What is Logic-Based/Logicist AI?; Its History

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-based AI
9/5/2024

Logic-and-AI in the news

• • •

Bit of Background: The OpenAI Soap Opera

Bit of Background: The OpenAI Soap Opera

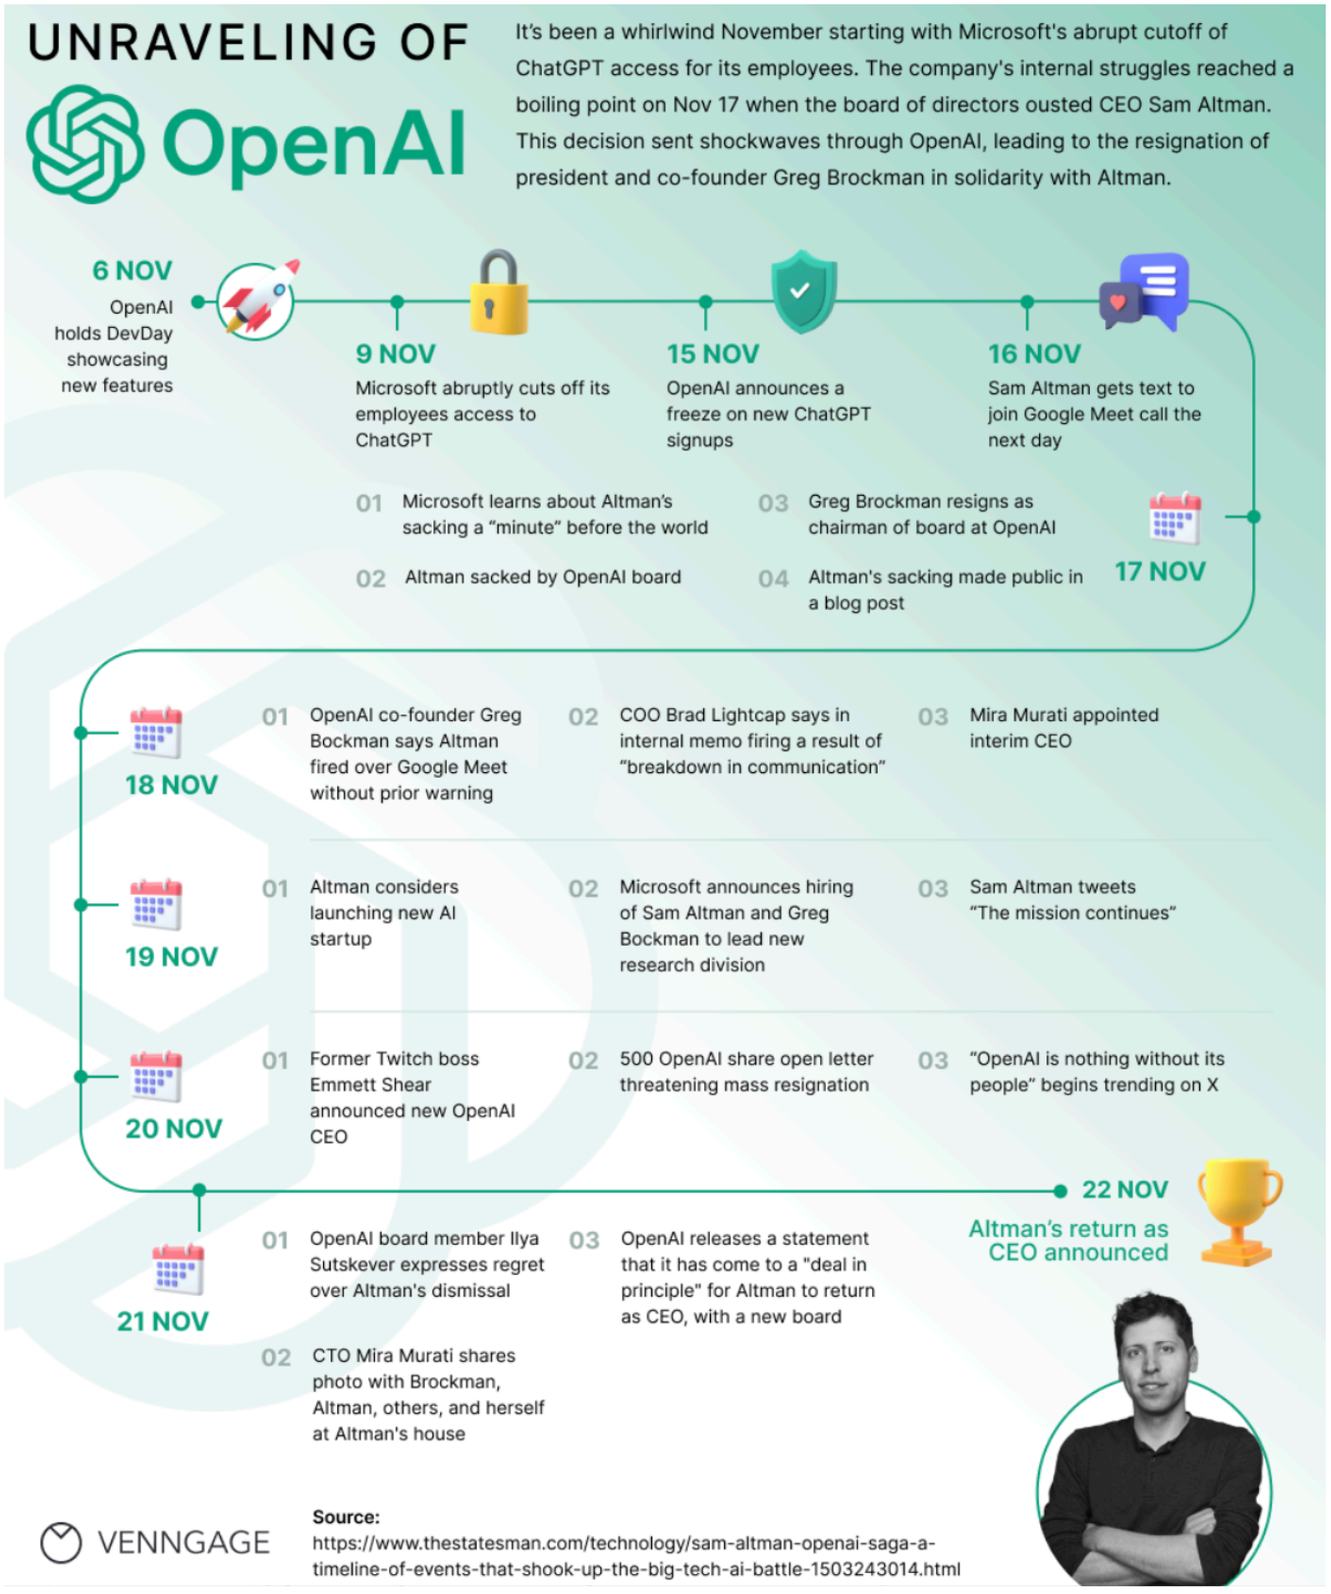

OpenAI Co-Founder Ilya Sutskever's New AI Firm Raises $1 Billion

Sutskever's AI startup, said its investors include Andreessen Horowitz, Sequoia Capital, DST Global and SV Angel

Ilya Sutskever 's new AI firm has raised $1 billion from investors. (PHOTO: AMIR COHEN/REUTERS)

By Kimberley Kao + Follow

Sep 05, 2024 12:26 a.m. ET

OpenAI Co-Founder Ilya Sutskever's New AI Firm Raises $1 Billion

Sutskever's AI startup, said its investors include Andreessen Horowitz, Sequoia Capital, DST Global and SV Angel

Ilya Sutskever's new AI firm has raised $1 billion from investors. (PHOTO: AMIR COHEN/REUTERS)

By Kimberley Kao + Follow

Sep 05, 2024 12:26 a.m. ET

OpenAI co-founder Ilya Sutskever's new firm, aimed at building "safe" artificial intelligence models, has raised $1 billion from investors.

Safe Superintelligence (SSI), Sutskever's AI startup, said in a post on X, formerly Twitter, that its investors include Andreessen Horowitz, Sequoia Capital, DST Global and SV Angel.

NFDG also participated in the fundraising round. NFDG is a venture capital partnership between Nat Friedman, former chief executive of Github, and SSI's co-founder Daniel Gross.

WHAT'S NEWS

OpenAI Rival Anthropic Sets Sights on Enterprise Market

Listen 12 min

SSI was founded in June with a focus on building safe AI models, shortly after Sutskever left OpenAI, where he was chief scientist and led a team focused on developing safety systems to control AI and keep it in line with a set of human values.

OpenAI Co-Founder Ilya Sutskever's New AI Firm Raises $1 Billion

Sutskever's AI startup, said its investors include Andreessen Horowitz, Sequoia Capital, DST Global and SV Angel

Ilya Sutskever's new AI firm has raised $1 billion from investors. (PHOTO: AMIR COHEN/REUTERS)

By Kimberley Kao + Follow

Sep 05, 2024 12:26 a.m. ET

OpenAI co-founder Ilya Sutskever's new firm, aimed at building "safe" artificial intelligence models, has raised $1 billion from investors.

Safe Superintelligence (SSI), Sutskever's AI startup, said in a post on X, formerly Twitter, that its investors include Andreessen Horowitz, Sequoia Capital, DST Global and SV Angel.

NFDG also participated in the fundraising round. NFDG is a venture capital partnership between Nat Friedman, former chief executive of Github, and SSI's co-founder Daniel Gross.

WHAT'S NEWS

OpenAI Rival Anthropic Sets Sights on Enterprise Market

Listen 12 min

SSI was founded in June with a focus on building safe AI models, shortly after Sutskever left OpenAI, where he was chief scientist and led a team focused on developing safety systems to control AI and keep it in line with a set of human values.

December 2 & 5!!

Re. *Blade Runner* …

Re. *Blade Runner* . . .

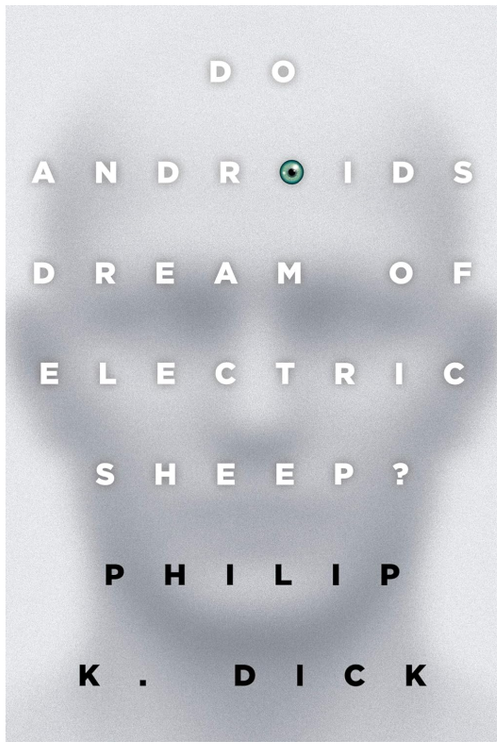

Re. *Blade Runner* . . .

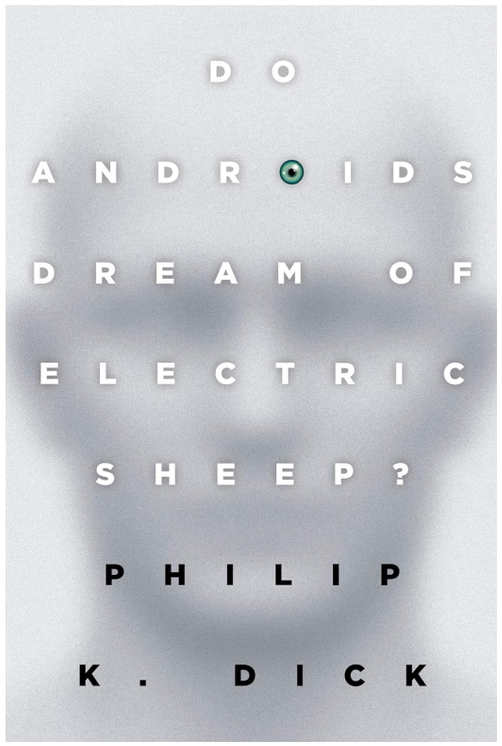

https://www.brmovie.com/Downloads/Docs/BR_Scripts.htm

Re. *Blade Runner* . . .

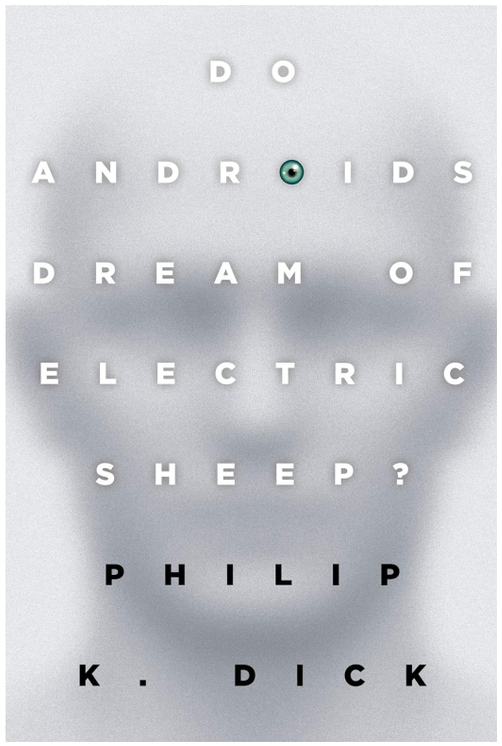

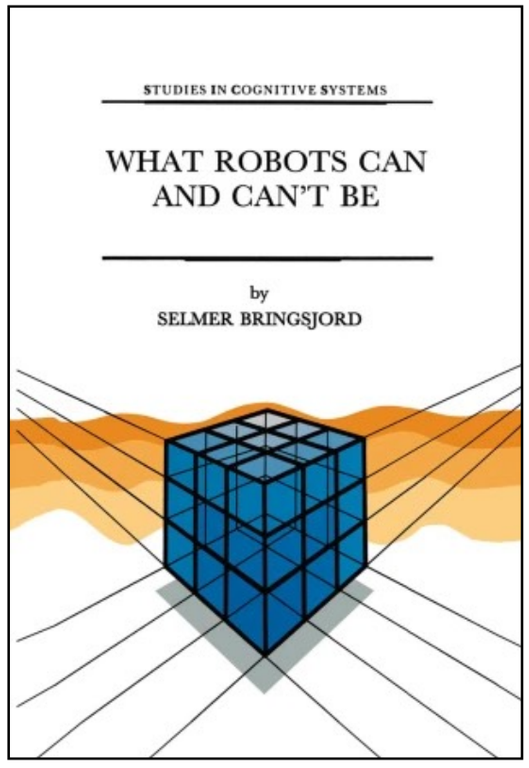

https://www.brmovie.com/Downloads/Docs/BR_Scripts.htm

Re. *Blade Runner* . . .

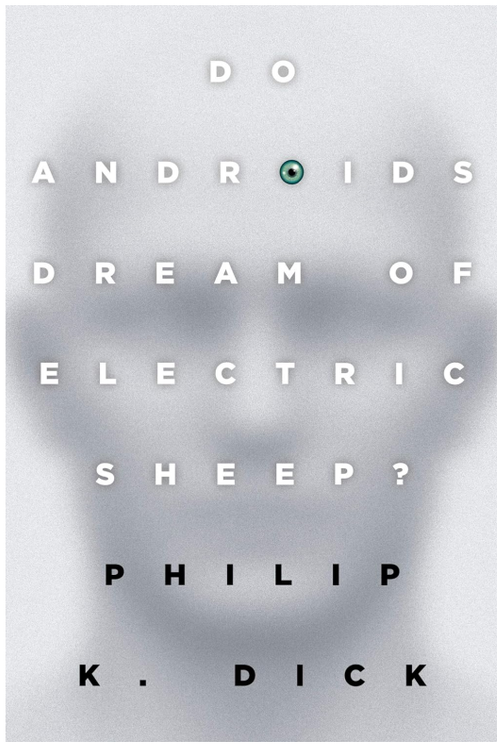

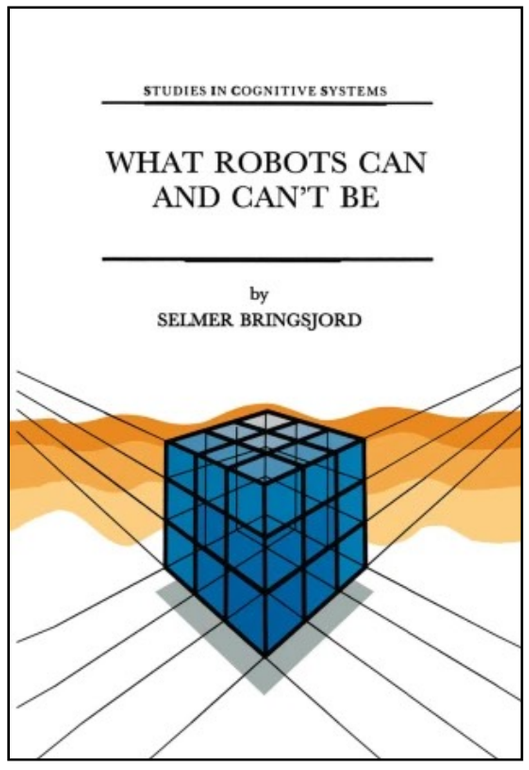

https://www.brmovie.com/Downloads/Docs/BR_Scripts.htm

Re. *Blade Runner* …

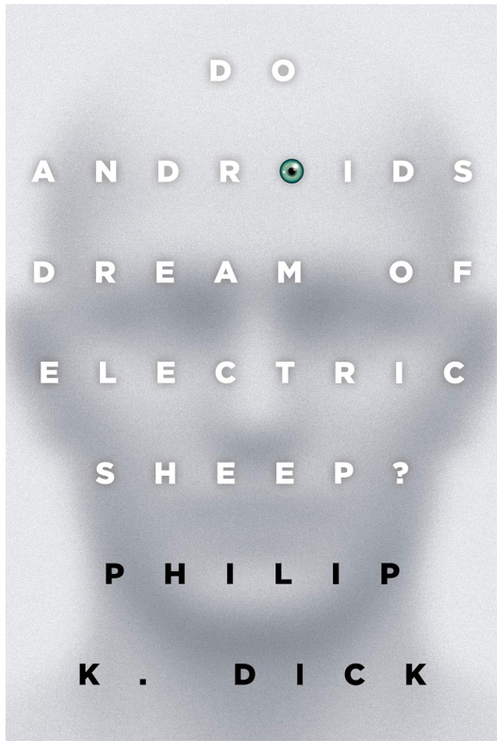

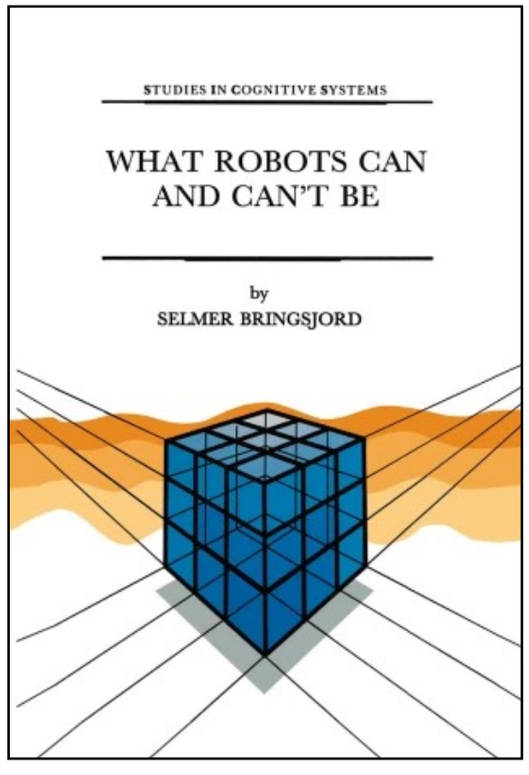

Fake persons

$(t_1 - t_n)$

https://www.brmovie.com/Downloads/Docs/BR_Scripts.htm

Re. *Blade Runner* …

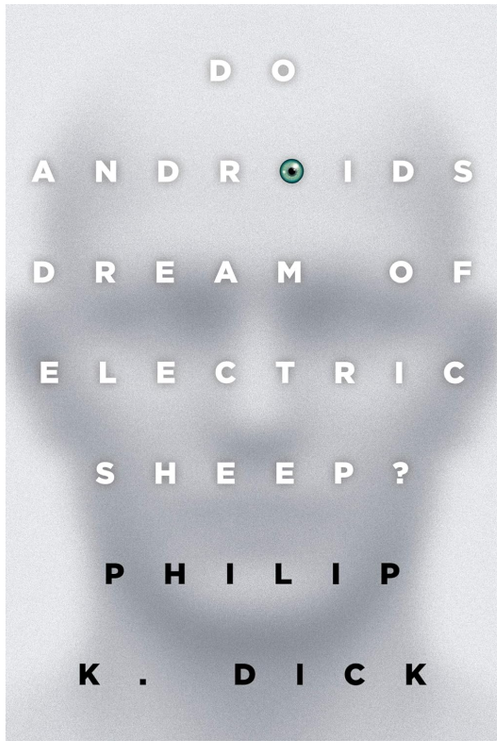

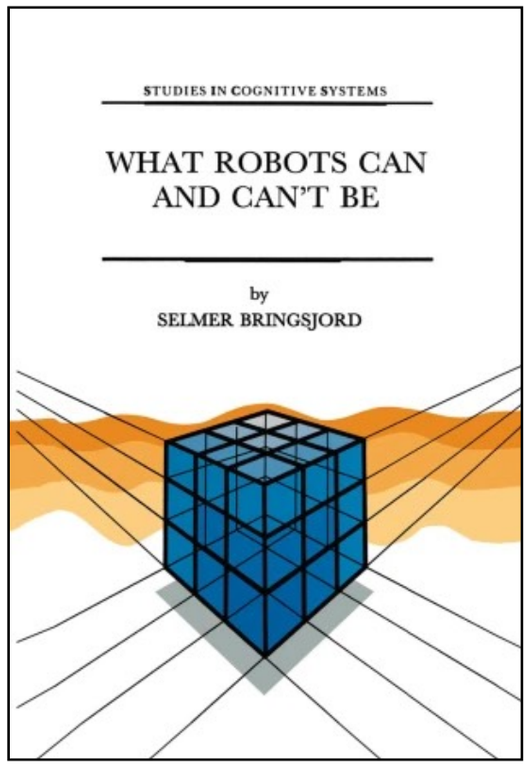

Fake persons

$(t_1 - t_n)$

Persons

https://www.brmovie.com/Downloads/Docs/BR_Scripts.htm

The Future

The Future

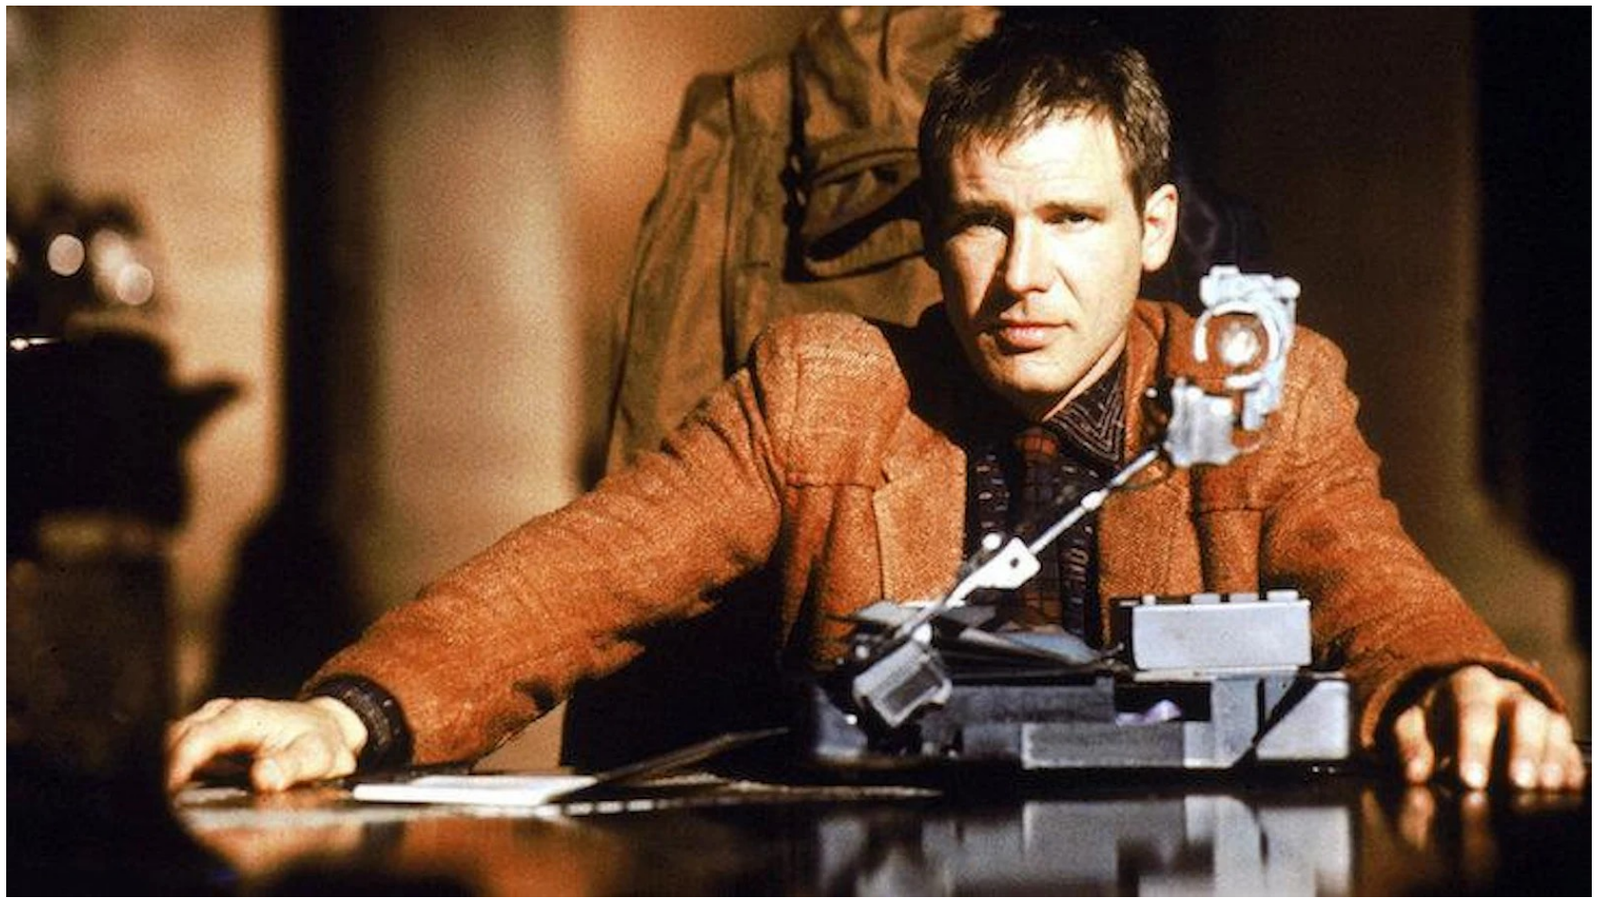

The Future

The Future

What is logic-based (= logicist = logical) AI …

L-BAI is the science and engineering of artificial (intelligent) agents whose computation mediating percepts and actions if reasoning in one or more logics/meta-logics.

L-BAI is the science and engineering of artificial (intelligent) agents whose computation mediating percepts and actions if reasoning in one or more logics/meta-logics.

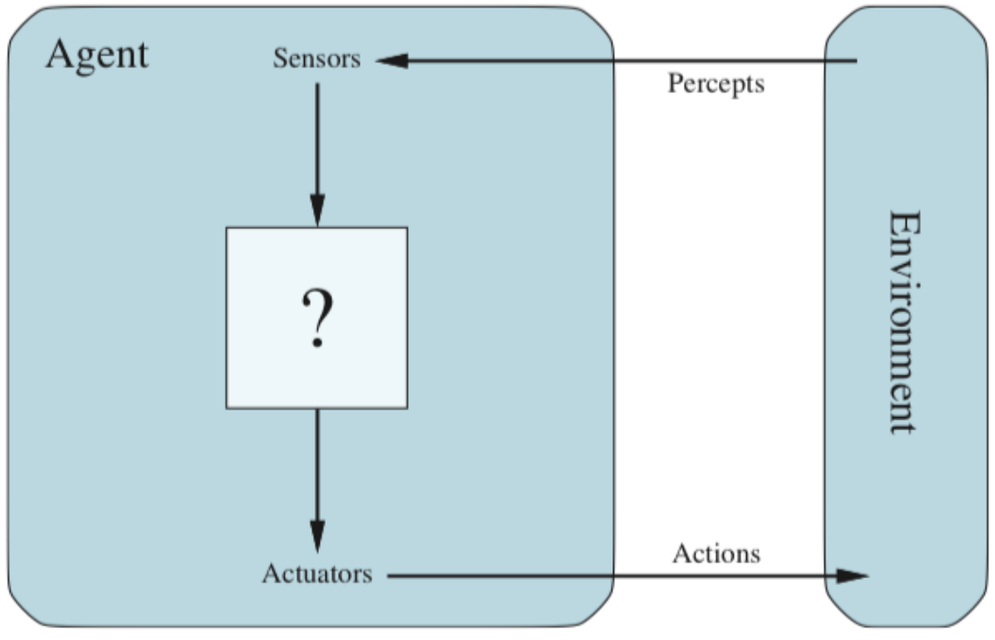

Figure 2.1 Agents interact with environments through sensors and actuators.

L-BAI is the science and engineering of artificial (intelligent) agents whose computation mediating percepts and actions if reasoning in one or more logics/meta-logics.

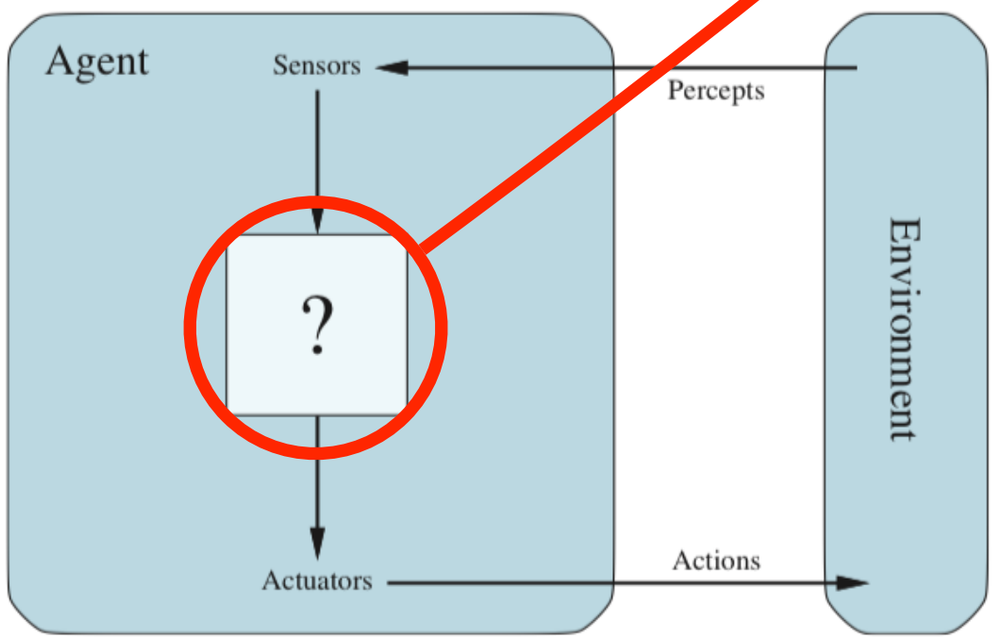

Figure 2.1 Agents interact with environments through sensors and actuators.

And now the whirlwind history …

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1956

Birth of Modern AI

LogicTheorist

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

1956

Birth of Modern AI

LogicTheorist

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

Neural Networks for Logic

McCulloch & Pitts

1956

Birth of Modern AI

LogicTheorist

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

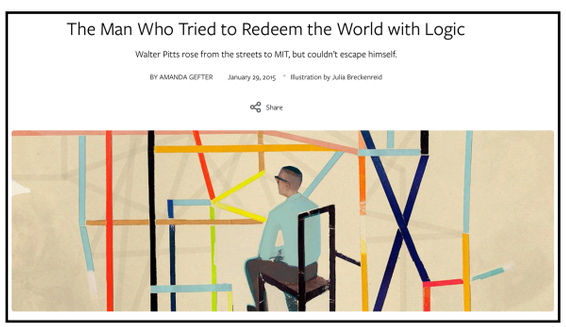
The Man Who Tried to Redeem the World with Logic

Walter Pitts rose from the streets to MIT, but couldn't escape himself.

BY AMANDA GEFTER January 29, 2015 Illustration by Julia Breckenreid

Share

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

Neural Networks for Logic

McCulloch & Pitts

1956

Birth of Modern AI

LogicTheorist

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

1956

Birth of Modern AI

LogicTheorist

Neural Networks for Logic

McCulloch & Pitts

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

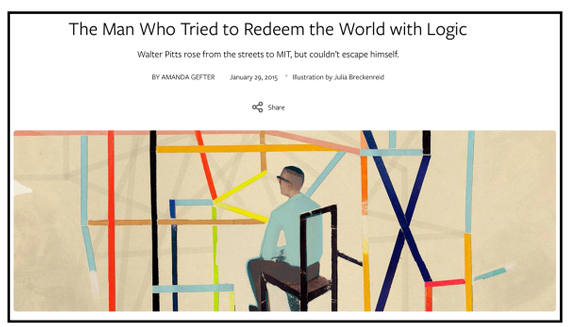

The Man Who Tried to Redeem the World with Logic

Walter Pitts rose from the streets to MIT, but couldn't escape himself.

BY AMANDA GEFTER January 29, 2015 Illustration by Julia Breckenreid

Share

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

1956

Birth of Modern AI

LogicTheorist

Neural Networks for Logic

McCulloch & Pitts

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

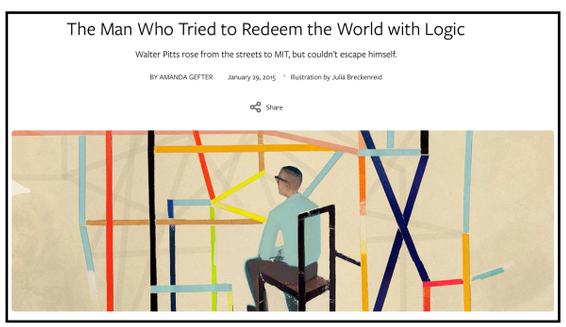
The Man Who Tried to Redeem the World with Logic

Walter Pitts rose from the streets to MIT, but couldn't escape himself.

BY AMANDA GEFTER January 29, 2015 · Illustration by Julia Breckenreid

Share

Unfortunately, AI's Great Divorce …

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

1956

Birth of Modern AI

LogicTheorist

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

Neural Networks for Logic

McCulloch & Pitts

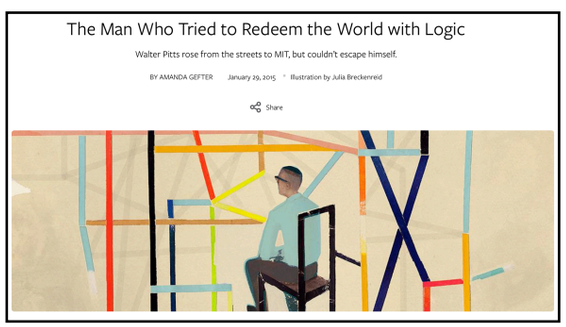

The Man Who Tried to Redeem the World with Logic

Walter Pitts rose from the streets to MIT, but couldn't escape himself.

BY AMANDA GEFTER January 29, 2015 Illustration by Julia Breckenreid

Share

Unfortunately, AI's Great Divorce …

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

1956

Birth of Modern AI

LogicTheorist

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

Neural Networks for Logic

McCulloch & Pitts

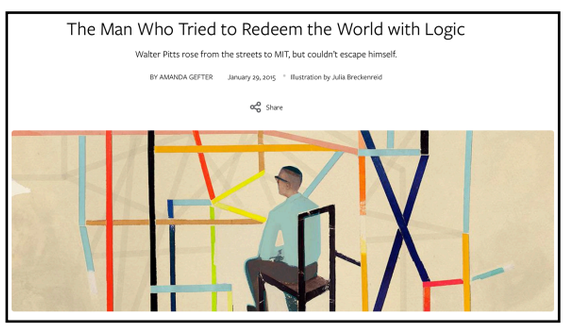
The Man Who Tried to Redeem the World with Logic

Walter Pitts rose from the streets to MIT, but couldn't escape himself.

BY AMANDA GEFTER January 29, 2015 Illustration by Julia Breckenreid

Share

Unfortunately, AI's Great Divorce …

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

1956

Birth of Modern AI

LogicTheorist

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

Neural Networks for Logic

McCulloch & Pitts

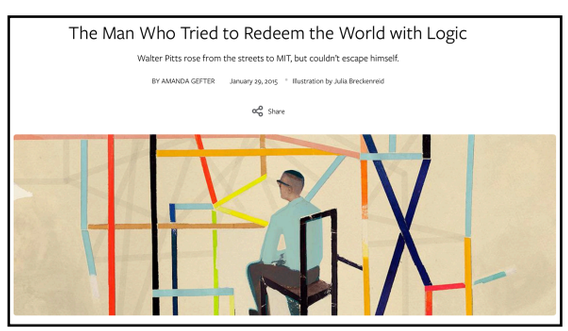
The Man Who Tried to Redeem the World with Logic

Walter Pitts rose from the streets to MIT, but couldn't escape himself.

BY AMANDA GEFTER January 29, 2015 Illustration by Julia Breckenreid

Share

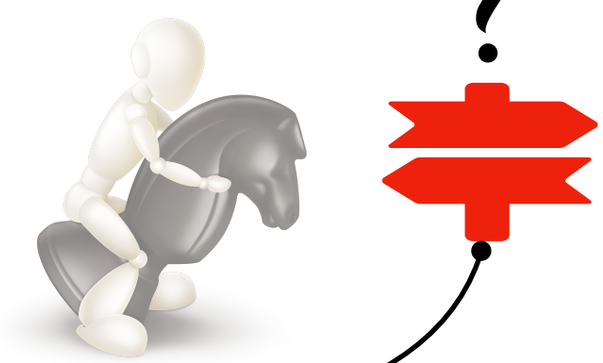

Unfortunately, AI's Great Divorce …

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

1956

Birth of Modern AI

LogicTheorist

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

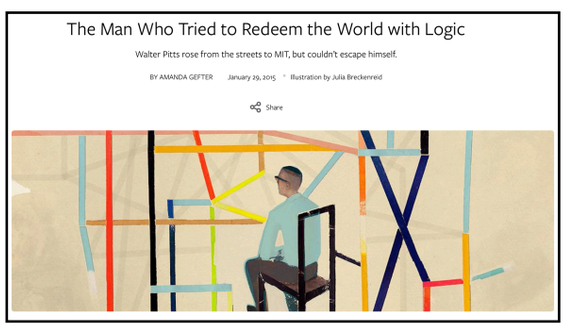
The Man Who Tried to Redeem the World with Logic

Walter Pitts rose from the streets to MIT, but couldn't escape himself.

BY AMANDA GEFTER January 29, 2015 Illustration by Julia Breckenreid

Share

Neural Networks for Logic

McCulloch & Pitts

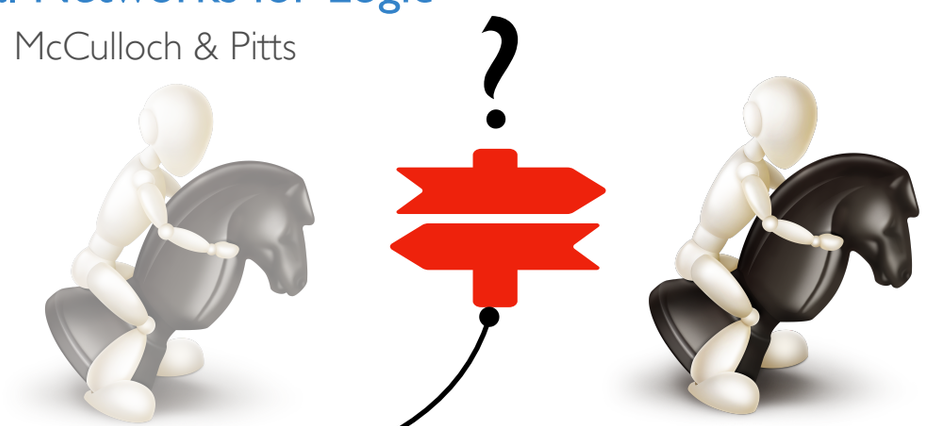

Unfortunately, AI's Great Divorce …

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

1956

Birth of Modern AI

LogicTheorist

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

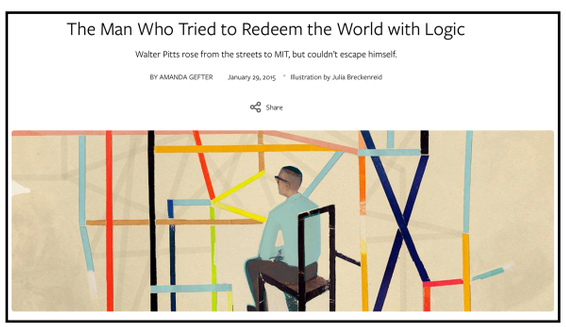
The Man Who Tried to Redeem the World with Logic

Walter Pitts rose from the streets to MIT, but couldn't escape himself.

BY AMANDA GEFTER January 29, 2015 Illustration by Julia Breckenreid

Share

Neural Networks for Logic

McCulloch & Pitts

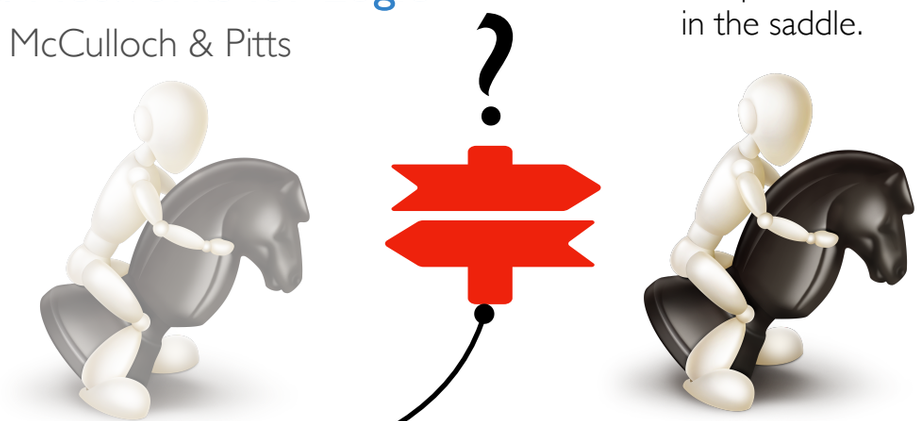

Deep Learner in the saddle.

Unfortunately, AI's Great Divorce …

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

1956

Birth of Modern AI

LogicTheorist

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

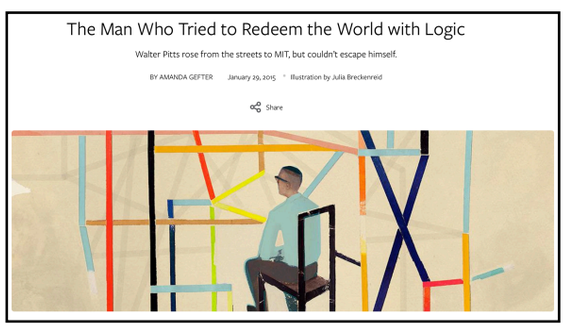
The Man Who Tried to Redeem the World with Logic

Walter Pitts rose from the streets to MIT, but couldn't escape himself.

BY AMANDA GEFTER January 29, 2015 Illustration by Julia Breckenreid

Share

Neural Networks for Logic

McCulloch & Pitts

?

Deep Learner in the saddle.

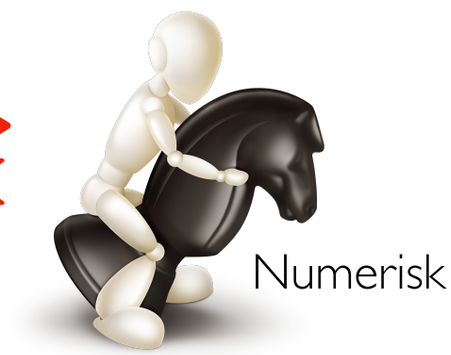

Unfortunately, AI's Great Divorce …

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

Birth of Modern AI

LogicTheorist

1956

Neural Networks for Logic

McCulloch & Pitts

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

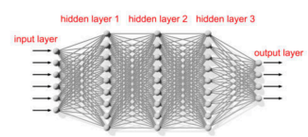

Deep Learner in the saddle.

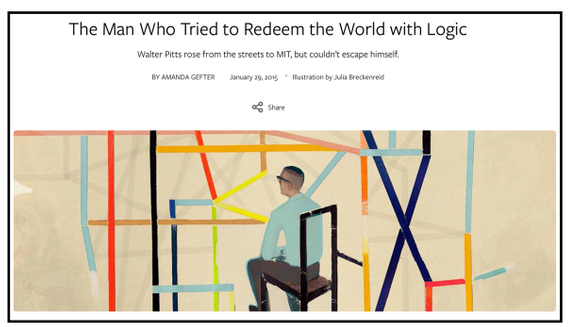

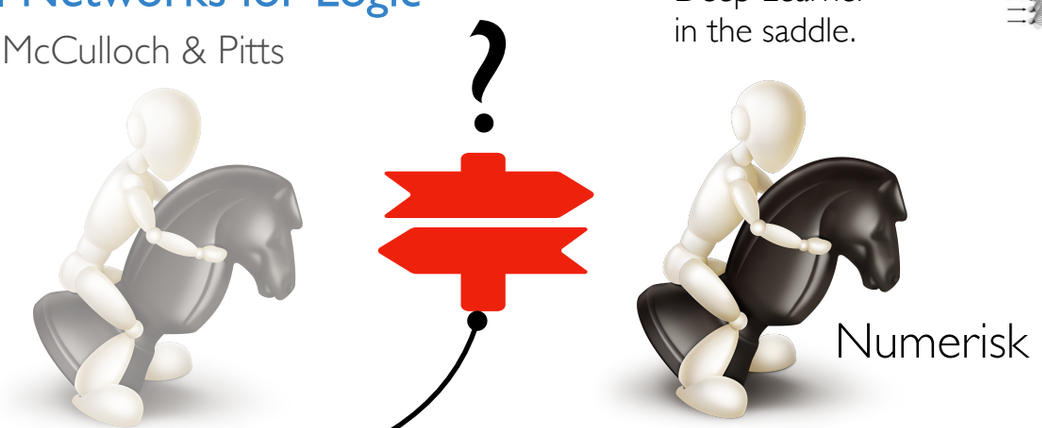

Unfortunately, AI's Great Divorce …

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\mathscr{L}_1$ discovered, and modal logic as well.

1943

Birth of Modern AI

LogicTheorist

1956

Neural Networks for Logic

McCulloch & Pitts

2024

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

"Danger, Will Robinson!"

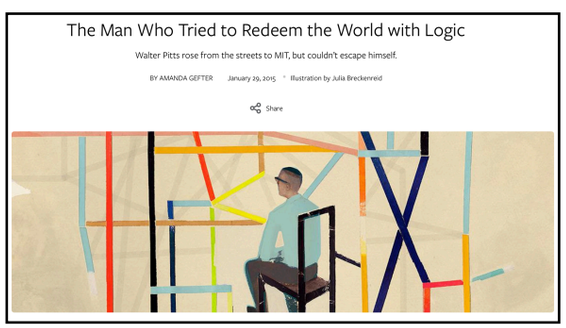

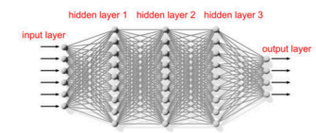

Deep Learner in the saddle.

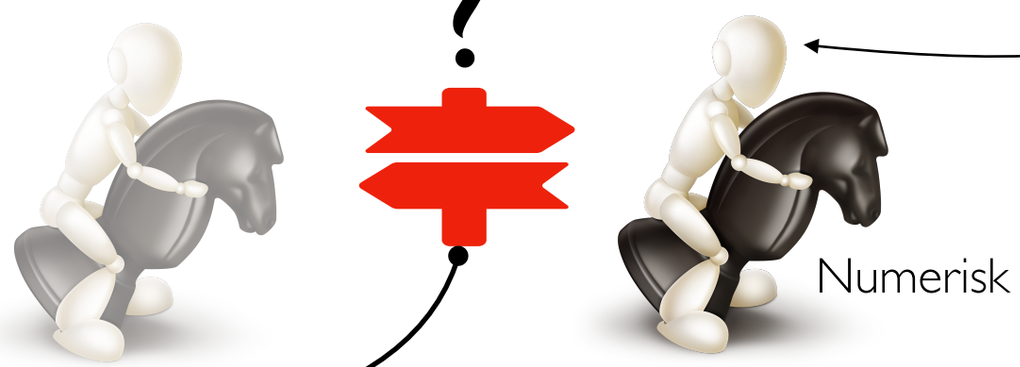

Numerisk

Unfortunately, AI's Great Divorce …

325 BC

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

C. 1700

Leibniz

First-order logic = $\vdash \mathscr{L}_1$

$\mathscr{L}_1$ discovered, and modal logic as well.

1936

1943

Birth of Modern AI

LogicTheorist

1956

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

2024

"Danger, Will Robinson!"

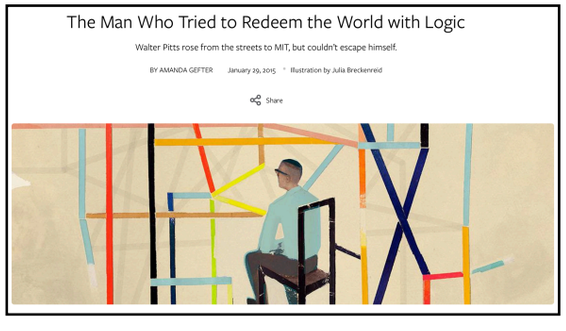

Neural Networks for Logic

McCulloch & Pitts

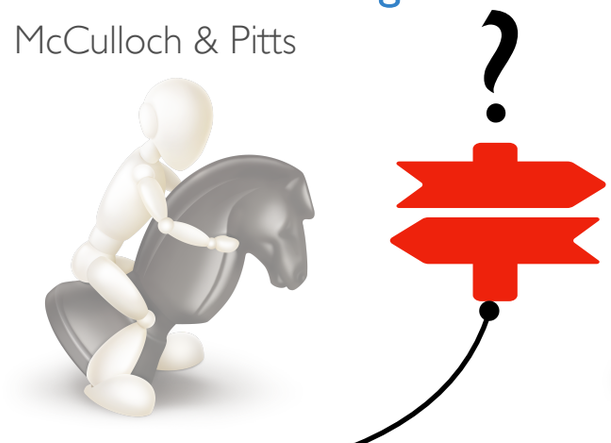

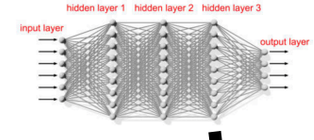

Deep Learner in the saddle.

Numerisk

Unfortunately, AI's Great Divorce …

325 BC

C. 1700

1943

1956

20006

2024

"Danger, Will Robinson!"

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

Leibniz

1936

First-order logic = $\vdash \mathscr{L}_1$

$\mathscr{L}_1$ discovered, and modal logic as well.

Birth of Modern AI

LogicTheorist

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

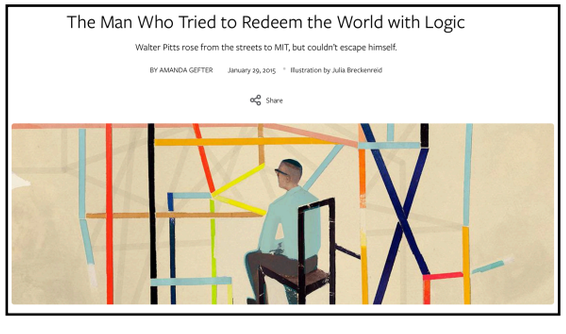

Neural Networks for Logic

McCulloch & Pitts

?

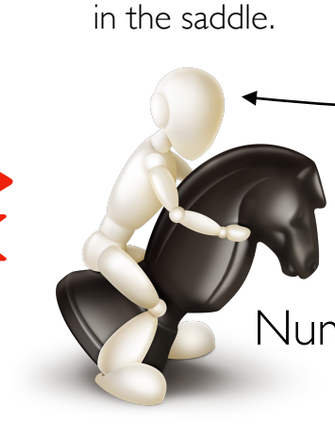

Numerisk

Deep Learner in the saddle.

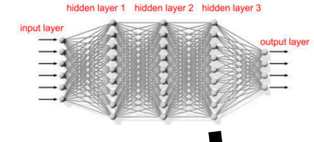

Unfortunately, AI's Great Divorce …

AI@50





Logisk

325 BC

C. 1700

1943

1956

20006

2024

"Danger, Will Robinson!"

Aristotle

"Wow Euclid, humans are really smart!"

A fragment of first-order logic = $\mathscr{L}_1$ introduced.

Leibniz

First-order logic = $\mathscr{L}_1$

$\mathscr{L}_1$ discovered, and modal logic as well.

1936 $\vdash$

Birth of Modern AI

LogicTheorist

OLCSU

Only Logic Can Save Us, i.e. only Logisk can save us.

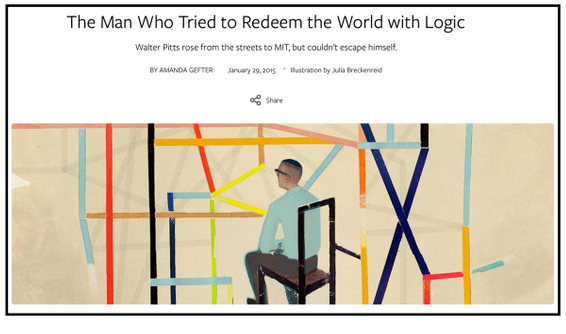

Neural Networks for Logic

McCulloch & Pitts

?

Deep Learner in the saddle.

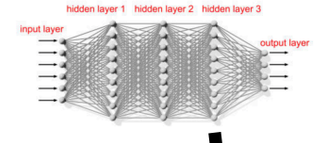

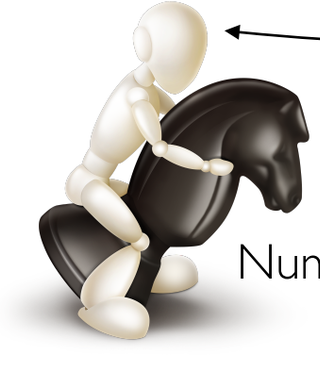

Numerisk

Unfortunately, AI's Great Divorce …

2024

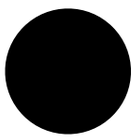

2024

Intro to Logic-Based AI @ RPI

$\mathcal{DCEC}^*$

Syntax

$$
\begin{aligned}
S ::= {} & \text{Object} \mid \text{Agent} \mid \text{Self} \sqsubset \text{Agent} \mid \text{ActionType} \mid \text{Action} \sqsubseteq \text{Event} \mid \\
& \text{Moment} \mid \text{Boolean} \mid \text{Fluent} \mid \text{Numeric} \\
f ::= {} & \mathit{action} : \text{Agent} \times \text{ActionType} \rightarrow \text{Action} \\
& \mathit{initially} : \text{Fluent} \rightarrow \text{Boolean} \\
& \mathit{holds} : \text{Fluent} \times \text{Moment} \rightarrow \text{Boolean} \\
& \mathit{happens} : \text{Event} \times \text{Moment} \rightarrow \text{Boolean} \\
& \mathit{clipped} : \text{Moment} \times \text{Fluent} \times \text{Moment} \rightarrow \mathit{Boolean} \\
& \mathit{initiates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Boolean} \\
& \mathit{terminates} : \text{Event} \times \text{Fluent} \times \text{Moment} \rightarrow \text{Boolean} \\
& \mathit{prior} : \text{Moment} \times \text{Moment} \rightarrow \text{Boolean} \\
& \mathit{interval} : \text{Moment} \times \text{Boolean} \\
& * : \text{Agent} \rightarrow \text{Self} \\
& \mathit{payoff} : \text{Agent} \times \text{ActionType} \times \text{Moment} \rightarrow \text{Numeric} \\
t ::= {} & x : S \mid c : S \mid f(t_1, \ldots, t_n) \\
\phi ::= {} & t : \text{Boolean} \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \\
& \mathbf{P}(a,t,\phi) \mid \mathbf{K}(a,t,\phi) \mid \mathbf{C}(t,\phi) \mid \mathbf{S}(a,b,t,\phi) \mid \mathbf{S}(a,t,\phi) \\
& \mathbf{B}(a,t,\phi) \mid \mathbf{D}(a,t,holds(f,t')) \mid \mathbf{I}(a,t,happens(action(a^*,\alpha),t')) \\
& \mathbf{O}(a,t,\phi,happens(action(a^*,\alpha),t'))
\end{aligned}
$$

Rules of Inference

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

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

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

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

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

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

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

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

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

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

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

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

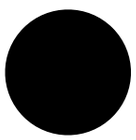

2024

Intro to Logic-Based AI @ RPI

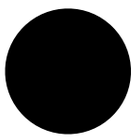

2024

Intro to Logic-Based AI @ RPI

2024

Intro to Logic-Based AI @ RPI

2024

Intro to Logic-Based AI @ RPI

350 BC

2024

Intro to Logic-Based AI @ RPI

350 BC

Euclid

2024

350 BC

2024

Euclid

350 BC

300 BC

2024

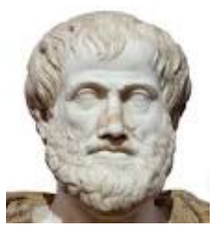

Euclid

350 BC

300 BC

2024

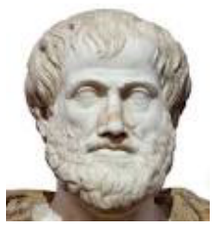

I don’t believe in magic! Why exactly is that so convincing? What exactly is he doing?!?

Euclid

350 BC

300 BC

2024

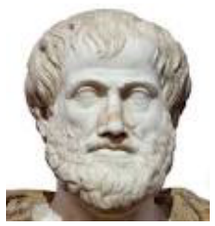

Euclid

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

2024

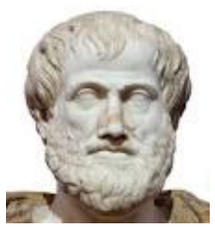

I don't believe in magic! Why exactly is that so convincing? What exactly is he doing?!?

Euclid

Organon

He's using syllogisms!

E.g.,

All As are Bs.
All Bs are Cs.

All As are Cs.

350 BC

Euclid

300 BC

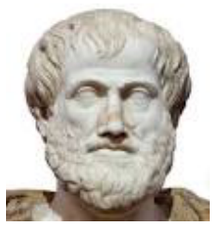

Organon

2024

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

Balderdash!

He's using syllogisms!

E.g.,

All As are Bs.
All Bs are Cs.

All As are Cs.

350 BC

Euclid

300 BC

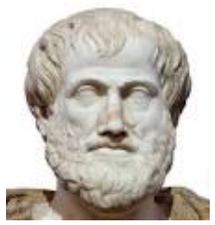

Organon

2024

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

Balderdash!

350 BC

Euclid

300 BC

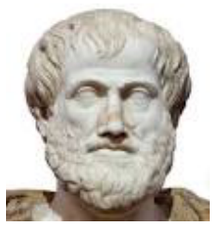

Organon

2024

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

350 BC

300 BC

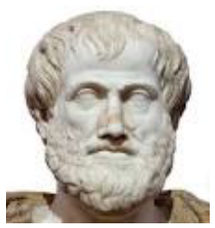

Euclid

Organon

350 BC

300 BC

1666

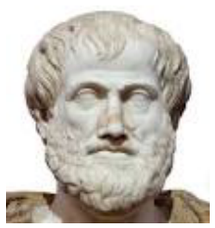

Euclid

Organon

350 BC

Euclid

300 BC

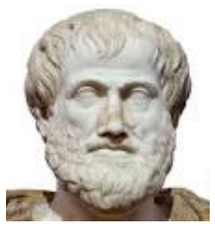

Organon

1666

Leibniz

350 BC

Euclid

300 BC

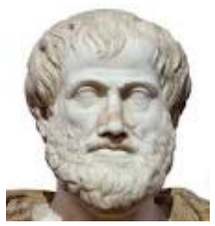

Organon

1666

Leibniz

$$\int$$

"Universal Computational Logic"

350 BC

Euclid

300 BC

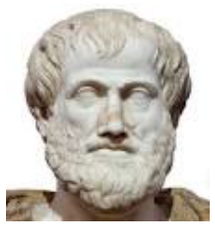

Organon

1666

Leibniz

$$\int$$

"Universal Computational Logic"

350 BC

Euclid

300 BC

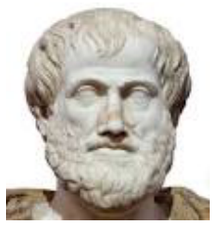

Organon

1666

Leibniz

$$\int$$

1854

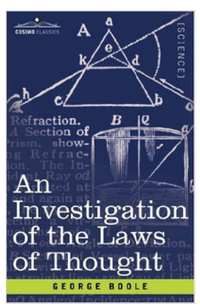

"Universal Computational Logic"

350 BC

Euclid

300 BC

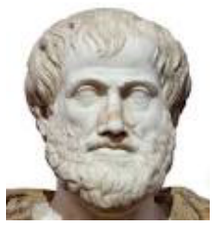

Organon

1666

Leibniz

$$\int$$

1854

“Universal Computational Logic”

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

Organon

1666

Leibniz

$$\int$$

1854

1956

Simon

Intro to Logic-Based AI @ RPI

“Astonishing” Logic Theorist Proof @ Dawn of AI

"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

"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	$(\phi \vee \phi) \rightarrow \phi$	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®).

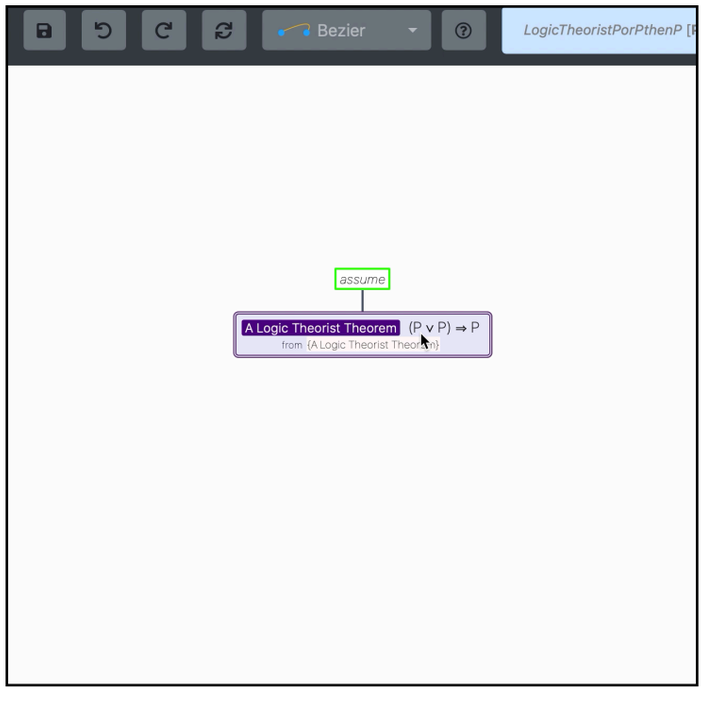

“Astonishing” Logic Theorist Proof @ Dawn of AI

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®).

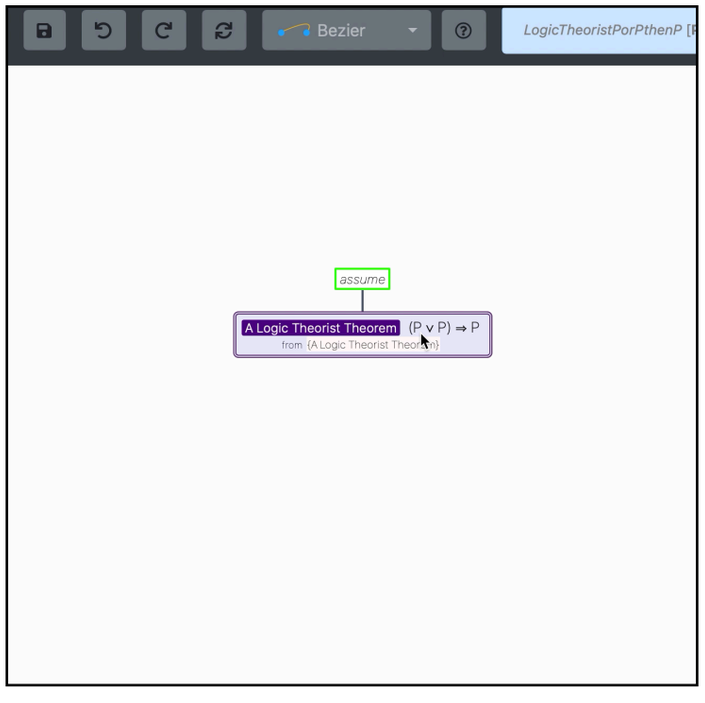

"Universal Computational Logic"

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

Organon

Leibniz

$\int$

1854

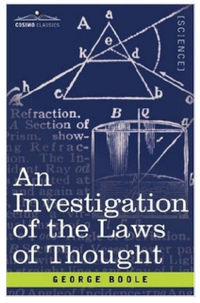

1956

Simon

2024

Intro to Logic-Based AI @ RPI

“Universal Computational Logic”

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

Organon

1666

Leibniz

$\int$

1854

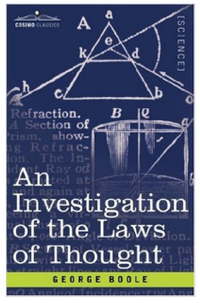

1956

Simon

2024

Intro to Logic-Based 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 Logic-Based 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

2025

Intro to Logic-Based 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 Logic-Based AI @ RPI

2025

Thesingularity?

Entscheidungsproblem

"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 Logic-Based AI @ RPI

2025

Thesingularity?

?

Entscheidungsproblem

"Universal Computational Logic"

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

Organon

1666

Leibniz

$\int$

1854

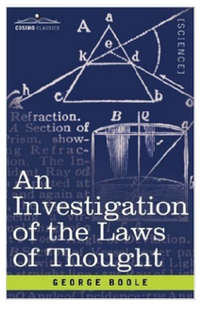

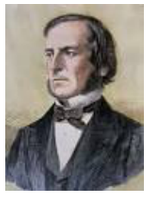

1956

Simon

2023

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

2024

Thesingularity?

?

Entscheidungsproblem

"Universal Computational Logic"

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

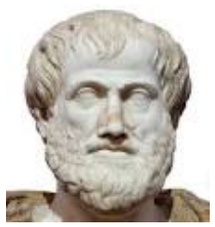

Organon

1666

Leibniz

$\int$

1854

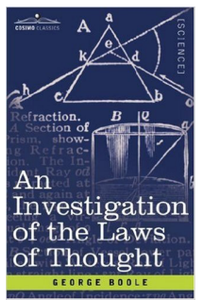

1956

Simon

2023

2024

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

Thesingularity?

?

Entscheidungsproblem

"Universal Computational Logic"

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

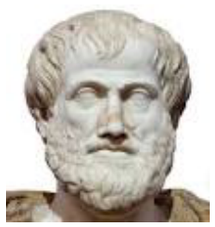

Organon

1666

Leibniz

$\int$

1854

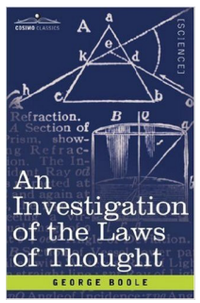

1956

Simon

2023

2024

Frege

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

Thesingularity?

?

Entscheidungsproblem

"Universal Computational Logic"

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

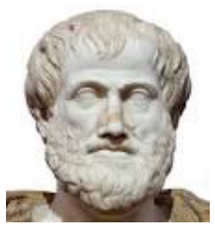

Organon

1666

Leibniz

$\int$

1854

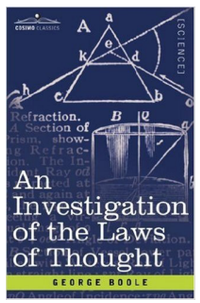

1956

Simon

2023

2024

Frege

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

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

Thesingularity?

?

Entscheidungsproblem

"Universal Computational Logic"

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

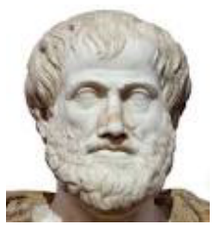

Organon

1666

Leibniz

$\int$

1854

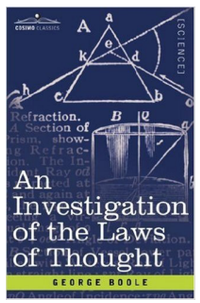

1956

Simon

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

2023

2024

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

Thesingularity?

?

Entscheidungsproblem

"Universal Computational Logic"

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

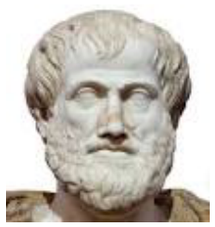

Organon

1666

Leibniz

$\int$

1854

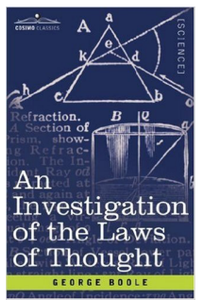

1956

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

2023

2024

The Singularity?

?

Entscheidungsproblem

"Universal Computational Logic"

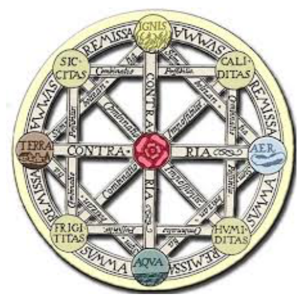

Logic Theorist

(birth of modern logicist AI)

350 BC — 300 BC — 1666 — 1854 — 1956 — 2023 — 2024

Euclid

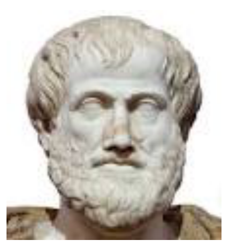

Organon

Leibniz

$\int$

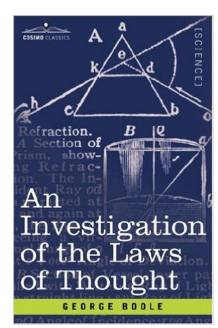

Simon

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

Turing

Post

The Singularity?

?

Entscheidungsproblem

"Universal Computational Logic"

Logic Theorist
(birth of modern logicist AI)

350 BC

Euclid

300 BC

Organon

1666

Leibniz

$\int$

1854

1956

Simon

2023

2024

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

Post

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

Thesingularity?

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 (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

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 ($\text{THEOREM}_{\text{PC}}$) for the Propositional Calculus

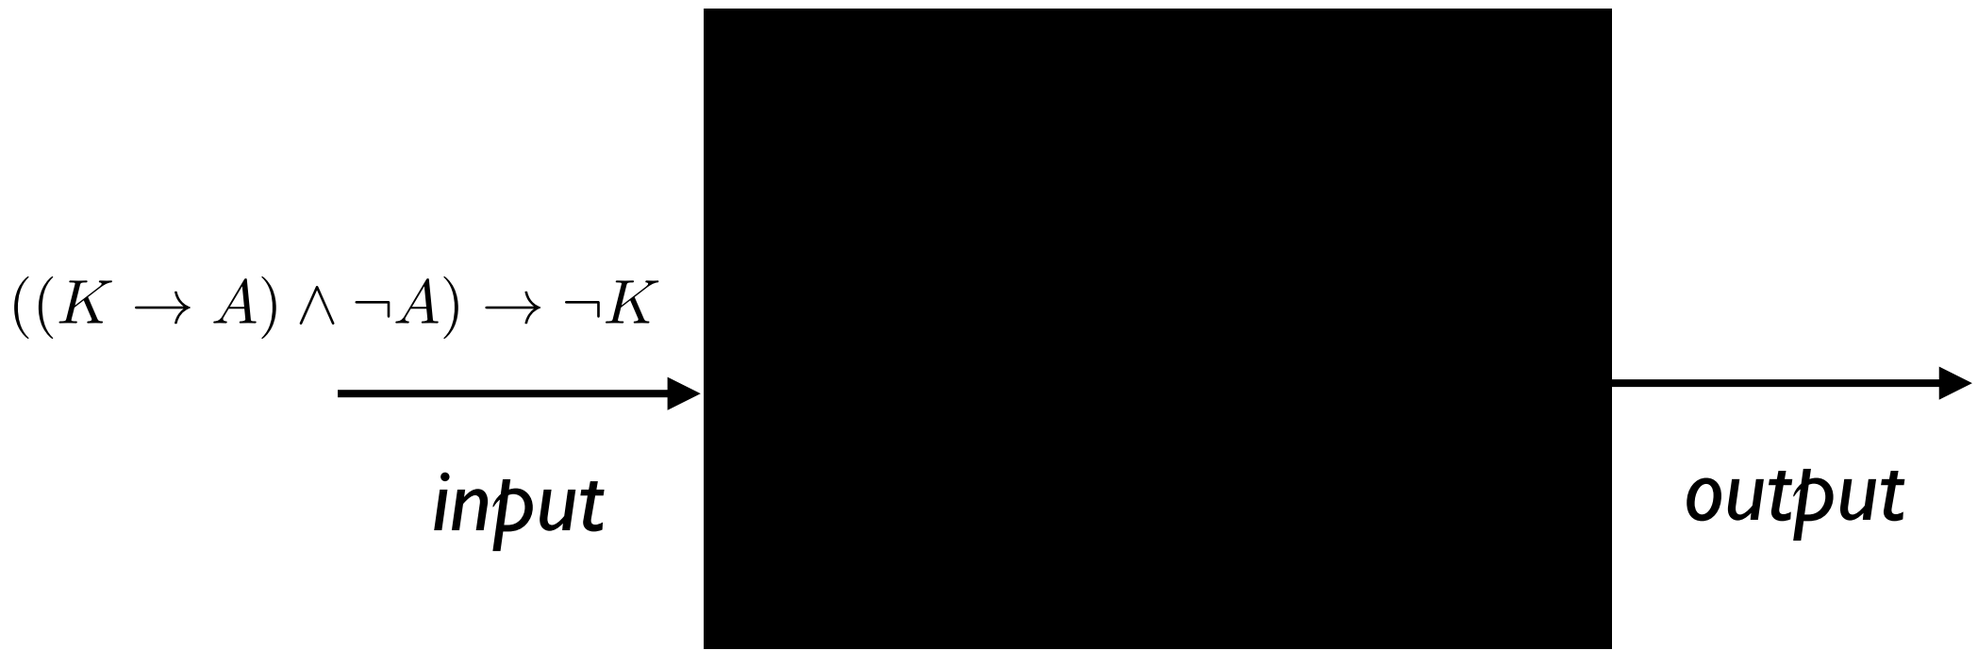

Hard!! — for apparently no polynomial-time algorithm for this!

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

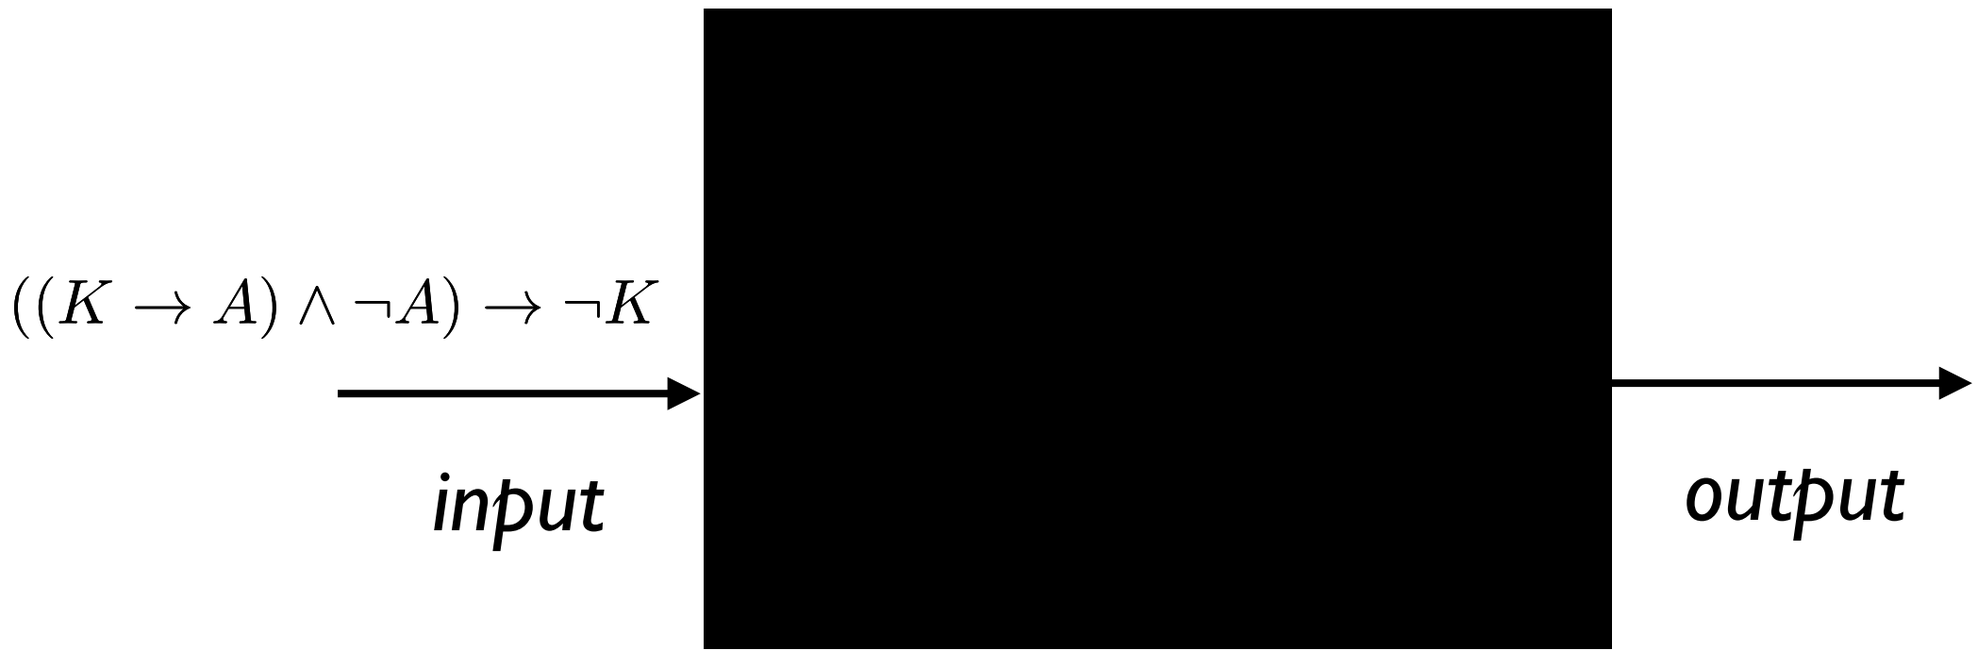

Hard!! — for apparently no polynomial-time algorithm for this!

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

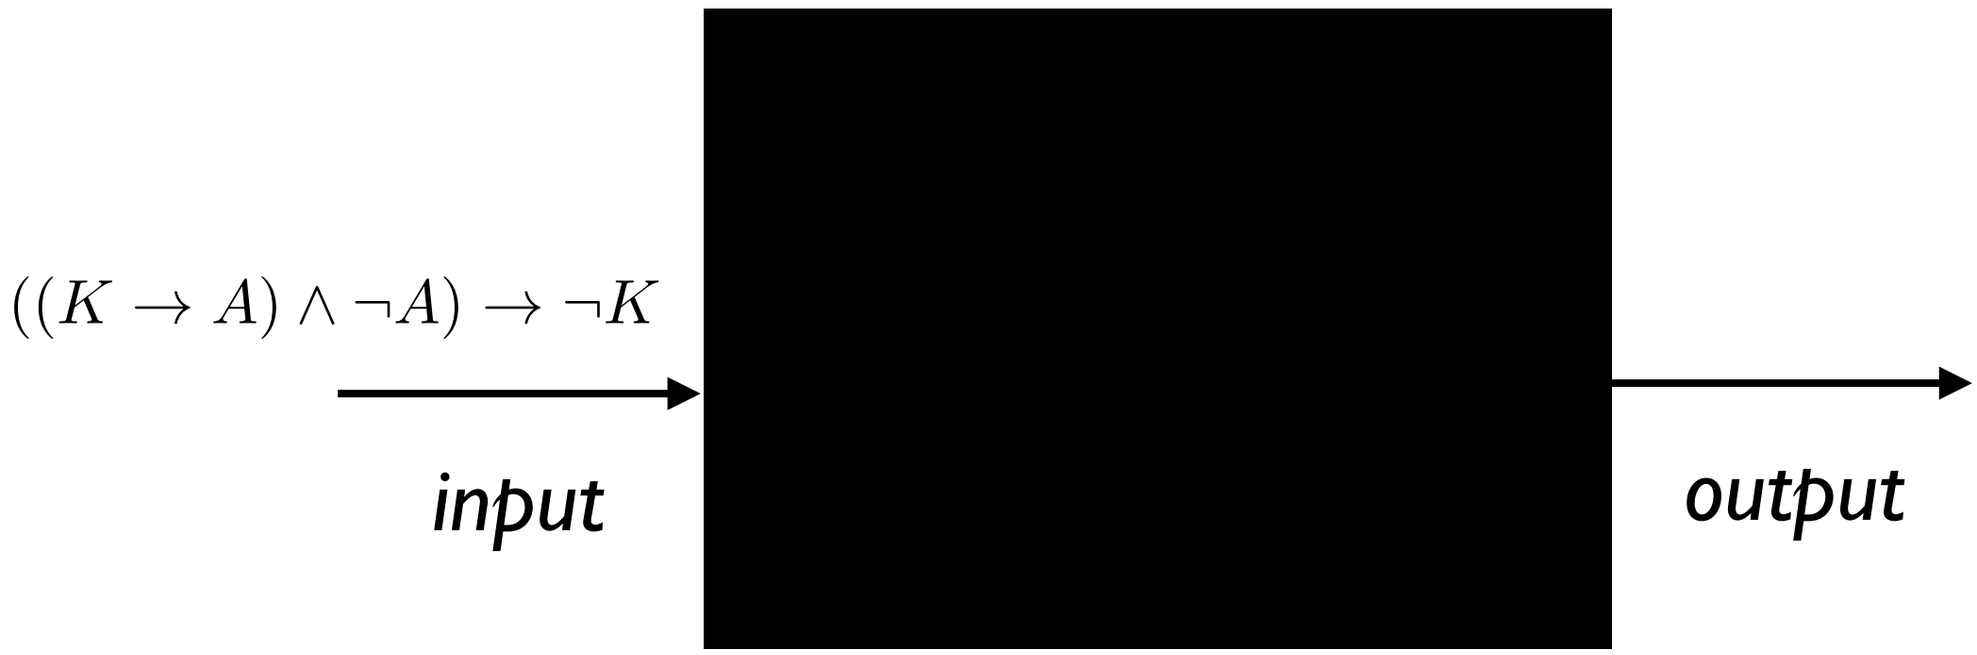

Hard!! — for apparently no polynomial-time algorithm for this!

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

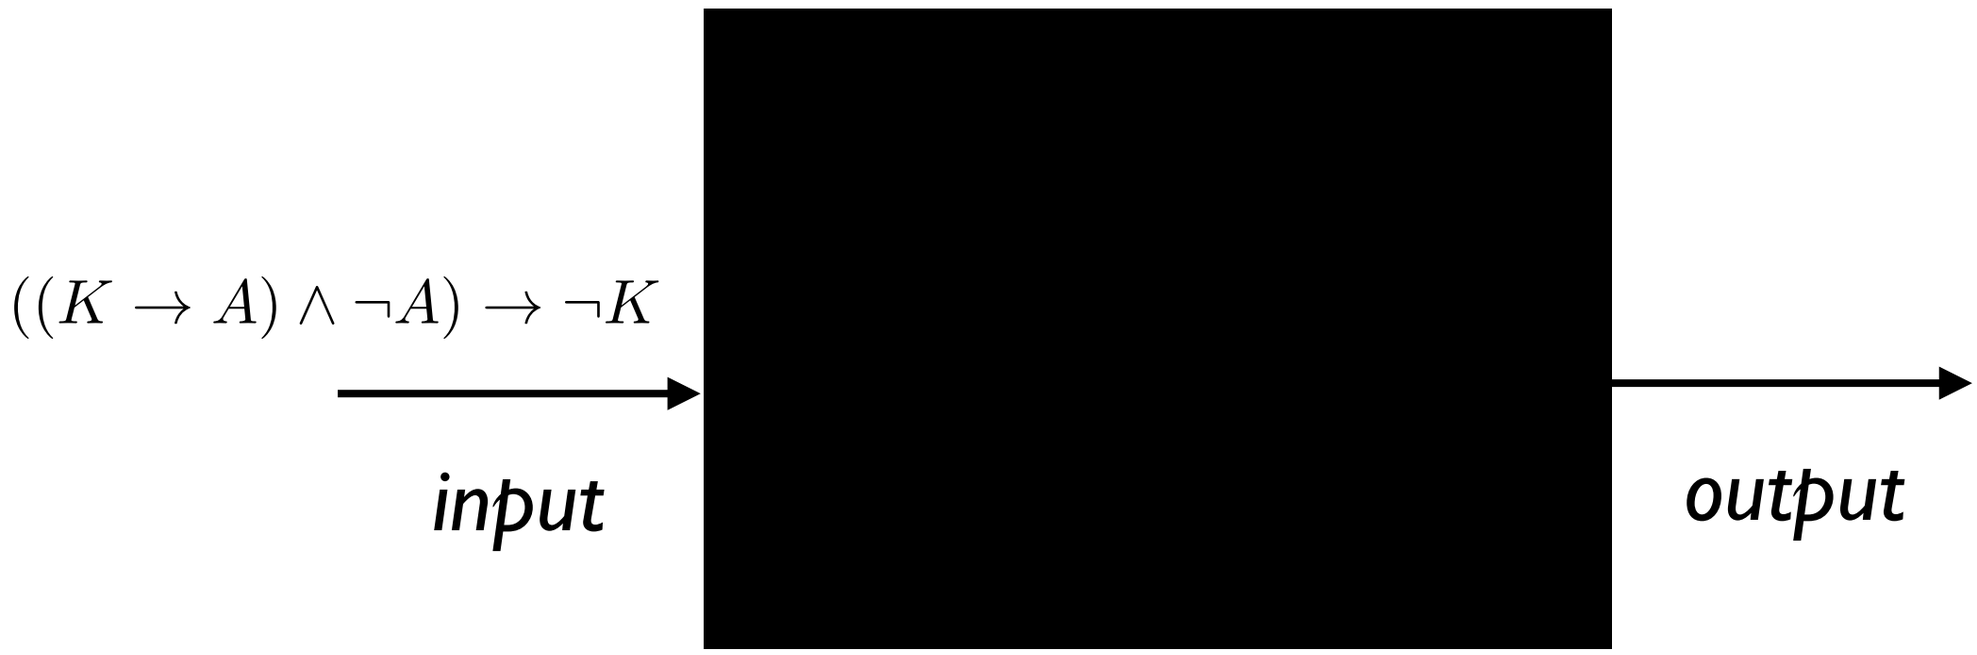

Hard!! — for apparently no polynomial-time algorithm for this!

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

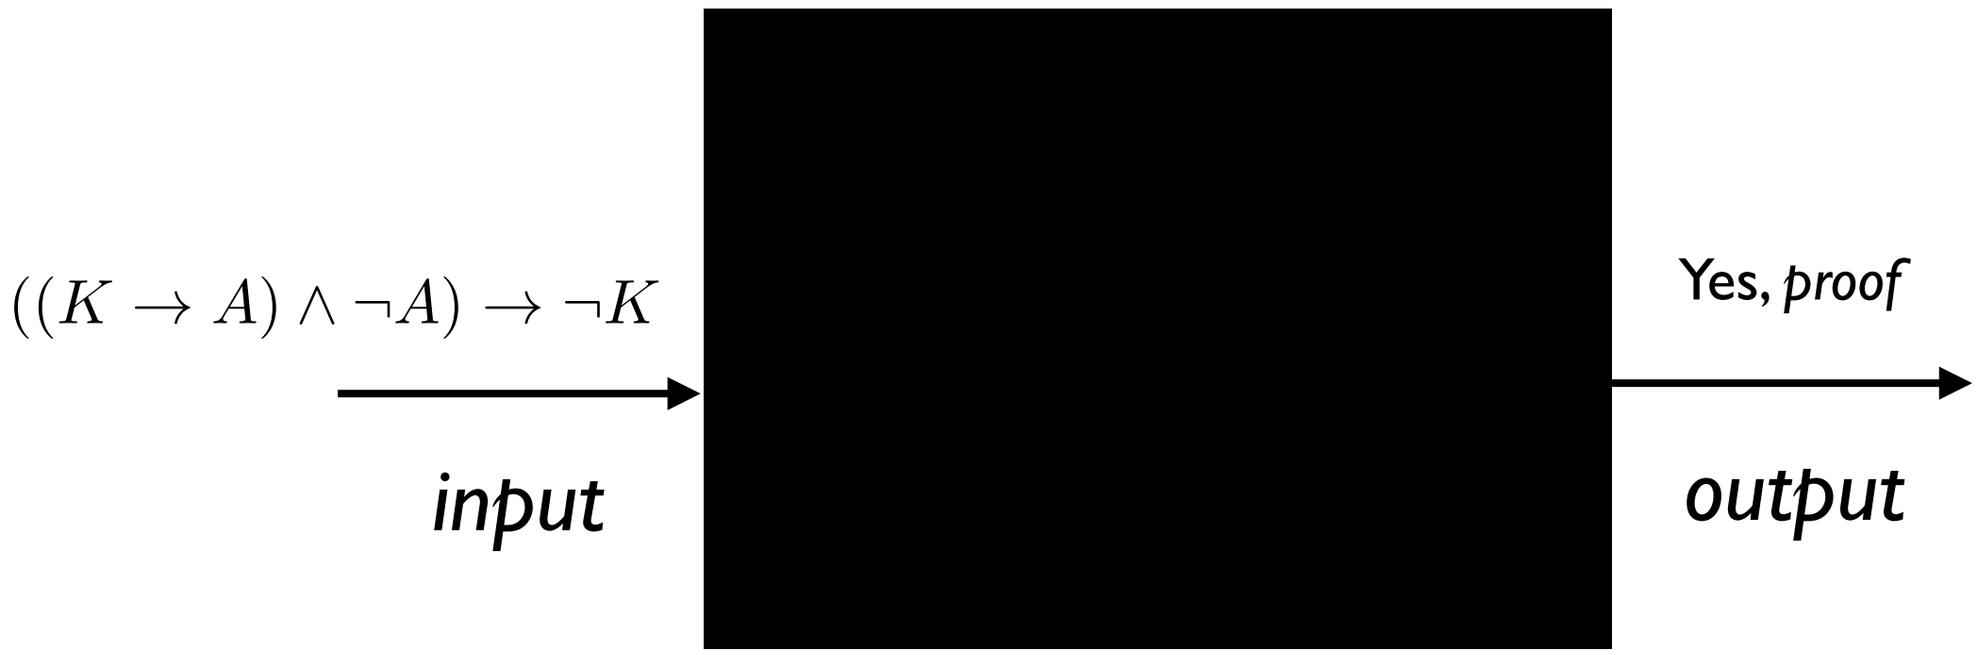

Hard!! — for apparently no polynomial-time algorithm for this!

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*,
($\text{THEOREM}_{\text{FOL}}$)
for First-Order Logic (FOL)

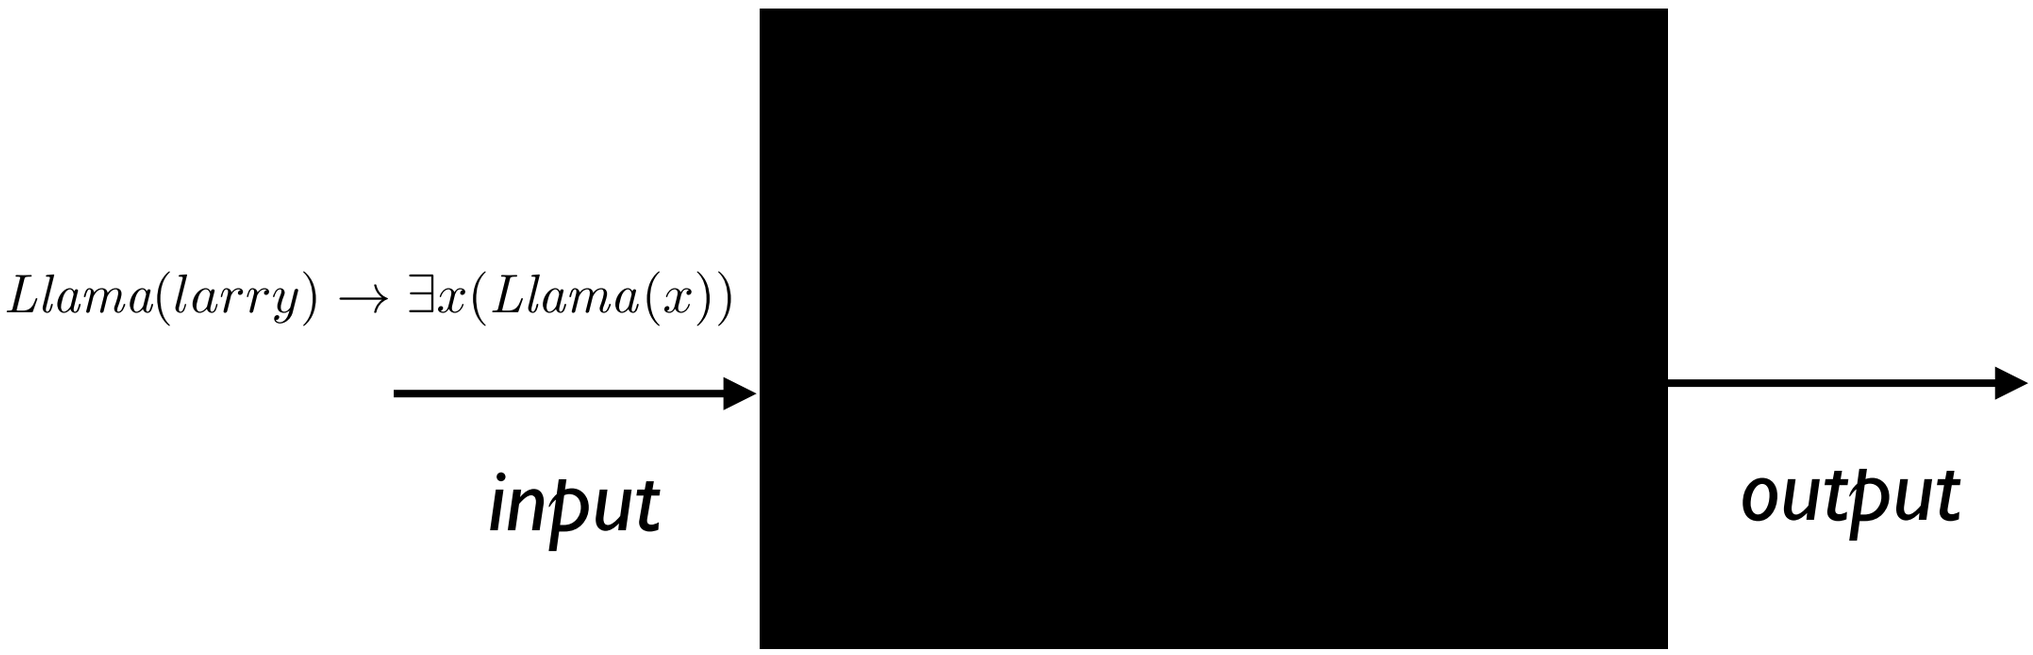

And now, the Theoremhood Decision Problem, i.e., the *Entscheidungsproblem*, ($\text{THEOREM}_{\text{FOL}}$) for First-Order Logic (FOL)

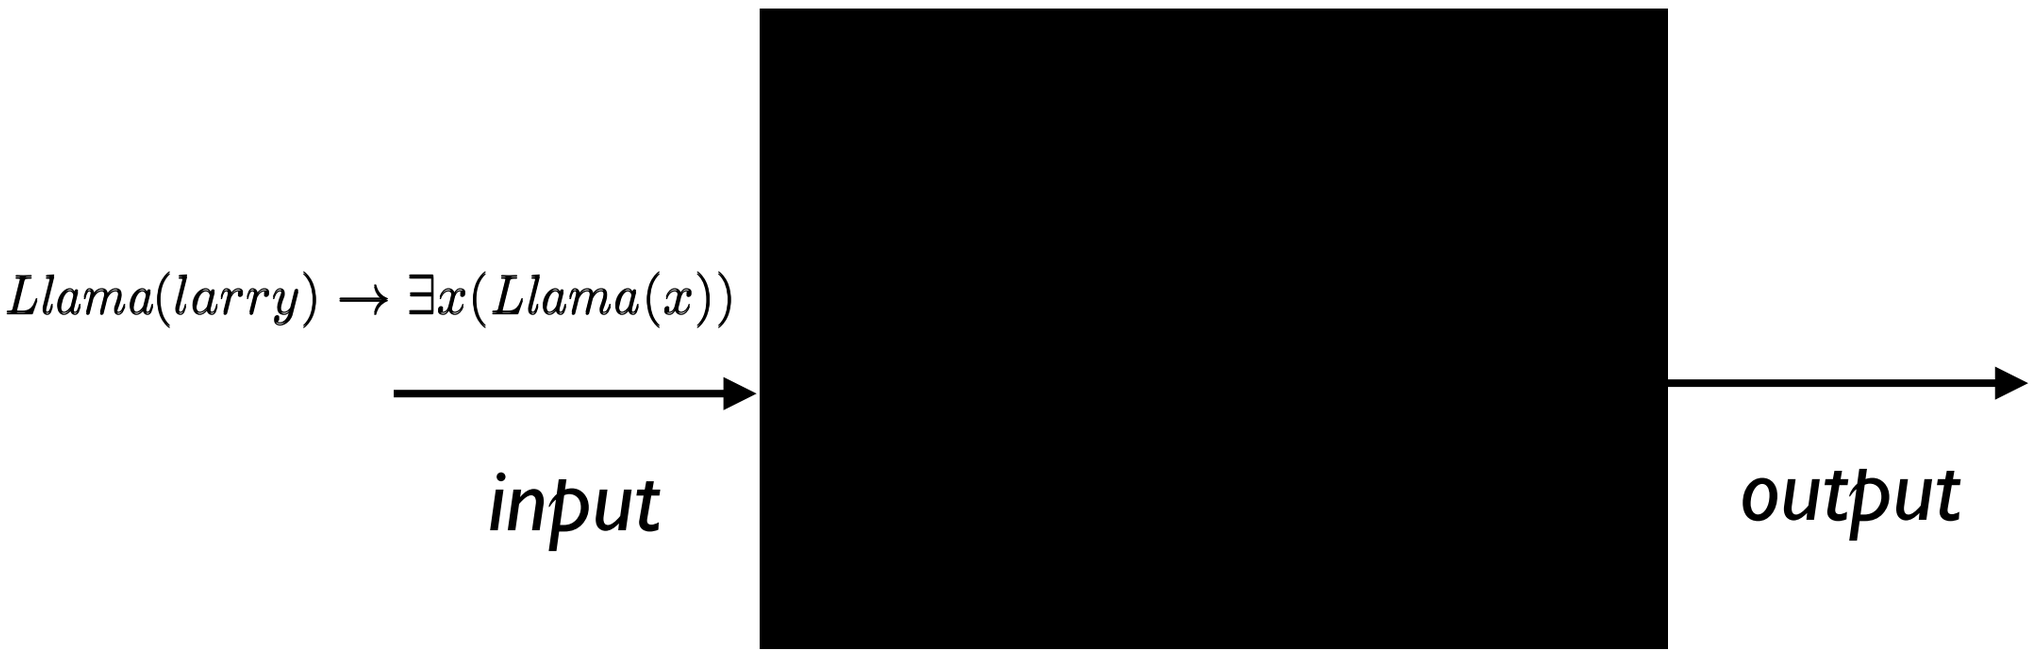

And now, the Theoremhood Decision Problem,
i.e., the *Entscheidungsproblem*,
($\text{THEOREM}_{\text{FOL}}$)
for First-Order Logic (FOL)

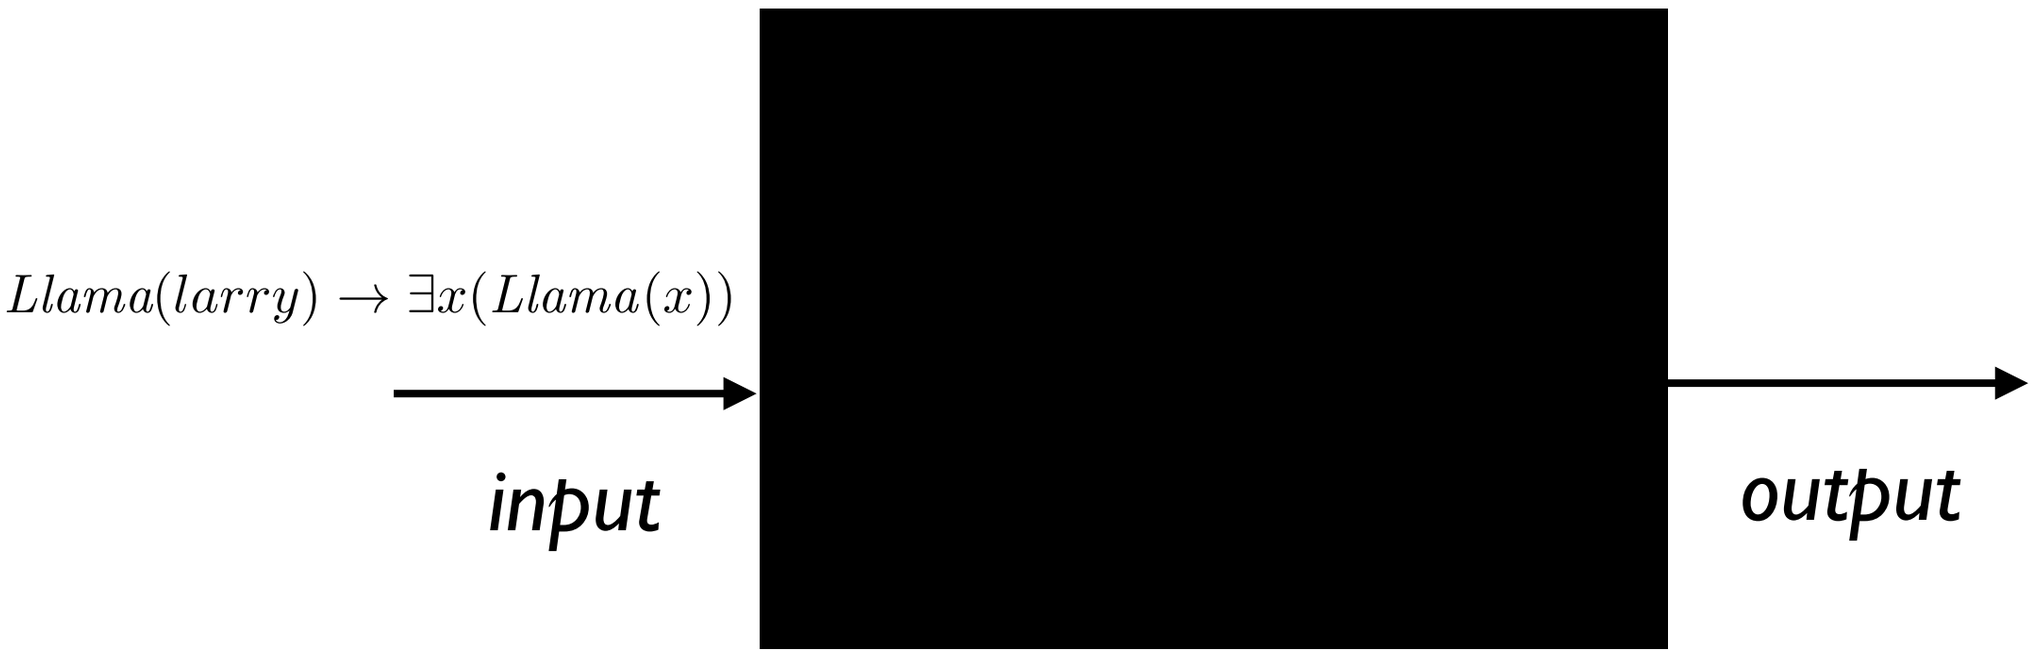

And now, the Theoremhood Decision Problem,
i.e., the *Entscheidungsproblem*,
($\text{THEOREM}_{\text{FOL}}$)
for First-Order Logic (FOL)

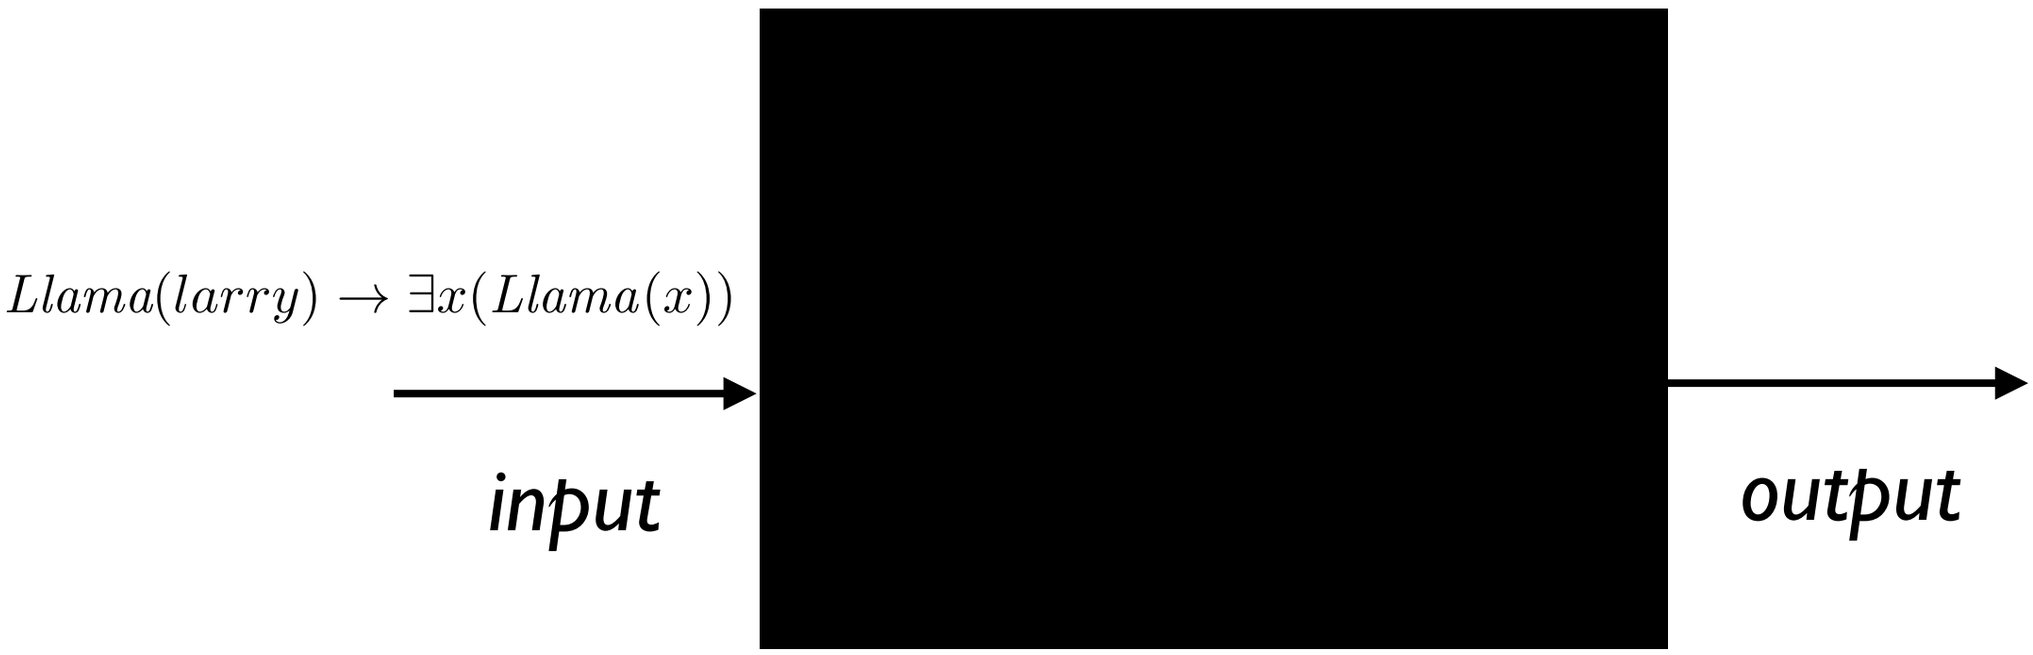

And now, the Theoremhood Decision Problem, i.e., the *Entscheidungsproblem*, ($\text{THEOREM}_{\text{FOL}}$) for First-Order Logic (FOL)

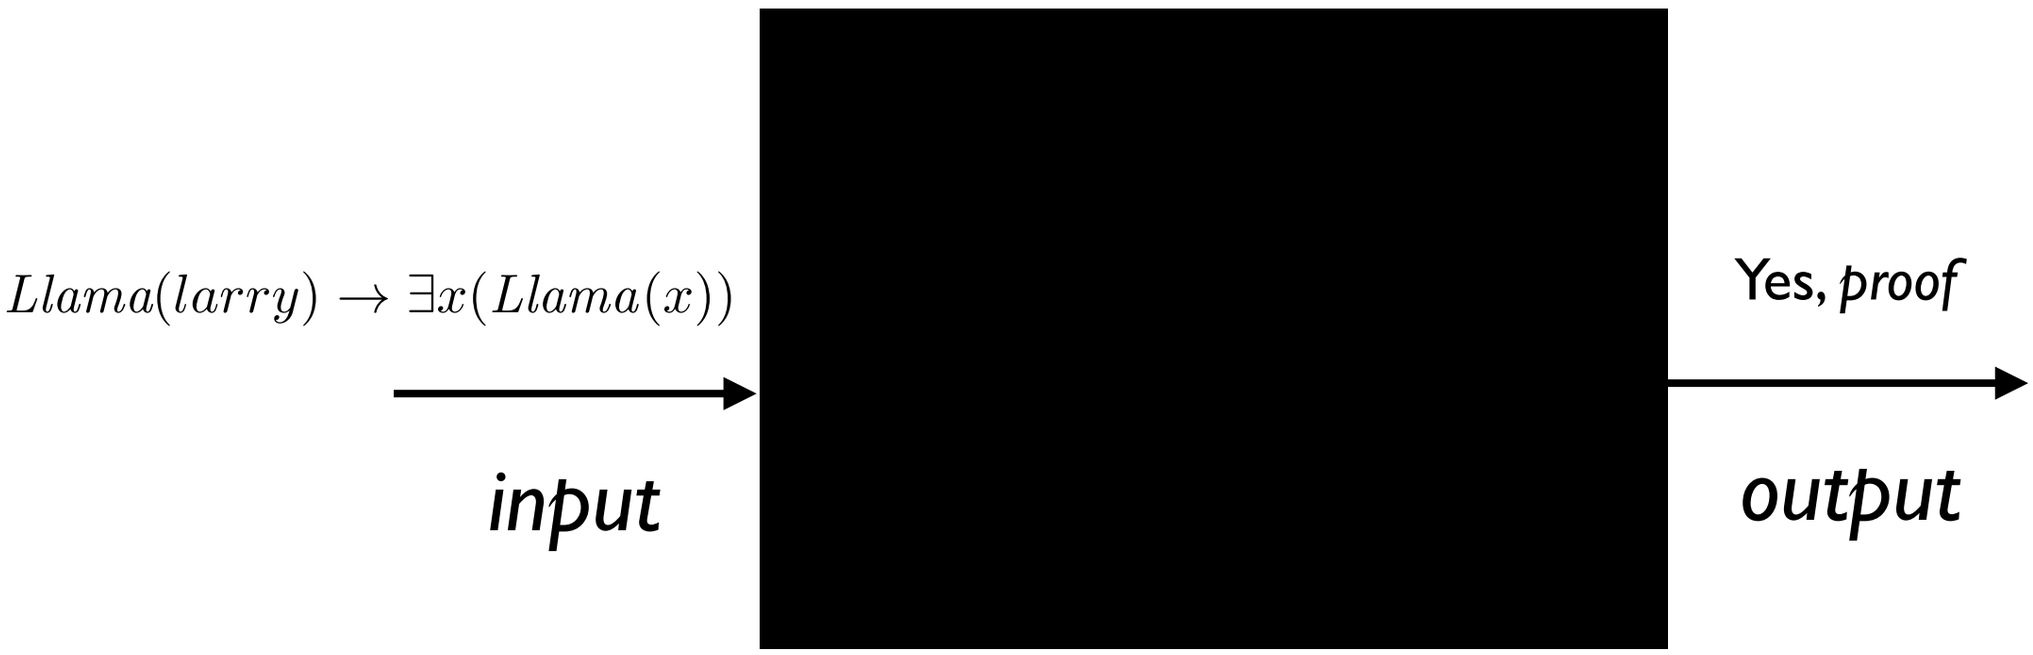

And now, the Theoremhood Decision Problem,
i.e., the *Entscheidungsproblem*,
($\text{THEOREM}_{\text{FOL}}$)
for First-Order Logic (FOL)

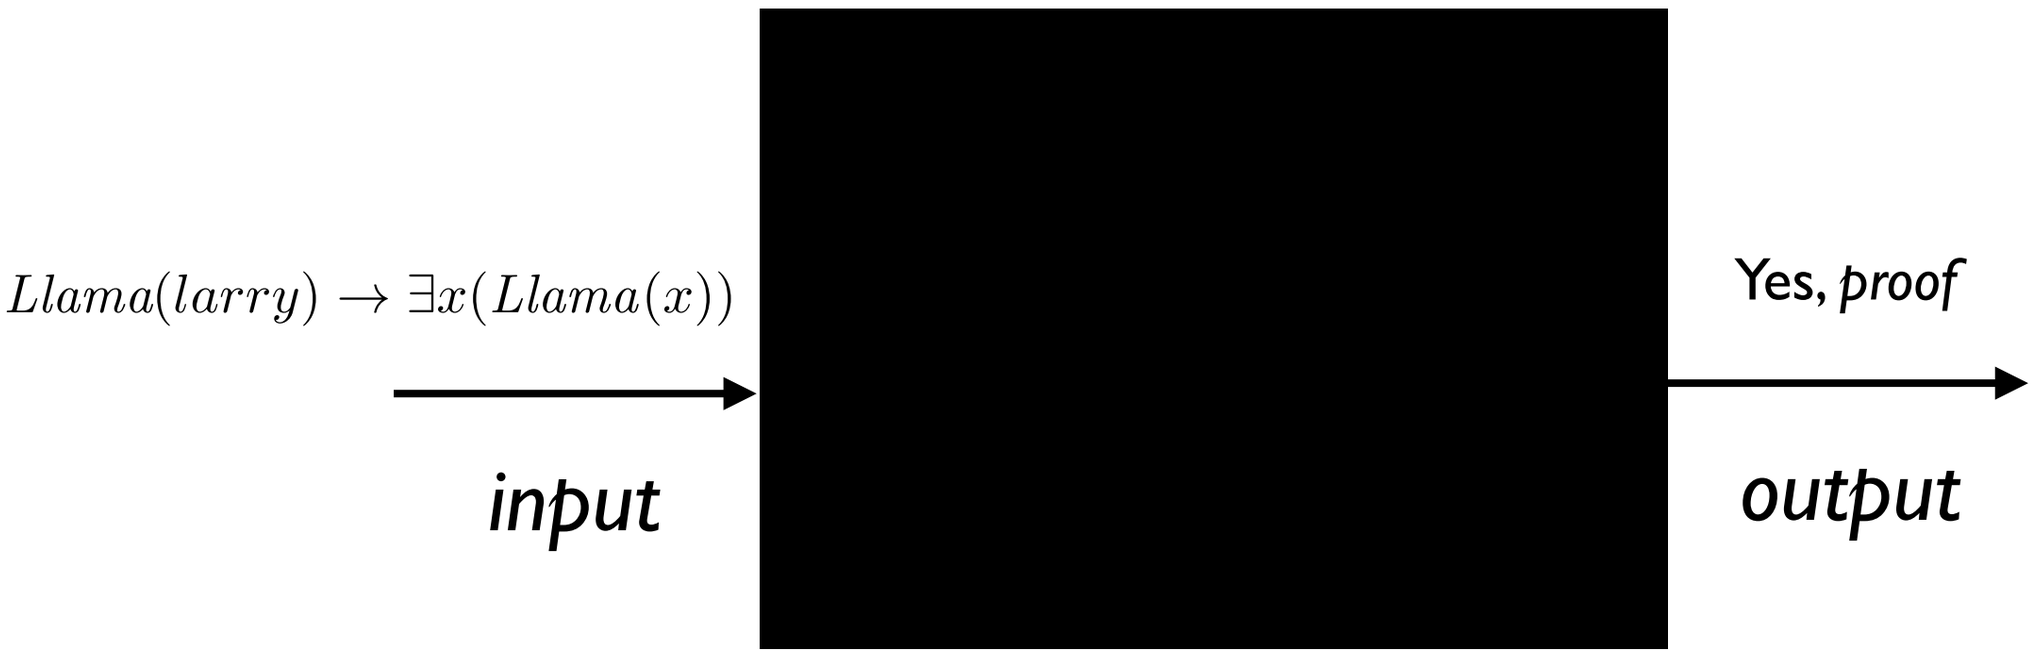

Not just hard: *impossible* for a (and this needed to be *invented* in the course of clarifying and solving the problem) standard computing machine.

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*?

Logikk kan gi dyp glede!