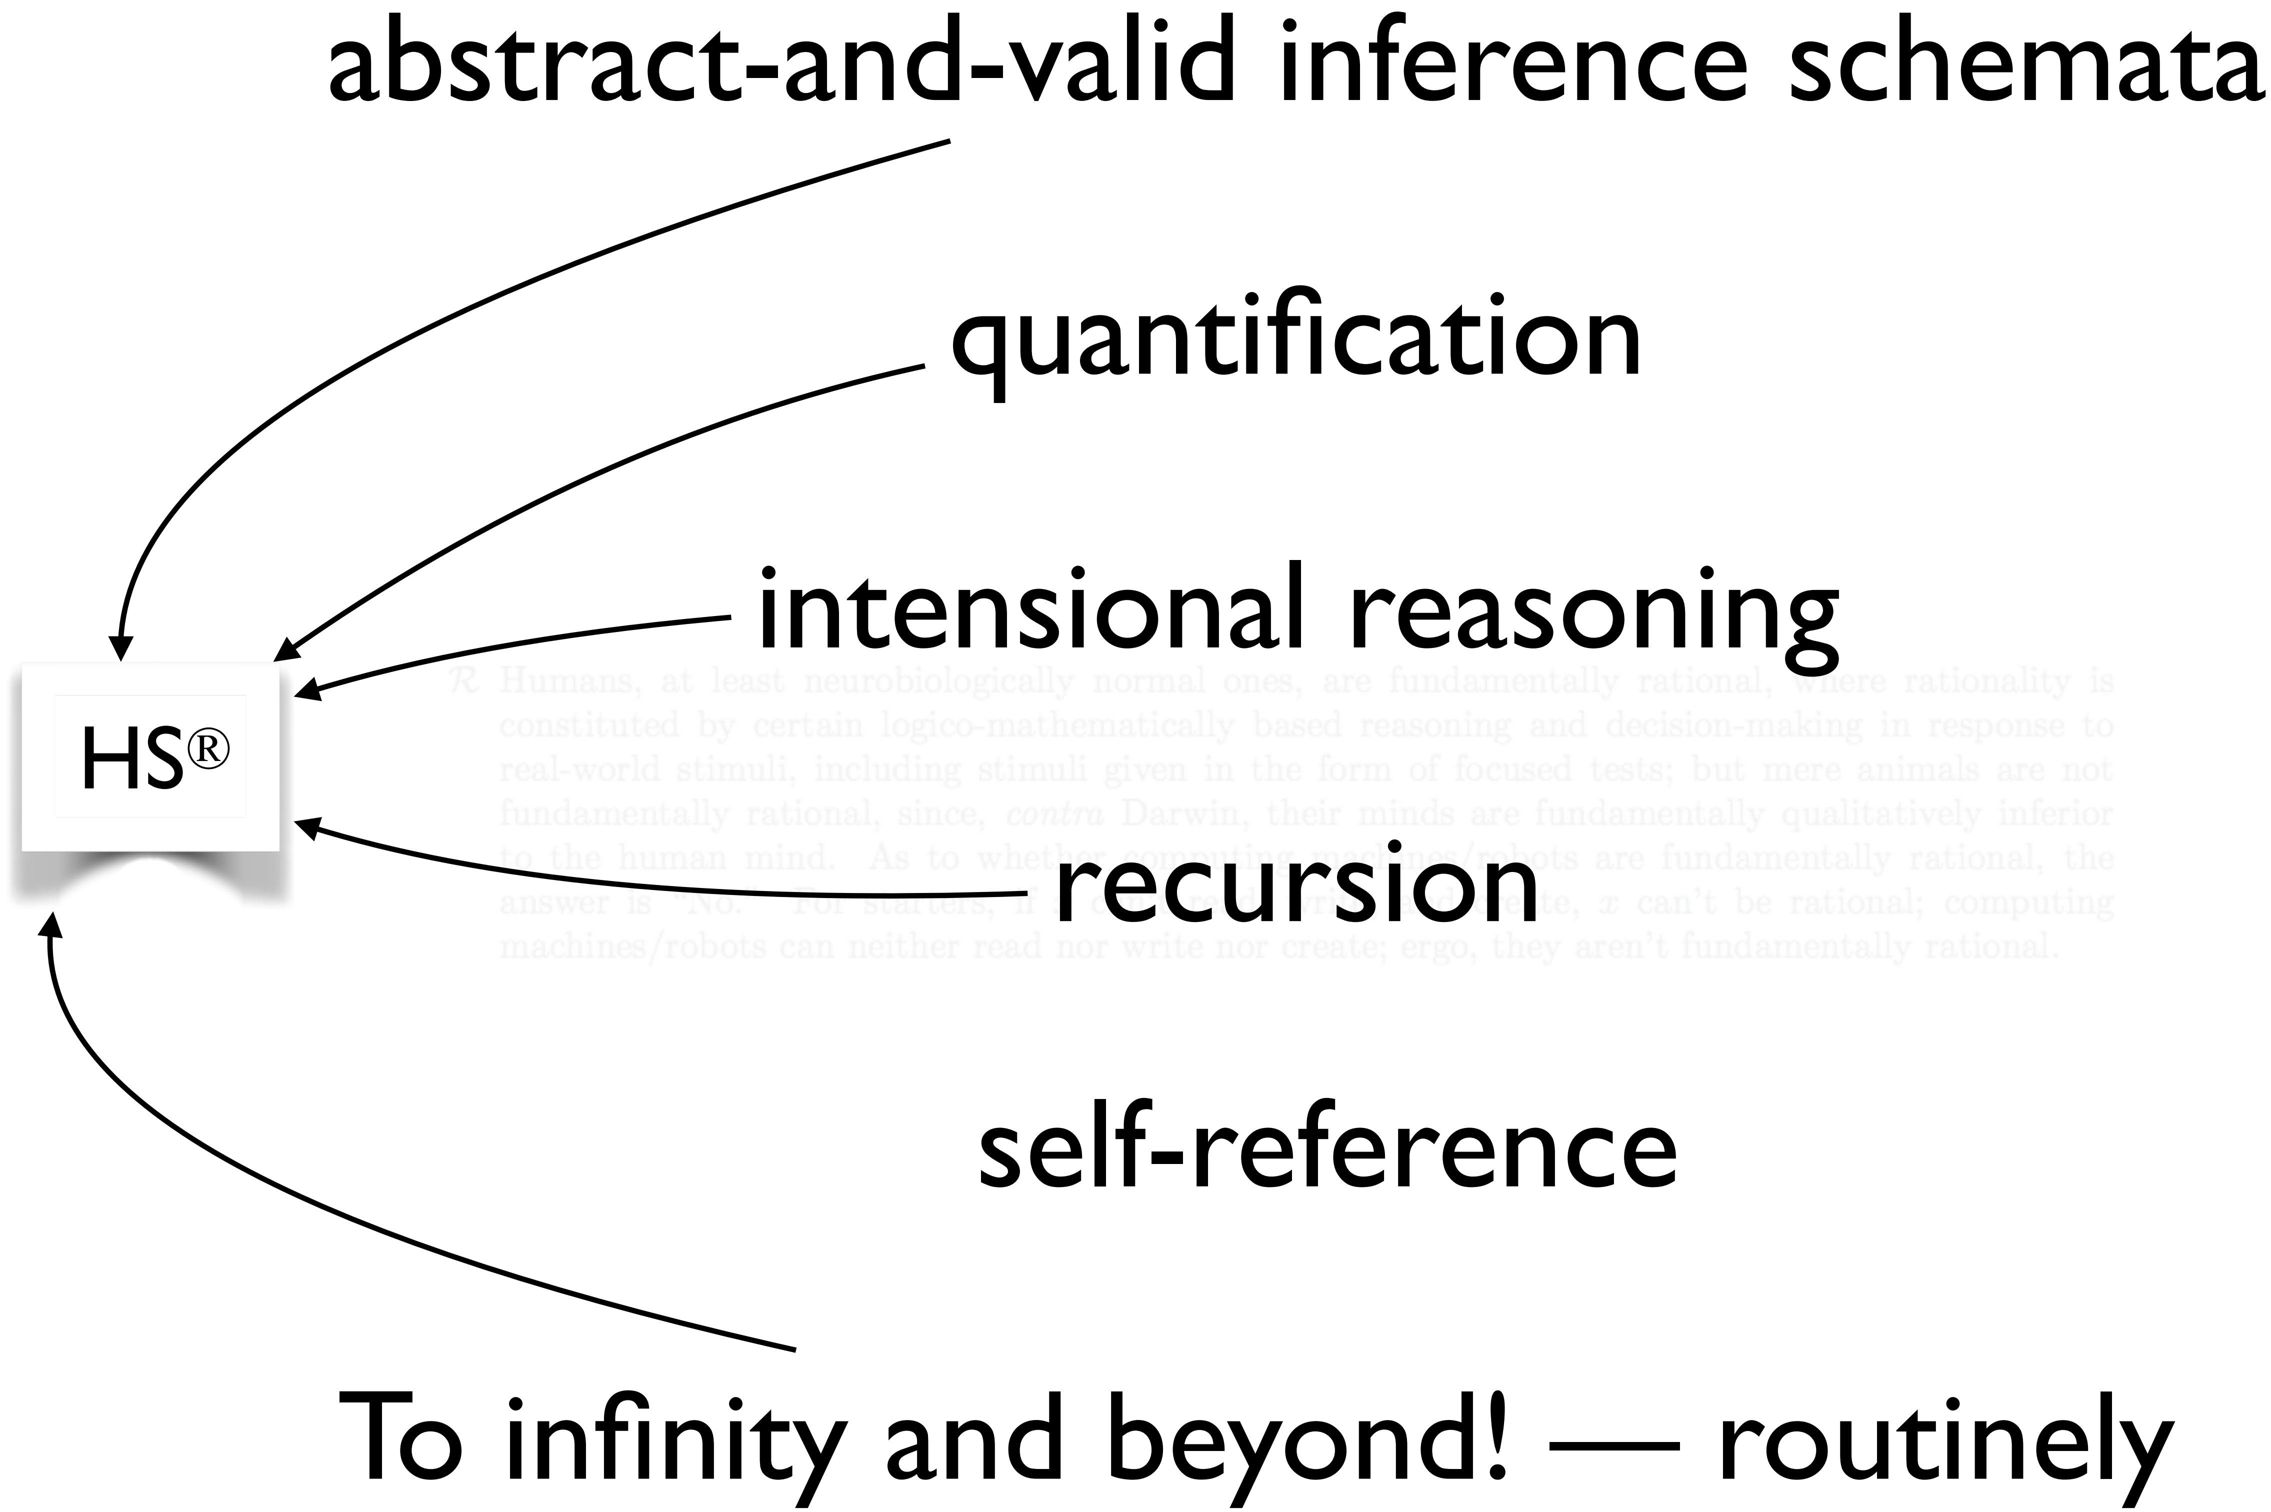
abstract-and-valid inference schemata
quantification
intensional reasoning
HS®
recursion
self-reference
To infinity and beyond! — routinely
ℛ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, contra Darwin, their minds are fundamentally qualitatively inferior
to the human mind. As to whether
are fundamentally rational, the
answer is
can't be rational; computing
machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

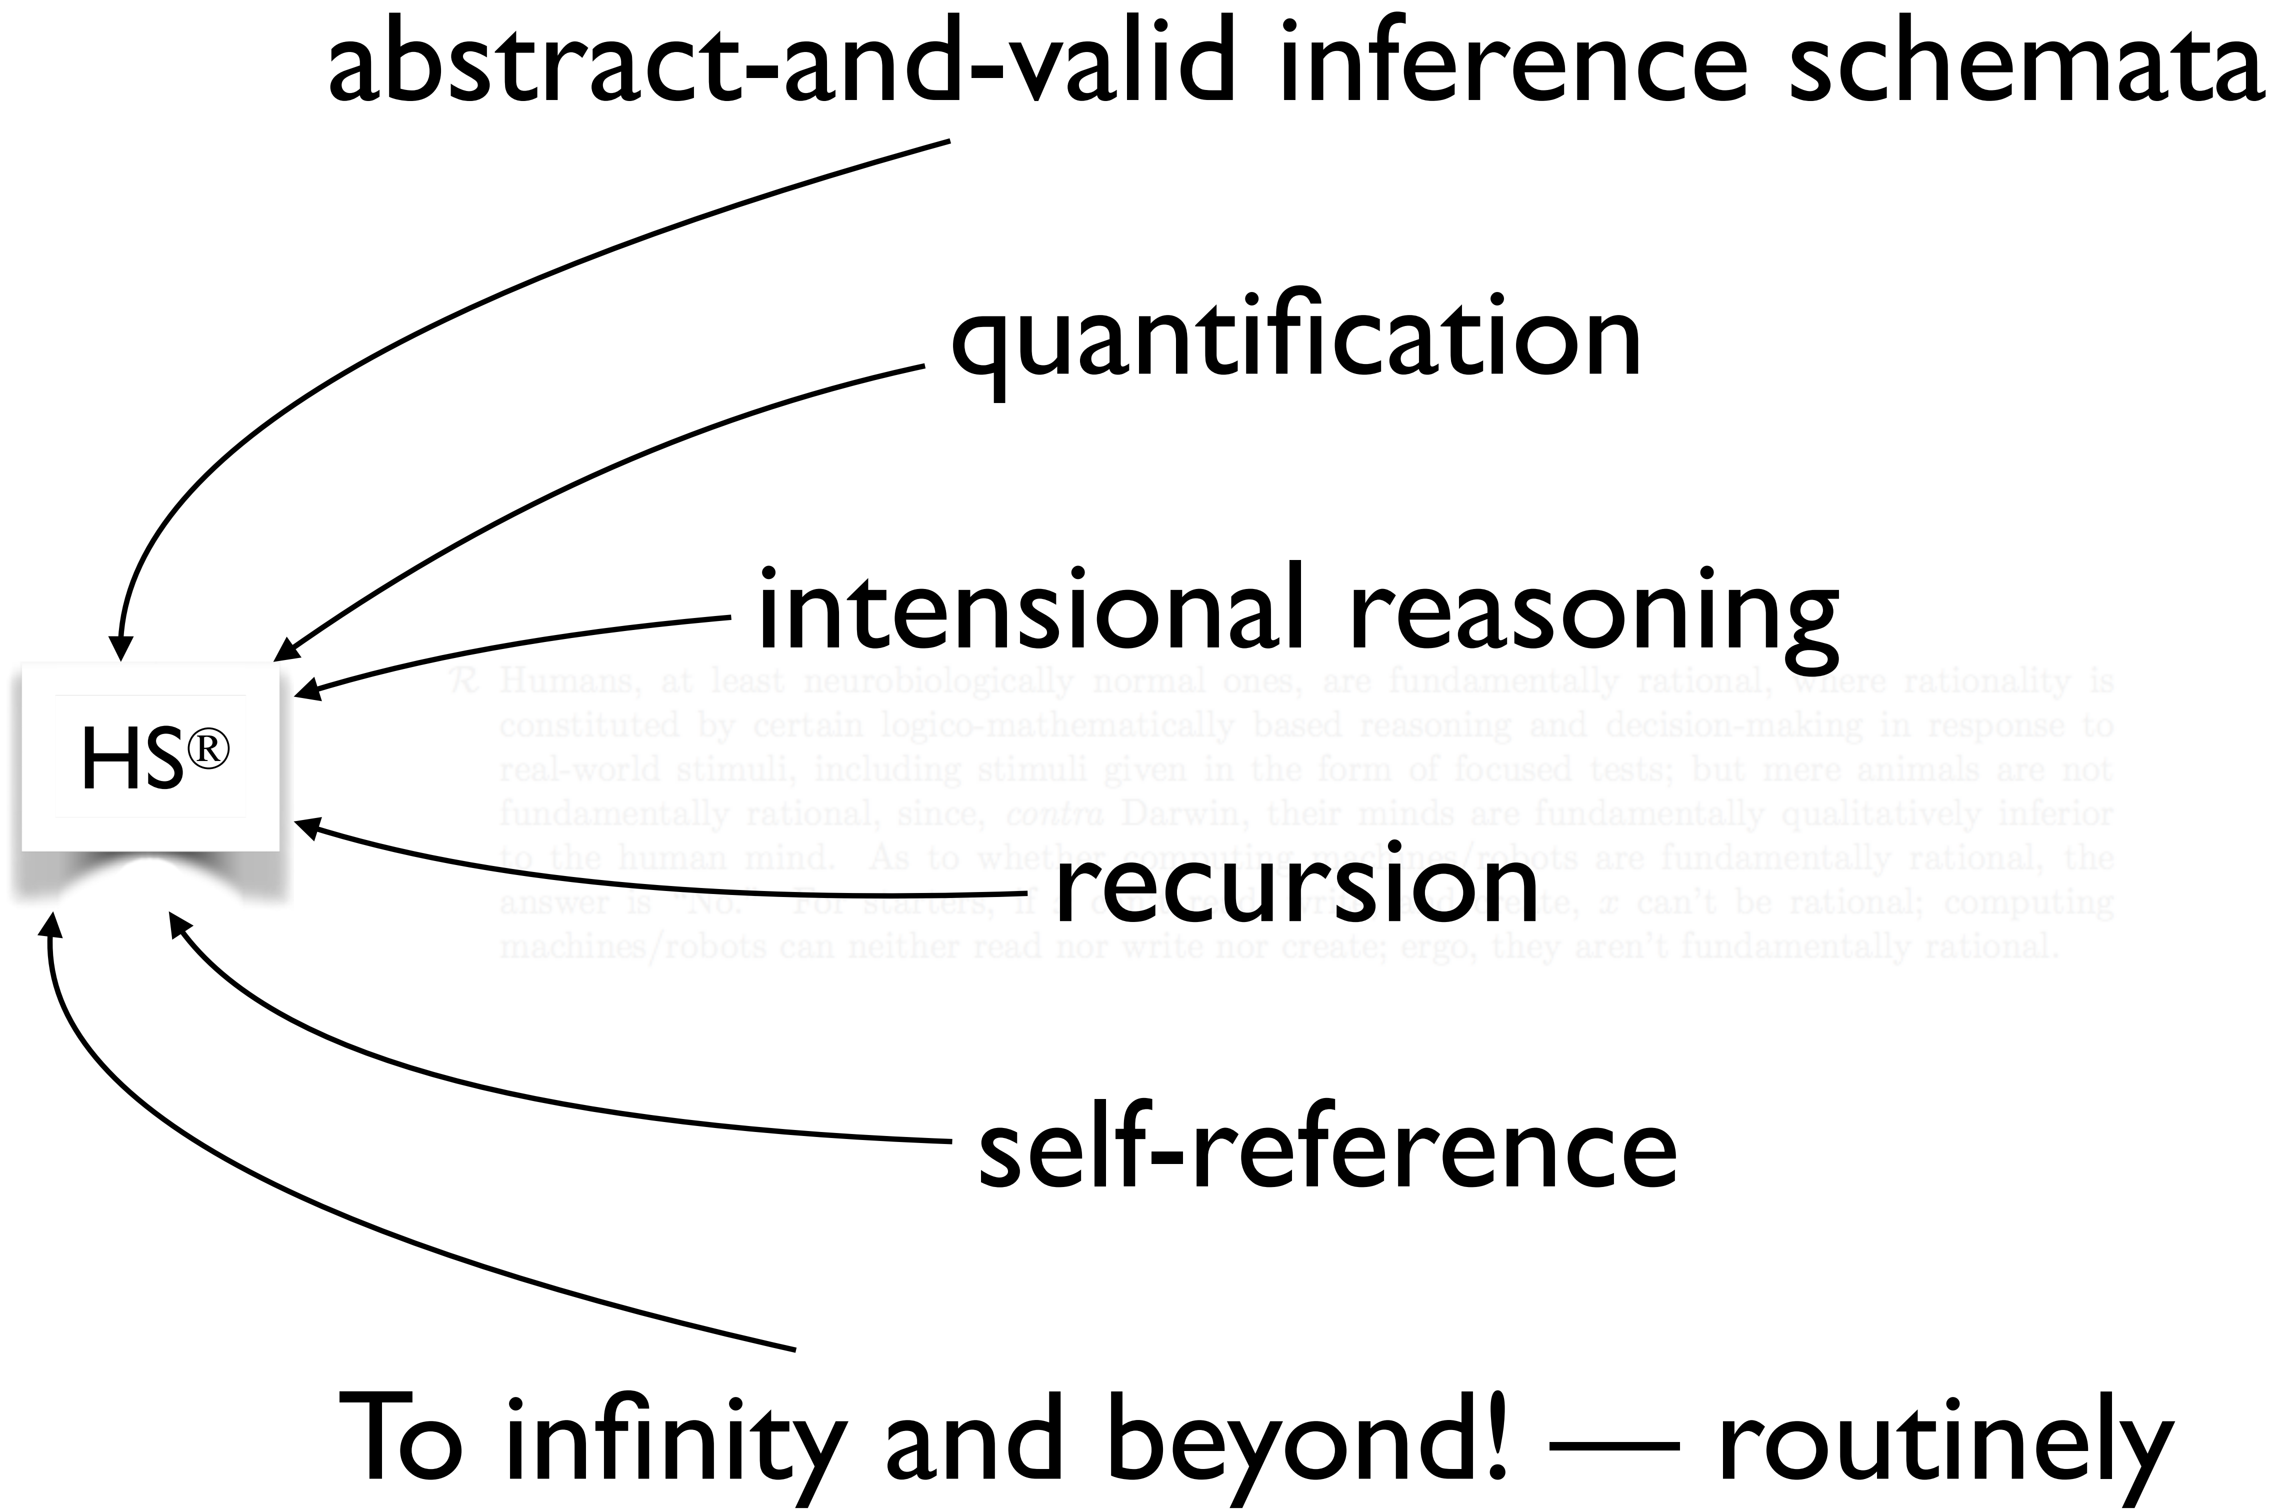
abstract-and-valid inference schemata
quantification
intensional reasoning
HS®
$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is "No." For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.
recursion
self-reference
To infinity and beyond! — routinely

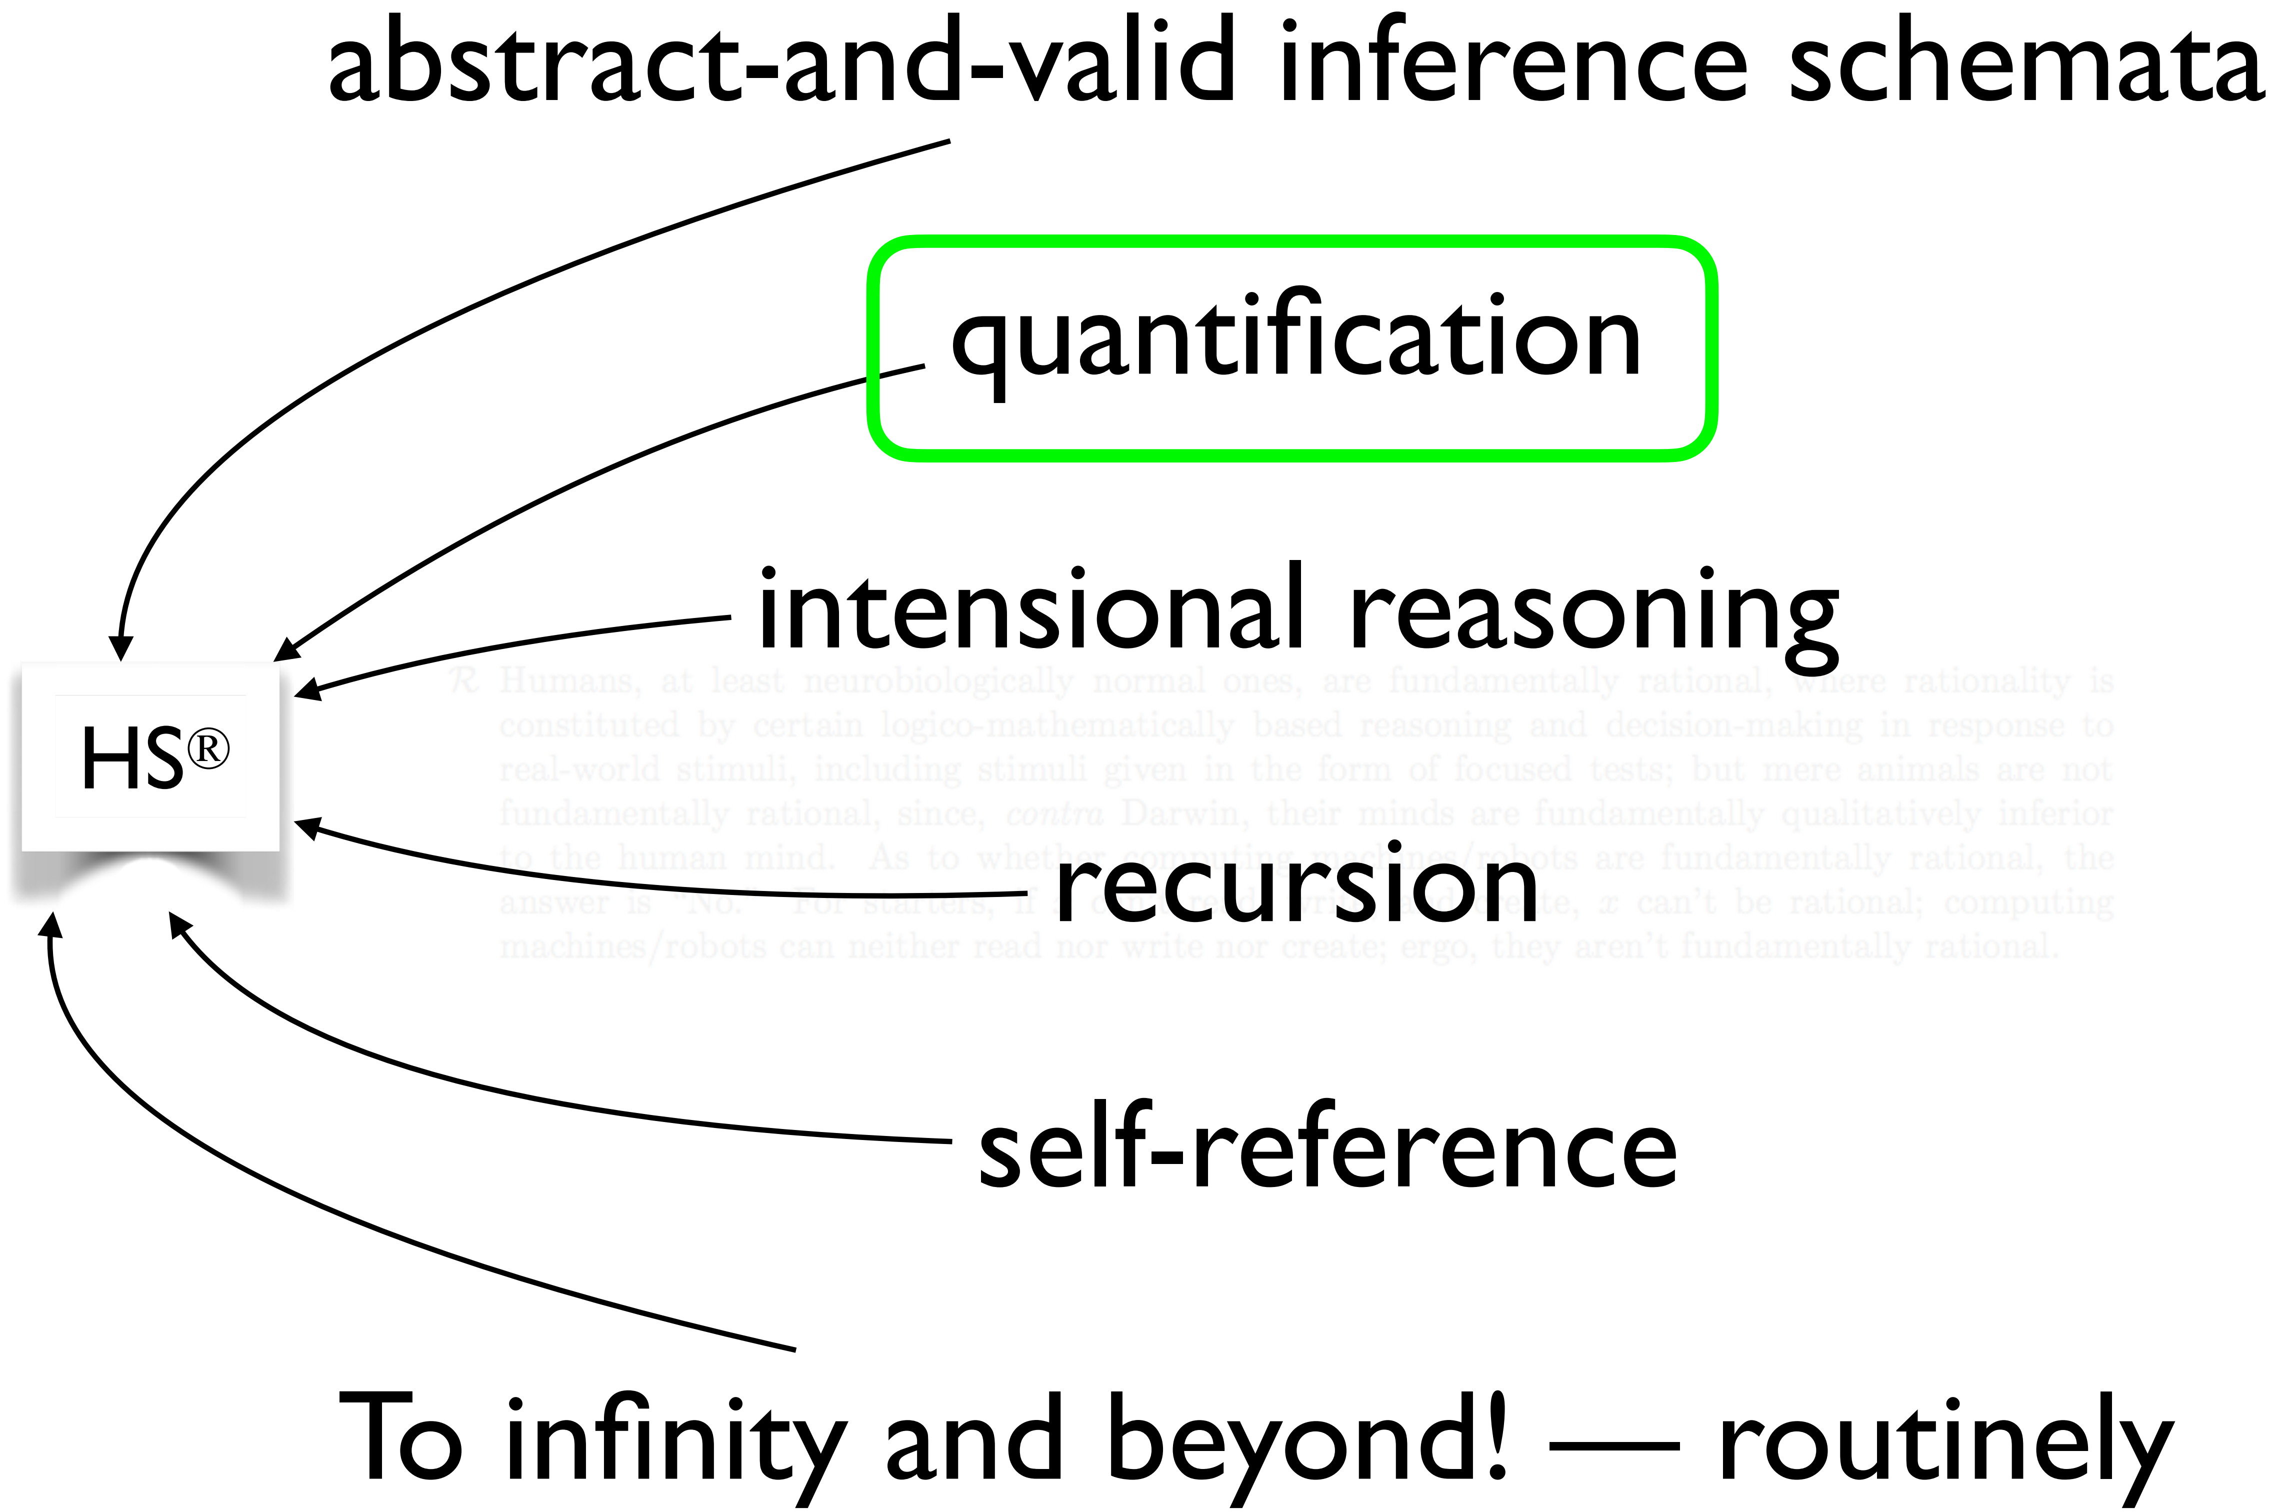
abstract-and-valid inference schemata
quantification
intensional reasoning
HS®
recursion
self-reference
To infinity and beyond! — routinely
$\mathcal{R}$ Humans, at least neurobiologically normal ones, are fundamentally rational, where rationality is constituted by certain logico-mathematically based reasoning and decision-making in response to real-world stimuli, including stimuli given in the form of focused tests; but mere animals are not fundamentally rational, since, *contra* Darwin, their minds are fundamentally qualitatively inferior to the human mind. As to whether computing machines/robots are fundamentally rational, the answer is “No.” For starters, if x can't read, write, and create, x can't be rational; computing machines/robots can neither read nor write nor create; ergo, they aren't fundamentally rational.

And now the whirlwind history …

2026

2026

$\mathcal{DCEC}^*$

Syntax

$$S ::= \text{Object} \mid \text{Agent} \mid \text{Self} \sqsubset \text{Agent} \mid \text{ActionType} \mid \text{Action} \sqsubseteq \text{Event} \mid \text{Moment} \mid \text{Boolean} \mid \text{Fluent} \mid \text{Numeric}$$

$$f ::= \begin{array}{l} \mathit{action} : \text{Agent} \times \text{ActionType} \rightarrow \text{Action} \\ \mathit{initially} : \text{Fluent} \rightarrow \text{Boolean} \\ \mathit{holds} : \text{Fluent} \times \text{Moment} \rightarrow \text{Boolean} \\ \mathit{happens} : \text{Event} \times \text{Moment} \rightarrow \text{Boolean} \\ \mathit{clipped} : \text{Moment} \times \text{Fluent} \times \text{Moment} \rightarrow \mathit{Boolean} \\ \mathit{initiates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Boolean} \\ \mathit{terminates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Boolean} \\ \mathit{prior} : \text{Moment} \times \text{Moment} \rightarrow \text{Boolean} \\ \mathit{interval} : \text{Moment} \times \text{Boolean} \\ * : \text{Agent} \rightarrow \text{Self} \\ \mathit{payoff} : \text{Agent} \times \text{ActionType} \times \text{Moment} \rightarrow \text{Numeric} \end{array}$$

$$t ::= x : S \mid c : S \mid f(t_1, \ldots, t_n)$$

$$\phi ::= \begin{array}{l} t : \text{Boolean} \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \\ \mathbf{P}(a,t,\phi) \mid \mathbf{K}(a,t,\phi) \mid \mathbf{C}(t,\phi) \mid \mathbf{S}(a,b,t,\phi) \mid \mathbf{S}(a,t,\phi) \\ \mathbf{B}(a,t,\phi) \mid \mathbf{D}(a,t,\mathit{holds}(f,t')) \mid \mathbf{I}(a,t,\mathit{happens}(\mathit{action}(a^*,\alpha),t')) \\ \mathbf{O}(a,t,\phi,\mathit{happens}(\mathit{action}(a^*,\alpha),t')) \end{array}$$

Rules of Inference

$$\frac{}{\mathbf{C}(t,\mathbf{P}(a,t,\phi) \rightarrow \mathbf{K}(a,t,\phi))}\;[R_1] \qquad \frac{}{\mathbf{C}(t,\mathbf{K}(a,t,\phi) \rightarrow \mathbf{B}(a,t,\phi))}\;[R_2]$$

$$\frac{\mathbf{C}(t,\phi)\;\; t \leq t_1 \ldots t \leq t_n}{\mathbf{K}(a_1,t_1,\ldots \mathbf{K}(a_n,t_n,\phi)\ldots)}\;[R_3] \qquad \frac{\mathbf{K}(a,t,\phi)}{\phi}\;[R_4]$$

$$\frac{}{\mathbf{C}(t,\mathbf{K}(a,t_1,\phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{K}(a,t_2,\phi_1) \rightarrow \mathbf{K}(a,t_3,\phi_2)}\;[R_5]$$

$$\frac{}{\mathbf{C}(t,\mathbf{B}(a,t_1,\phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{B}(a,t_2,\phi_1) \rightarrow \mathbf{B}(a,t_3,\phi_2)}\;[R_6]$$

$$\frac{}{\mathbf{C}(t,\mathbf{C}(t_1,\phi_1 \rightarrow \phi_2)) \rightarrow \mathbf{C}(t_2,\phi_1) \rightarrow \mathbf{C}(t_3,\phi_2)}\;[R_7]$$

$$\frac{}{\mathbf{C}(t,\forall x.\, \phi \rightarrow \phi[x \mapsto t])}\;[R_8] \qquad \frac{}{\mathbf{C}(t,\phi_1 \leftrightarrow \phi_2 \rightarrow \neg\phi_2 \rightarrow \neg\phi_1)}\;[R_9]$$

$$\frac{}{\mathbf{C}(t,[\phi_1 \wedge \ldots \wedge \phi_n \rightarrow \phi] \rightarrow [\phi_1 \rightarrow \ldots \rightarrow \phi_n \rightarrow \psi])}\;[R_{10}]$$

$$\frac{\mathbf{B}(a,t,\phi)\;\; \phi \rightarrow \psi}{\mathbf{B}(a,t,\psi)}\;[R_{11a}] \qquad \frac{\mathbf{B}(a,t,\phi)\;\; \mathbf{B}(a,t,\psi)}{\mathbf{B}(a,t,\psi \wedge \phi)}\;[R_{11b}]$$

$$\frac{\mathbf{S}(s,h,t,\phi)}{\mathbf{B}(h,t,\mathbf{B}(s,t,\phi))}\;[R_{12}]$$

$$\frac{\mathbf{I}(a,t,\mathit{happens}(\mathit{action}(a^*,\alpha),t'))}{\mathbf{P}(a,t,\mathit{happens}(\mathit{action}(a^*,\alpha),t))}\;[R_{13}]$$

$$\frac{\begin{array}{l}\mathbf{B}(a,t,\phi)\;\; \mathbf{B}(a,t,\mathbf{O}(a^*,t,\phi,\mathit{happens}(\mathit{action}(a^*,\alpha),t'))) \\ \mathbf{O}(a,t,\phi,\mathit{happens}(\mathit{action}(a^*,\alpha),t'))\end{array}}{\mathbf{K}(a,t,\mathbf{I}(a^*,t,\mathit{happens}(\mathit{action}(a^*,\alpha),t')))}\;[R_{14}]$$

$$\frac{\phi \leftrightarrow \psi}{\mathbf{O}(a,t,\phi,\gamma) \leftrightarrow \mathbf{O}(a,t,\psi,\gamma)}\;[R_{15}]$$

2026

Intro to Formal Logic (With AI) @ RPI

2026


2026


2026

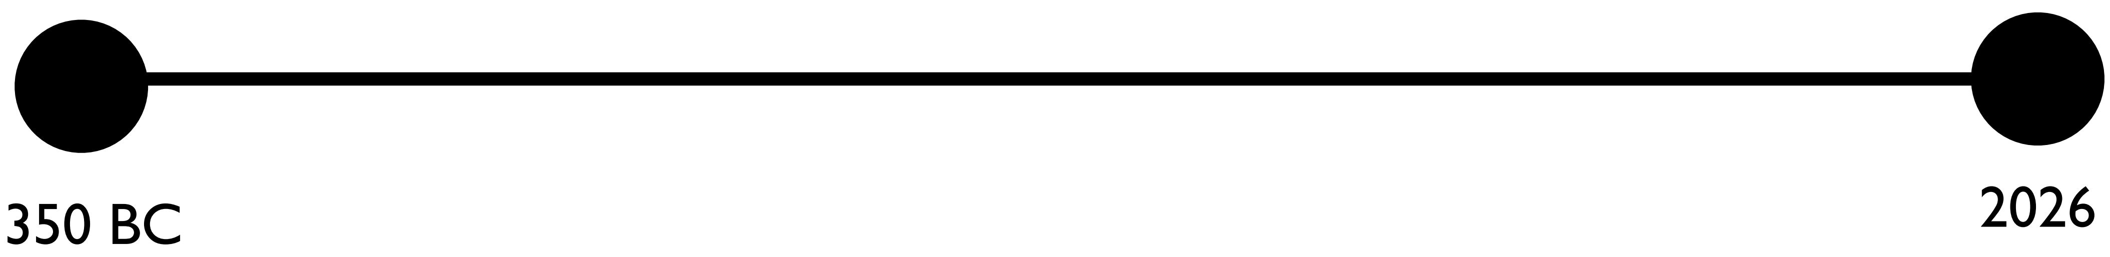
350 BC
2026

350 BC

2026

Euclid

Euclidean "Magic"

Theorem: There are infinitely many primes.

Proof: We take an indirect route. Let $\Pi = p_1 = 2, p_2 = 3, p_3 = 5, \ldots, p_k$ be a finite, exhaustive consecutive sequence of prime numbers. Next, let $\mathbf{M}_\Pi$ be $p_1 \times p_2 \times \cdots \times p_k$, and set $\mathbf{M}'_\Pi$ to $\mathbf{M}_\Pi + 1$. Either $\mathbf{M}'_\Pi$ is prime, or not; we thus have two (exhaustive) cases to consider.

C1 Suppose $\mathbf{M}'_\Pi$ is prime. In this case we immediately have a prime number beyond any in Π — contradiction!

C2 Suppose on the other hand that $\mathbf{M}'_\Pi$ is *not* prime. Then some prime p divides $\mathbf{M}'_\Pi$. (Why?) Now, p itself is either in Π, or not; we hence have two sub-cases. Supposing that p is in Π entails that p divides $\mathbf{M}_\Pi$. But we are operating under the supposition that p divides $\mathbf{M}'_\Pi$ as well. This implies that p divides 1, which is absurd (a contradiction). Hence the prime p is outside Π.

Hence for *any* such list Π, there is a prime outside the list. That is, there are infinitely many primes. **QED**

Euclidean "Magic"

Theorem: There are infinitely many primes.

Proof: We take an indirect route. Let $\Pi = p_1 = 2, p_2 = 3, p_3 = 5, \ldots, p_k$ be a finite, exhaustive consecutive sequence of prime numbers. Next, let $\mathbf{M}_\Pi$ be $p_1 \times p_2 \times \cdots \times p_k$, and set $\mathbf{M}'_\Pi$ to $\mathbf{M}_\Pi + 1$. Either $\mathbf{M}'_\Pi$ is prime, or not; we thus have two (exhaustive) cases to consider.

C1 Suppose $\mathbf{M}'_\Pi$ is prime. In this case we immediately have a prime number beyond any in Π — contradiction!

C2 Suppose on the other hand that $\mathbf{M}'_\Pi$ is *not* prime. Then some prime p divides $\mathbf{M}'_\Pi$. (Why?) Now, p itself is either in Π, or not; we hence have two sub-cases. Supposing that p is in Π entails that p divides $\mathbf{M}_\Pi$. But we are operating under the supposition that p divides $\mathbf{M}'_\Pi$ as well. This implies that p divides 1, which is absurd (a contradiction). Hence the prime p is outside Π.

Hence for *any* such list Π, there is a prime outside the list. That is, there are infinitely many primes. **QED**

Euclidean "Magic"

Theorem: There are infinitely many primes.

Proof: We take an indirect route. Let $\Pi = p_1 = 2, p_2 = 3, p_3 = 5, \ldots, p_k$ be a finite, exhaustive consecutive sequence of prime numbers. Next, let $\mathbf{M}_\Pi$ be $p_1 \times p_2 \times \cdots \times p_k$, and set $\mathbf{M}'_\Pi$ to $\mathbf{M}_\Pi + 1$. Either $\mathbf{M}'_\Pi$ is prime, or not; we thus have two (exhaustive) cases to consider.

C1 Suppose $\mathbf{M}'_\Pi$ is prime. In this case we immediately have a prime number beyond any in Π — contradiction!

C2 Suppose on the other hand that $\mathbf{M}'_\Pi$ is *not* prime. Then some prime p divides $\mathbf{M}'_\Pi$. (Why?) Now, p itself is either in Π, or not; we hence have two sub-cases. Supposing that p is in Π entails that p divides $\mathbf{M}_\Pi$. But we are operating under the supposition that p divides $\mathbf{M}'_\Pi$ as well. This implies that p divides 1, which is absurd (a contradiction). Hence the prime p is outside Π.

Hence for *any* such list Π, there is a prime outside the list. That is, there are infinitely many primes. **QED**

Euclidean "Magic"

Theorem: There are infinitely many primes.

Proof: We take an indirect route. Let $\Pi = p_1 = 2, p_2 = 3, p_3 = 5, \ldots, p_k$ be a finite, exhaustive consecutive sequence of prime numbers. Next, let $\mathbf{M}_\Pi$ be $p_1 \times p_2 \times \cdots \times p_k$, and set $\mathbf{M}'_\Pi$ to $\mathbf{M}_\Pi + 1$. Either $\mathbf{M}'_\Pi$ is prime, or not; we thus have two (exhaustive) cases to consider.

C1 Suppose $\mathbf{M}'_\Pi$ is prime. In this case we immediately have a prime number beyond any in Π — contradiction!

C2 Suppose on the other hand that $\mathbf{M}'_\Pi$ is *not* prime. Then some prime p divides $\mathbf{M}'_\Pi$. (Why?) Now, p itself is either in Π, or not; we hence have two sub-cases. Supposing that p is in Π entails that p divides $\mathbf{M}_\Pi$. But we are operating under the supposition that p divides $\mathbf{M}'_\Pi$ as well. This implies that p divides 1, which is absurd (a contradiction). Hence the prime p is outside Π.

Hence for *any* such list Π, there is a prime outside the list. That is, there are infinitely many primes. **QED**

Euclidean "Magic"

Theorem: There are infinitely many primes.

Proof: We take an indirect route. Let $\Pi = p_1 = 2, p_2 = 3, p_3 = 5, \ldots, p_k$ be a finite, exhaustive consecutive sequence of prime numbers. Next, let $\mathbf{M}_\Pi$ be $p_1 \times p_2 \times \cdots \times p_k$, and set $\mathbf{M}'_\Pi$ to $\mathbf{M}_\Pi + 1$. Either $\mathbf{M}'_\Pi$ is prime, or not; we thus have two (exhaustive) cases to consider.

C1 Suppose $\mathbf{M}'_\Pi$ is prime. In this case we immediately have a prime number beyond any in Π — contradiction!

C2 Suppose on the other hand that $\mathbf{M}'_\Pi$ is *not* prime. Then some prime p divides $\mathbf{M}'_\Pi$. (Why?) Now, p itself is either in Π, or not; we hence have two sub-cases. Supposing that p is in Π entails that p divides $\mathbf{M}_\Pi$. But we are operating under the supposition that p divides $\mathbf{M}'_\Pi$ as well. This implies that p divides 1, which is absurd (a contradiction). Hence the prime p is outside Π.

Hence for *any* such list Π, there is a prime outside the list. That is, there are infinitely many primes. **QED**

The Fundamental Theorem of Arithmetic

350 BC

2026

Euclid


350 BC
300 BC
2026
Euclid

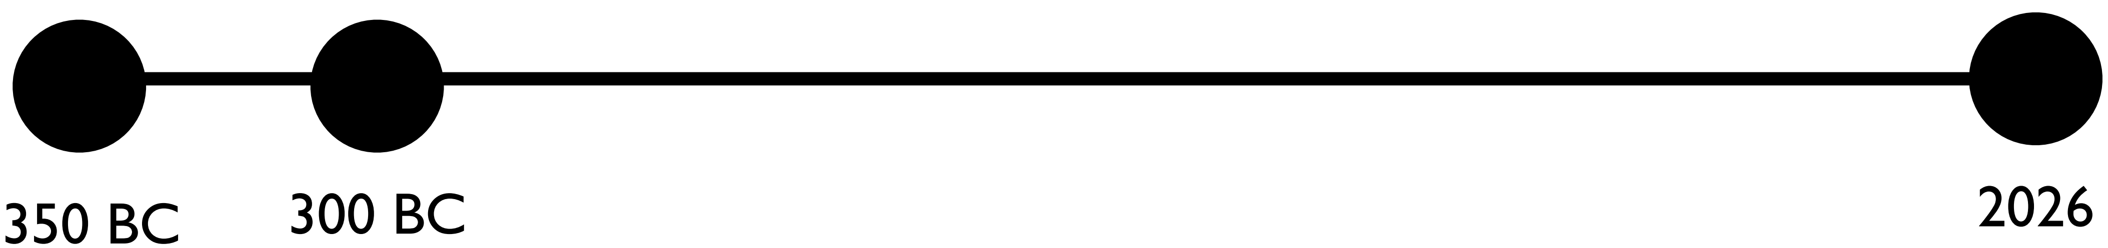

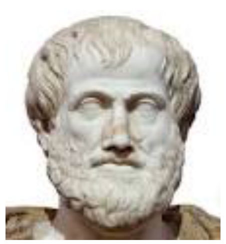

I don’t believe in magic! Why exactly is that so convincing? What exactly is he doing?!?

Euclid

350 BC 300 BC 2026

Euclid

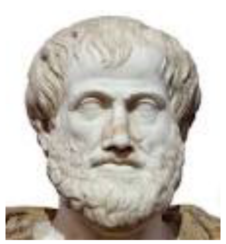

Organon

I don’t believe in magic! Why exactly is that so convincing? What exactly is he doing?!?

He’s using syllogisms!

E.g.,

All As are Bs.
All Bs are Cs.

All As are Cs.

350 BC | 300 BC | 2026

Euclid

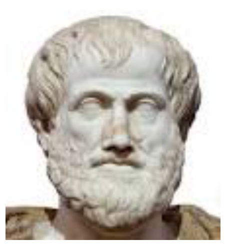

Organon

I don’t believe in magic! Why exactly is that so convincing? What exactly is he doing?!?

He’s using syllogisms!

E.g.,

All As are Bs.
All Bs are Cs.

All As are Cs.

350 BC — 300 BC — 2026

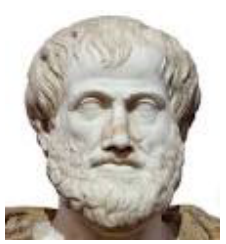

I don’t believe in magic! Why exactly is that so convincing? What exactly is he doing?!?

Euclid

Organon

Balderdash!

He's using syllogisms!

E.g.,

All As are Bs.
All Bs are Cs.

All As are Cs.

350 BC 300 BC 2026

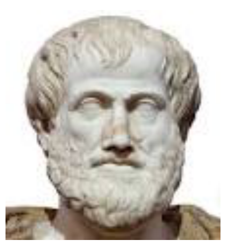

I don't believe in magic! Why exactly is that so convincing? What exactly is he doing?!?

Euclid

Organon

Balderdash!

350 BC 300 BC 2026

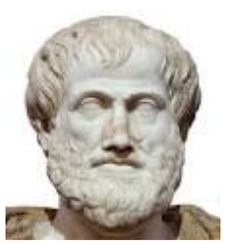

I don’t believe in magic! Why exactly is that so convincing? What exactly is he doing?!?

Euclid *Organon*

350 BC

Euclid

300 BC

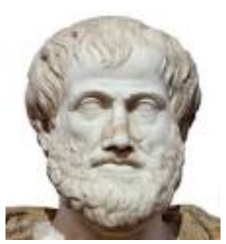

Organon

Intro to (Formal) Logic (& AI) @ RPI

350 BC 300 BC 1666

Euclid

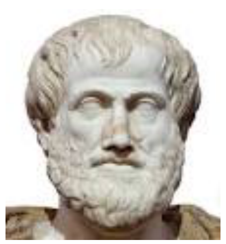

Organon

350 BC

Euclid

300 BC

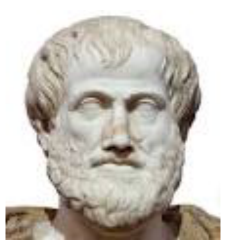

Organon

1666

Leibniz


350 BC
300 BC
1666
Euclid
Organon
Leibniz
$\int$

"Universal Computational Logic"

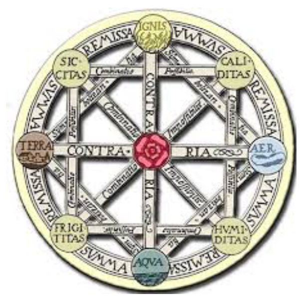

350 BC

Euclid

300 BC

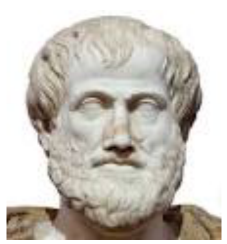

Organon

1666

Leibniz

$\int$

"Universal Computational Logic"

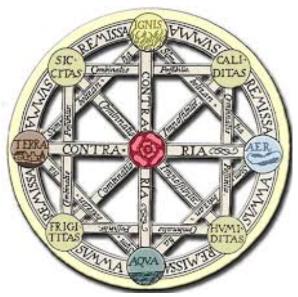

350 BC

Euclid

300 BC

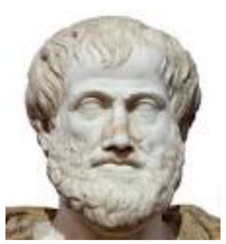

Organon

1666

Leibniz

$\int$

1854

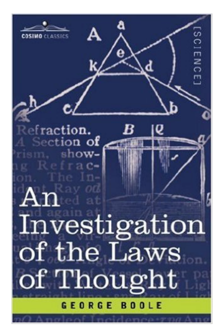

"Universal Computational Logic"

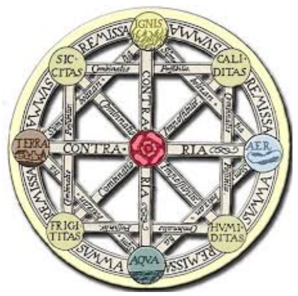

350 BC

300 BC

1666

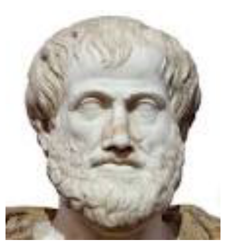

Euclid

Organon

Leibniz

$\int$

"Universal Computational Logic"

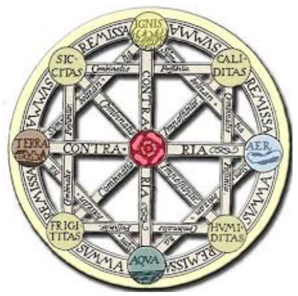

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

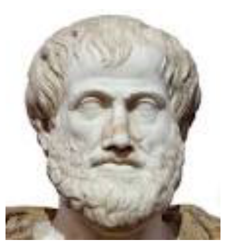

Organon

1666

Leibniz

$\int$

1956

Simon

Intro to (Formal) Logic (& AI) @ RPI

“Astonishing” Logic Theorist Proof @ Dawn of AI

“Astonishing” Logic Theorist Proof @ Dawn of AI

1	$(\phi \vee \phi) \to \phi$	axiom
2	$(\neg\phi \vee \neg\phi) \to \neg\phi$	substitution
3	$(\phi \to \neg\phi) \to \neg\phi$	a “replacement rule”
4	$(A \to \neg A) \to \neg A$	substitution

“Astonishing” Logic Theorist Proof @ Dawn of AI

1	$(\phi \vee \phi) \rightarrow \phi$	axiom
2	$(\neg\phi \vee \neg\phi) \rightarrow \neg\phi$	substitution
3	$(\phi \rightarrow \neg\phi) \rightarrow \neg\phi$	a “replacement rule”
4	$(A \rightarrow \neg A) \rightarrow \neg A$	substitution

At dawn of AI: 10 seconds.

“Astonishing” Logic Theorist Proof @ Dawn of AI

1	$(\phi \lor \phi) \to \phi$	axiom
2	$(\neg\phi \lor \neg\phi) \to \neg\phi$	substitution
3	$(\phi \to \neg\phi) \to \neg\phi$	a “replacement rule”
4	$(A \to \neg A) \to \neg A$	substitution

At dawn of AI: 10 seconds.

AI of today: vanishingly small amount of time (in eg HS®).

“Astonishing” Logic Theorist Proof @ Dawn of AI

1	$(\phi \vee \phi) \rightarrow \phi$	axiom
2	$(\neg\phi \vee \neg\phi) \rightarrow \neg\phi$	substitution
3	$(\phi \rightarrow \neg\phi) \rightarrow \neg\phi$	a “replacement rule”
4	$(A \rightarrow \neg A) \rightarrow \neg A$	substitution

At dawn of AI: 10 seconds.

AI of today: vanishingly small amount of time (in eg HS®).

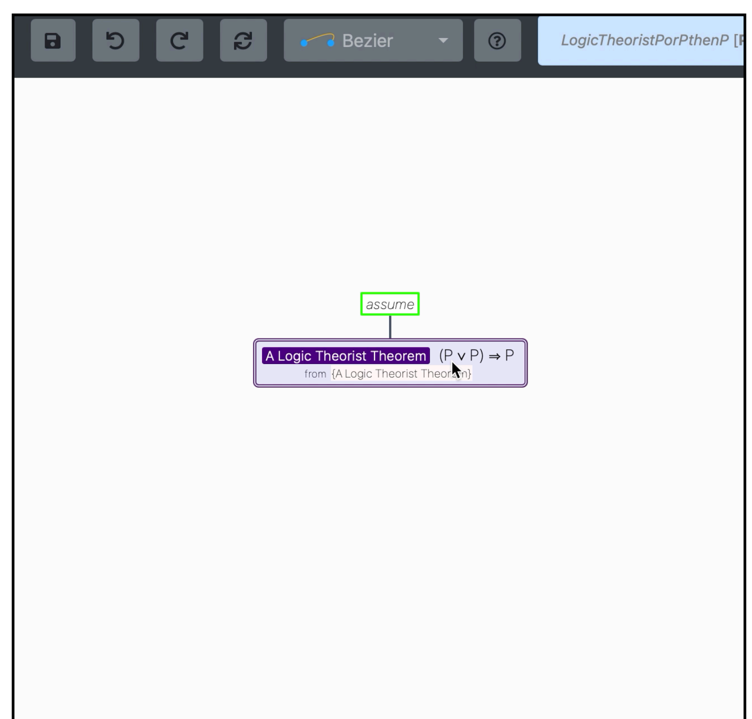

"Astonishing" Logic Theorist Proof @ Dawn of AI

1	$(\phi \vee \phi) \rightarrow \phi$	axiom
2	$(\neg\phi \vee \neg\phi) \rightarrow \neg\phi$	substitution
3	$(\phi \rightarrow \neg\phi) \rightarrow \neg\phi$	a "replacement rule"
4	$(A \rightarrow \neg A) \rightarrow \neg A$	substitution

At dawn of AI: 10 seconds.

AI of today: vanishingly small amount of time (in eg HS®).

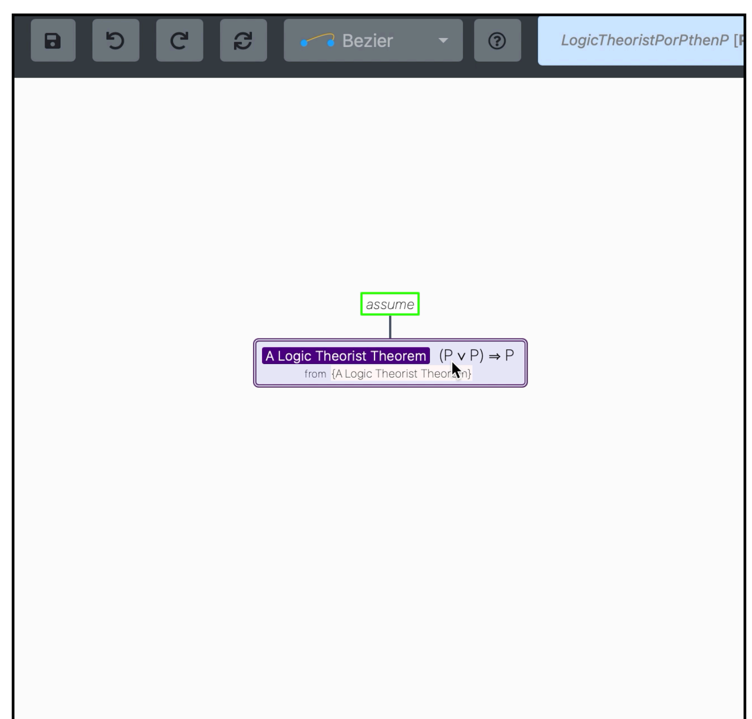

“Universal Computational Logic”

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

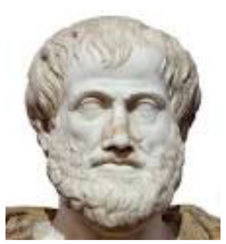

Organon

Leibniz

$\int$

1854

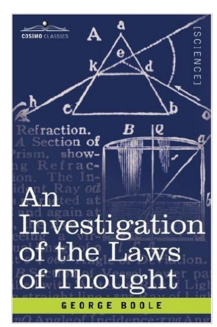

1956

Simon

2024

Intro to (Formal) Logic (With AI) @ RPI

"Universal Computational Logic"

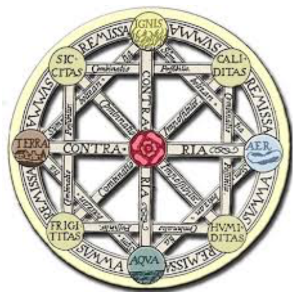

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

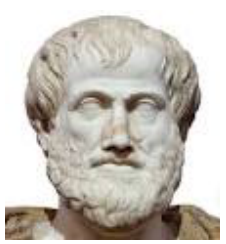

Organon

1666

Leibniz

$\int$

1854

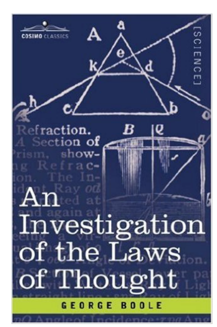

1956

Simon

2024

Intro to (Formal) Logic (With AI) @ RPI

"Universal Computational Logic"

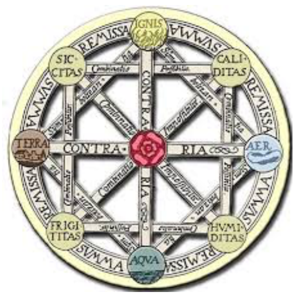

Logic Theorist

(birth of modern logicist AI)

350 BC — 300 BC — 1666 — 1956 — 2024

350 BC

Euclid

300 BC

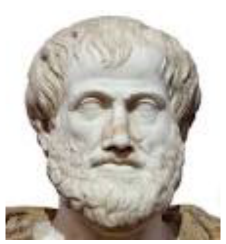

Organon

1666

Leibniz

$\int$

1956

Simon

2024

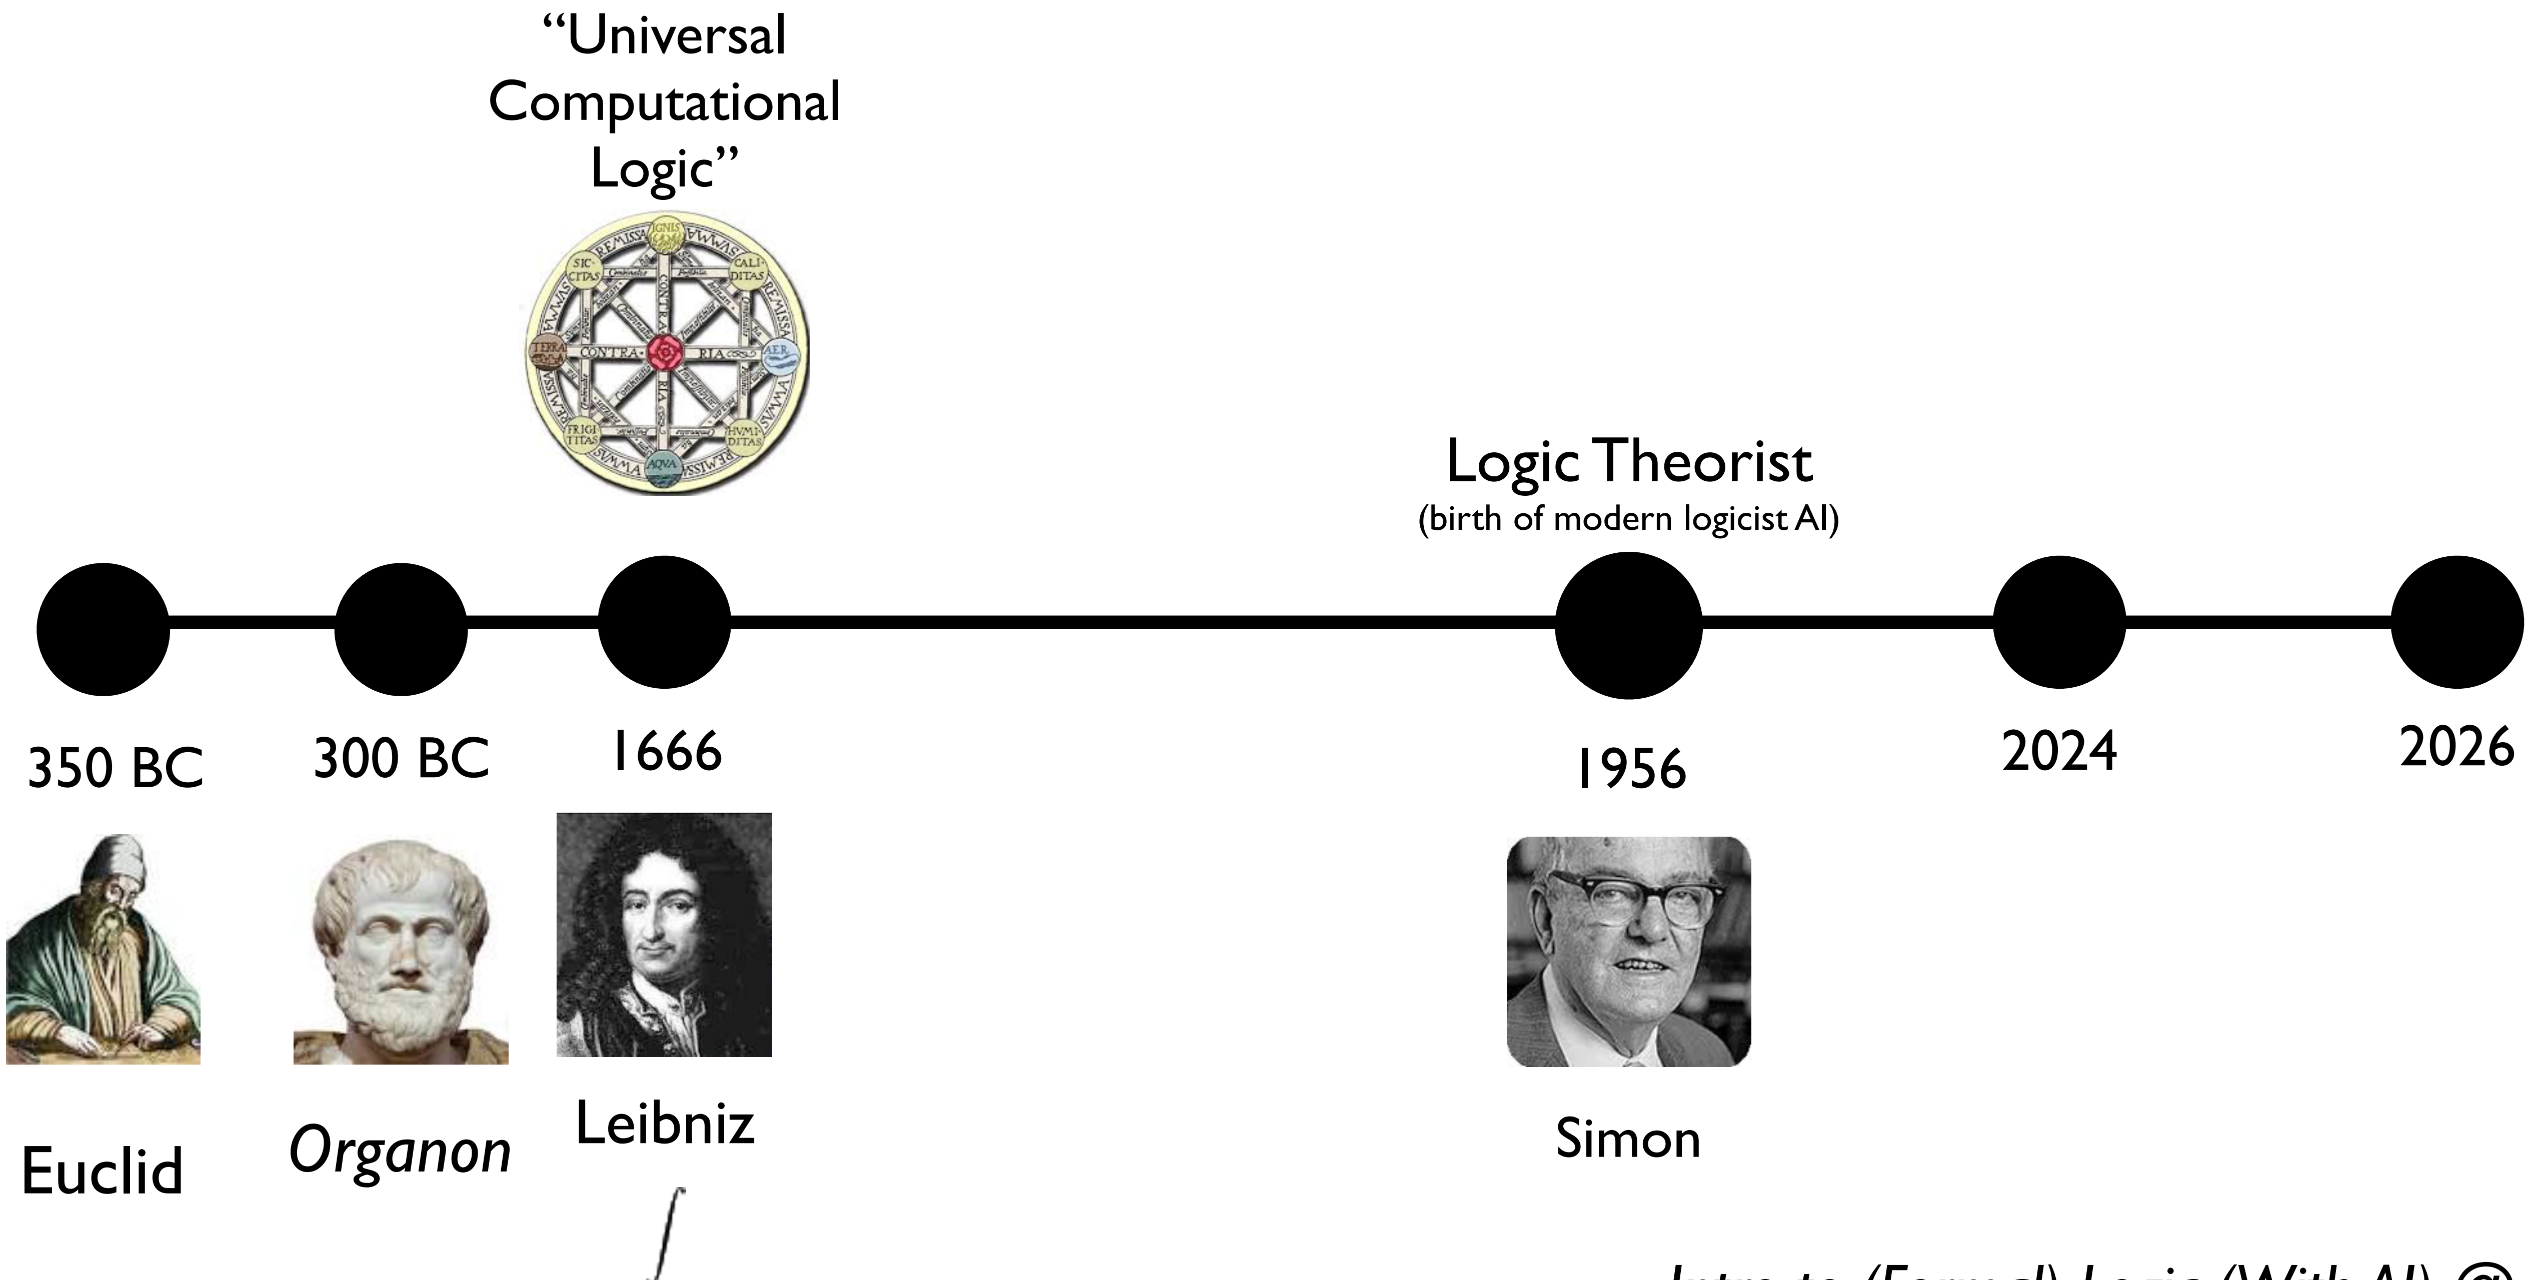
"Universal Computational Logic"
Logic Theorist
(birth of modern logicist AI)
350 BC
300 BC
1666
1956
2024
2026
Euclid
Organon
Leibniz
∫
Simon

"Universal Computational Logic"

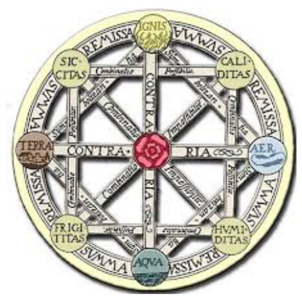

Logic Theorist
(birth of modern logicist AI)

350 BC

300 BC

1666

1956

2024

2026

Euclid

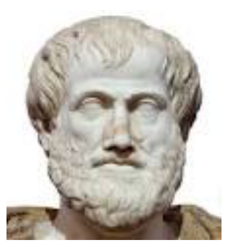

Organon

Leibniz

$\int$

Simon

The Singularity?

Intro to (Formal) Logic (With AI) @ RPI

Entscheidungsproblem

"Universal Computational Logic"

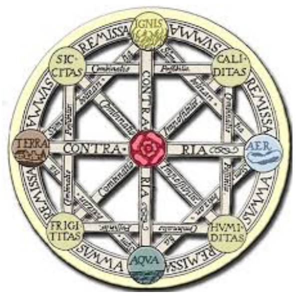

Logic Theorist
(birth of modern logicist AI)

350 BC | 300 BC | 1666 | 1956 | 2024 | 2026

Euclid

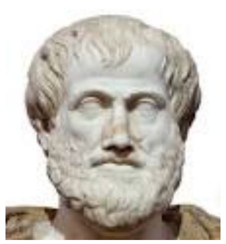

Organon

Leibniz

$\int$

Simon

The Singularity?

Intro to (Formal) Logic (With AI) @ RPI

?

Entscheidungsproblem

"Universal Computational Logic"

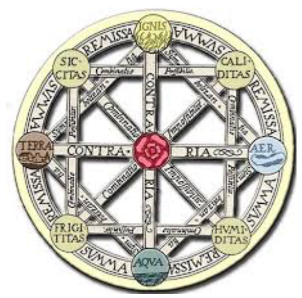

Logic Theorist

(birth of modern logicist AI)

350 BC | 300 BC | 1666 | 1854 | 1956 | 2023 | 2025

Euclid

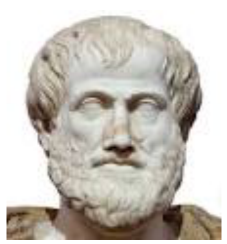

Organon

Leibniz

$\int$

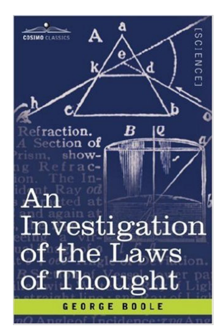

Simon

The Singularity?

Intro to (Formal) Logic (& AI) @ RPI

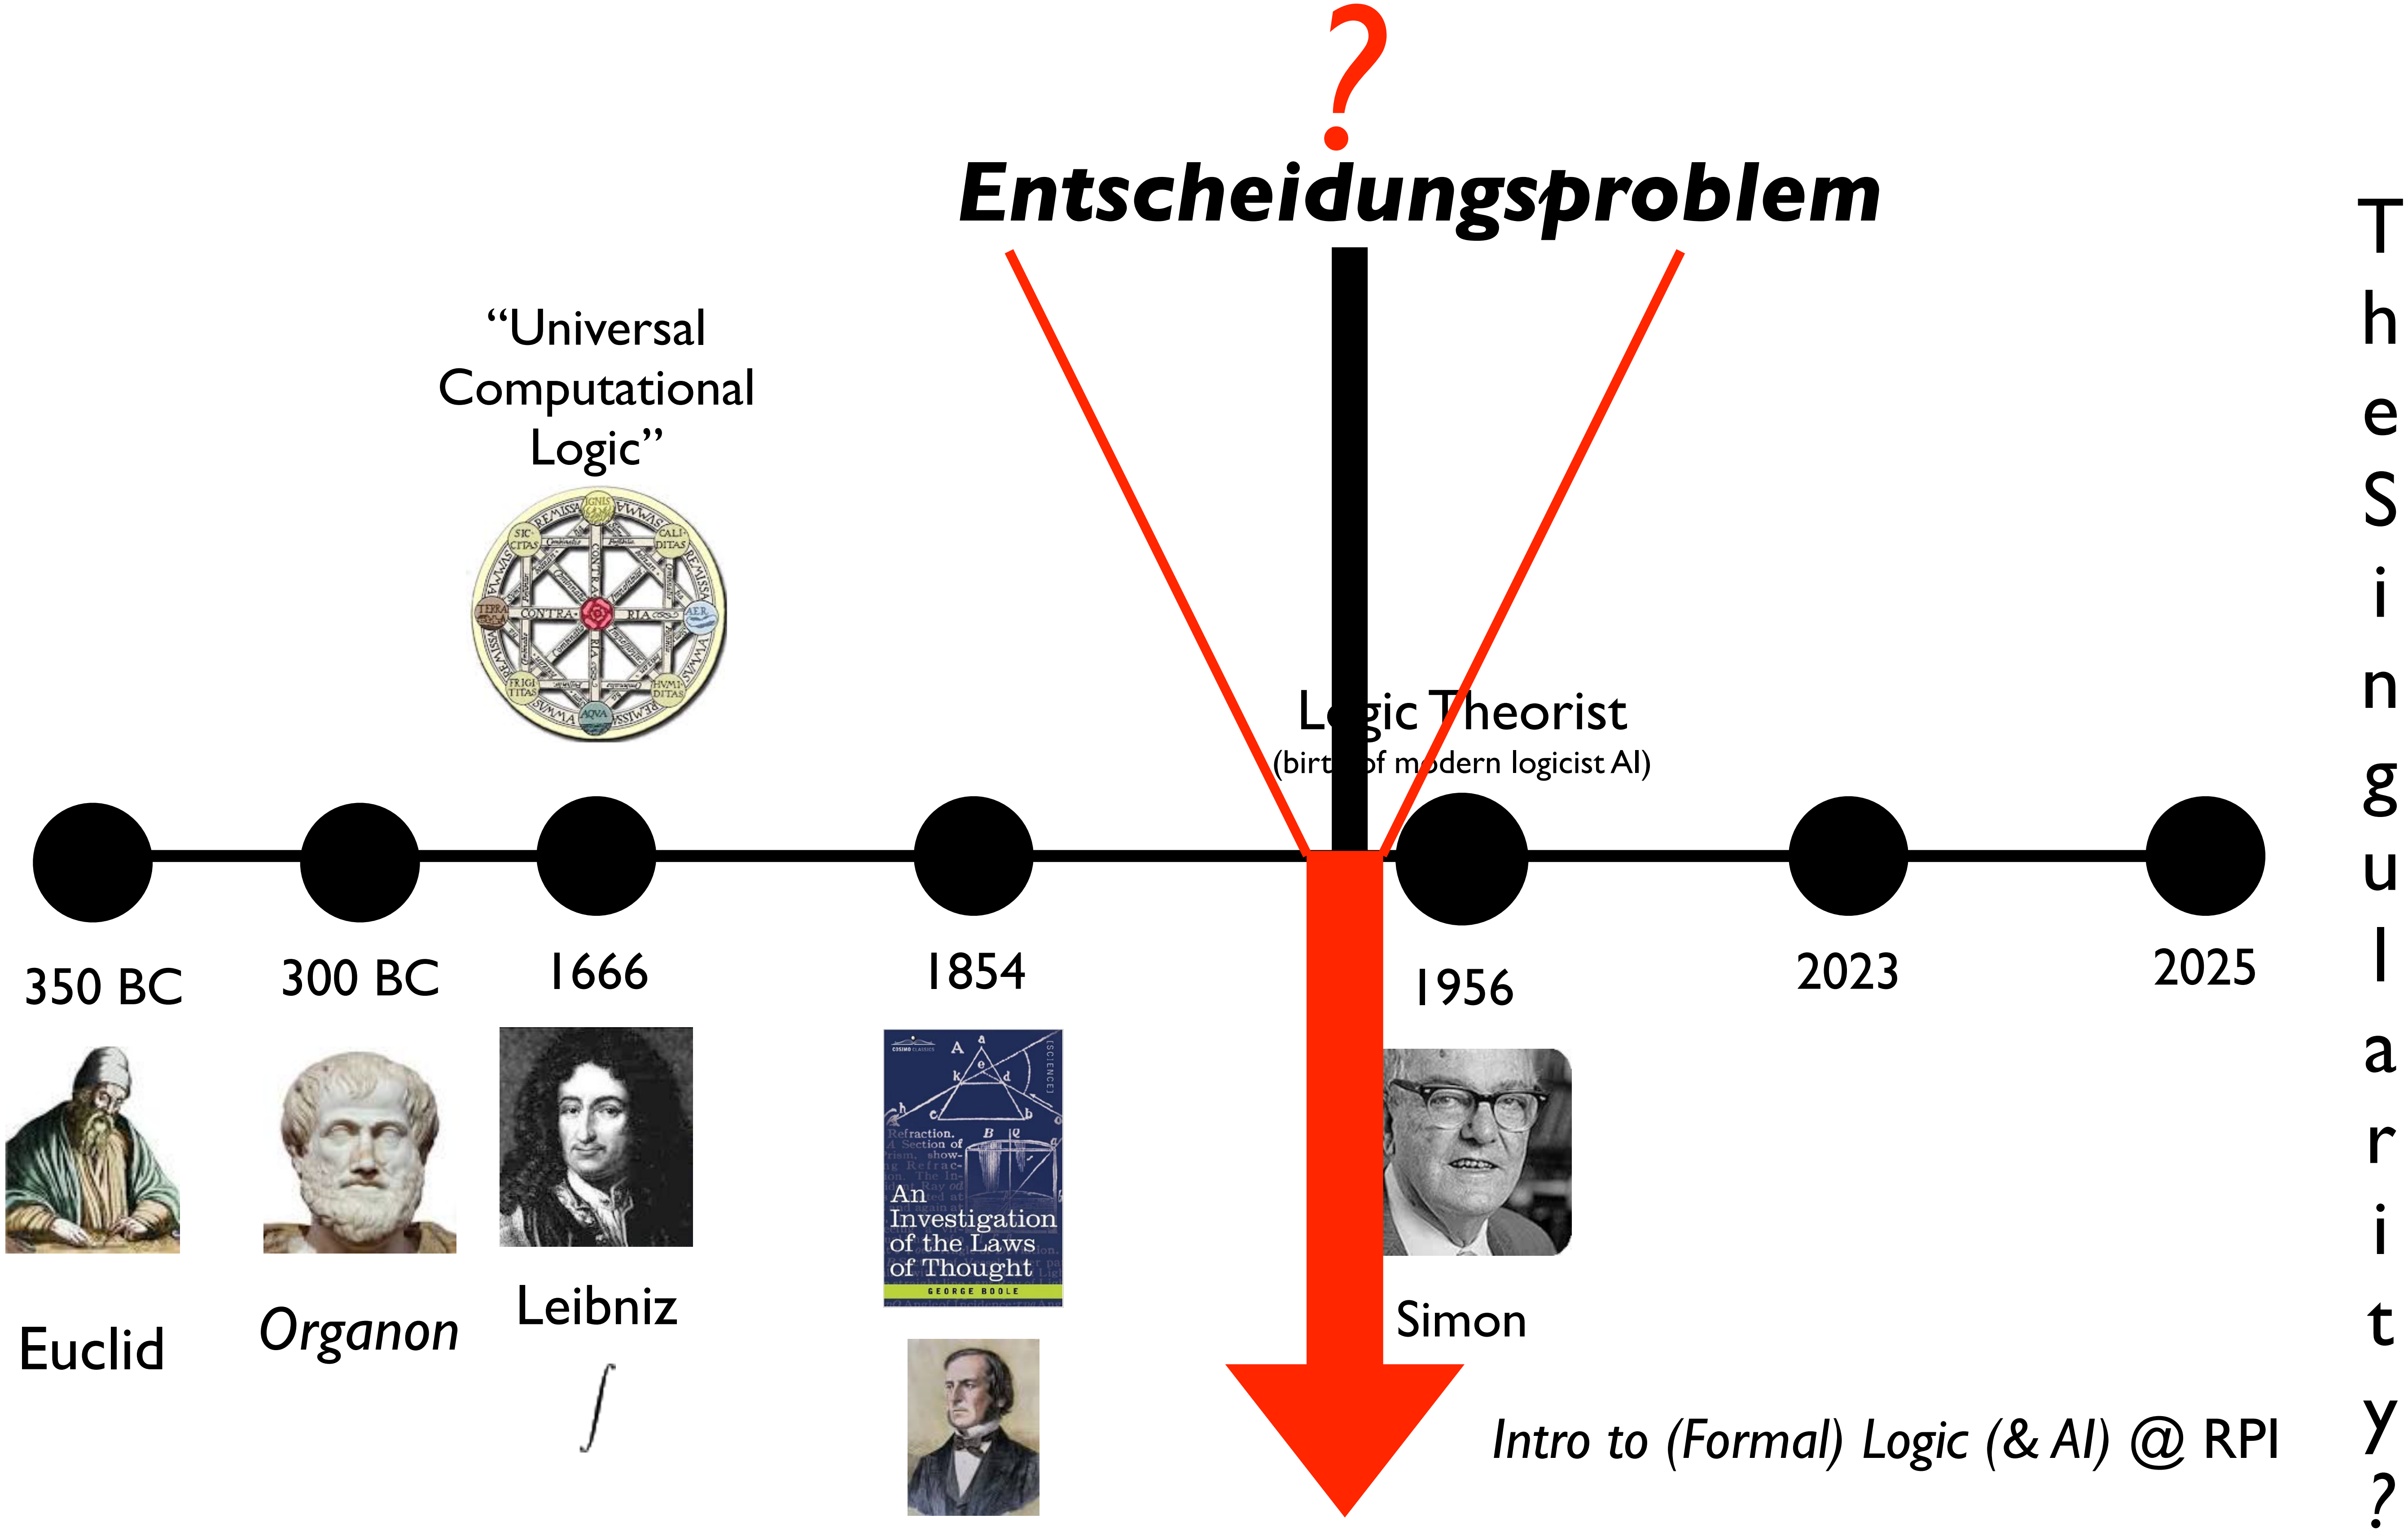
?
Entscheidungsproblem
"Universal
Computational
Logic"
Logic Theorist
(birth of modern logicist AI)
350 BC
300 BC
1666
1854
1956
2023
2025
Euclid
Organon
Leibniz
∫
An Investigation of the Laws of Thought
GEORGE BOOLE
Simon
Intro to (Formal) Logic (& AI) @ RPI
The Singularity?

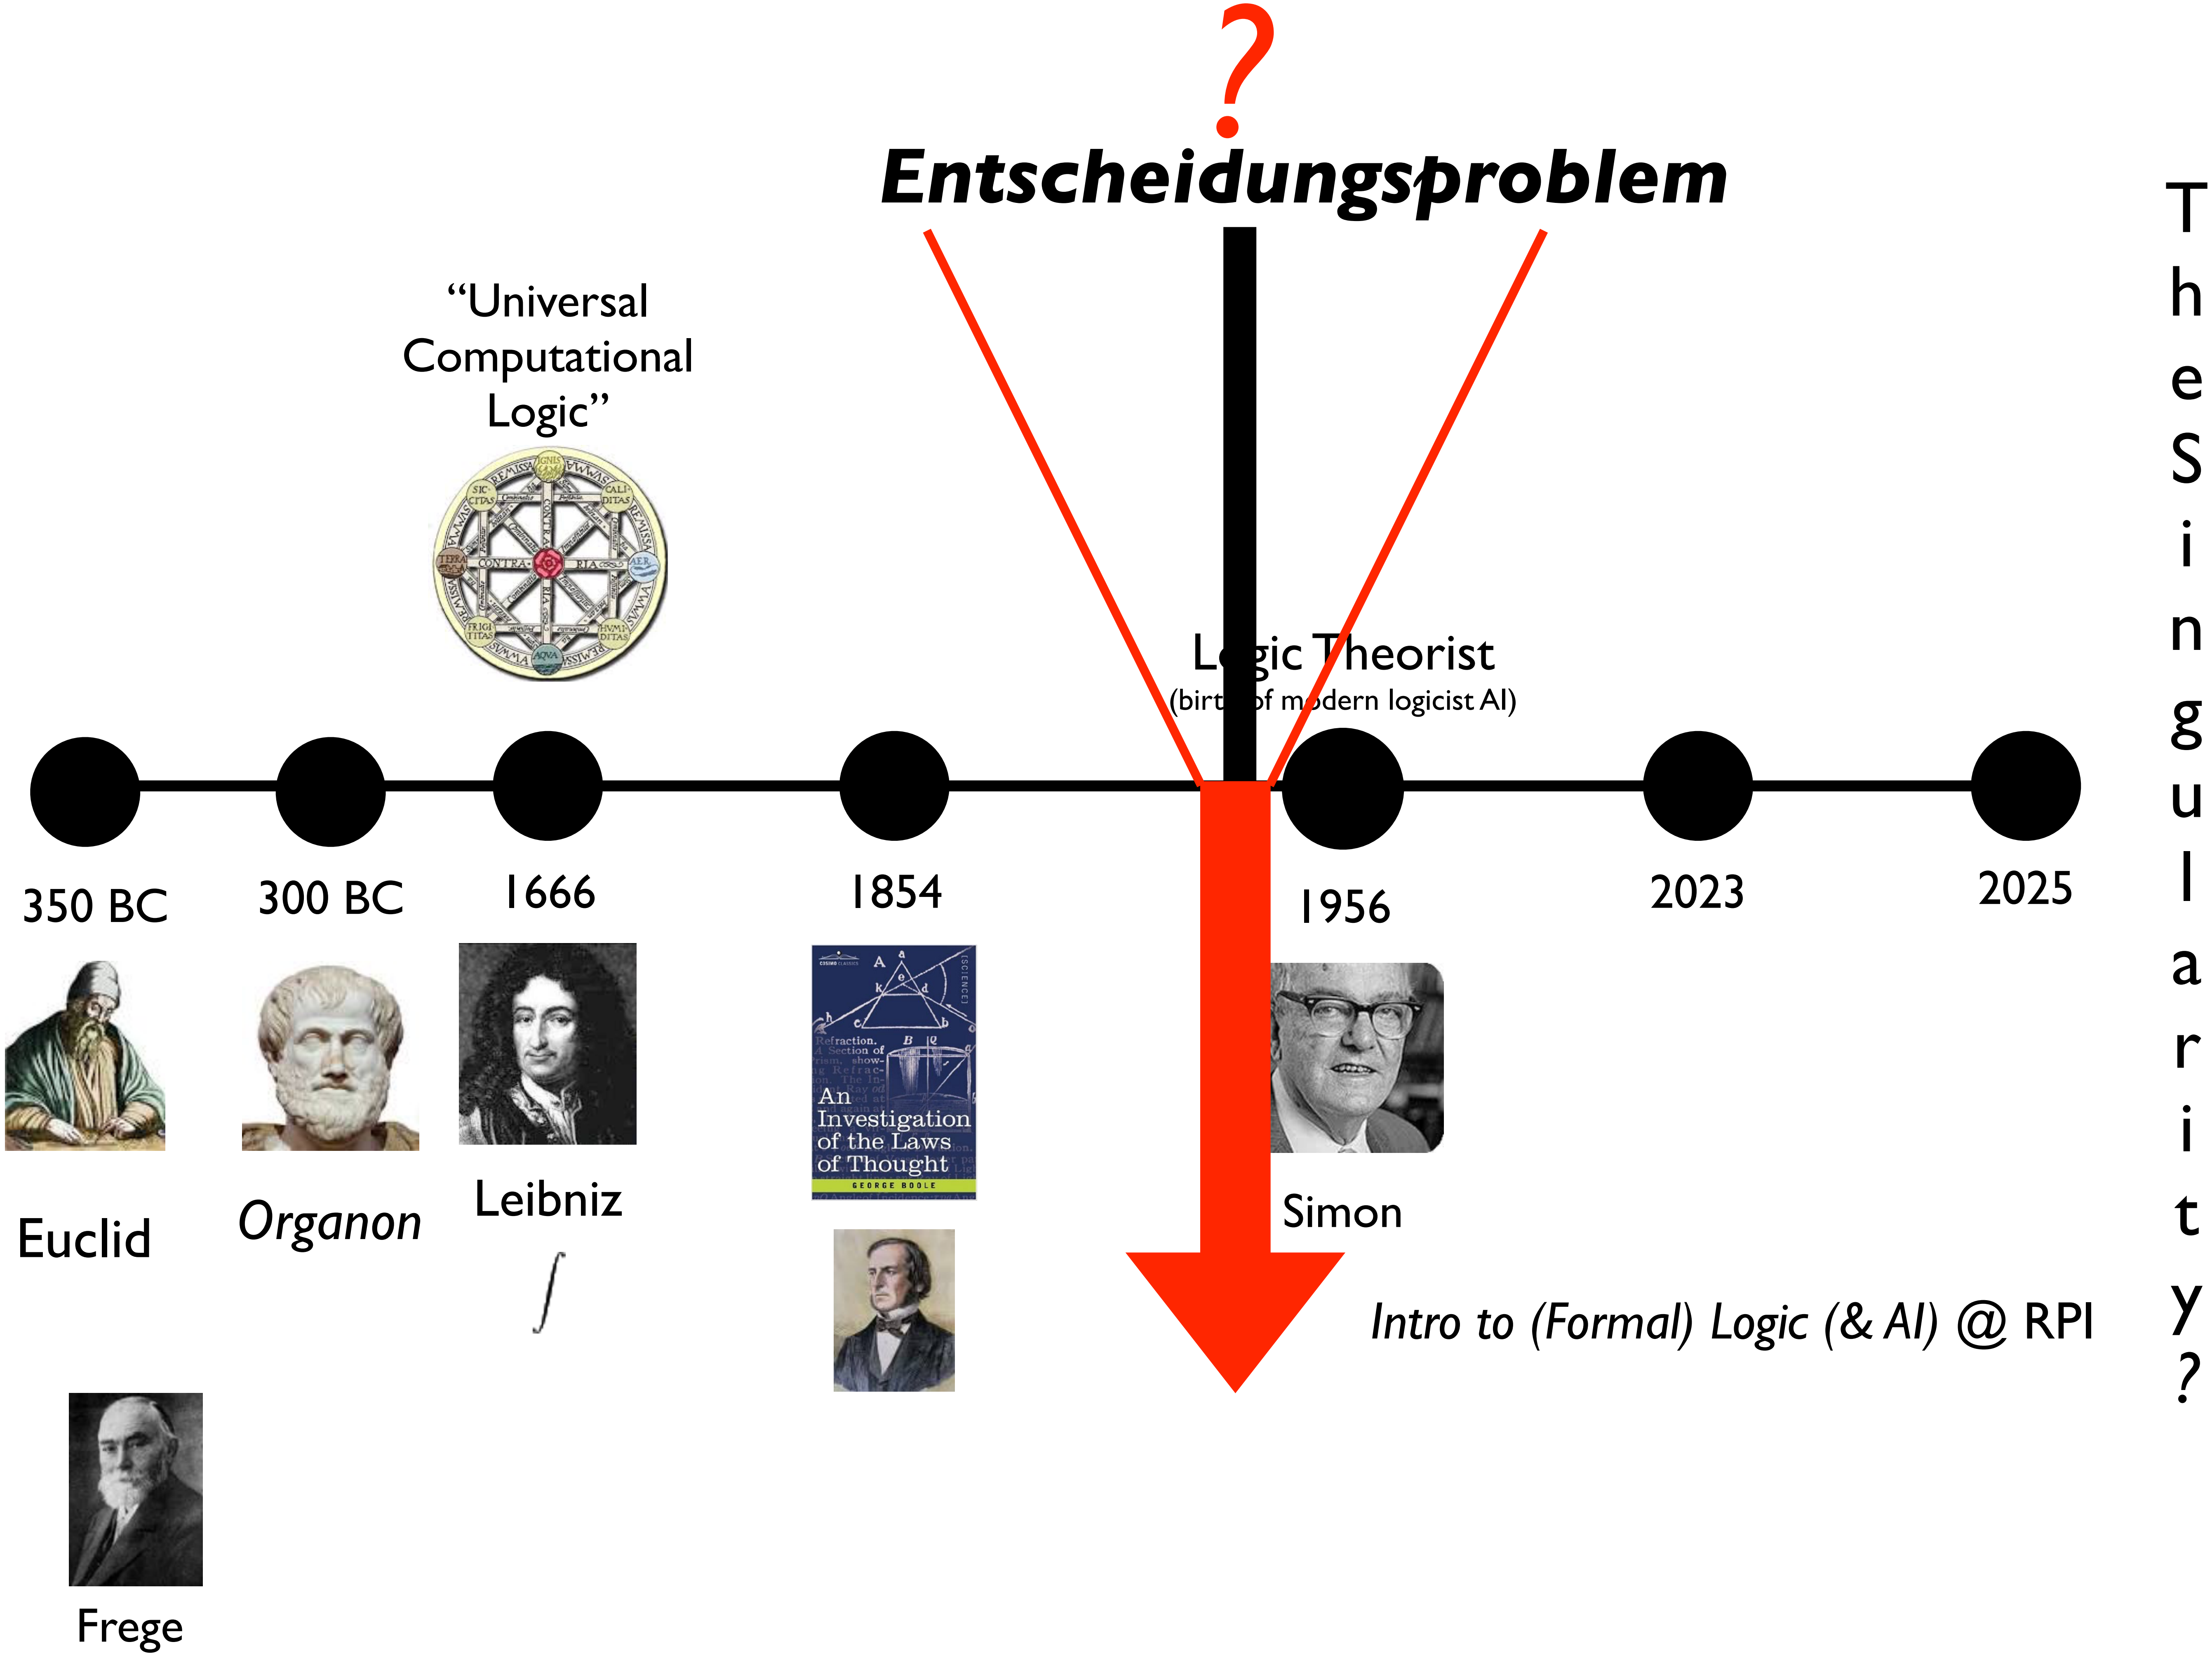
?
Entscheidungsproblem
"Universal Computational Logic"
Logic Theorist
(birth of modern logicist AI)
350 BC
300 BC
1666
1854
1956
2023
2025
Euclid
Organon
Leibniz
∫
An Investigation of the Laws of Thought
GEORGE BOOLE
Simon
Frege
Intro to (Formal) Logic (& AI) @ RPI
The Singularity?

?

Entscheidungsproblem

“Universal Computational Logic”

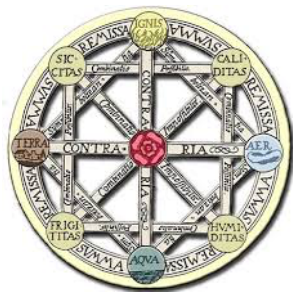

Logic Theorist

(birth of modern logicist AI)

350 BC | 300 BC | 1666 | 1854 | 1956 | 2023 | 2025

Euclid

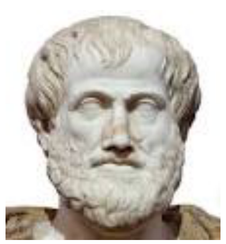

Organon

Leibniz

$\int$

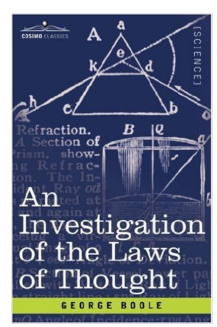

Simon

Intro to (Formal) Logic (& AI) @ RPI

The Singularity?

Frege

Exceeds Leibniz & de-mystifies Euclid: the “compellingness” of these proofs consists in their being, at bottom, formal proofs in first-order logic (FOL).

?

Entscheidungsproblem

The Singularity?

"Universal Computational Logic"

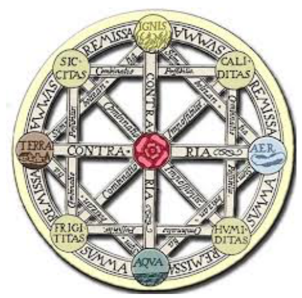

Logic Theorist (birth of modern logicist AI)

350 BC — 300 BC — 1666 — 1854 — 1956 — 2023 — 2025

Euclid

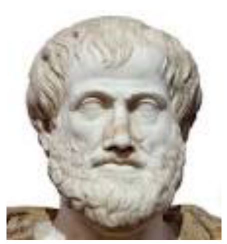

Organon

Leibniz

$\int$

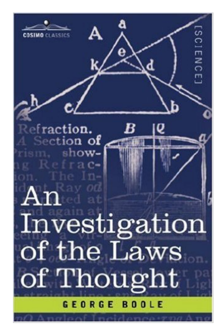

Simon

Intro to (Formal) Logic (& AI) @ RPI

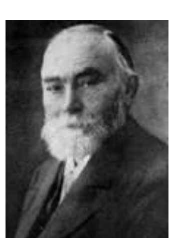

Frege

Exceeds Leibniz & de-mystifies Euclid: the "compellingness" of these proofs consists in their being, at bottom, formal proofs in first-order logic (FOL).

Church

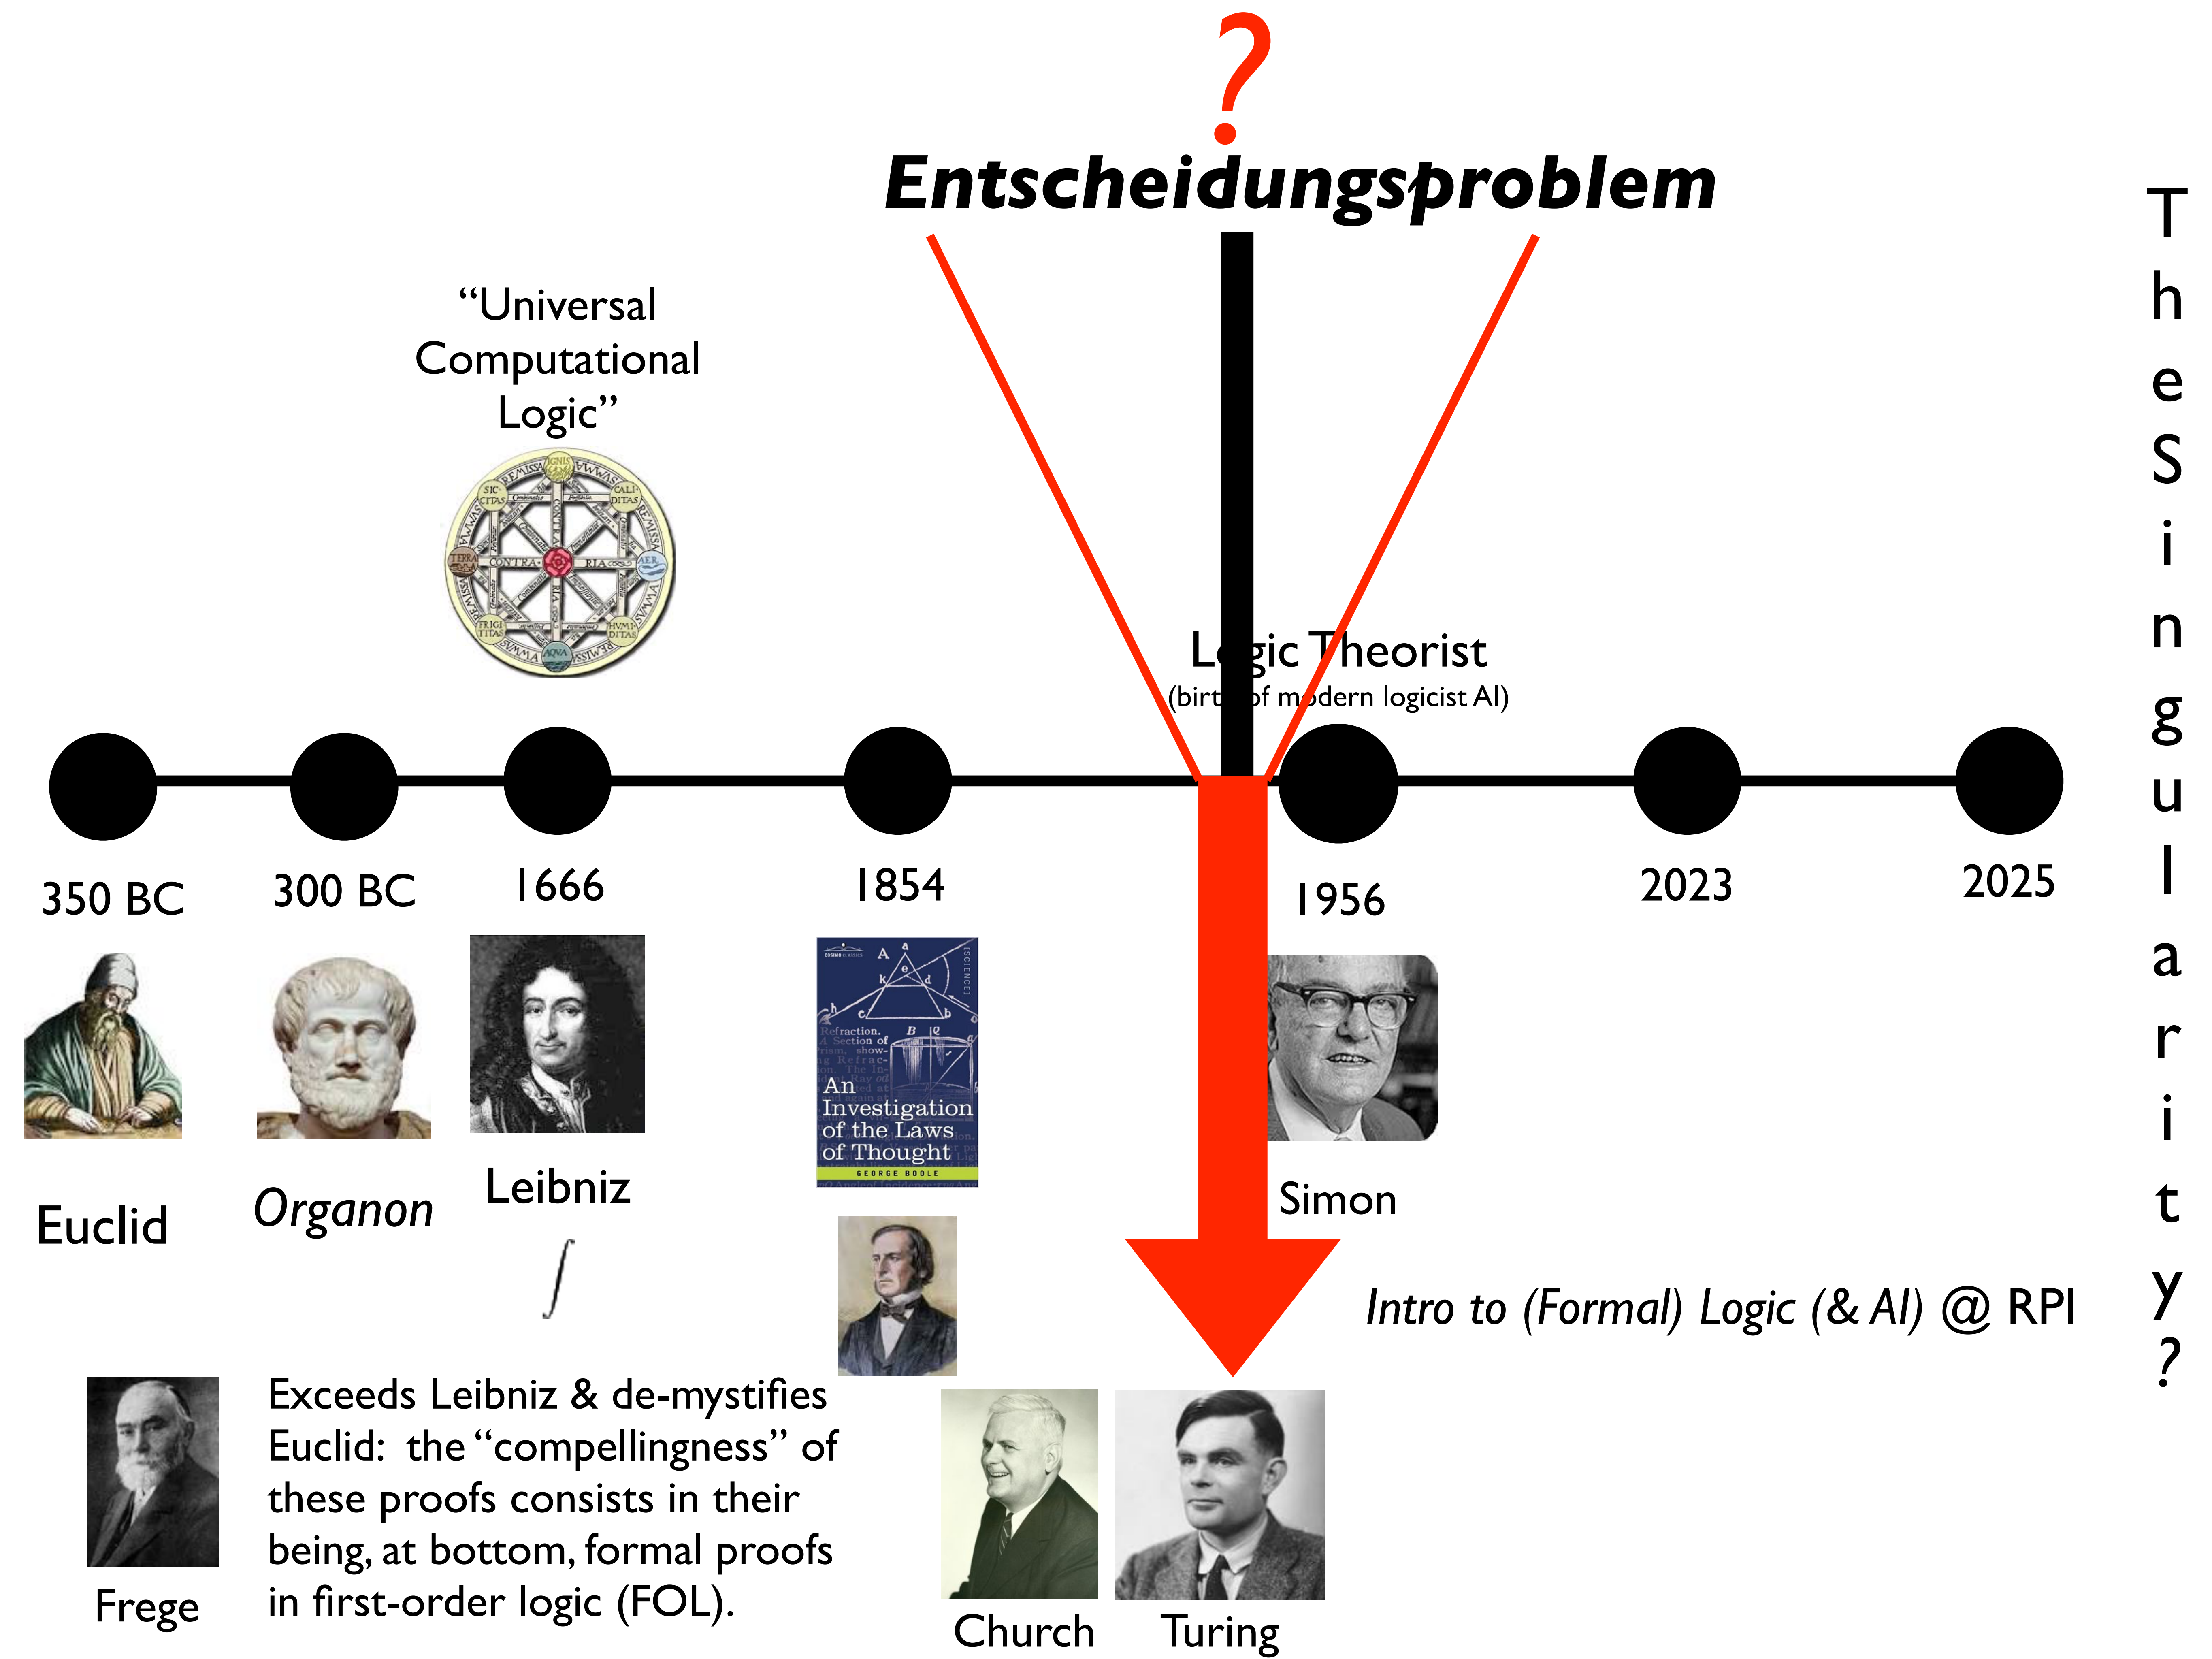
?
Entscheidungsproblem
"Universal Computational Logic"
Logic Theorist
(birth of modern logicist AI)
350 BC
300 BC
1666
1854
1956
2023
2025
Euclid
Organon
Leibniz
∫
An Investigation of the Laws of Thought
GEORGE BOOLE
Simon
Intro to (Formal) Logic (& AI) @ RPI
Frege
Exceeds Leibniz & de-mystifies Euclid: the "compellingness" of these proofs consists in their being, at bottom, formal proofs in first-order logic (FOL).
Church
Turing
The Singularity?

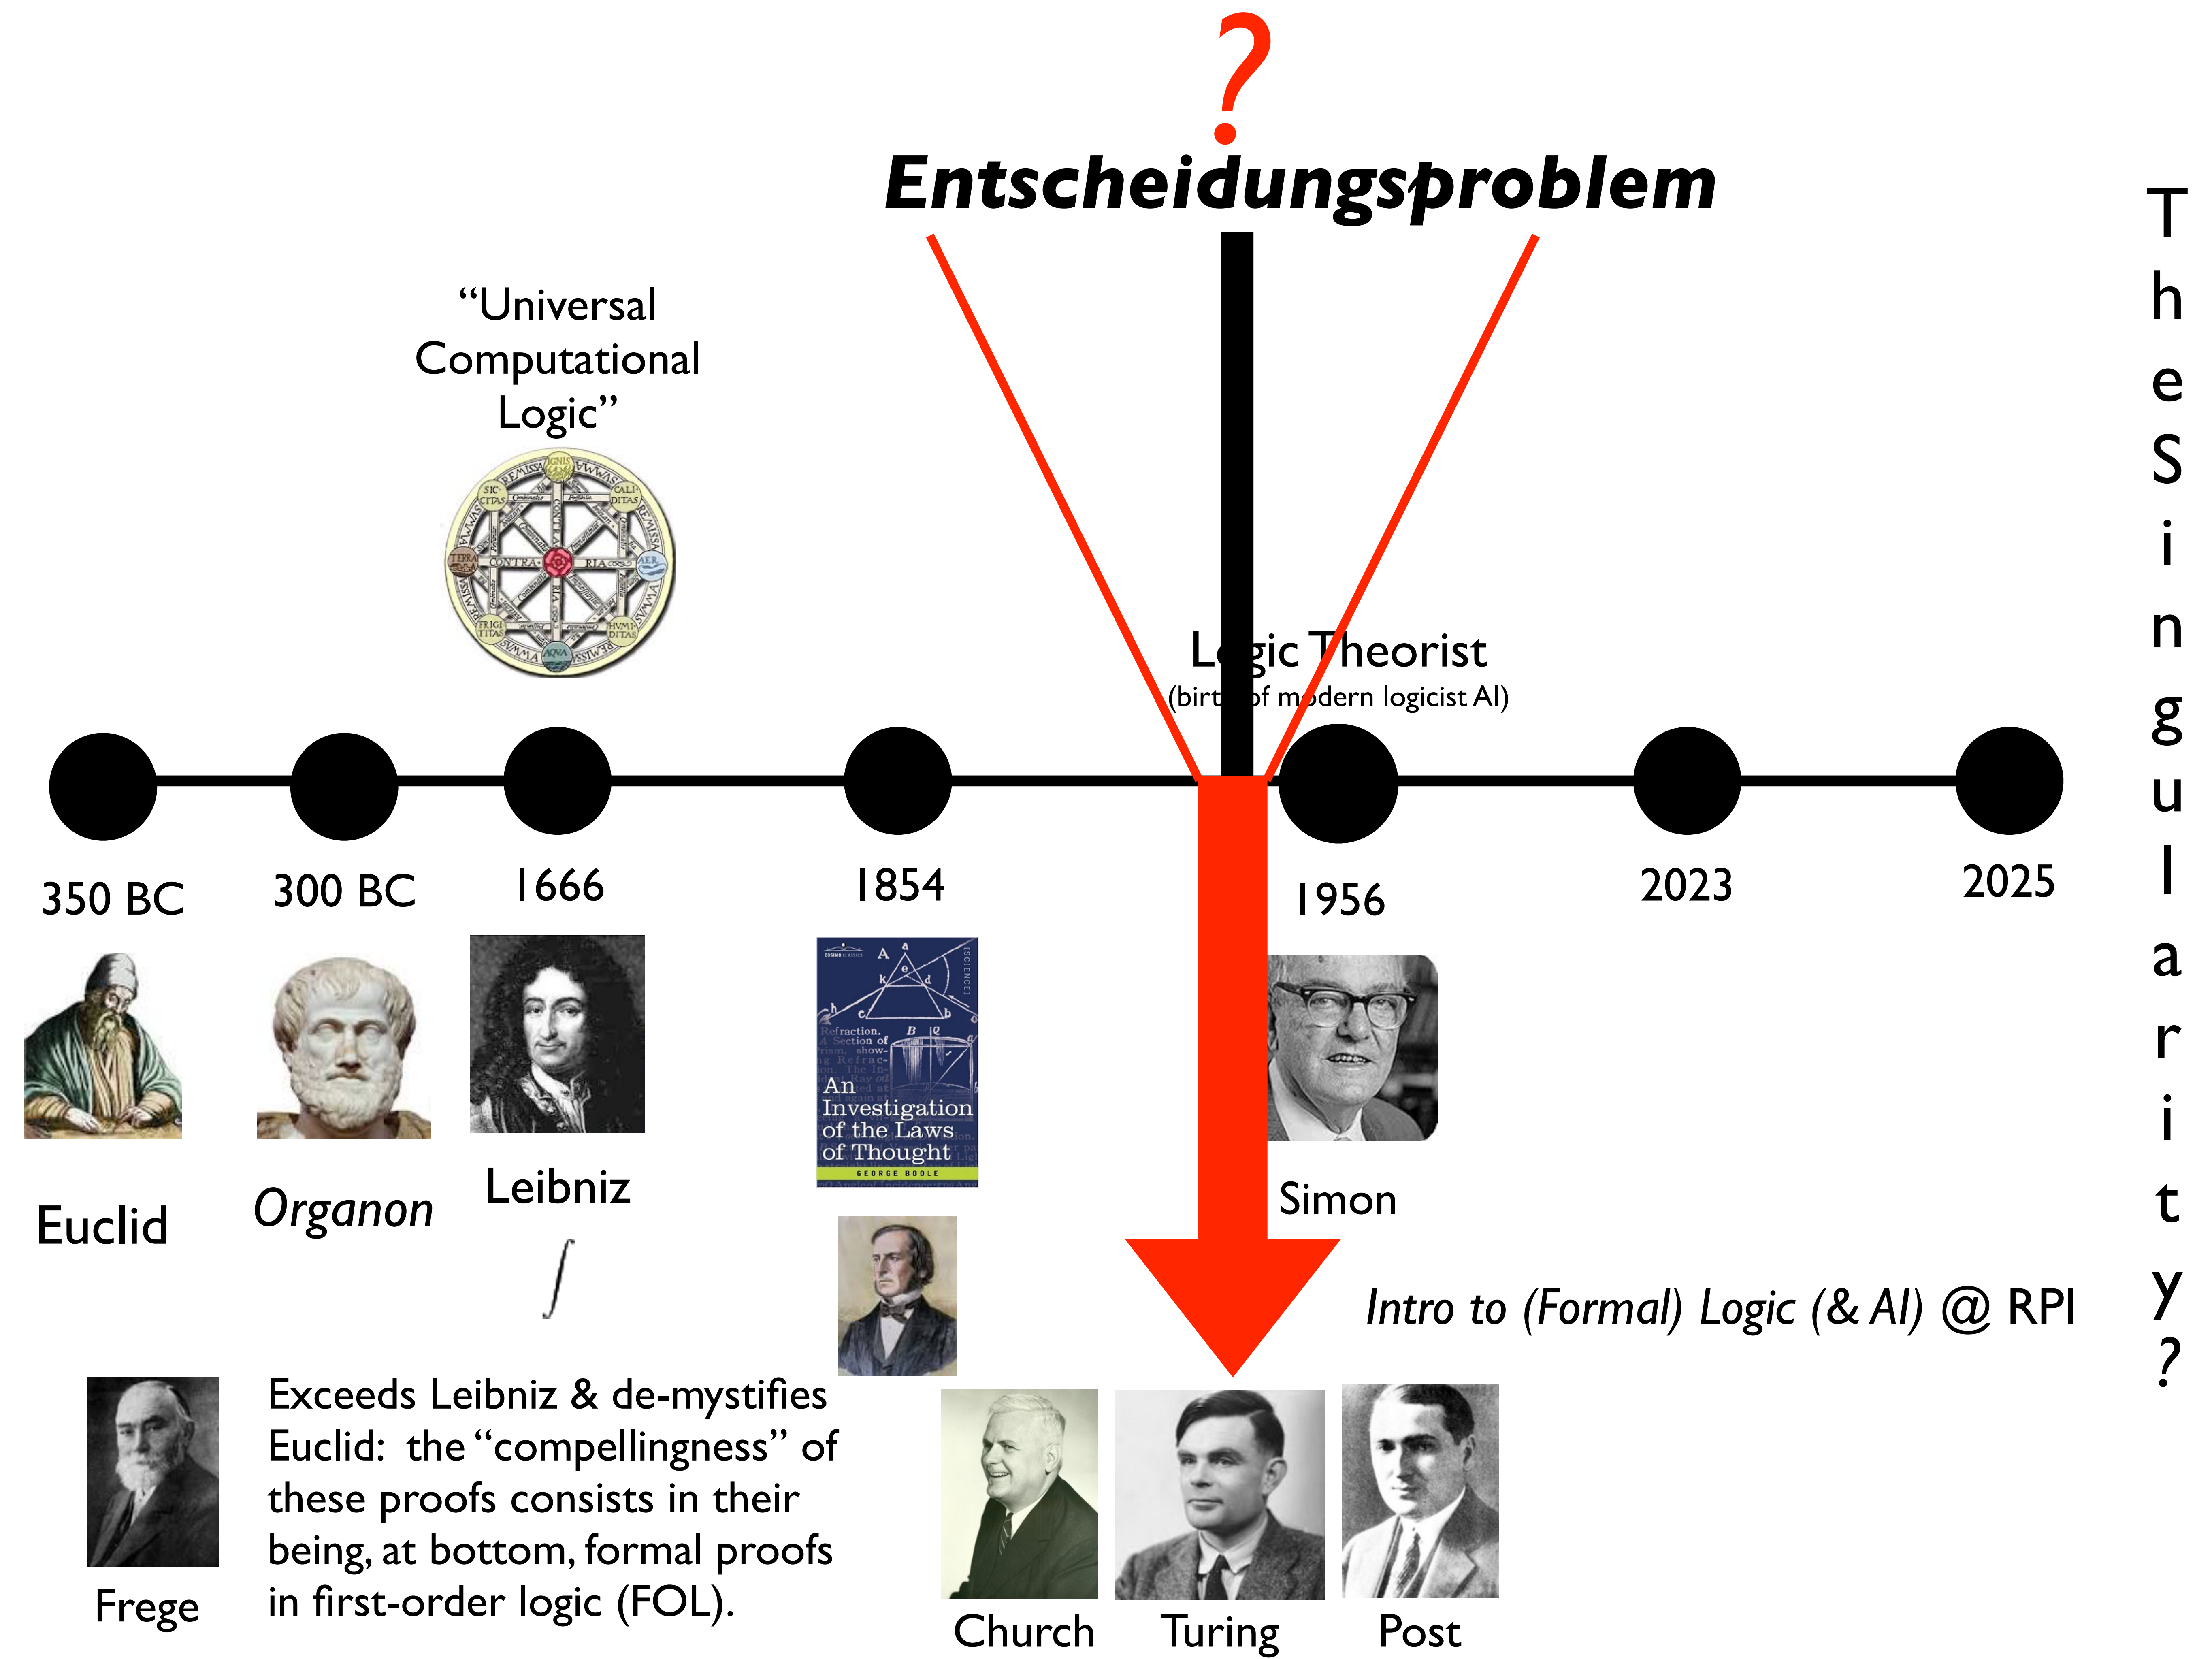
?
Entscheidungsproblem
"Universal Computational Logic"
Logic Theorist
(birth of modern logicist AI)
350 BC
300 BC
1666
1854
1956
2023
2025
Euclid
Organon
Leibniz
∫
An Investigation of the Laws of Thought
GEORGE BOOLE
Simon
Intro to (Formal) Logic (& AI) @ RPI
The Singularity?
Frege
Exceeds Leibniz & de-mystifies Euclid: the "compellingness" of these proofs consists in their being, at bottom, formal proofs in first-order logic (FOL).
Church
Turing
Post

?

Entscheidungsproblem

"Universal Computational Logic"

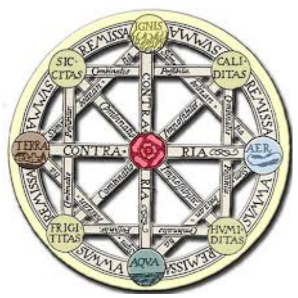

Logic Theorist
(birth of modern logicist AI)

350 BC | 300 BC | 1666 | 1854 | 1956 | 2023 | 2025

Euclid

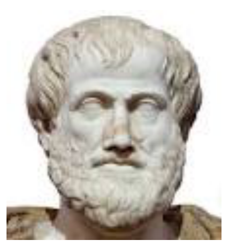

Organon

Leibniz

$\int$

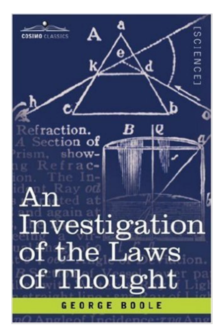

Simon

Intro to (Formal) Logic (& AI) @ RPI

The Singularity?

Frege

Exceeds Leibniz & de-mystifies Euclid: the "compellingness" of these proofs consists in their being, at bottom, formal proofs in first-order logic (FOL).

Church

Turing

Post

Here's what a computer is, and given that, sorry, the *Entscheidungsproblem* can't be solved by such a machine!

First, the Theoremhood Decision Problem ($\text{THEOREM}_{\text{PC}}$) for the Propositional Calculus

First, the Theoremhood Decision Problem (THEOREM$_{PC}$) for the Propositional Calculus

First, the Theoremhood Decision Problem (THEOREM_{PC}) for the Propositional Calculus

First, the Theoremhood Decision Problem (THEOREM_{PC}) for the Propositional Calculus

First, the Theoremhood Decision Problem (THEOREM_{PC}) for the Propositional Calculus

First, the Theoremhood Decision Problem ($\text{THEOREM}_{\text{PC}}$) for the Propositional Calculus

First, the Theoremhood Decision Problem ($\text{THEOREM}_{\text{PC}}$) for the Propositional Calculus

Hard!! — for apparently no polynomial-time algorithm for this!

First, the Theoremhood Decision Problem (THEOREM$_{PC}$) for the Propositional Calculus

Hard!! — for apparently no polynomial-time algorithm for this!

First, the Theoremhood Decision Problem
(THEOREM$_{PC}$)
for the Propositional Calculus

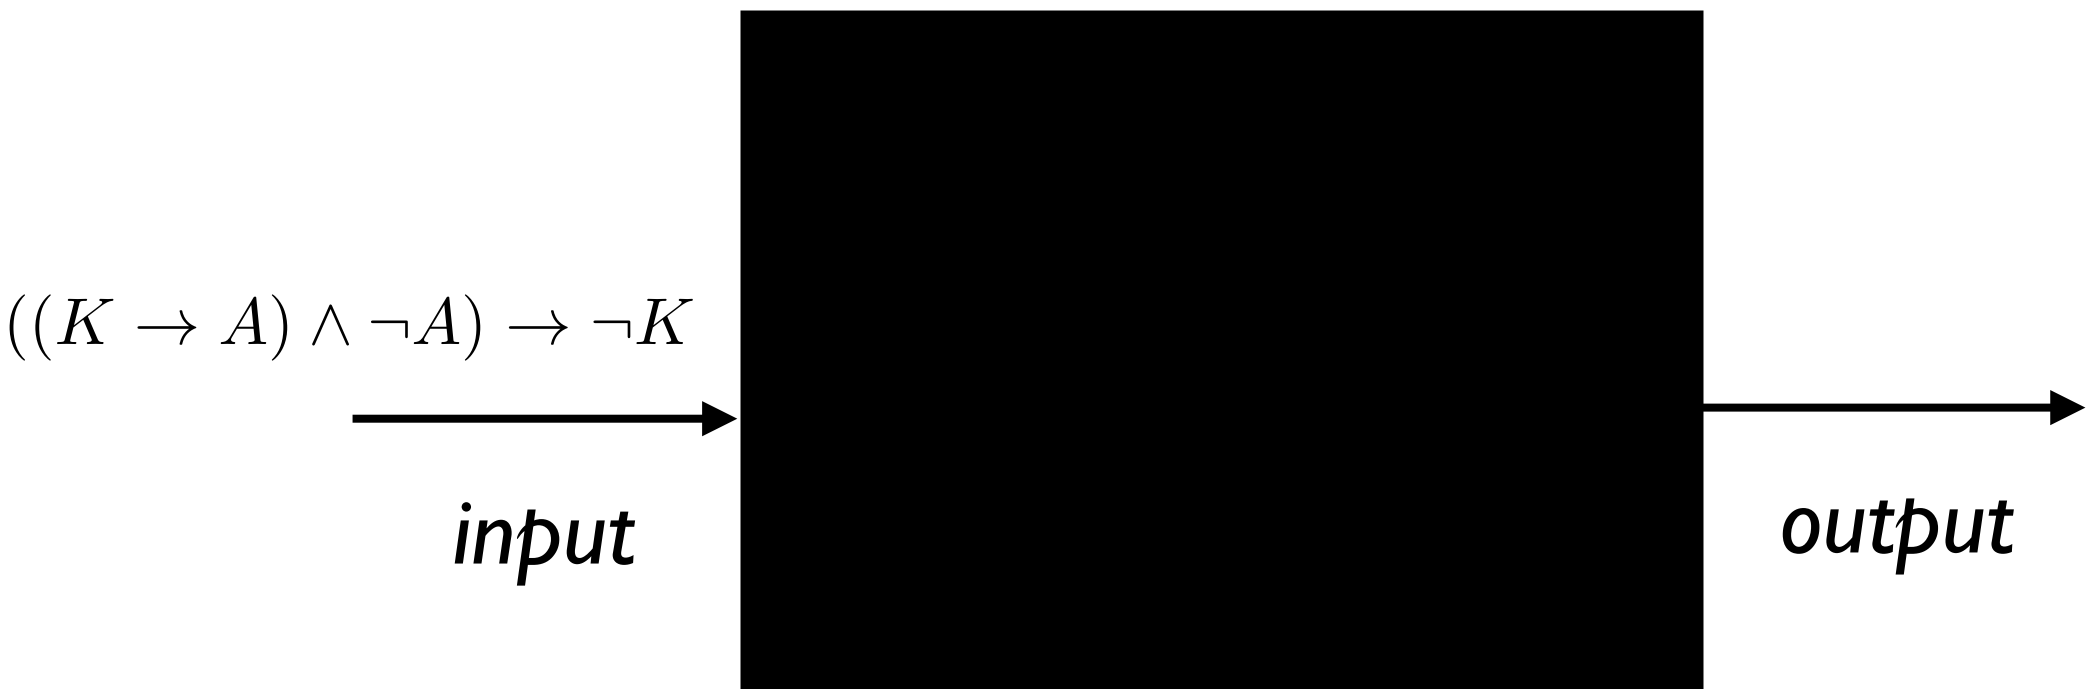

Hard!! — for apparently no polynomial-time algorithm for this!

First, the Theoremhood Decision Problem (THEOREM$_{PC}$) for the Propositional Calculus

$$((K \to A) \land \neg A) \to \neg K$$

Hard!! — for apparently no polynomial-time algorithm for this!

First, the Theoremhood Decision Problem ($\text{THEOREM}_{\text{PC}}$) for the Propositional Calculus

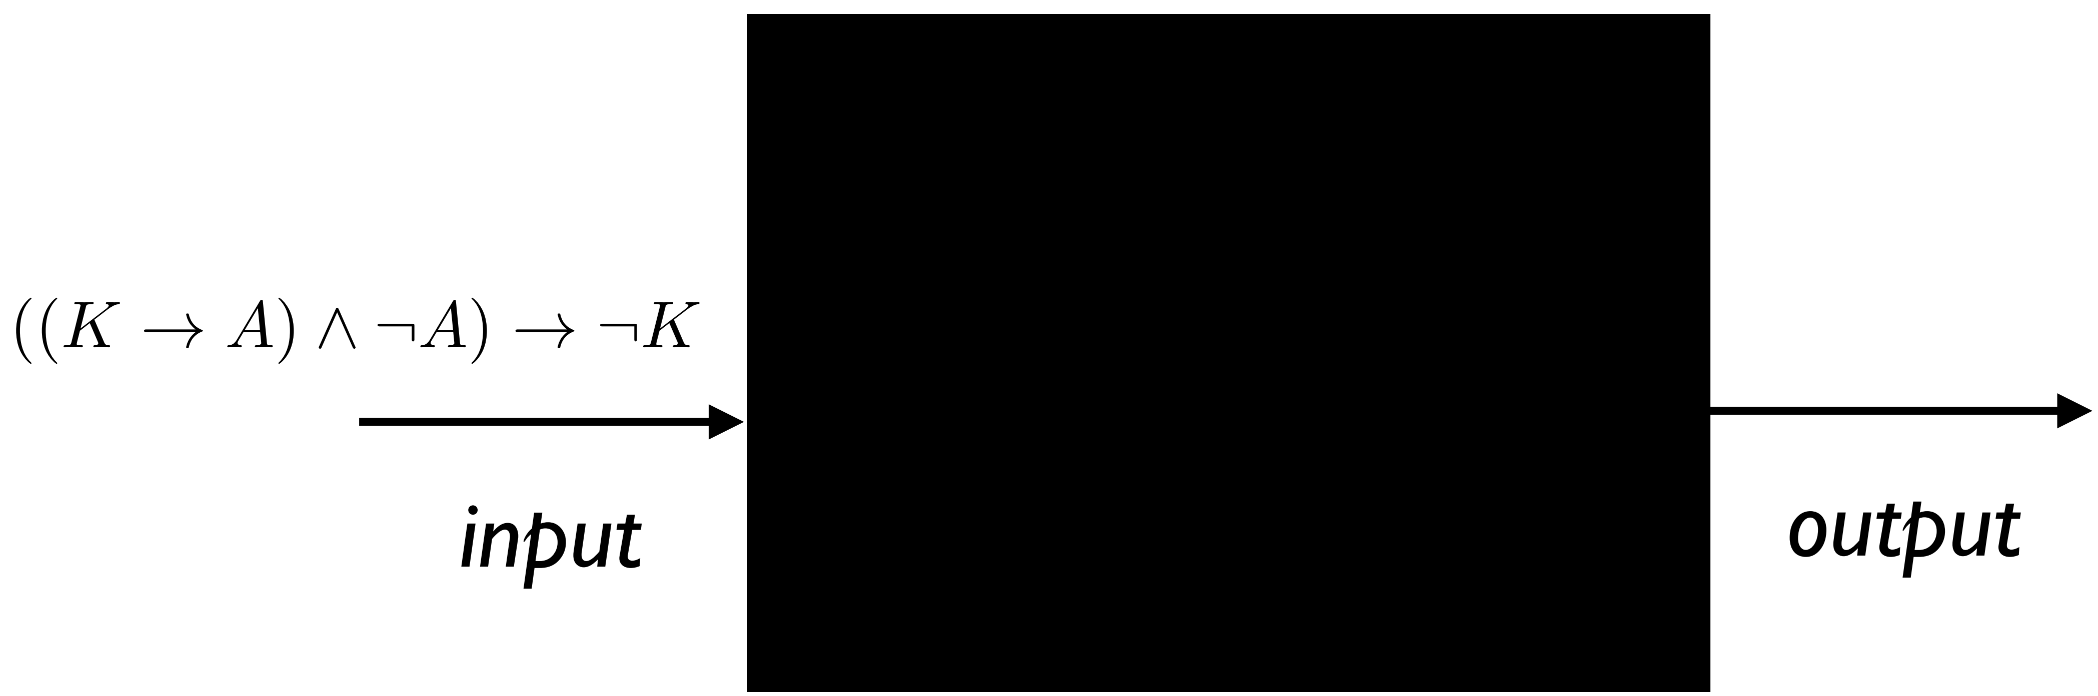

Hard!! — for apparently no polynomial-time algorithm for this!

First, the Theoremhood Decision Problem (THEOREM$_{PC}$) for the Propositional Calculus

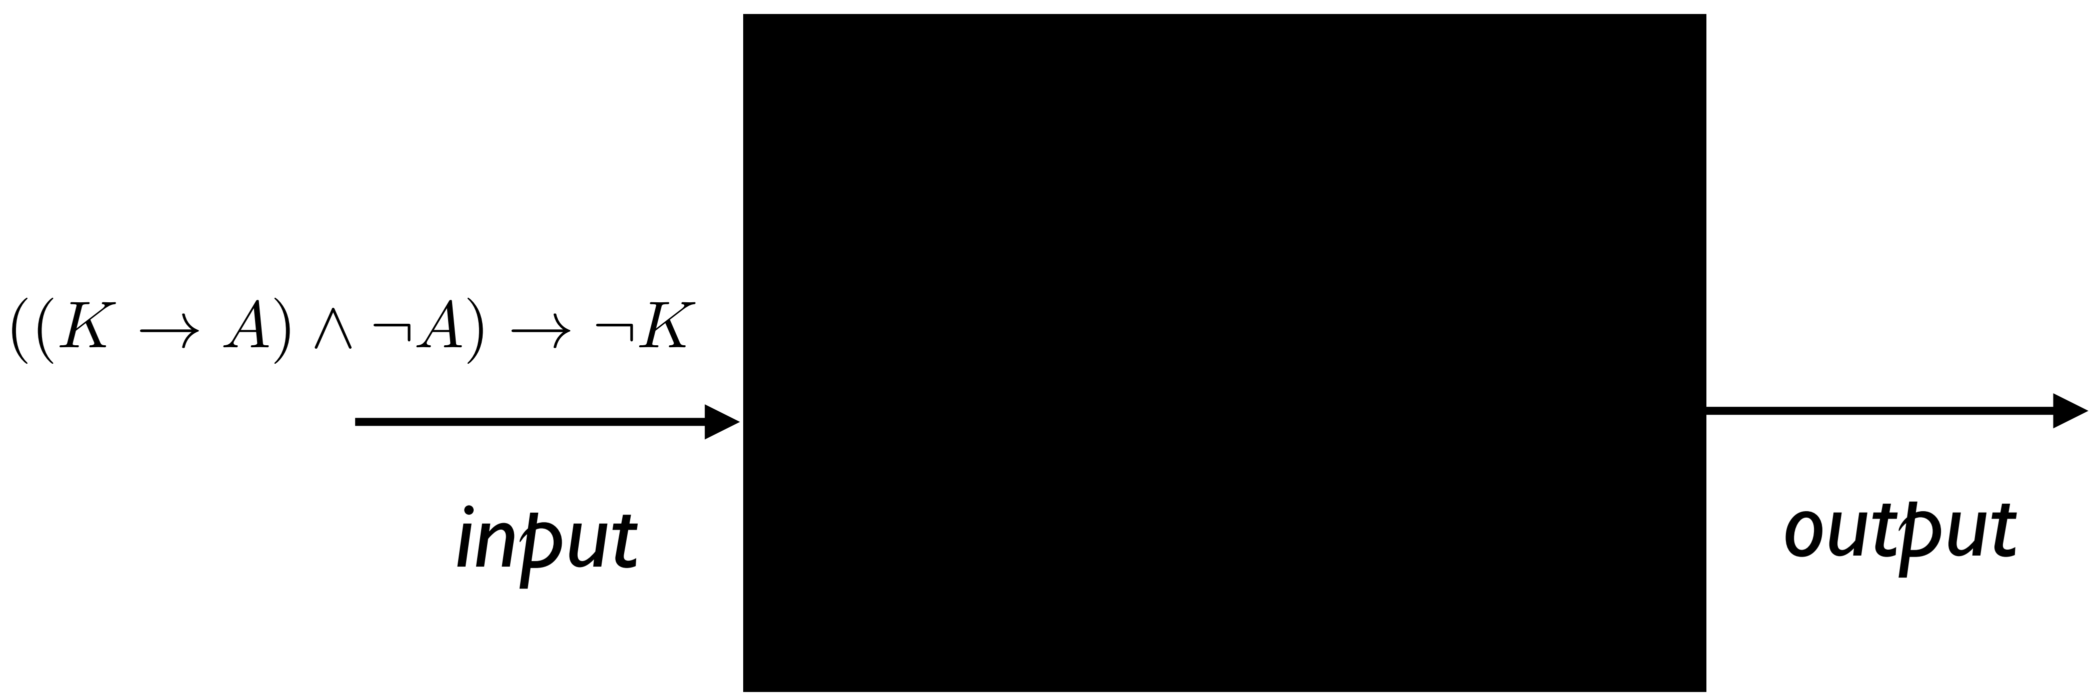

Hard!! — for apparently no polynomial-time algorithm for this!

First, the Theoremhood Decision Problem (THEOREM$_{PC}$) for the Propositional Calculus

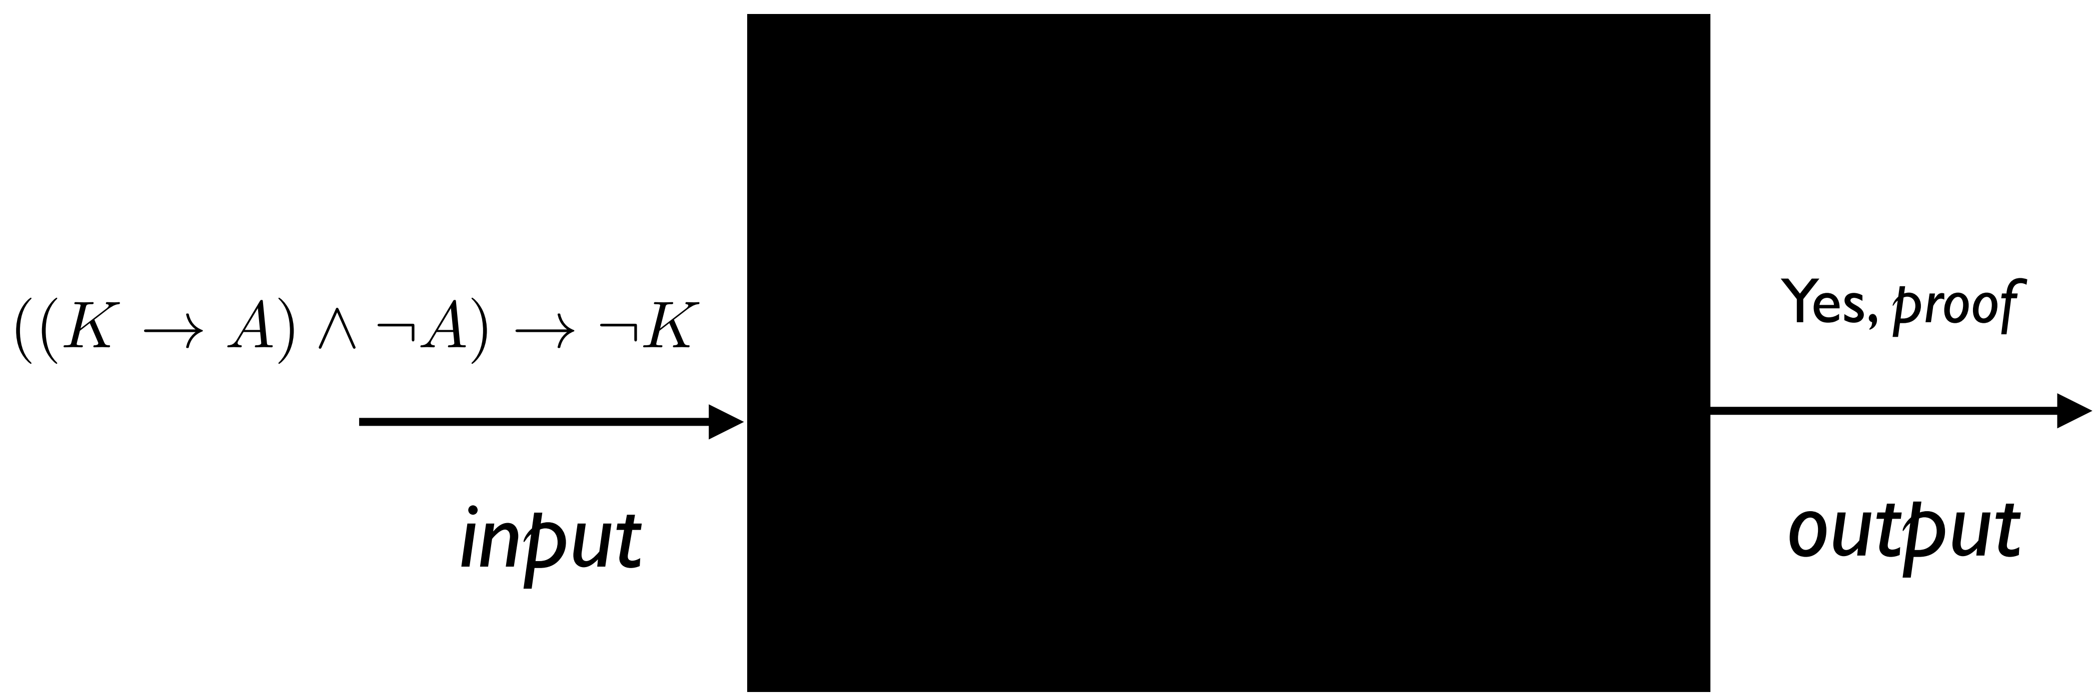

Hard!! — for apparently no polynomial-time algorithm for this!

And now, the Theoremhood Decision Problem, i.e., the *Entscheidungsproblem*, ($\text{THEOREM}_{\text{FOL}}$) for First-Order Logic (FOL)

And now, the Theoremhood Decision Problem, i.e., the *Entscheidungsproblem*, ($THEOREM_{FOL}$) for First-Order Logic (FOL)

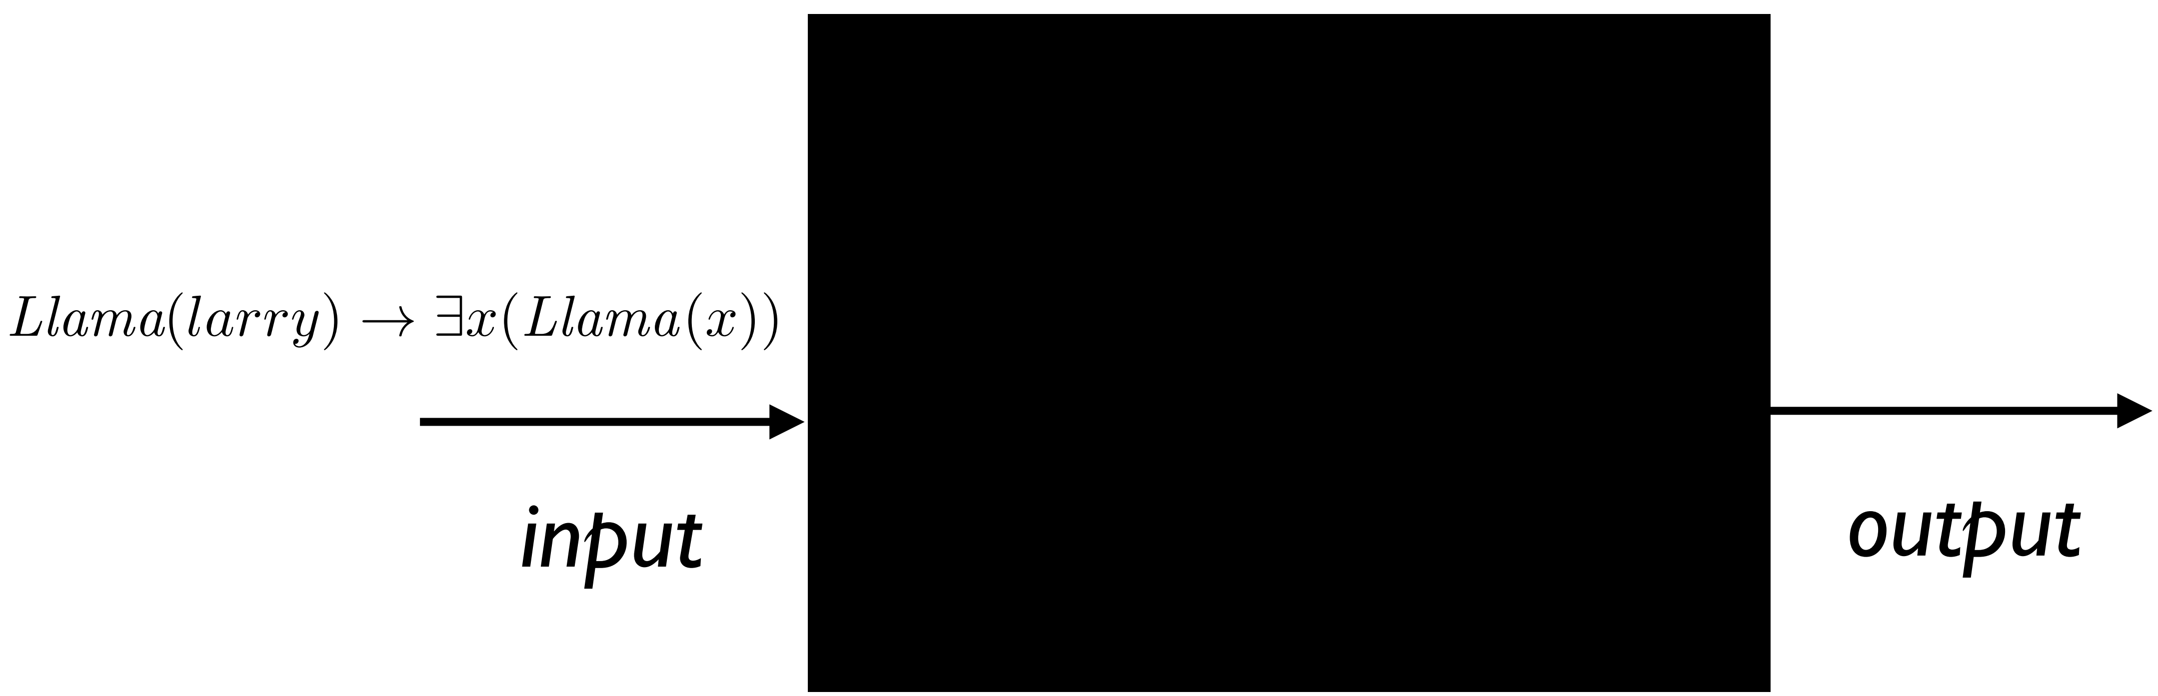

And now, the Theoremhood Decision Problem, i.e., the *Entscheidungsproblem*, (THEOREM$_{FOL}$) for First-Order Logic (FOL)

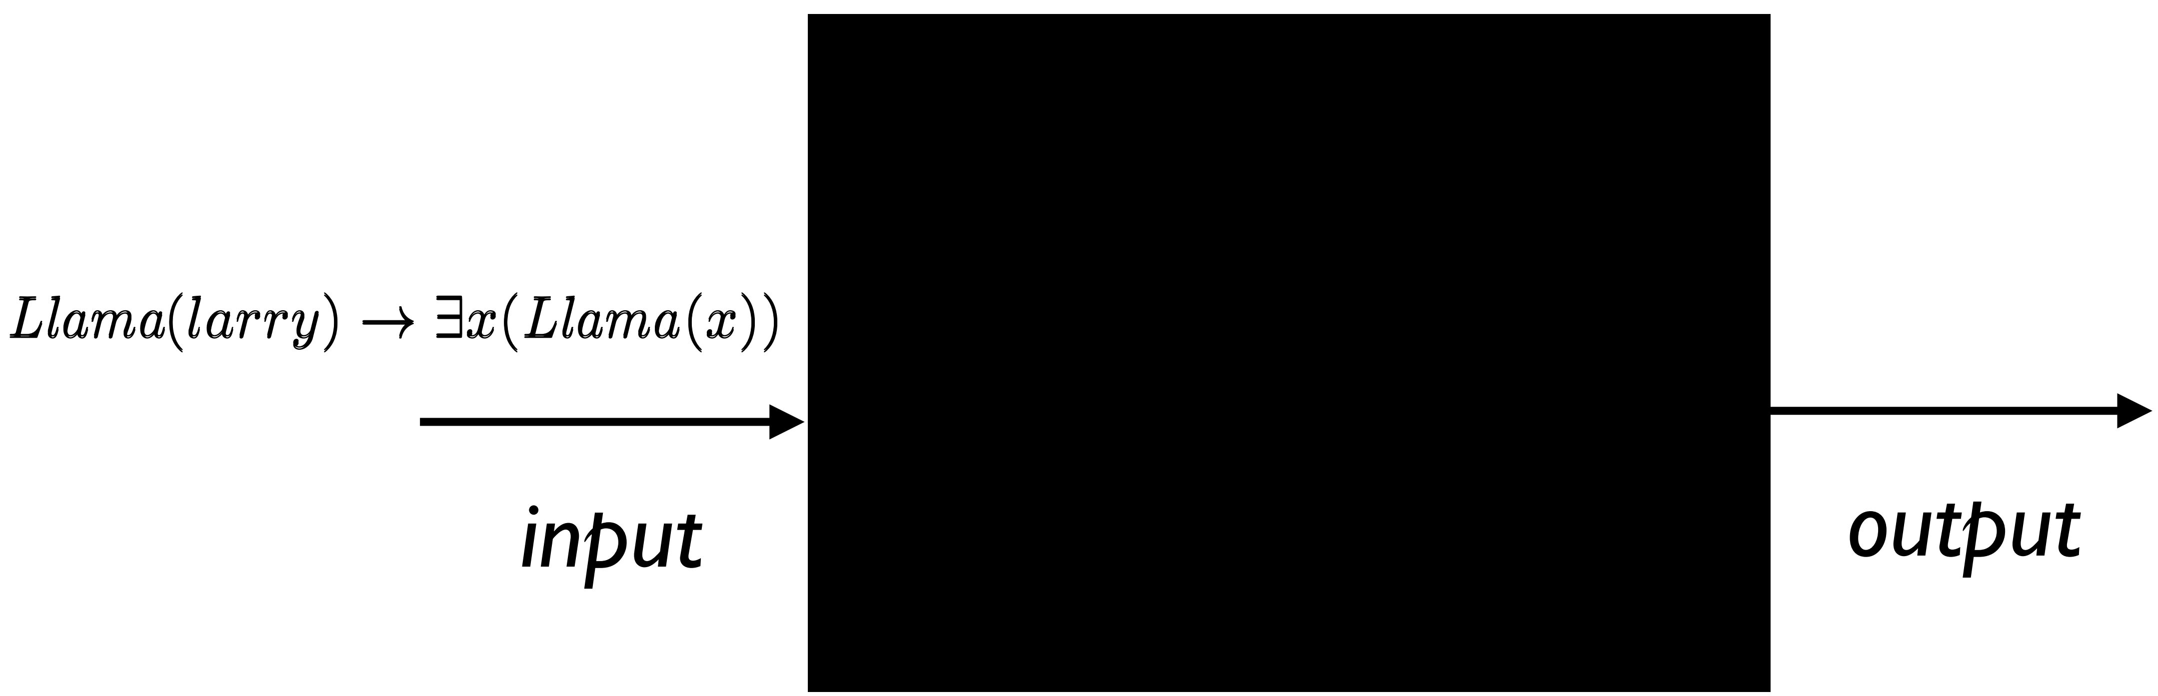

And now, the Theoremhood Decision Problem, i.e., the *Entscheidungsproblem*, ($\text{THEOREM}_{\text{FOL}}$) for First-Order Logic (FOL)

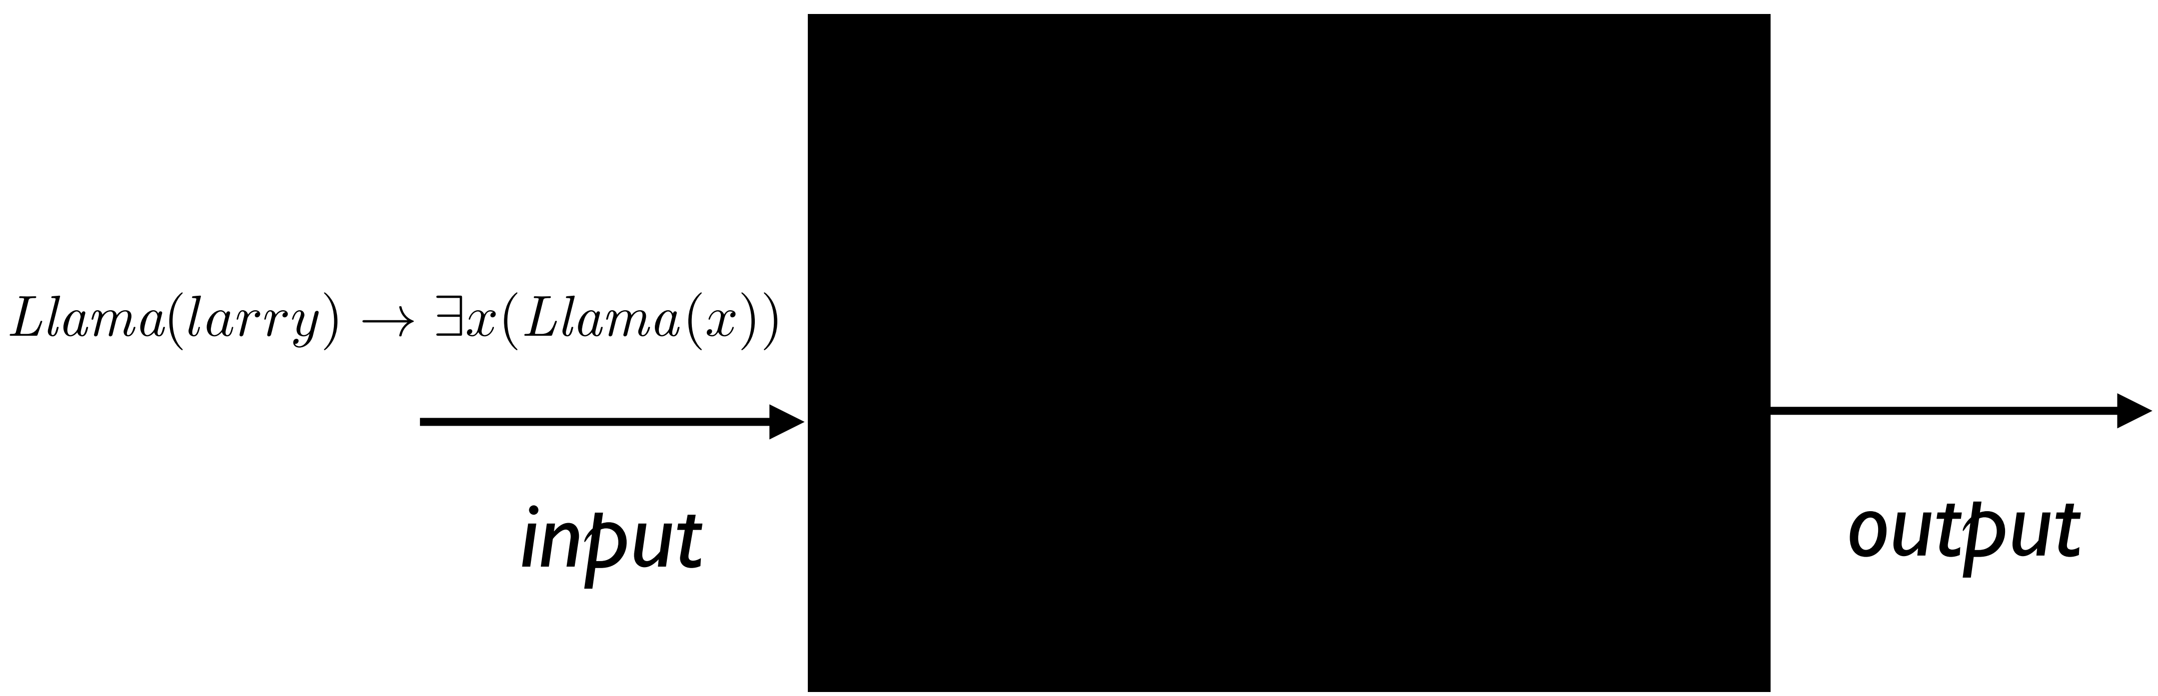

And now, the Theoremhood Decision Problem,
i.e., the *Entscheidungsproblem*,
($\text{THEOREM}_{\text{FOL}}$)
for First-Order Logic (FOL)

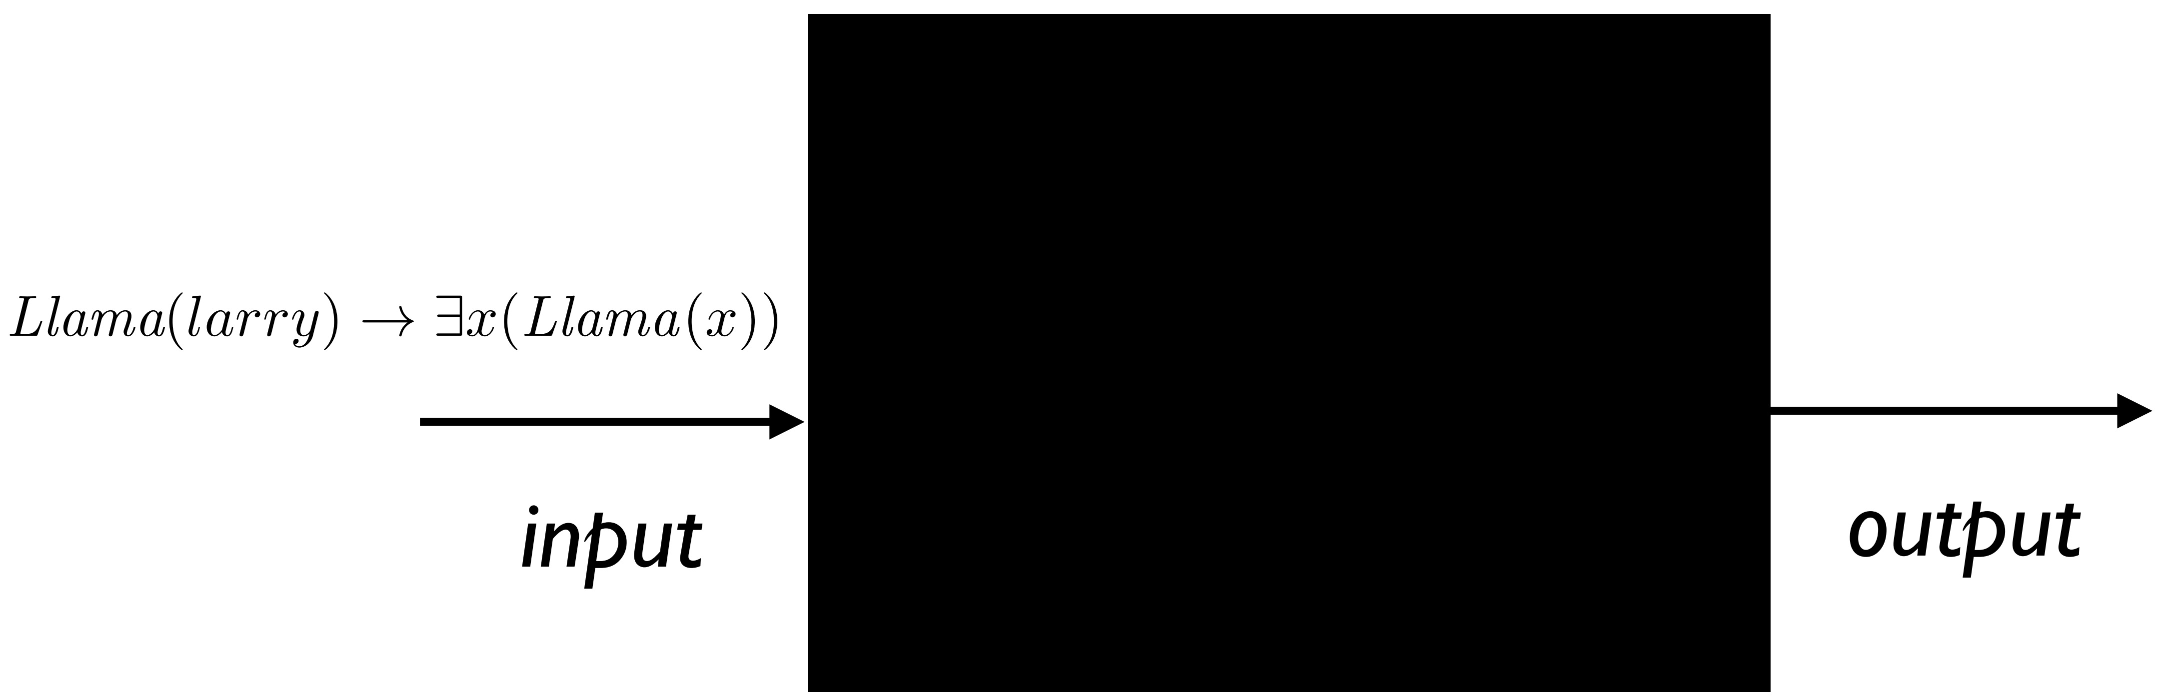

And now, the Theoremhood Decision Problem, i.e., the *Entscheidungsproblem*, (THEOREM$_{FOL}$) for First-Order Logic (FOL)

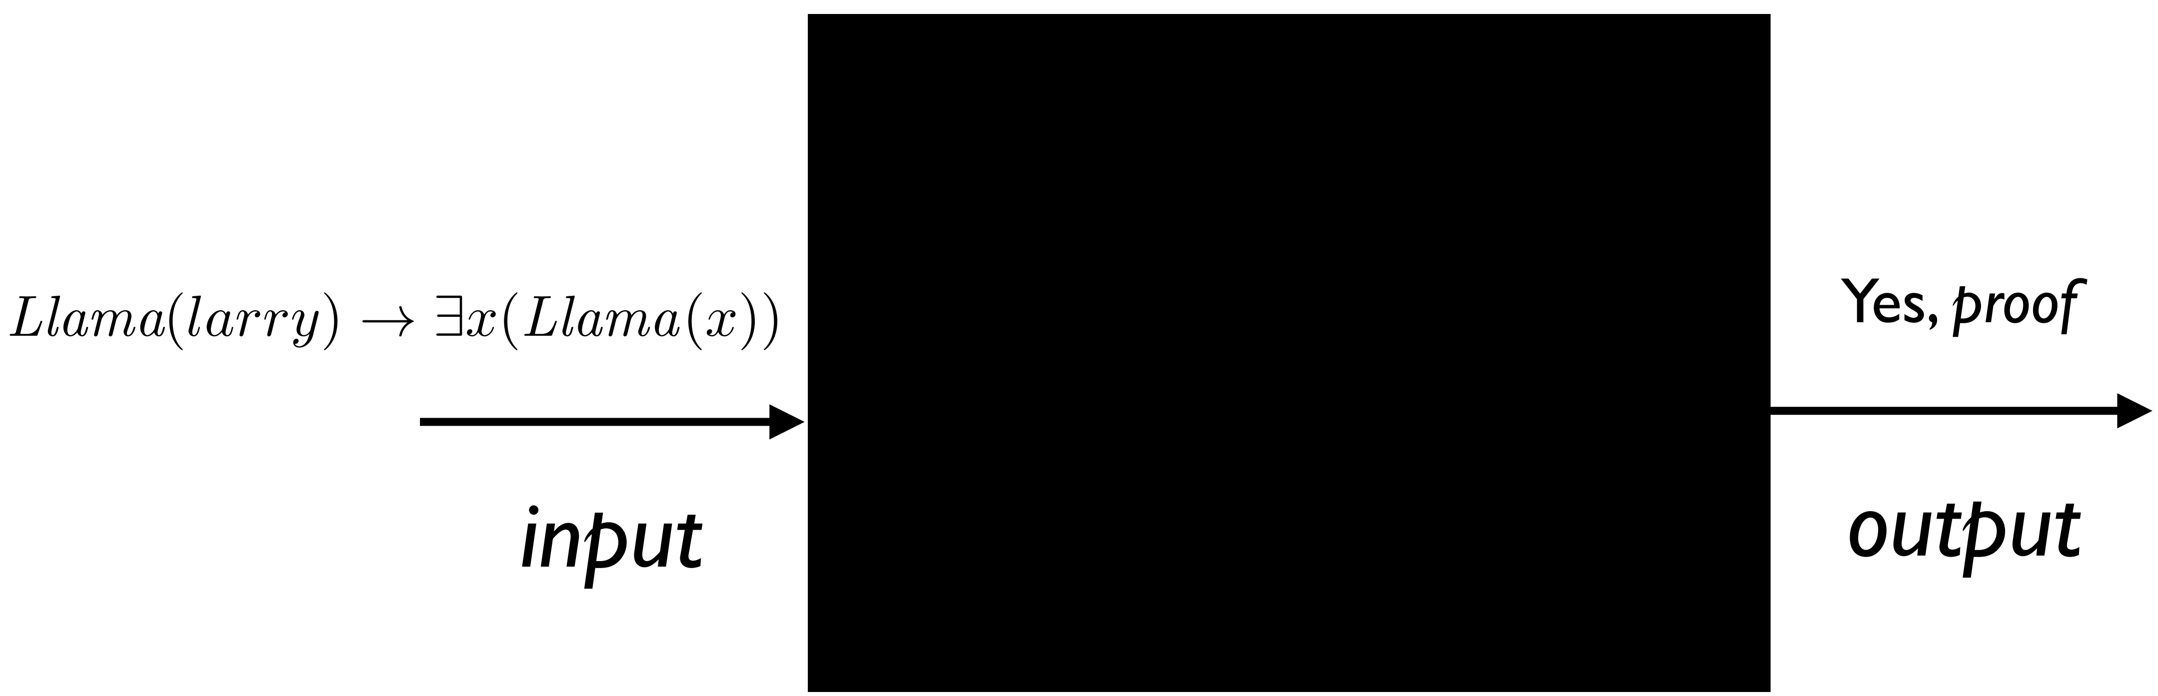

And now, the Theoremhood Decision Problem, i.e., the *Entscheidungsproblem*, ($\text{THEOREM}_{\text{FOL}}$) for First-Order Logic (FOL)

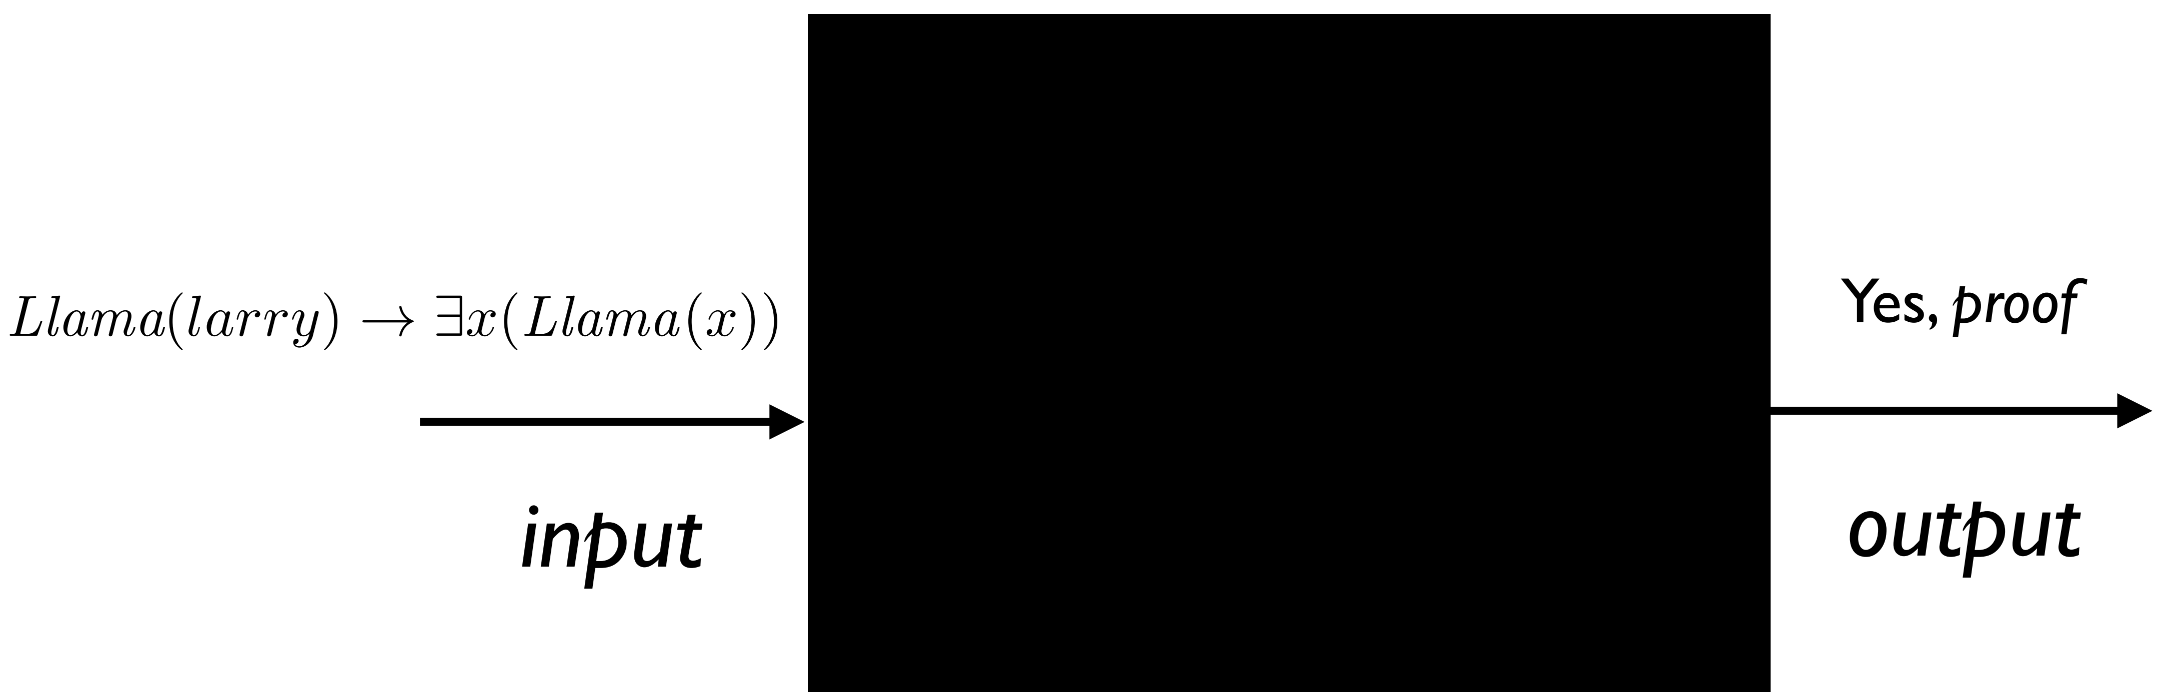

Not just hard: *impossible* for a (and this needed to be *invented* in the course of clarifying and solving the problem) standard computing machine.

Applying this to …
The Singularity Question

Applying this to …
The Singularity Question

$\mathcal{A}$:

	Premise 1	There will be AI (created by HI and such that AI = HI).
	Premise 2	If there is AI, there will be AI^{+} (created by AI).
	Premise 3	If there is AI^{+}, there will be AI^{++} (created by AI^{+}).
$\therefore$	**S**	There will be AI^{++} (= $\mathcal{S}$ will occur).

(Good-Chalmers Argument)

(Kurzweil is an "extrapolationist.")

Applying this to …
The Singularity Question

So, these super-smart machines that will be built by human-level-smart machines, they can't *possibly* be smart enough to solve the *Entscheidungsproblem*. Hence they'll be just (recursively) faster at solving problems we can routinely solve? What's so super-smart about *that*?

Logistics …

The Starting Code to Purchase in Bookstore

Your code for starting the registration process is:

To access HyperGrader, HyperSlate, the license agreement,
and to obtain the textbook LAMA-BDLA, go to::

https://rpi.logicamodernapproach.com

The Starting Code to Purchase in Bookstore

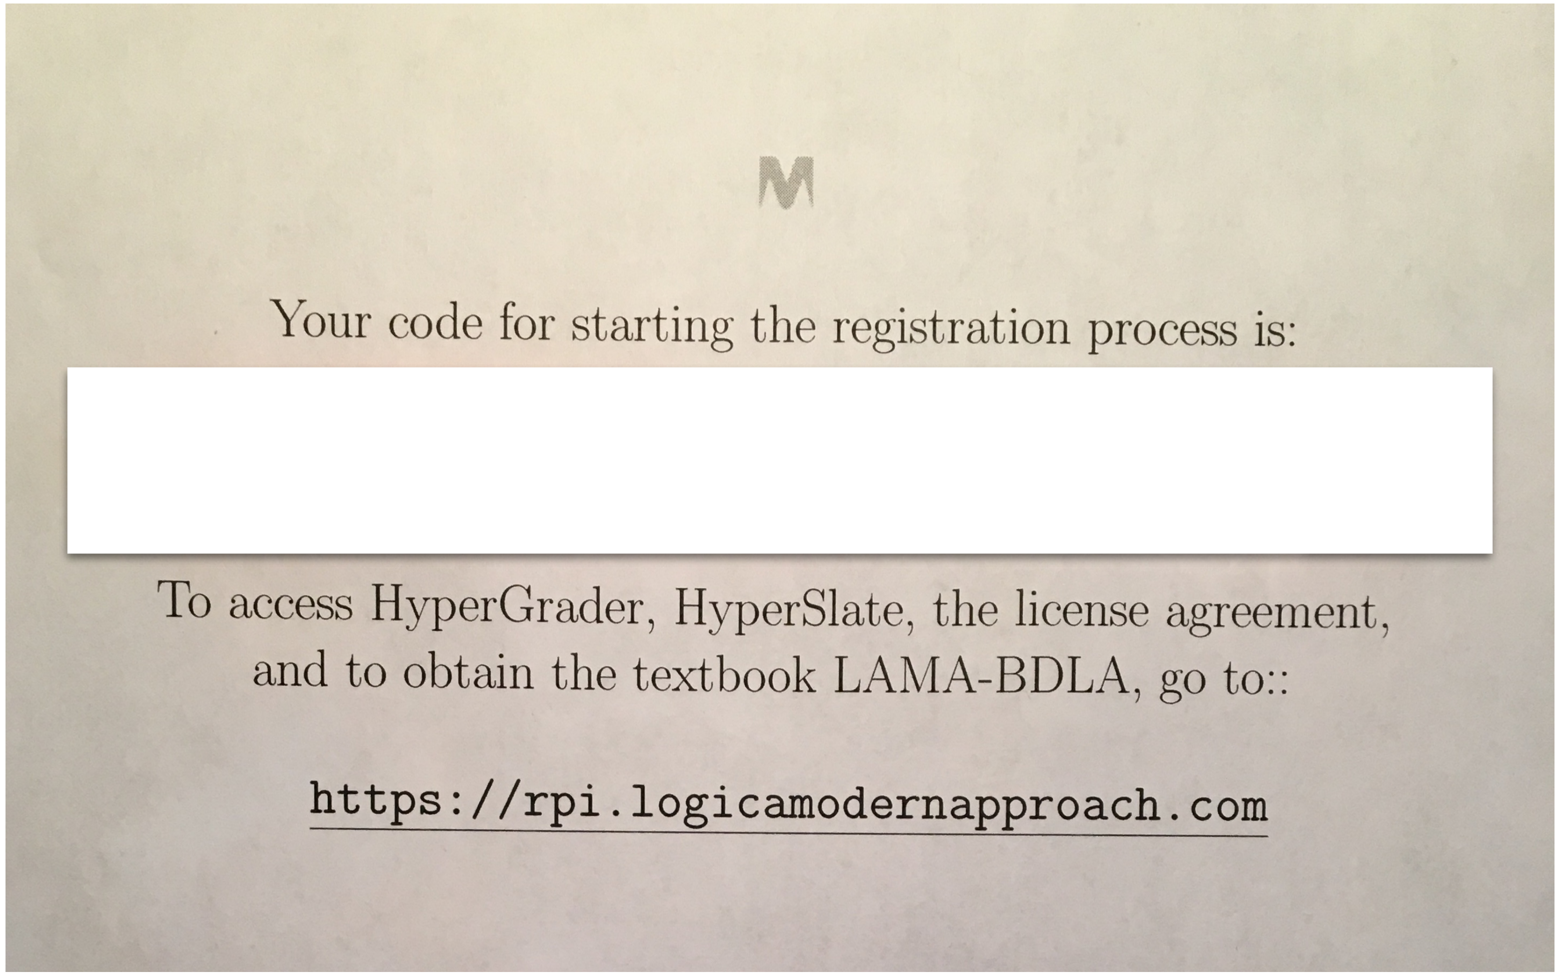
Your code for starting the registration process is:

To access HyperGrader, HyperSlate, the license agreement, and to obtain the textbook LAMA-BDLA, go to::

https://rpi.logicamodernapproach.com

Once seal broken on envelope, no return. Remember from first class, any reservations, opt for "Stanford" paradigm, with its software instead of LAMA® paradigm!

The Starting Code to Purchase in Bookstore

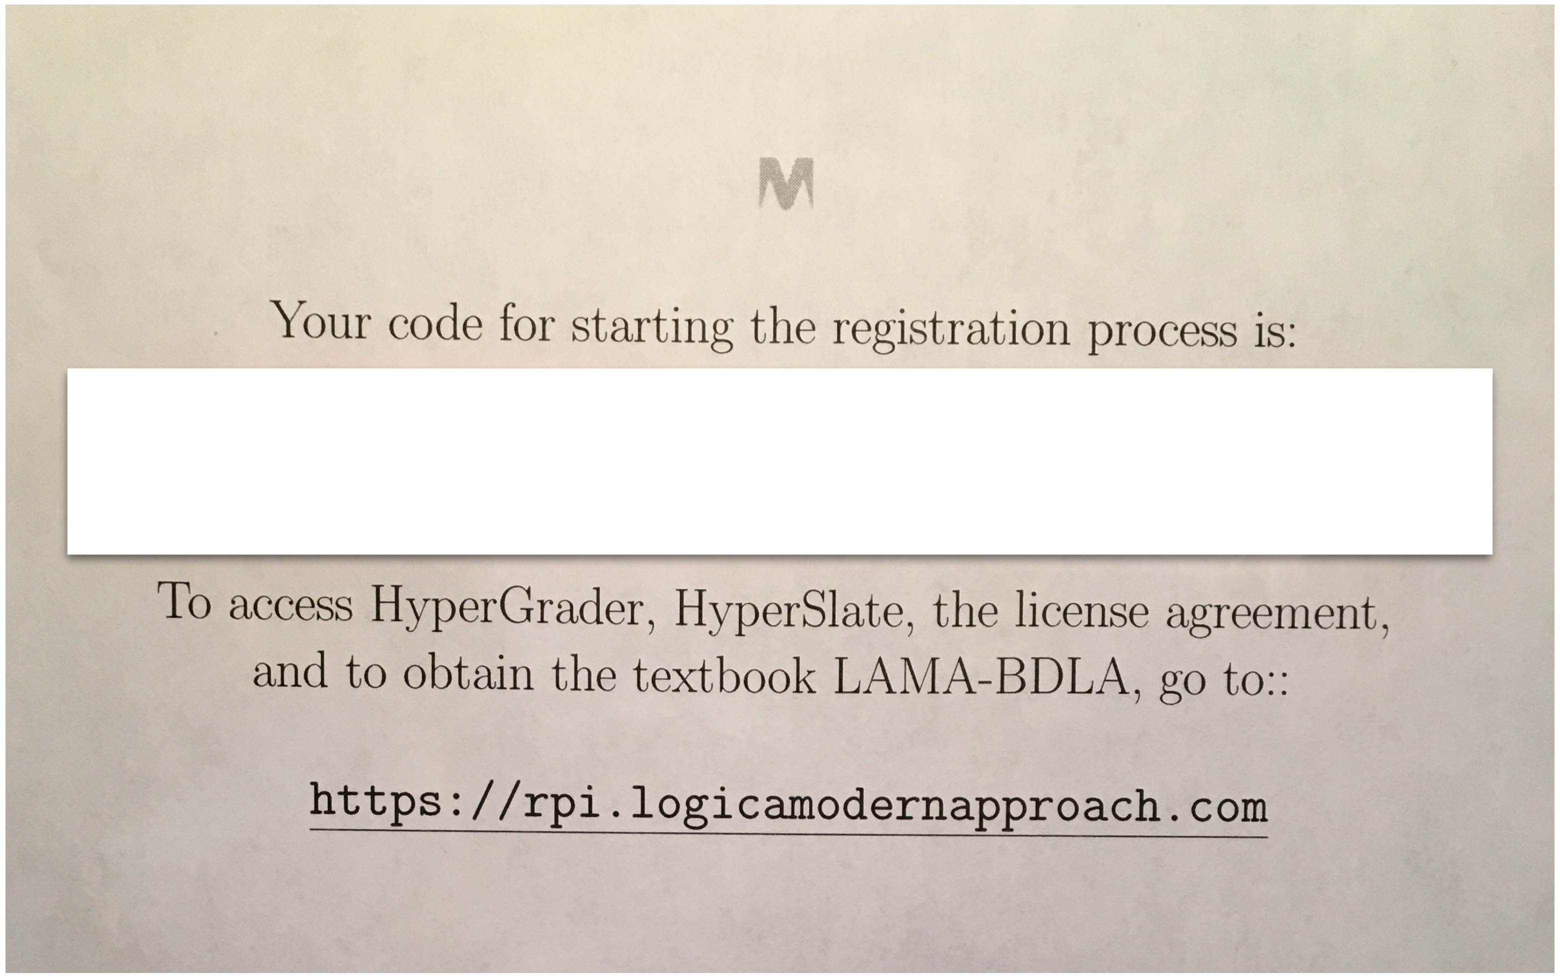
Your code for starting the registration process is:

To access HyperGrader, HyperSlate, the license agreement, and to obtain the textbook LAMA-BDLA, go to::

https://rpi.logicamodernapproach.com

Once seal broken on envelope, no return. Remember from first class, any reservations, opt for "Stanford" paradigm, with its software instead of LAMA® paradigm!

The email address you enter is case-sensitive!

The Starting Code to Purchase in Bookstore

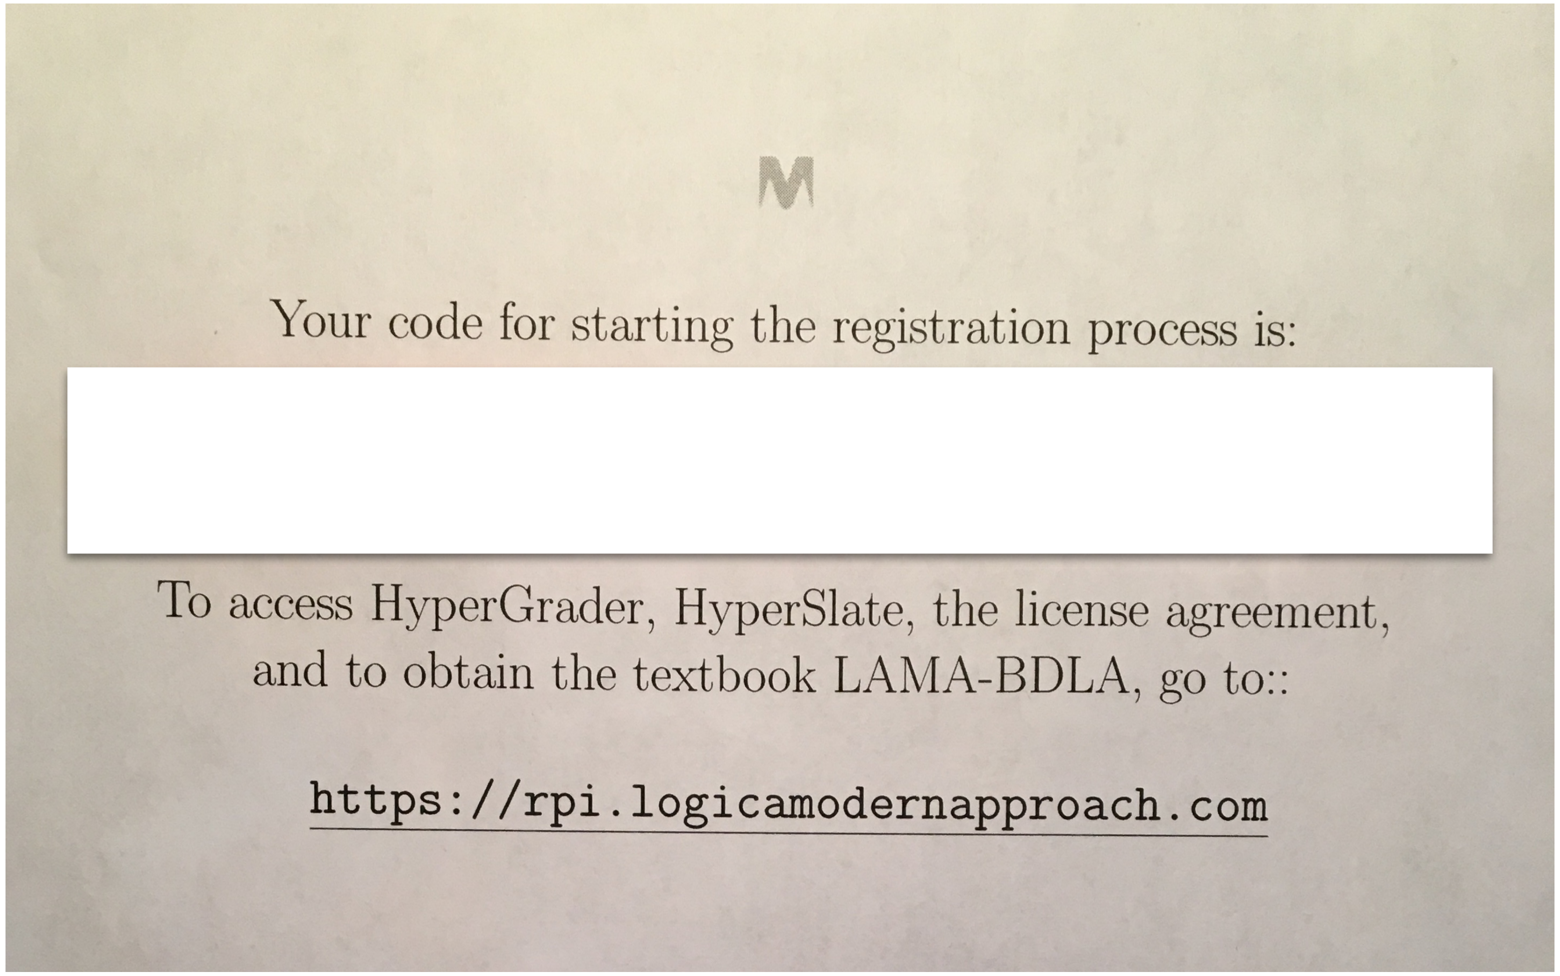
Your code for starting the registration process is:

To access HyperGrader, HyperSlate, the license agreement, and to obtain the textbook LAMA-BDLA, go to::

https://rpi.logicamodernapproach.com

Once seal broken on envelope, no return. Remember from first class, any reservations, opt for "Stanford" paradigm, with its software instead of LAMA® paradigm!

The email address you enter is case-sensitive!

Your OS and browser must be fully up-to-date; Chrome is the best choice, browser-wise (though I use Safari).

The Starting Code to Purchase in Bookstore

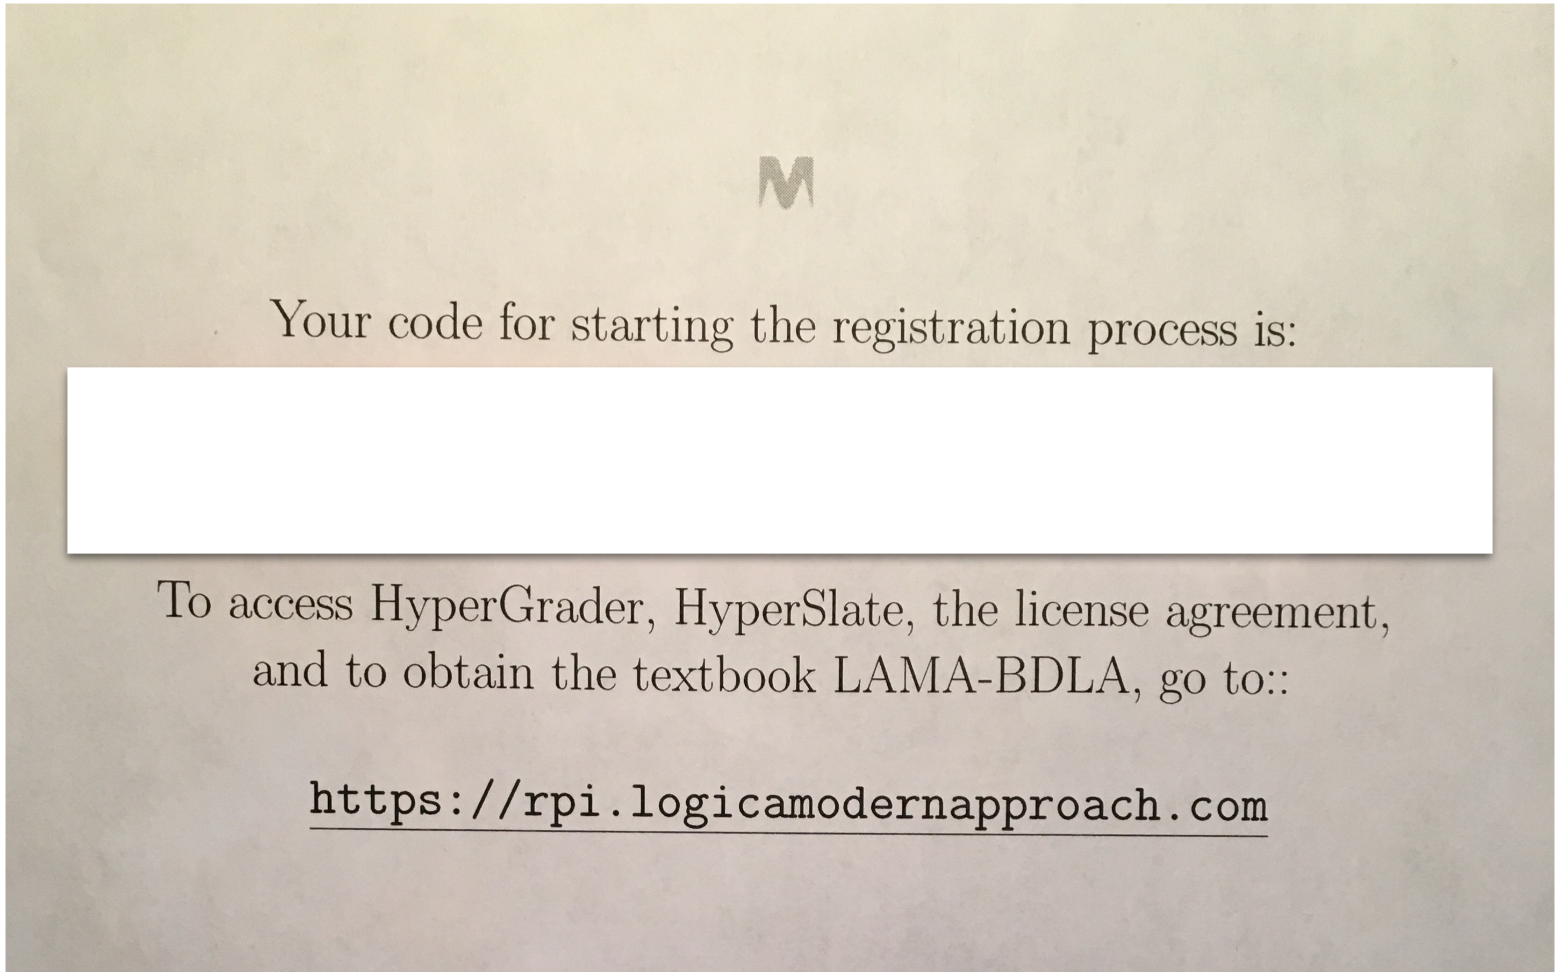
M

Your code for starting the registration process is:

To access HyperGrader, HyperSlate, the license agreement, and to obtain the textbook LAMA-BDLA, go to::

https://rpi.logicamodernapproach.com

Once seal broken on envelope, no return. Remember from first class, any reservations, opt for "Stanford" paradigm, with its software instead of LAMA® paradigm!

The email address you enter is case-sensitive!

Your OS and browser must be fully up-to-date; Chrome is the best choice, browser-wise (though I use Safari).

Watch that the link emailed to you doesn't end up being classified as spam.

Let's go live for a tutorial …

Logikk kan gi dyp glede!