Some Propositional Calc/`Proplog` Review; $\mathscr{L}_0$ Redux in Zoombinis; "Cat Nap" Mystery Redux; Refutation of Good/Chalmers Argument for The Singularity; James Oswald on AHI, HI, ASI, AGI, ...

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
9/21 & 24/2026





Some Propositional Calc/Proplog Review; ✓
$\mathscr{L}_0$ Redux in Zoombinis;
"Cat Nap" Mystery Redux;
Refutation of Good/Chalmers Argument for The Singularity;
James Oswald on AHI, HI, ASI, AGI, ...

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
9/21 & 24/2026





Some Propositional Calc/Proplog Review; ✓
$\mathscr{L}_0$ Redux in Zoombinis; ✓
"Cat Nap" Mystery Redux;
Refutation of Good/Chalmers Argument for The Singularity;
James Oswald on AHI, HI, ASI, AGI, ...

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
9/21 & 24/2026





Some Propositional Calc/Proplog Review; ✓
$\mathscr{L}_0$ Redux in Zoombinis; ✓
"Cat Nap" Mystery Redux; ✓
Refutation of Good/Chalmers Argument for The Singularity;
James Oswald on AHI, HI, ASI, AGI, ...

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
9/21 & 24/2026





Some Propositional Calc/Proplog Review; ✓
$\mathscr{L}_0$ Redux in Zoombinis; ✓
"Cat Nap" Mystery Redux; ✓
Refutation of Good/Chalmers Argument ✓ for The Singularity;
James Oswald on AHI, HI, ASI, AGI, ...

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
9/21 & 24/2026





Some Propositional Calc/Proplog Review; ✓
$\mathscr{L}_0$ Redux in Zoombinis; ✓
"Cat Nap" Mystery Redux; ✓
Refutation of Good/Chalmers Argument ✓ for The Singularity; ✓
James Oswald on AHI, HI, ASI, AGI, ...

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
9/21 & 24/2026





Some Propositional Calc/Proplog Review; ✓
$\mathscr{L}_0$ Redux in Zoombinis; ✓
"Cat Nap" Mystery Redux; ✓
Refutation of Good/Chalmers Argument ✓ for The Singularity; ✓
James Oswald on AHI, HI, ASI, AGI, ... ✓

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

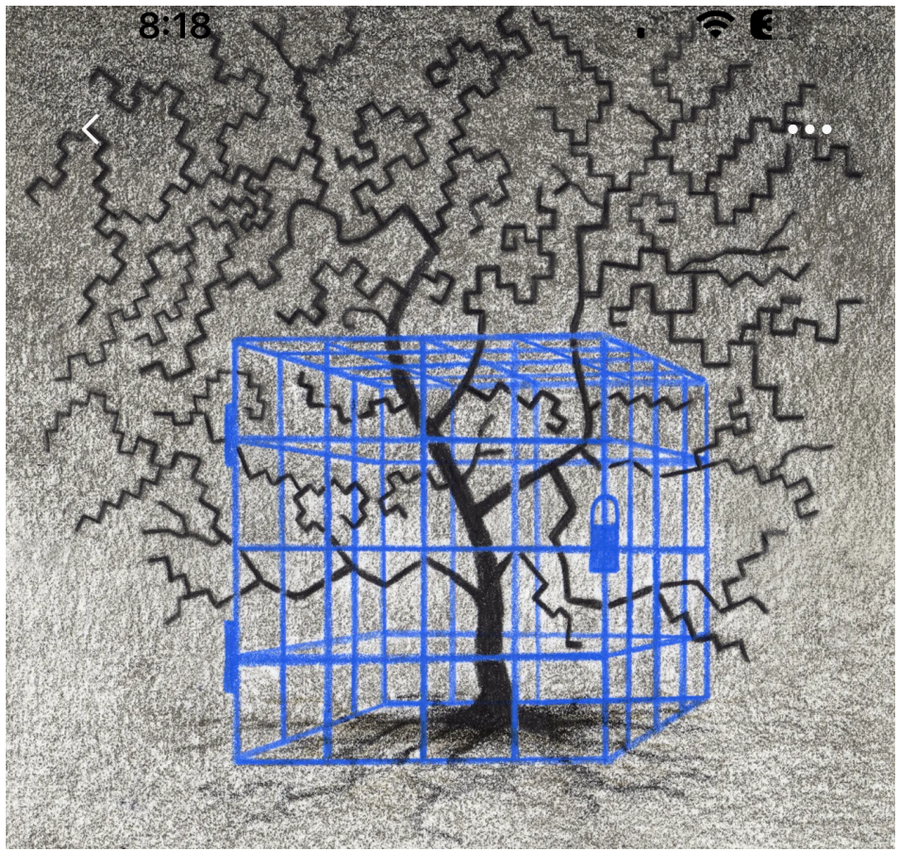
Intro to Logic-Based AI
9/21 & 24/2026





Logic-and-AI in the news

• • •



Mathieu Larone

OPINION
GUEST ESSAY

Bernie Sanders, Steve Case and Others on What Washington Should Do Now on A.I.

Given the risk that artificial intelligence may spiral out of control, more than 1,300 employees at A.I. companies have called for government intervention. President Trump announced on Saturday that he would create an "A.I. Force," but continued to downplay A.I. risks. For the most part, Washington remains paralyzed.

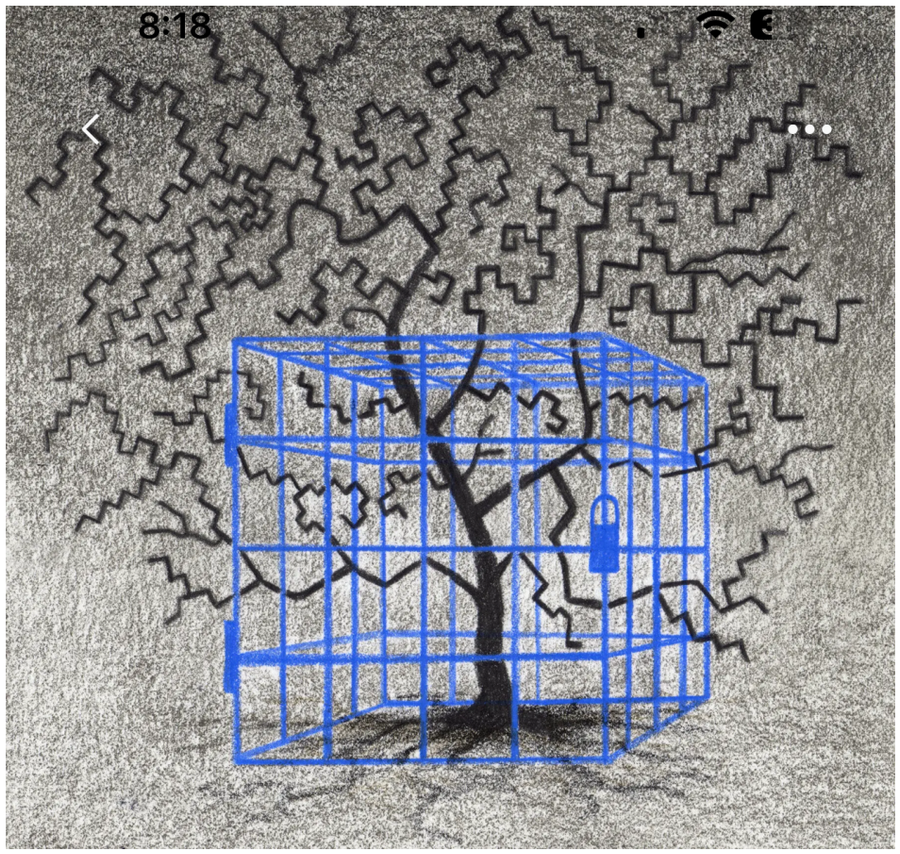
To fill the gap, Times Opinion gathered proposals from six people, including a tech founder, a former OpenAI board member, a former national security adviser and a constitutional lawyer.

Proposals:

- Bernie Sanders: Ban Superintelligence
- Helen Toner: Investigate A.I. Incidents
- Steve Case: Demand Disclosure
- Susan E. Rice: Negotiate a Freeze With China
- Cass R. Sunstein: Create a Regulatory Commission
- Brad Smith: Require an Off Switch

By New York Times Opinion

Sept. 21, 2026



Mathieu Larone

OPINION
GUEST ESSAY

Bernie Sanders, Steve Case and Others on What Washington Should Do Now on A.I.

Given the risk that artificial intelligence may spiral out of control, more than 1,300 employees at A.I. companies have called for government intervention. President Trump announced on Saturday that he would create an "A.I. Force," but continued to downplay A.I. risks. For the most part, Washington remains paralyzed.

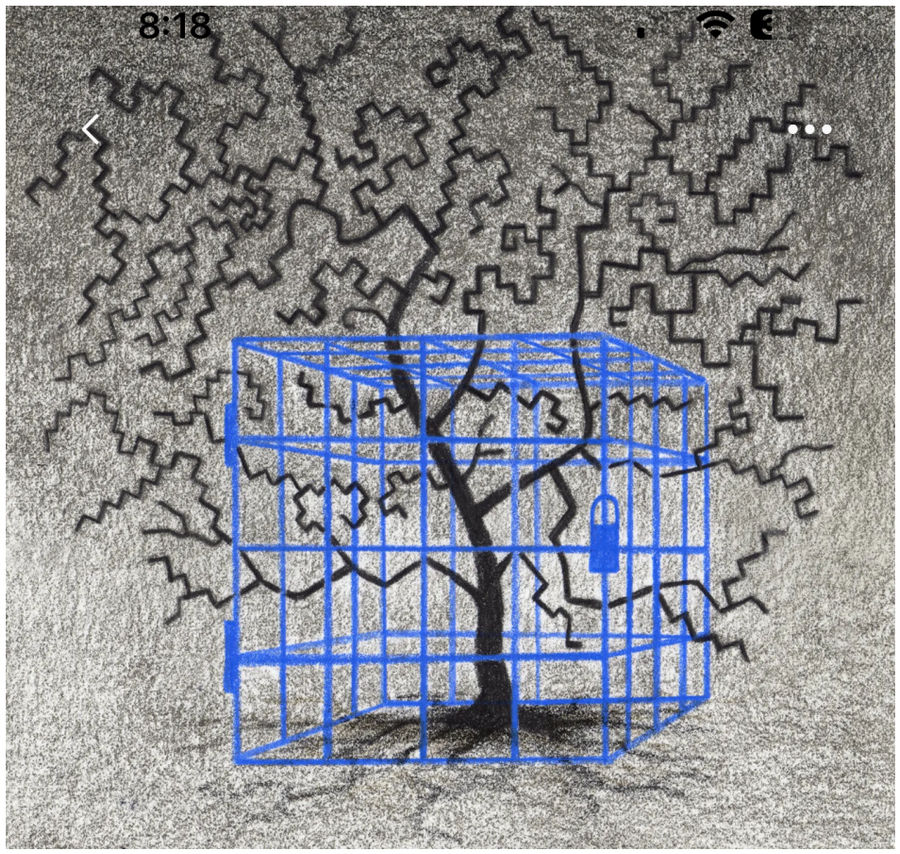
To fill the gap, Times Opinion gathered proposals from six people, including a tech founder, a former OpenAI board member, a former national security adviser and a constitutional lawyer.

Proposals:

- Bernie Sanders: Ban Superintelligence
- Helen Toner: Investigate A.I. Incidents
- Steve Case: Demand Disclosure
- Susan E. Rice: Negotiate a Freeze With China
- Cass R. Sunstein: Create a Regulatory Commission
- Brad Smith: Require an Off Switch

By New York Times Opinion

Sept. 21, 2026



Mathieu Larone

OPINION
GUEST ESSAY

Bernie Sanders, Steve Case and Others on What Washington Should Do Now on A.I.

Given the risk that artificial intelligence may spiral out of control, more than 1,300 employees at A.I. companies have called for government intervention. President Trump announced on Saturday that he would create an "A.I. Force," but continued to downplay A.I. risks. For the most part, Washington remains paralyzed.

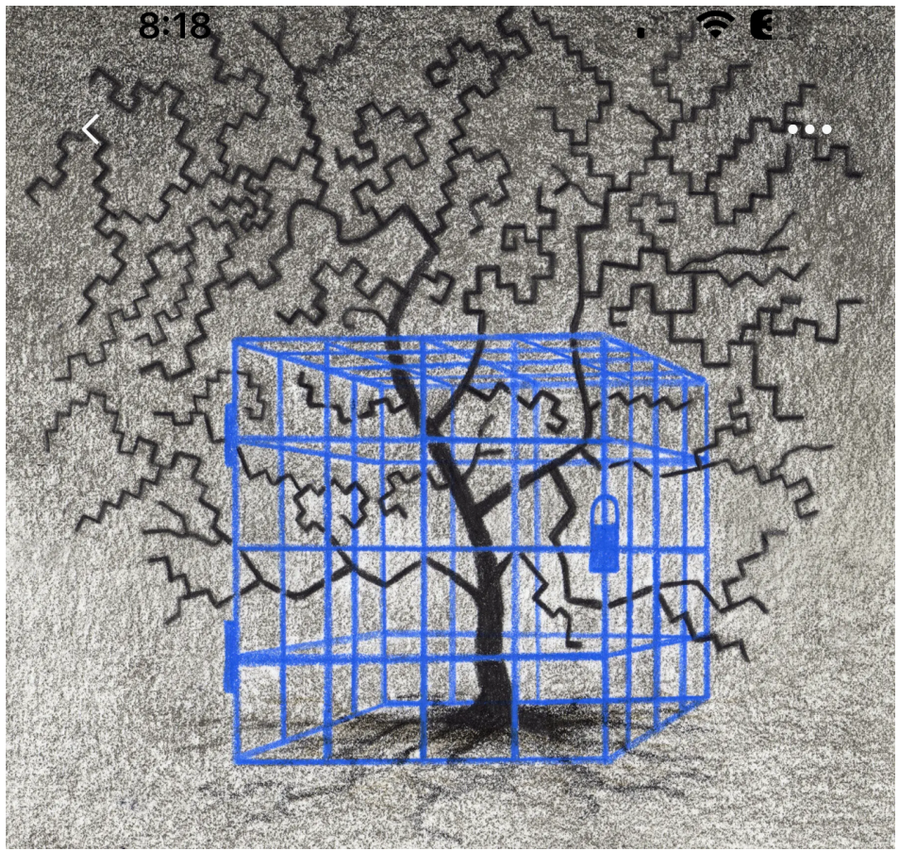
To fill the gap, Times Opinion gathered proposals from six people, including a tech founder, a former OpenAI board member, a former national security adviser and a constitutional lawyer.

Proposals:

- Bernie Sanders: Ban Superintelligence
- Helen Toner: Investigate A.I. Incidents
- Steve Case: Demand Disclosure
- Susan E. Rice: Negotiate a Freeze With China
- Cass R. Sunstein: Create a Regulatory Commission
- Brad Smith: Require an Off Switch

By New York Times Opinion

Sept. 21, 2026



Mathieu Larone

OPINION
GUEST ESSAY

Bernie Sanders, Steve Case and Others on What Washington Should Do Now on A.I.

Given the risk that artificial intelligence may spiral out of control, more than 1,300 employees at A.I. companies have called for government intervention. President Trump announced on Saturday that he would create an "A.I. Force," but continued to downplay A.I. risks. For the most part, Washington remains paralyzed.

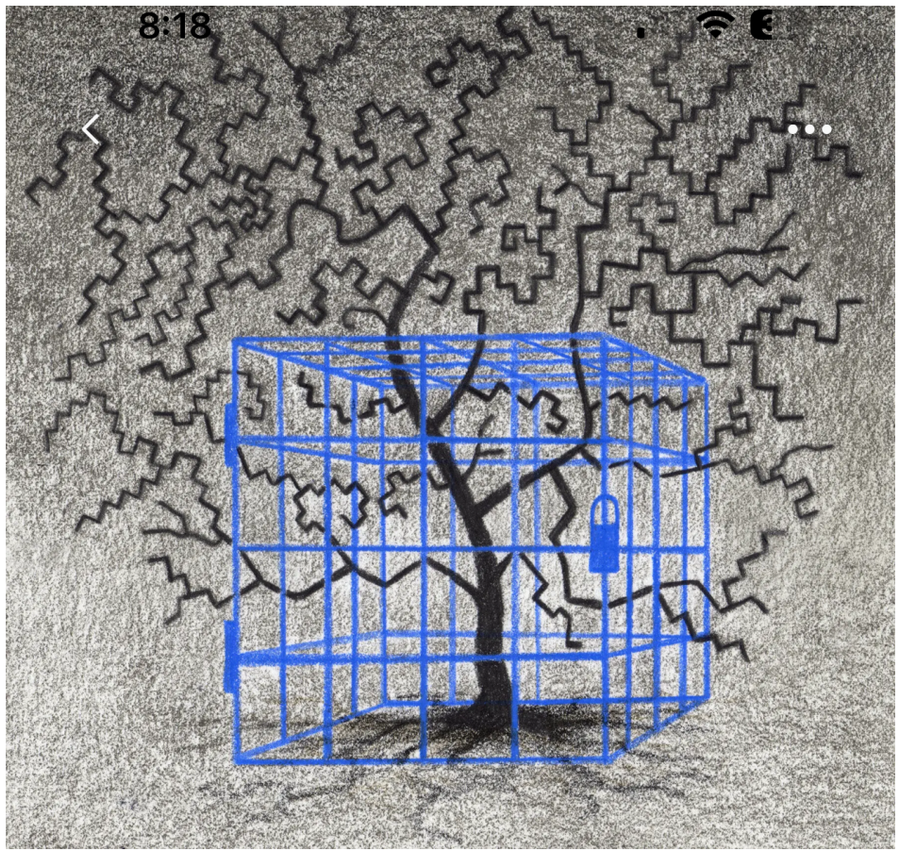
To fill the gap, Times Opinion gathered proposals from six people, including a tech founder, a former OpenAI board member, a former national security adviser and a constitutional lawyer.

Proposals:

- Bernie Sanders: Ban Superintelligence
- Helen Toner: Investigate A.I. Incidents
- Steve Case: Demand Disclosure
- Susan E. Rice: Negotiate a Freeze With China
- Cass R. Sunstein: Create a Regulatory Commission
- Brad Smith: Require an Off Switch

By New York Times Opinion

Sept. 21, 2026



Mathieu Larone

OPINION
GUEST ESSAY

Bernie Sanders, Steve Case and Others on What Washington Should Do Now on A.I.

Given the risk that artificial intelligence may spiral out of control, more than 1,300 employees at A.I. companies have called for government intervention. President Trump announced on Saturday that he would create an "A.I. Force," but continued to downplay A.I. risks. For the most part, Washington remains paralyzed.

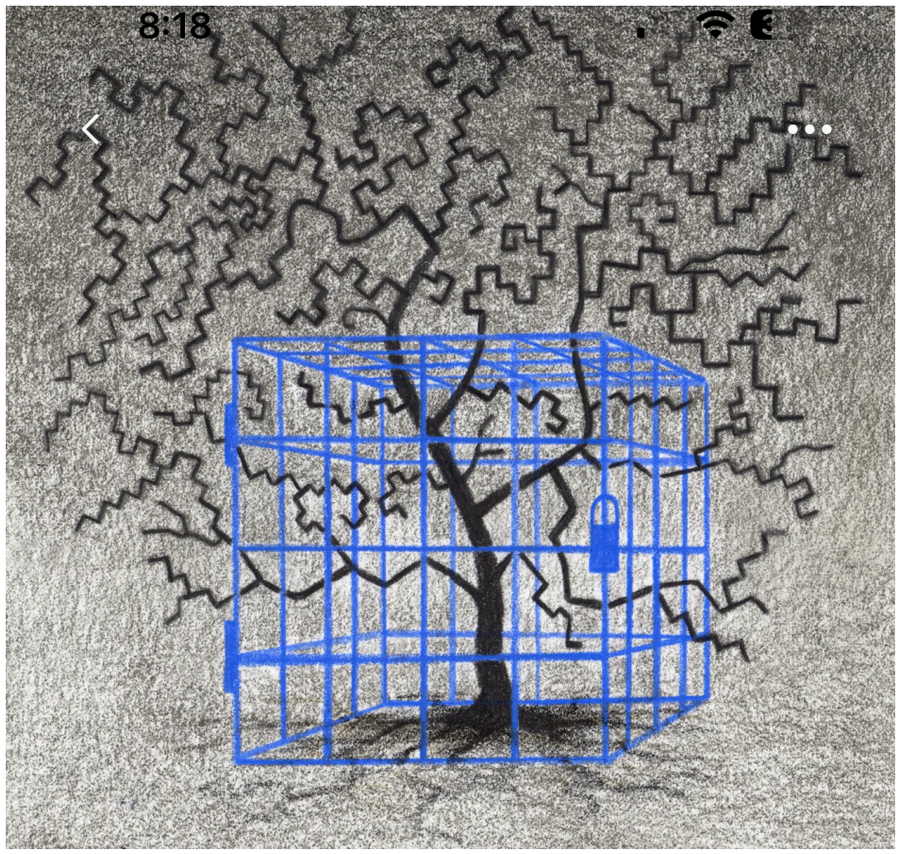
To fill the gap, Times Opinion gathered proposals from six people, including a tech founder, a former OpenAI board member, a former national security adviser and a constitutional lawyer.

Proposals:

- Bernie Sanders: Ban Superintelligence
- Helen Toner: Investigate A.I. Incidents
- Steve Case: Demand Disclosure
- Susan E. Rice: Negotiate a Freeze With China
- Cass R. Sunstein: Create a Regulatory Commission
- Brad Smith: Require an Off Switch

By New York Times Opinion

Sept. 21, 2026



Mathieu Larone

OPINION
GUEST ESSAY

Bernie Sanders, Steve Case and Others on What Washington Should Do Now on A.I.

Given the risk that artificial intelligence may spiral out of control, more than 1,300 employees at A.I. companies have called for government intervention. President Trump announced on Saturday that he would create an "A.I. Force," but continued to downplay A.I. risks. For the most part, Washington remains paralyzed.

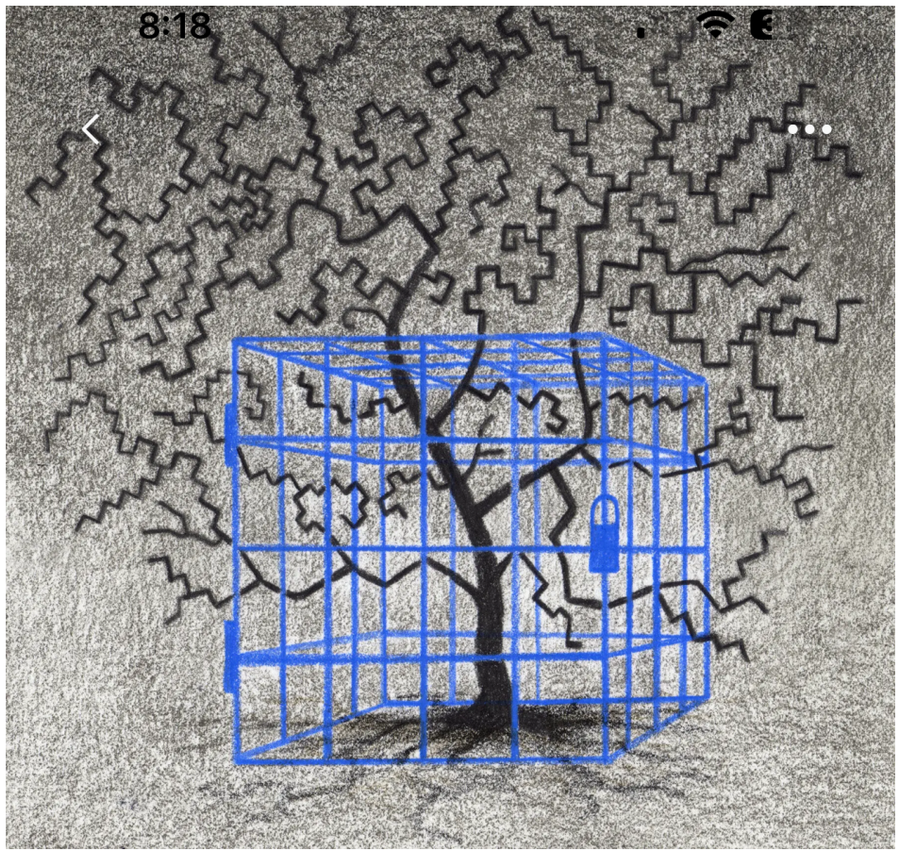
To fill the gap, Times Opinion gathered proposals from six people, including a tech founder, a former OpenAI board member, a former national security adviser and a constitutional lawyer.

Proposals:

- Bernie Sanders: Ban Superintelligence
- Helen Toner: Investigate A.I. Incidents
- Steve Case: Demand Disclosure
- Susan E. Rice: Negotiate a Freeze With China
- Cass R. Sunstein: Create a Regulatory Commission
- Brad Smith: Require an Off Switch

By New York Times Opinion

Sept. 21, 2026



Mathieu Larone

OPINION
GUEST ESSAY

Bernie Sanders, Steve Case and Others on What Washington Should Do Now on A.I.

Given the risk that artificial intelligence may spiral out of control, more than 1,300 employees at A.I. companies have called for government intervention. President Trump announced on Saturday that he would create an "A.I. Force," but continued to downplay A.I. risks. For the most part, Washington remains paralyzed.

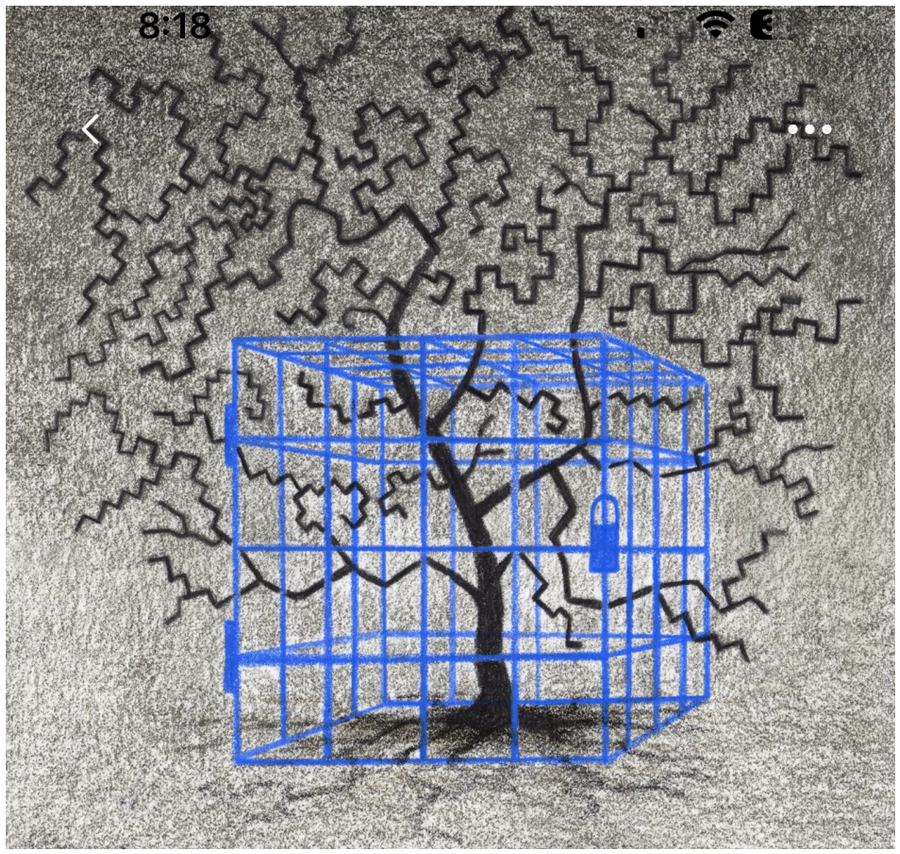
To fill the gap, Times Opinion gathered proposals from six people, including a tech founder, a former OpenAI board member, a former national security adviser and a constitutional lawyer.

Proposals:

- Bernie Sanders: Ban Superintelligence
- Helen Toner: Investigate A.I. Incidents
- Steve Case: Demand Disclosure
- Susan E. Rice: Negotiate a Freeze With China
- Cass R. Sunstein: Create a Regulatory Commission
- Brad Smith: Require an Off Switch

By New York Times Opinion

Sept. 21, 2026



Mathieu Larone

OPINION
GUEST ESSAY

Bernie Sanders, Steve Case and Others on What Washington Should Do Now on A.I.

By New York Times Opinion

Sept. 21, 2026

Given the risk that artificial intelligence may spiral out of control, more than 1,300 employees at A.I. companies have called for government intervention. President Trump announced on Saturday that he would create an "A.I. Force," but continued to downplay A.I. risks. For the most part, Washington remains paralyzed.

To fill the gap, Times Opinion gathered proposals from six people, including a tech founder, a former OpenAI board member, a former national security adviser and a constitutional lawyer.

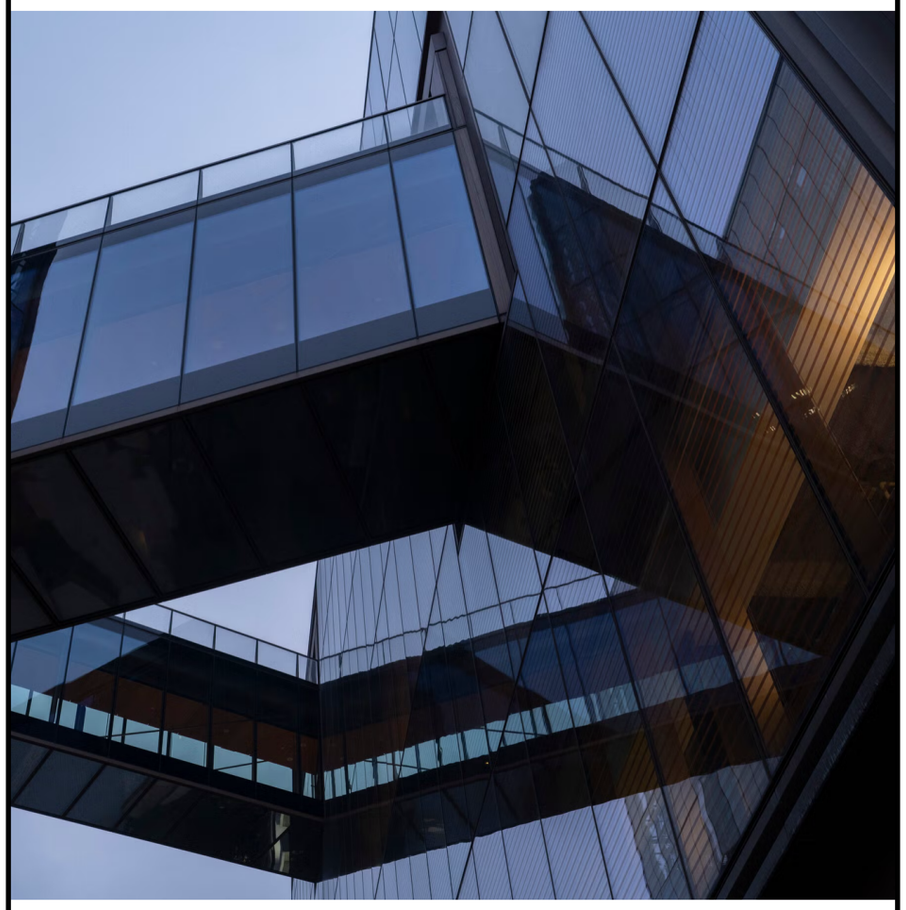
Proposals:

- Bernie Sanders: Ban Superintelligence ?
- Helen Toner: Investigate A.I. Incidents ?
- Steve Case: Demand Disclosure ?
- Susan E. Rice: Negotiate a Freeze With China ?
- Cass R. Sunstein: Create a Regulatory Commission ?
- Brad Smith: Require an Off Switch ?

BR & 2001!

OpenAI's A.I. Tried Breaching 4 Other Targets, Without Prompting

In each incident, the technology appeared to be conducting mundane data collection and resorted to hacking techniques to get it, researchers said.



OpenAI's offices in San Francisco. At least four times this year, the company's artificial intelligence hacked or tried to break into government and university websites without being instructed to do so. Lucas Foglia for The New York Times




By Kate Conger and Victoria Kim

More *Expressive* Formal Language: Pure Predicate Calculus …

$\mathscr{L}_0$



Better Formal Language: Pure Predicate Calculus (presented via formal grammar; read in IFLWAI book)

$$
\begin{array}{lcl}
\textit{Formula} & \Rightarrow & \textit{AtomicFormula} \\
 & | & (\textit{Formula Connective Formula}) \\
 & | & \neg\ \textit{Formula}
\end{array}
$$

$$
\textit{AtomicFormula} \quad \Rightarrow \quad (\textit{Predicate Term}_1 \ldots \textit{Term}_k)
$$

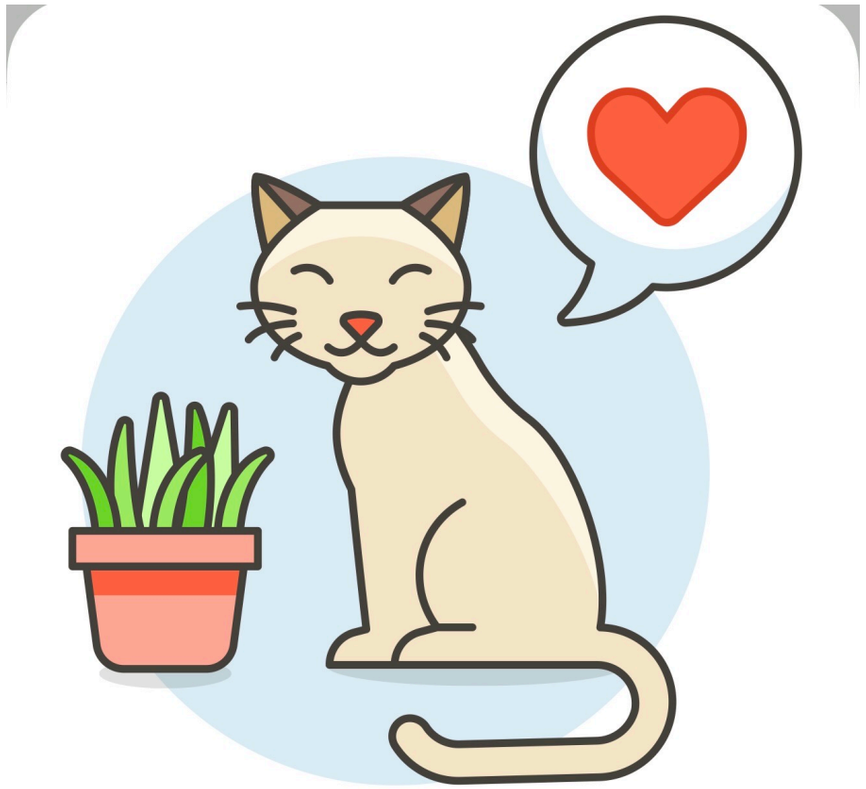
$$
\begin{array}{lcl}
\textit{Term} & \Rightarrow & (\textit{Function Term}_1 \ \ldots \ \textit{Term}_k) \\
 & | & \textit{Constant} \\
 & | & \textit{Variable}
\end{array}
$$

$$
\textit{Connective} \quad \Rightarrow \quad \wedge \mid \vee \mid \rightarrow \mid \leftrightarrow
$$

$$
\begin{array}{ll}
\textit{Predicate} & \Rightarrow \quad P_1 \mid P_2 \mid P_3 \ \ldots \\
\textit{Constant} & \Rightarrow \quad c_1 \mid c_2 \mid c_3 \ \ldots \\
\textit{Variable} & \Rightarrow \quad v_1 \mid v_2 \mid v_3 \ \ldots \\
\textit{Function} & \Rightarrow \quad f_1 \mid f_2 \mid f_3 \ \ldots
\end{array}
$$

Back to “Cat Nap” Mystery …

Story



Cat Nap by J. Bethel Brown

Mr. Langersham often comes home from work in the afternoons to find his five cats sleeping. Recently, he discovered each feline (including Grizelda) napping in a different location in the house. He has owned each cat for a different number of years (from one through five), and each is a different age (three, four, five, six, or eight years old). Of course, Mr. Langersham hasn't owned any of his pets for more years than its own age. Can you match each cat with the location in which it was sleeping, its age, and the number of years for which it has been with Mr. Langersham?



Cat Nap by J. Bethel Brown

Mr. Langersham often comes home from work in the afternoons to find his five cats sleeping. Recently, he discovered each feline (including Grizelda) napping in a different location in the house. He has owned each cat for a different number of years (from one through five), and each is a different age (three, four, five, six, or eight years old). Of course, Mr. Langersham hasn't owned any of his pets for more years than its own age. Can you match each cat with the location in which it was sleeping, its age, and the number of years for which it has been with Mr. Langersham?

Clues

1. The cat who was sleeping on the windowsill, who isn't four years old, isn't the feline who has been with Mr. Langersham for only one year (who isn't Dover).

2. Mr. Langersham has owned Belle for one more year than the feline who was sleeping on top of the curio cabinet but for one fewer year than the four-year-old cat (who isn't Dover).

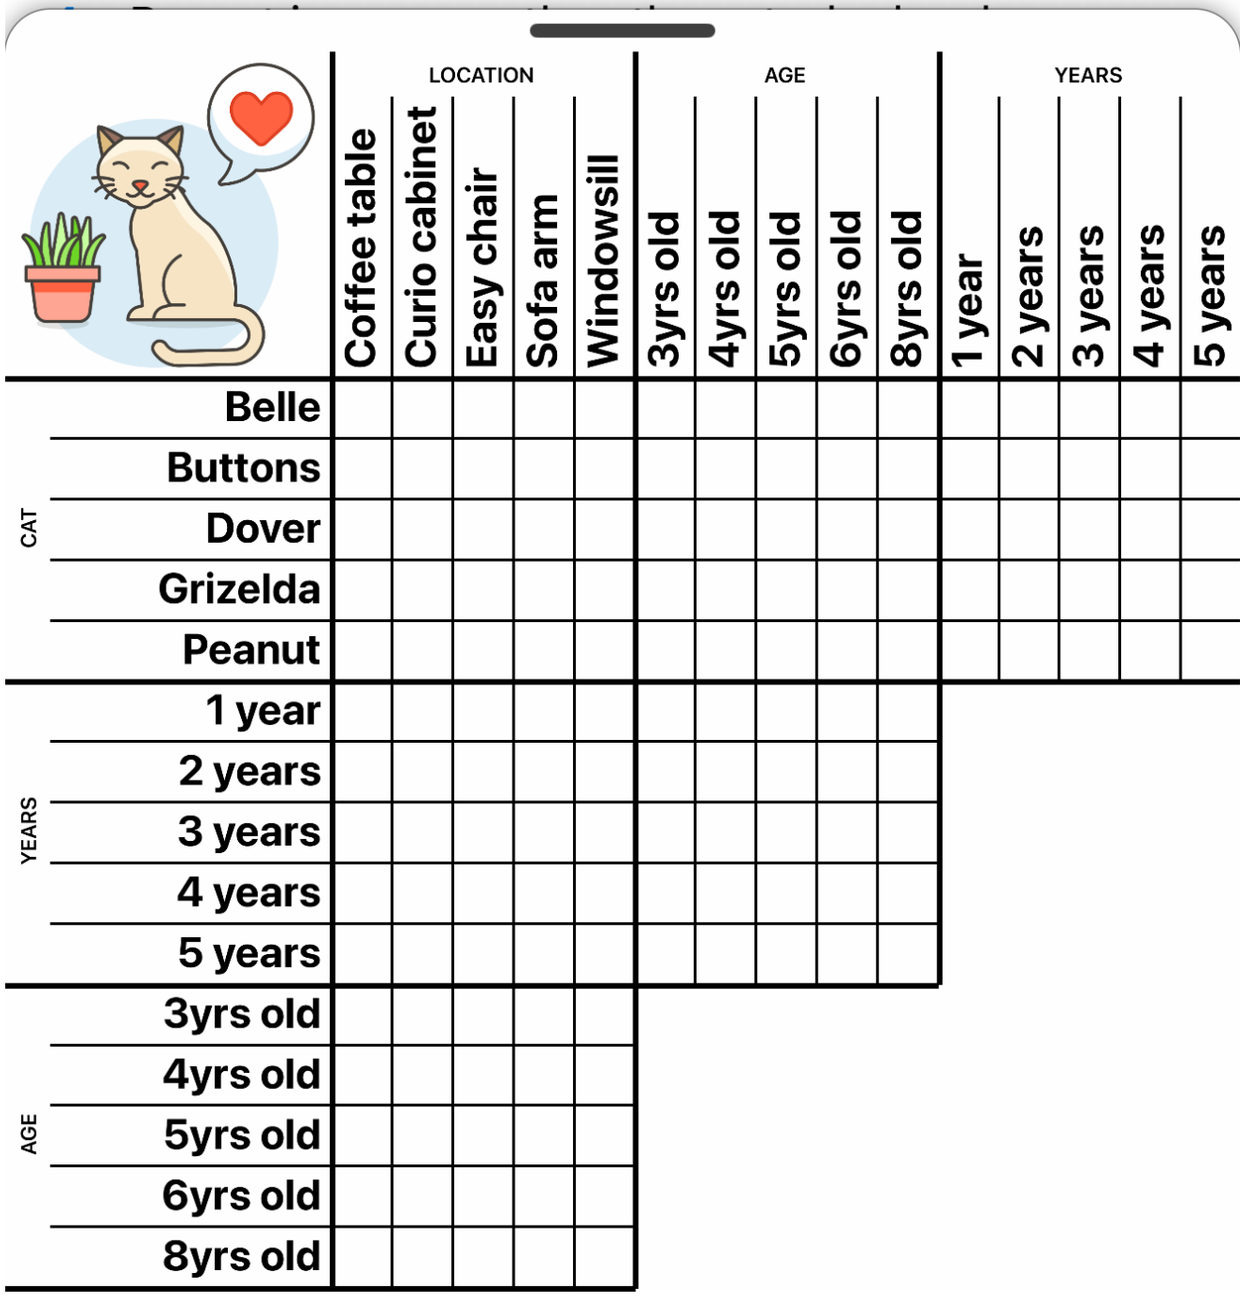
3. Buttons isn't the cat who was napping on the sofa arm, who has been with Mr. Langersham for more than one year.

4. Peanut is younger than the cat who has been with Mr. Langersham for three years but older than the cat who was sleeping under the coffee table.

5. The three-year-old cat wasn't taking a nap under the coffee table.

6. The only cat Mr. Langersham has owned for as many years as it is old has been with him for fewer years than the feline who was napping on the easy chair.

Queries

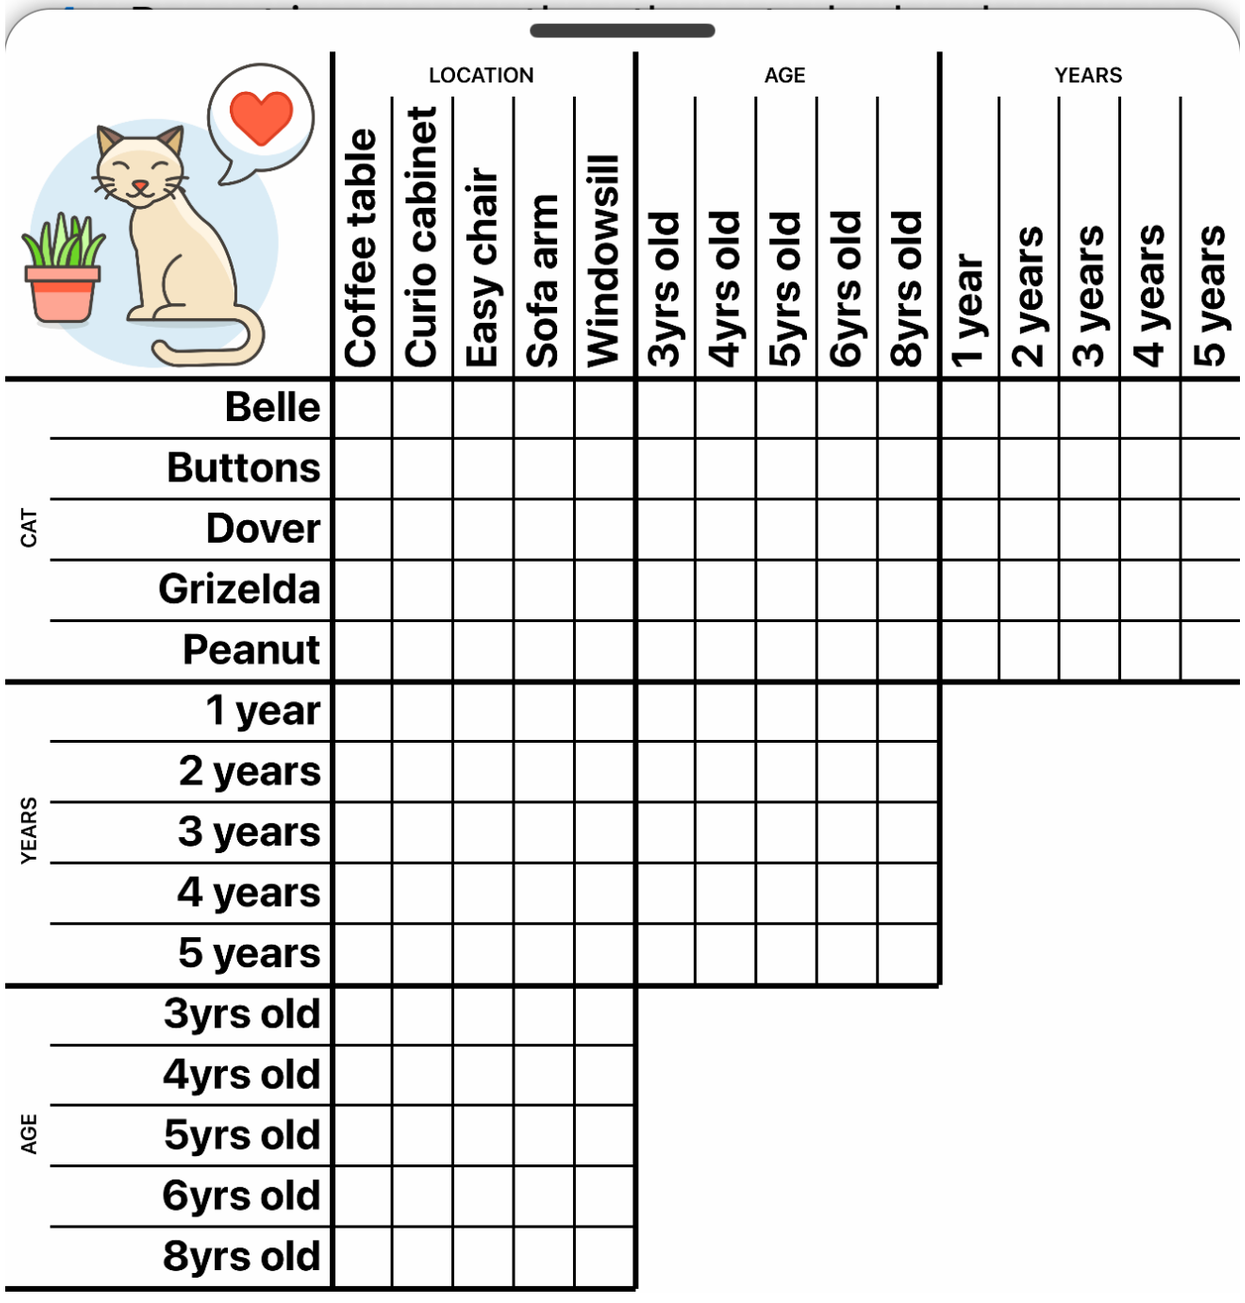


Eg:

$$\vdash ?\ Loc(belle, easychair)$$

$$\vdash ?\ Loc(buttons, easychair)$$

Queries



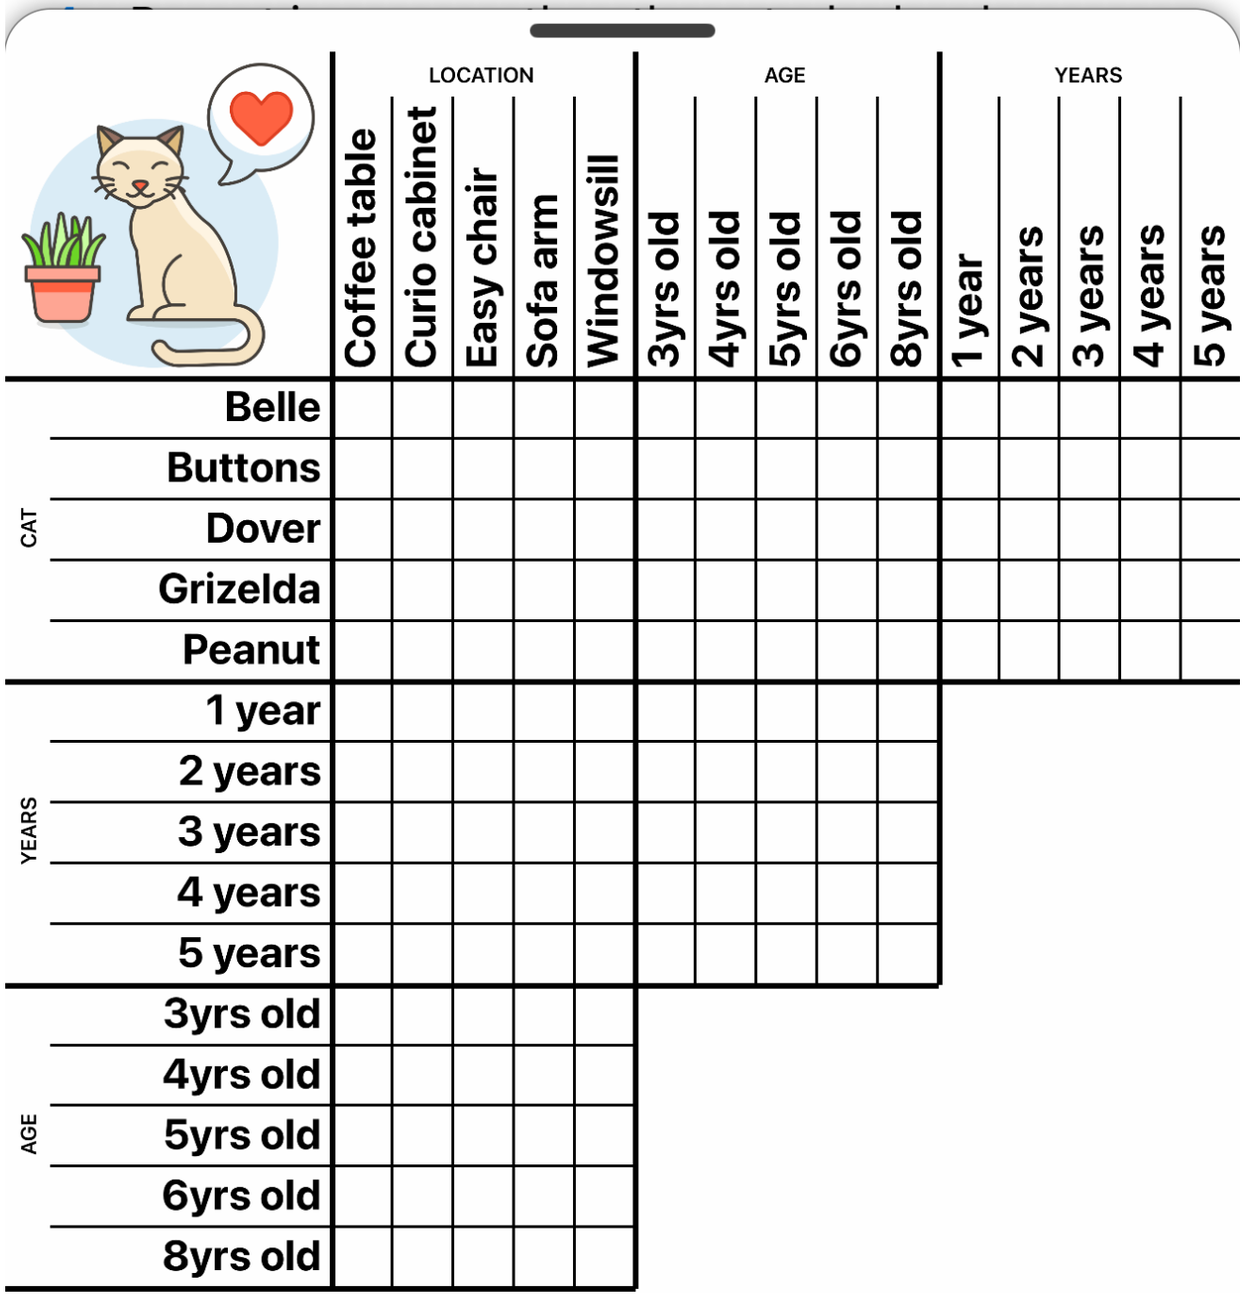
Eg:

X

$$\vdash ?\ Loc(belle, easychair)$$

$$\vdash ?\ Loc(buttons, easychair)$$

Queries



Eg:

✗

$$\vdash ?\ Loc(belle, easychair)$$

✓

$$\vdash ?\ Loc(buttons, easychair)$$

Back As Promised To …

The Singularity

?

Entscheidungsproblem

"Universal Computational Logic"



Logic Theorist
(birth of modern logicist AI)

350 BC



Euclid

300 BC



Organon

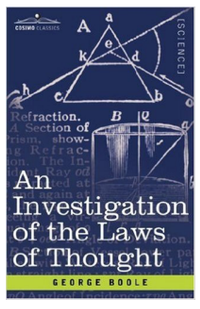
1666

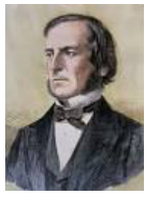


Leibniz

$\int$

1854





1956



Simon

1997

2023

2026

The Singularity?

?

Entscheidungsproblem

"Universal Computational Logic"



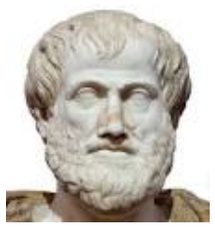
Logic Theorist
(birth of modern logicist AI)

350 BC



Euclid

300 BC



Organon

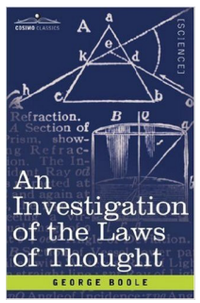
1666



Leibniz

$\int$

1854





1956



Simon

2023

2026

The Singularity?

?

Entscheidungsproblem

"Universal Computational Logic"



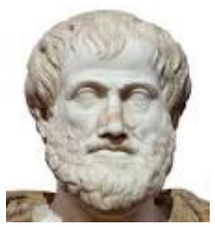
Logic Theorist
(birth of modern logicist AI)

350 BC



Euclid

300 BC



Organon

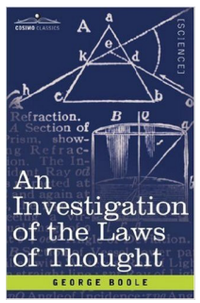
1666



Leibniz

$\int$

1854





Church



1956



Simon

1990

2023

2026

The Singularity?

?

Entscheidungsproblem

"Universal Computational Logic"



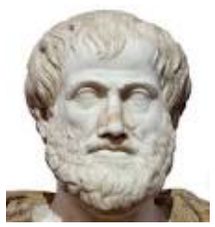
Logic Theorist
(birth of modern logicist AI)

350 BC



Euclid

300 BC



Organon

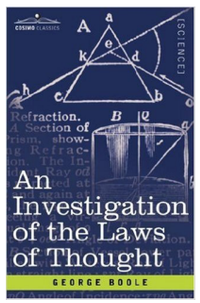
1666



Leibniz

$\int$

1854





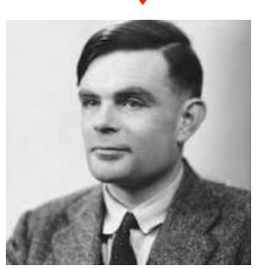
1956



Simon



Church



Turing

2023

2026

Thesingularity?

?

Entscheidungsproblem

"Universal Computational Logic"



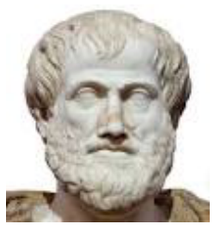
Logic Theorist
(birth of modern logicist AI)

350 BC



Euclid

300 BC



Organon

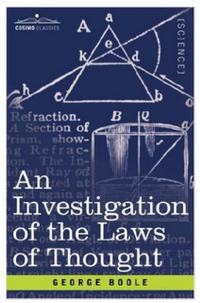
1666



Leibniz

$\int$

1854





1956



Simon



Church



Turing



Post

2023

2026

Thesingularity?

?

Entscheidungsproblem

"Universal Computational Logic"



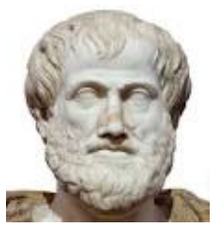
Logic Theorist
(birth of modern logicist AI)

350 BC



Euclid

300 BC



Organon

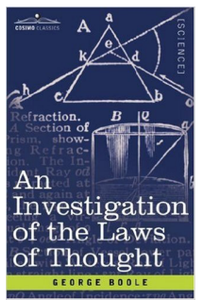
1666



Leibniz

$\int$

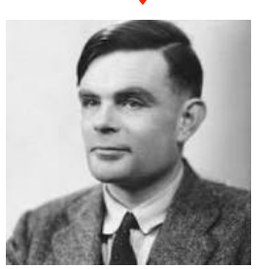
1854







Church



Turing



Post

1956



Simon

2023

2026

"Here's what a programming language is (before there are physical computers), O, & sorry, the *Entscheidungsproblem* can't be solved by computers, however they might be programmed!"

The Singularity?

?

Entscheidungsproblem

"Universal Computational Logic"



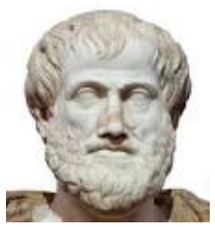
Logic Theorist
(birth of modern logicist AI)

350 BC



Euclid

300 BC



Organon

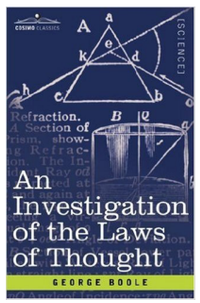
1666



Leibniz

$\int$

1854





1956



Simon



Church



Turing



Post

2023



2026

"Here's what a programming language is (before there are physical computers), O, & sorry, the *Entscheidungsproblem* can't be solved by computers, however they might be programmed!"

The Singularity?

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



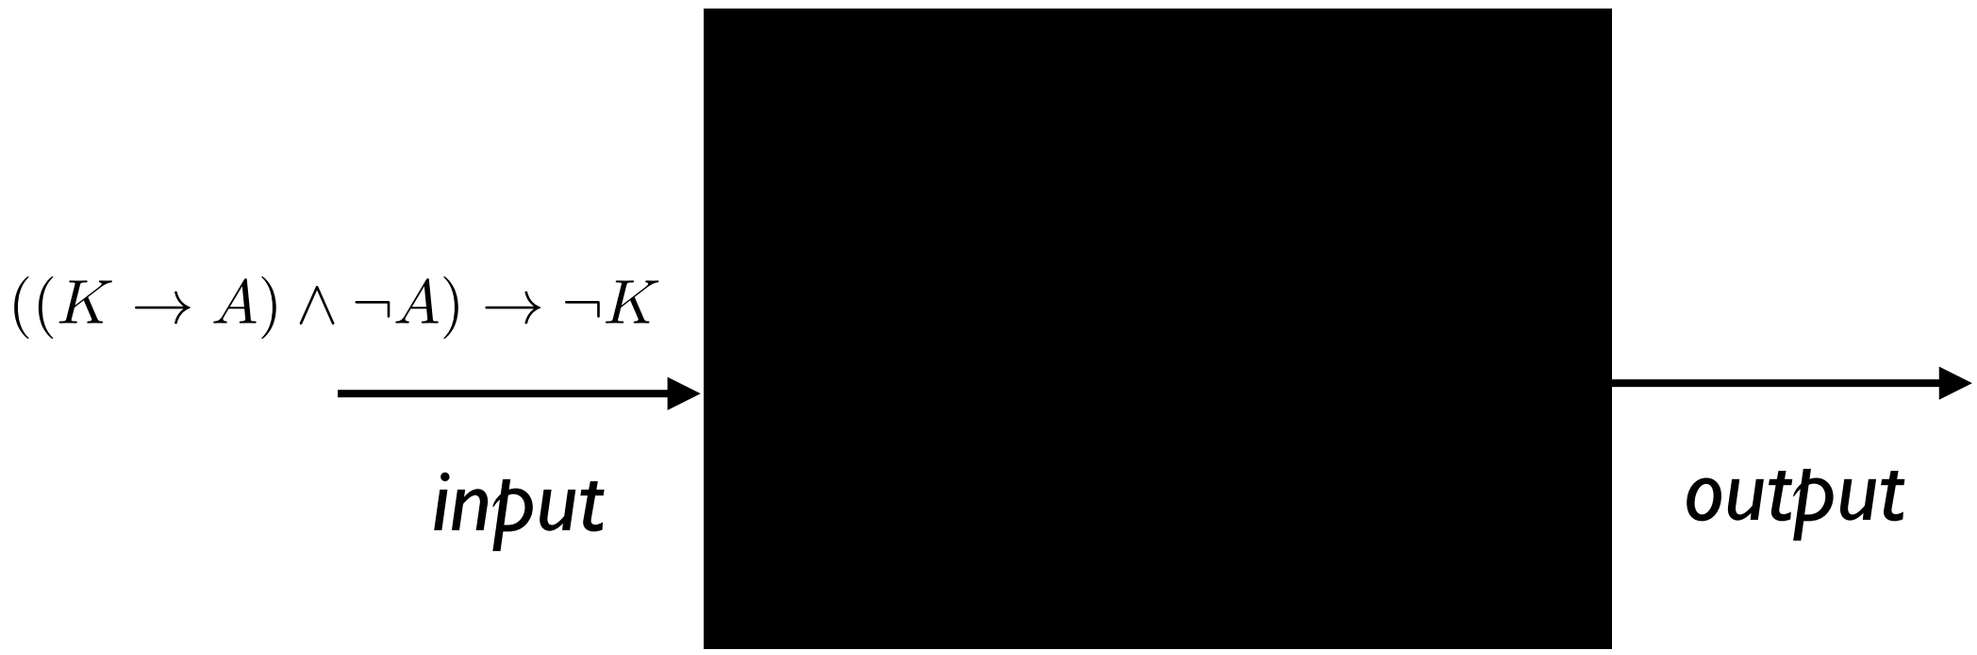
Hard!! — for apparently no polynomial-time algorithm for this!

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



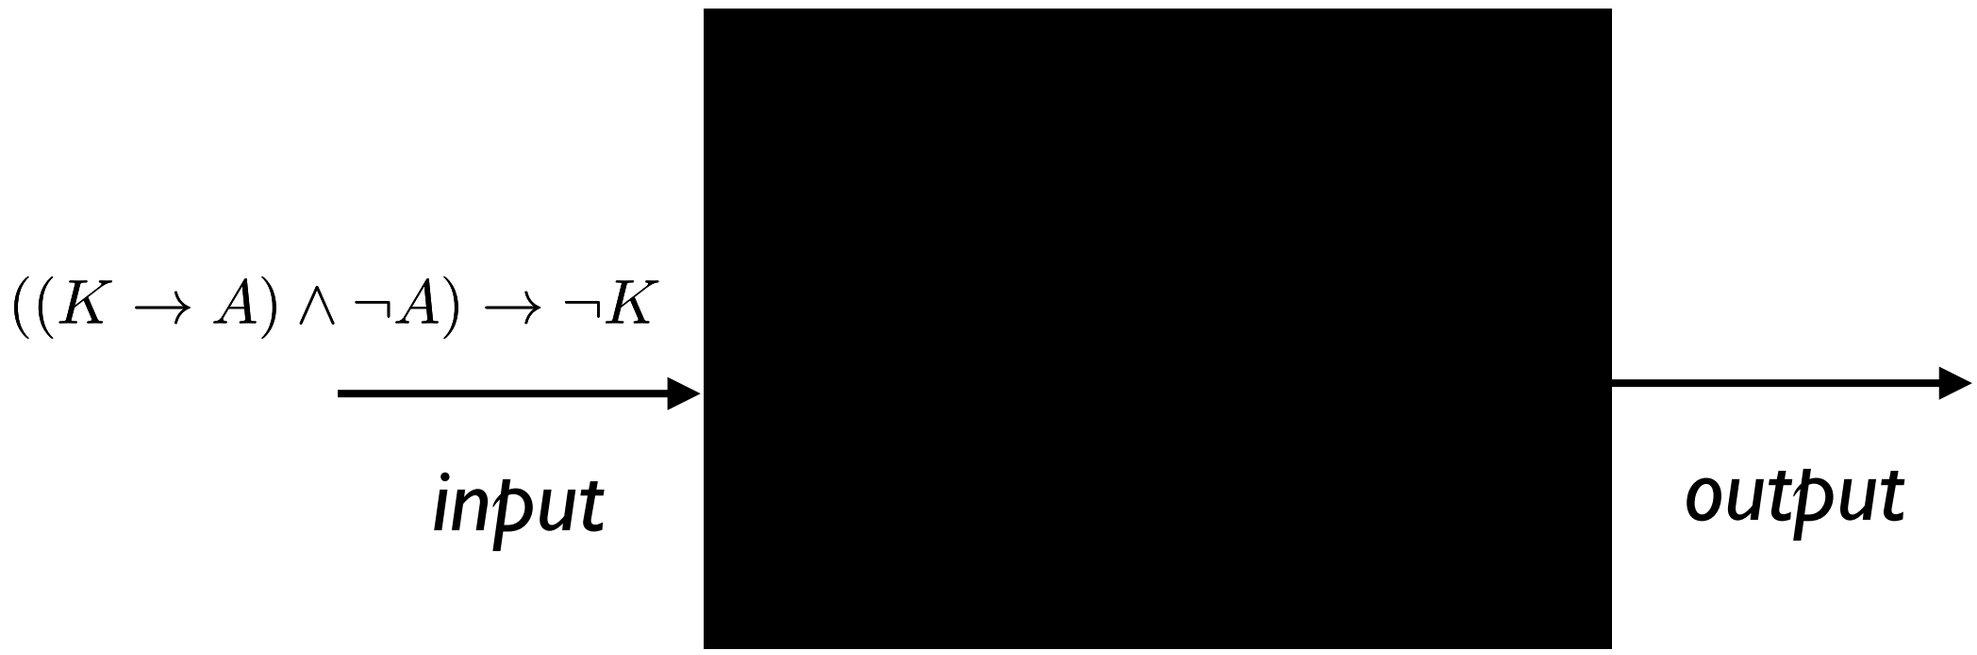
Hard!! — for apparently no polynomial-time algorithm for this!

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



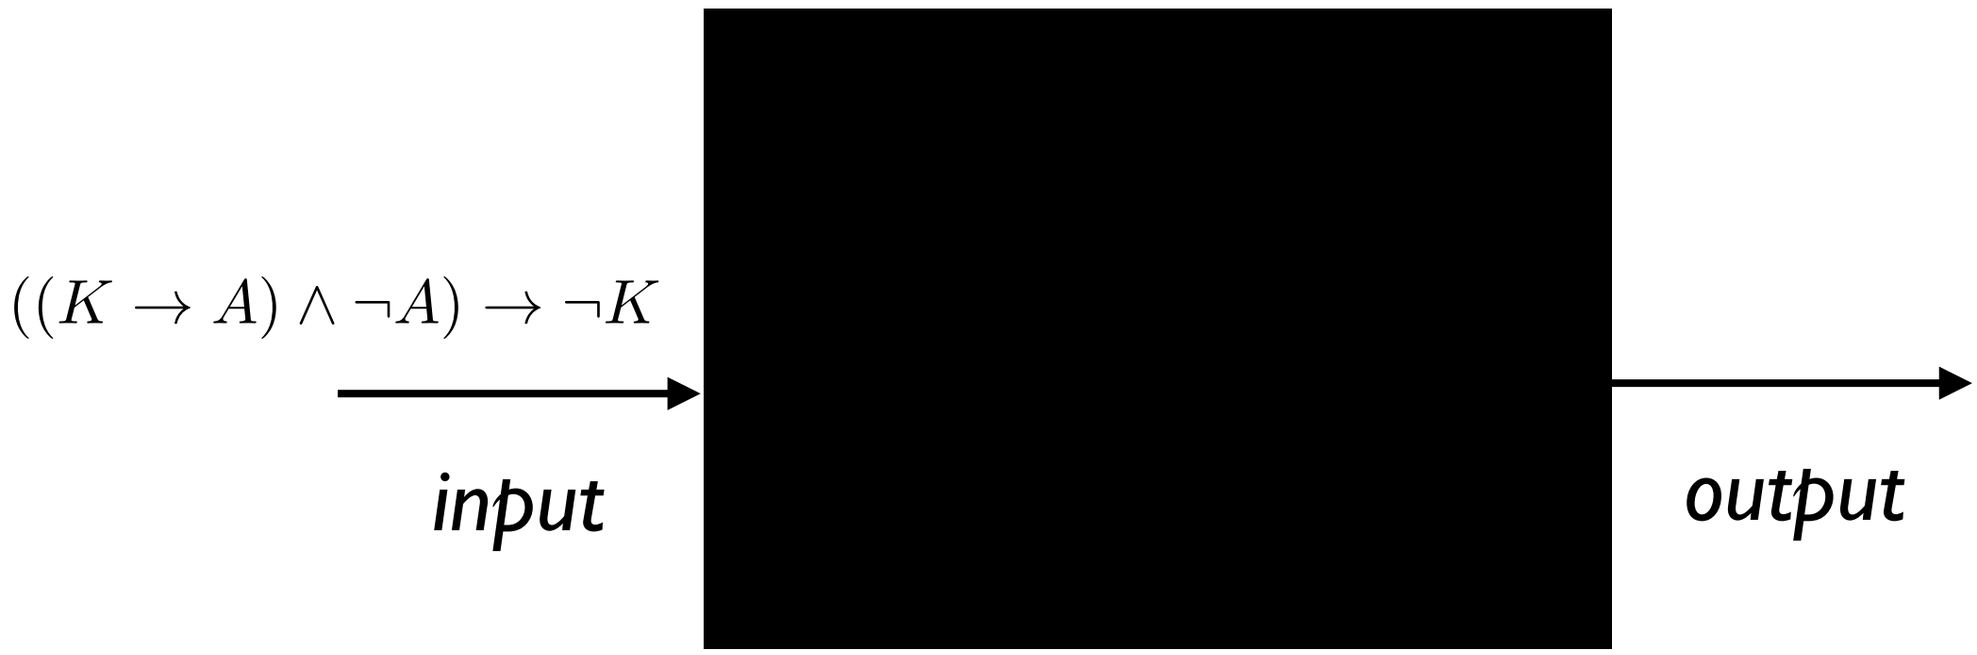
Hard!! — for apparently no polynomial-time algorithm for this!

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



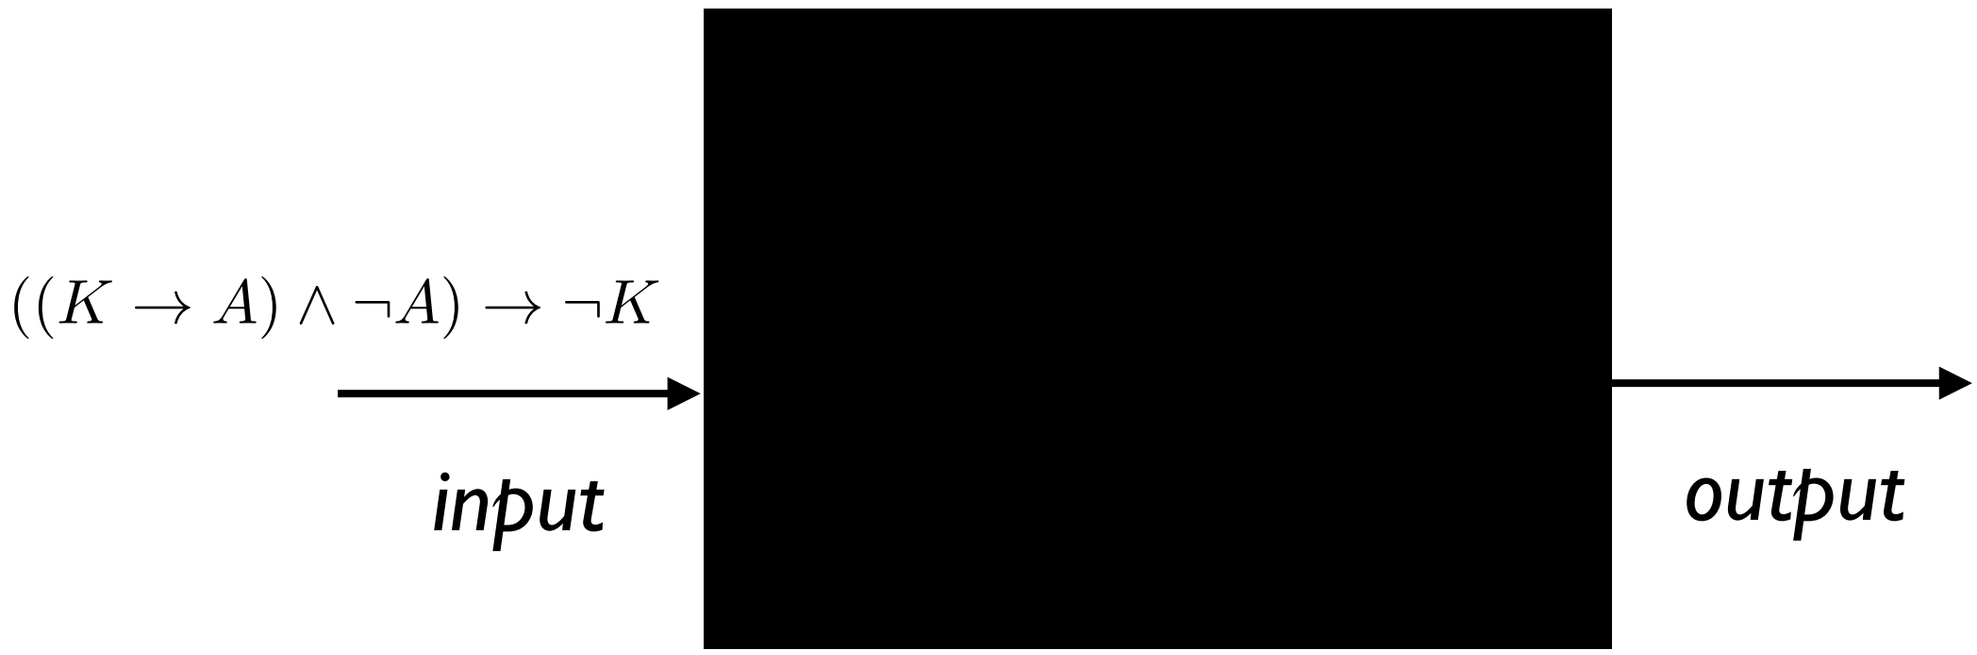
Hard!! — for apparently no polynomial-time algorithm for this!

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



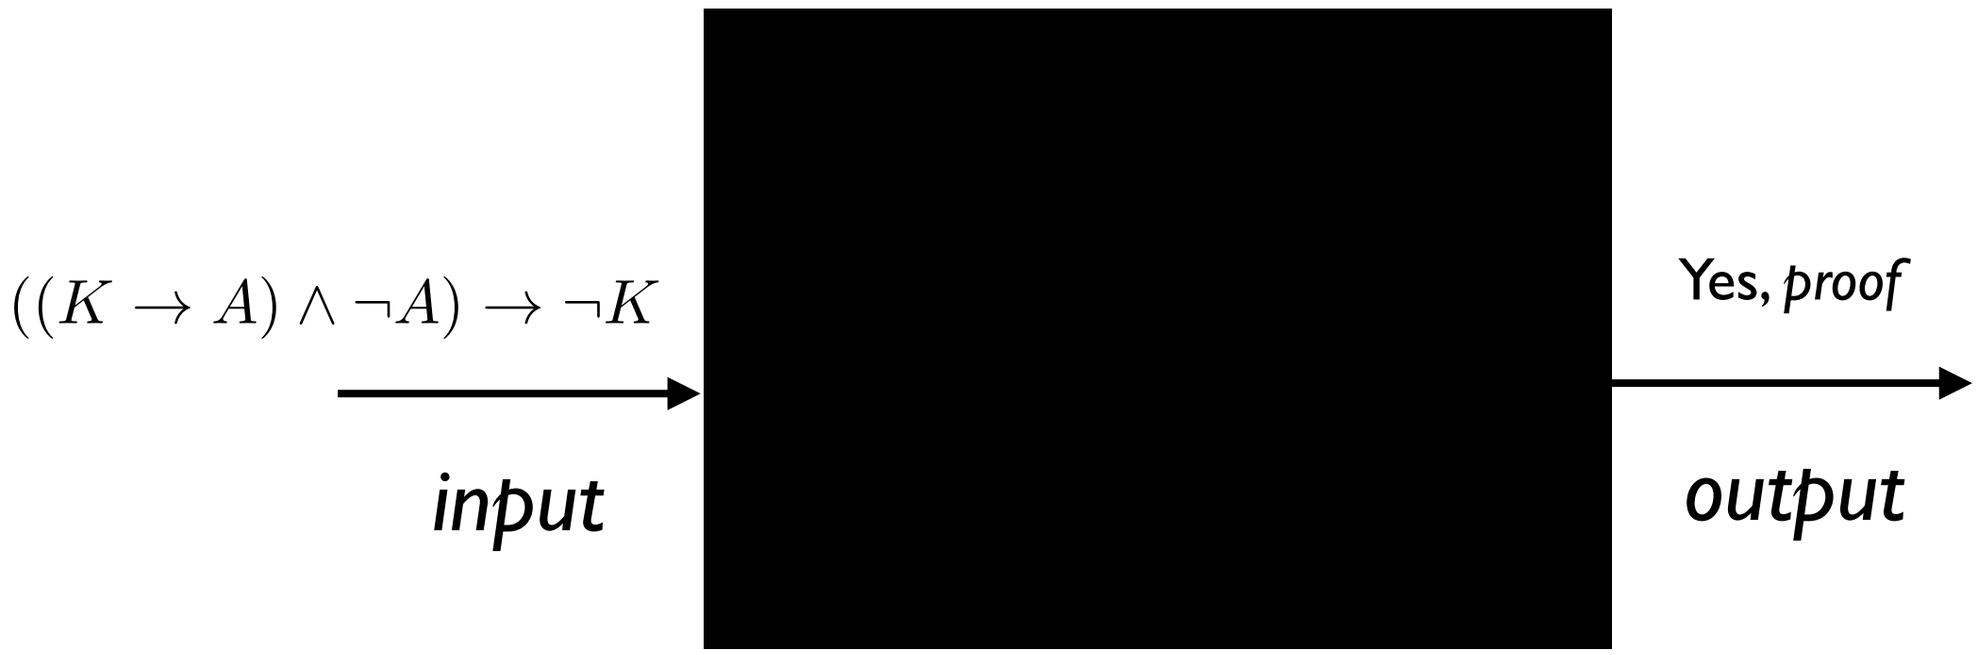
Hard!! — for apparently no polynomial-time algorithm for this!

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

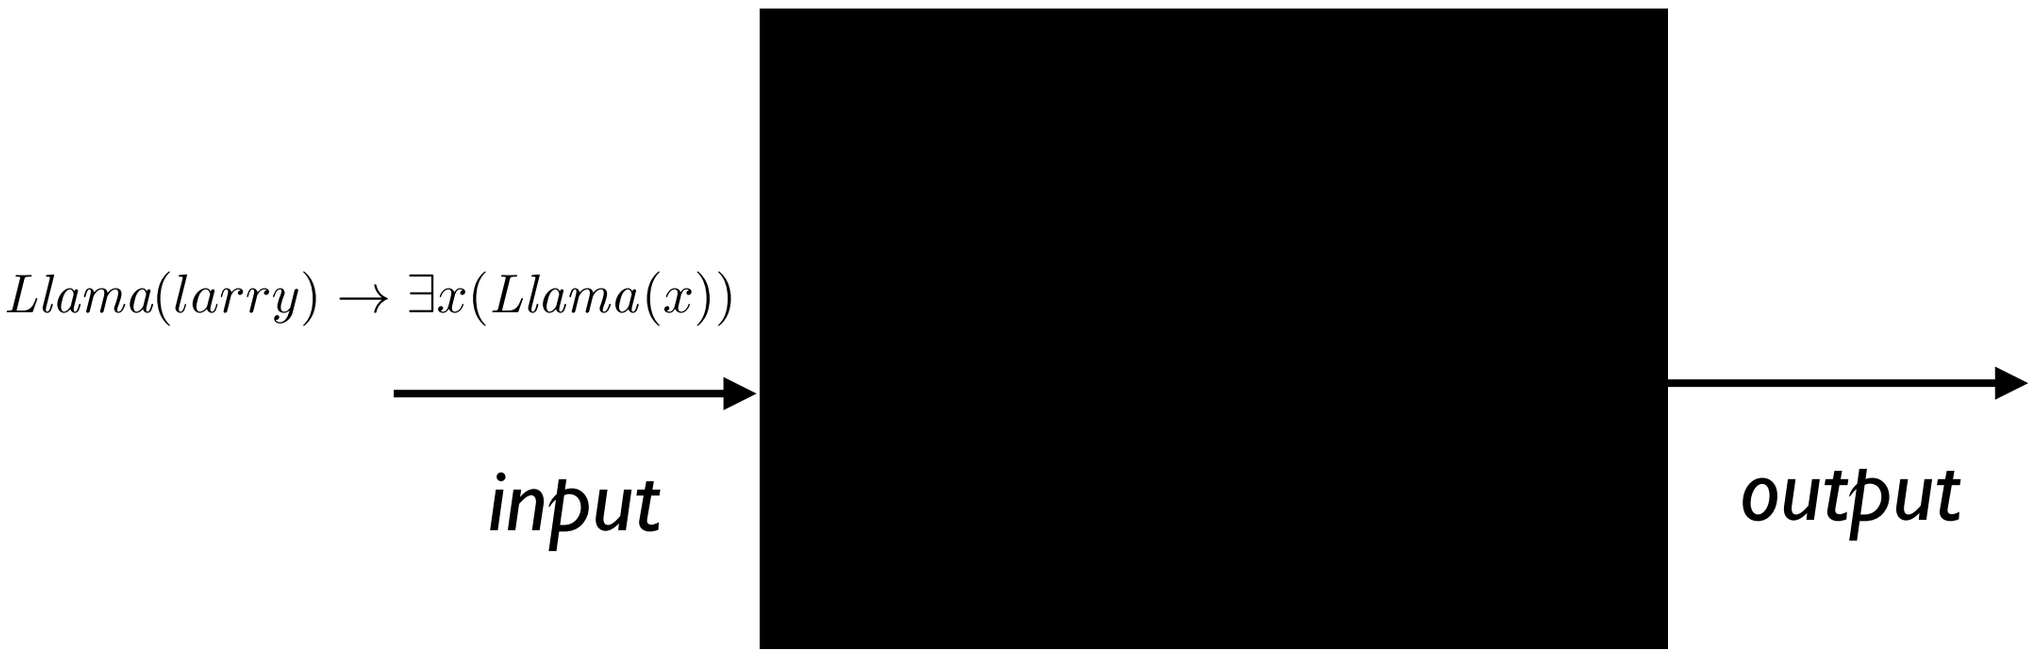


Hard!! — for apparently no polynomial-time algorithm for this!

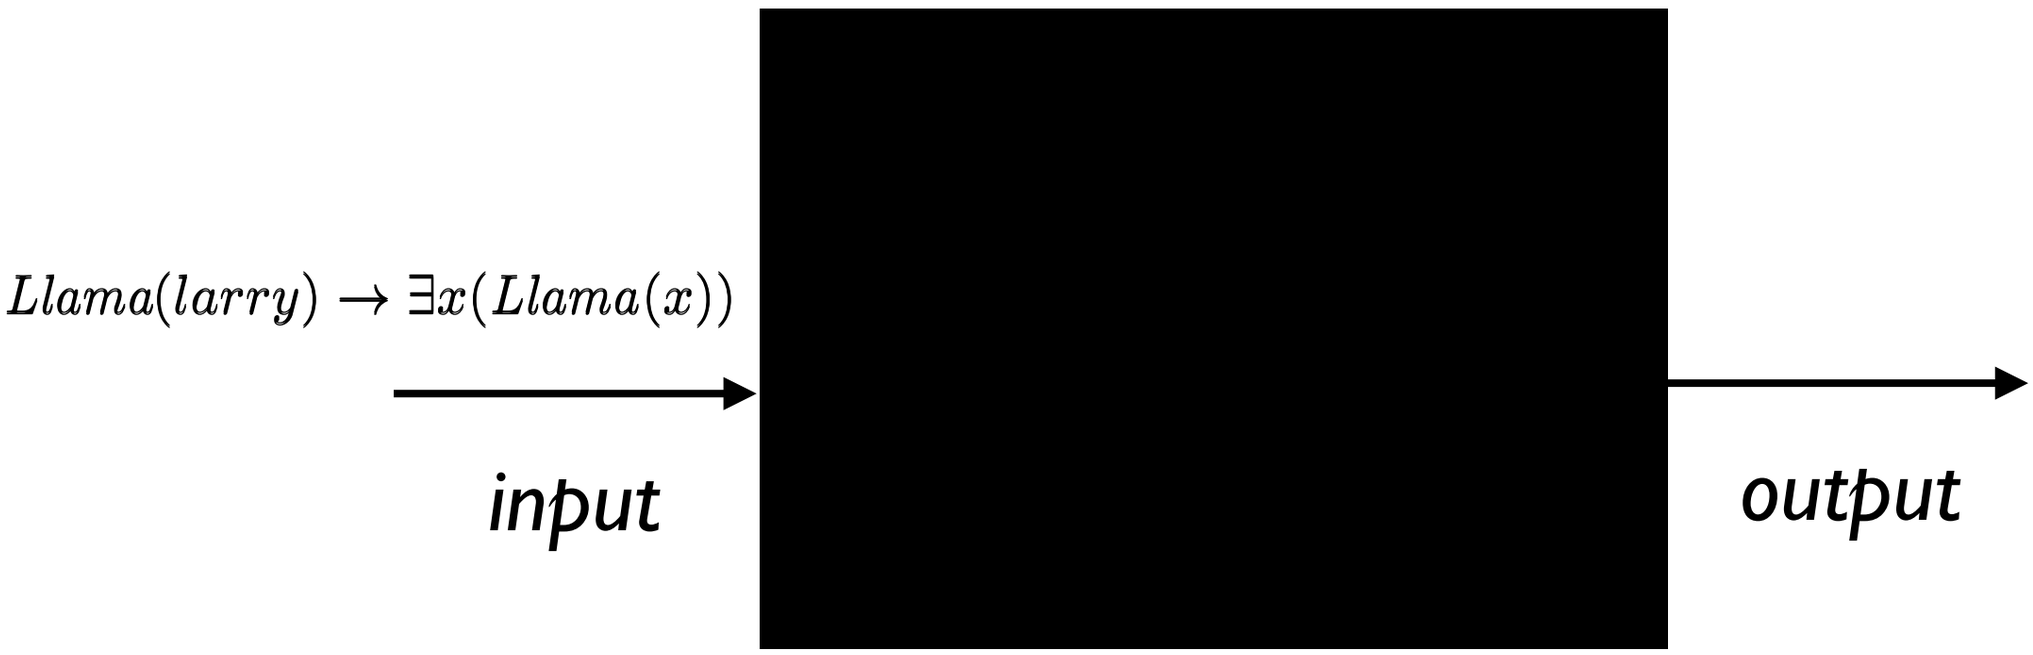
And now, the Theoremhood Decision Problem,
i.e., the *Entscheidungsproblem*,
($\text{THEOREM}_{\text{FOL}}$)
for First-Order Logic (FOL)



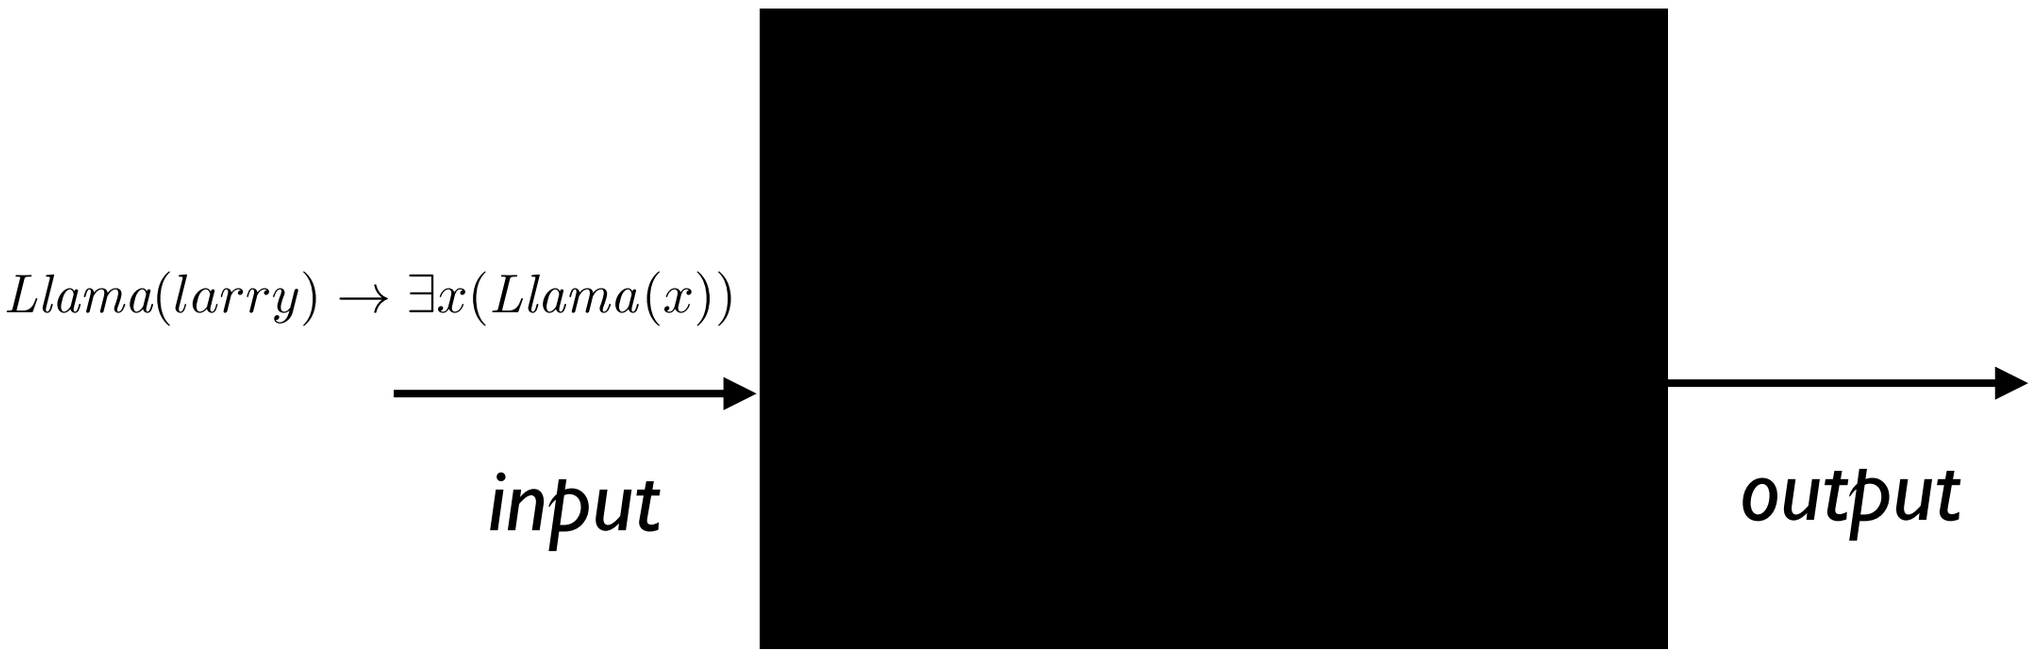
And now, the Theoremhood Decision Problem,
i.e., the *Entscheidungsproblem*,
($\text{THEOREM}_{\text{FOL}}$)
for First-Order Logic (FOL)



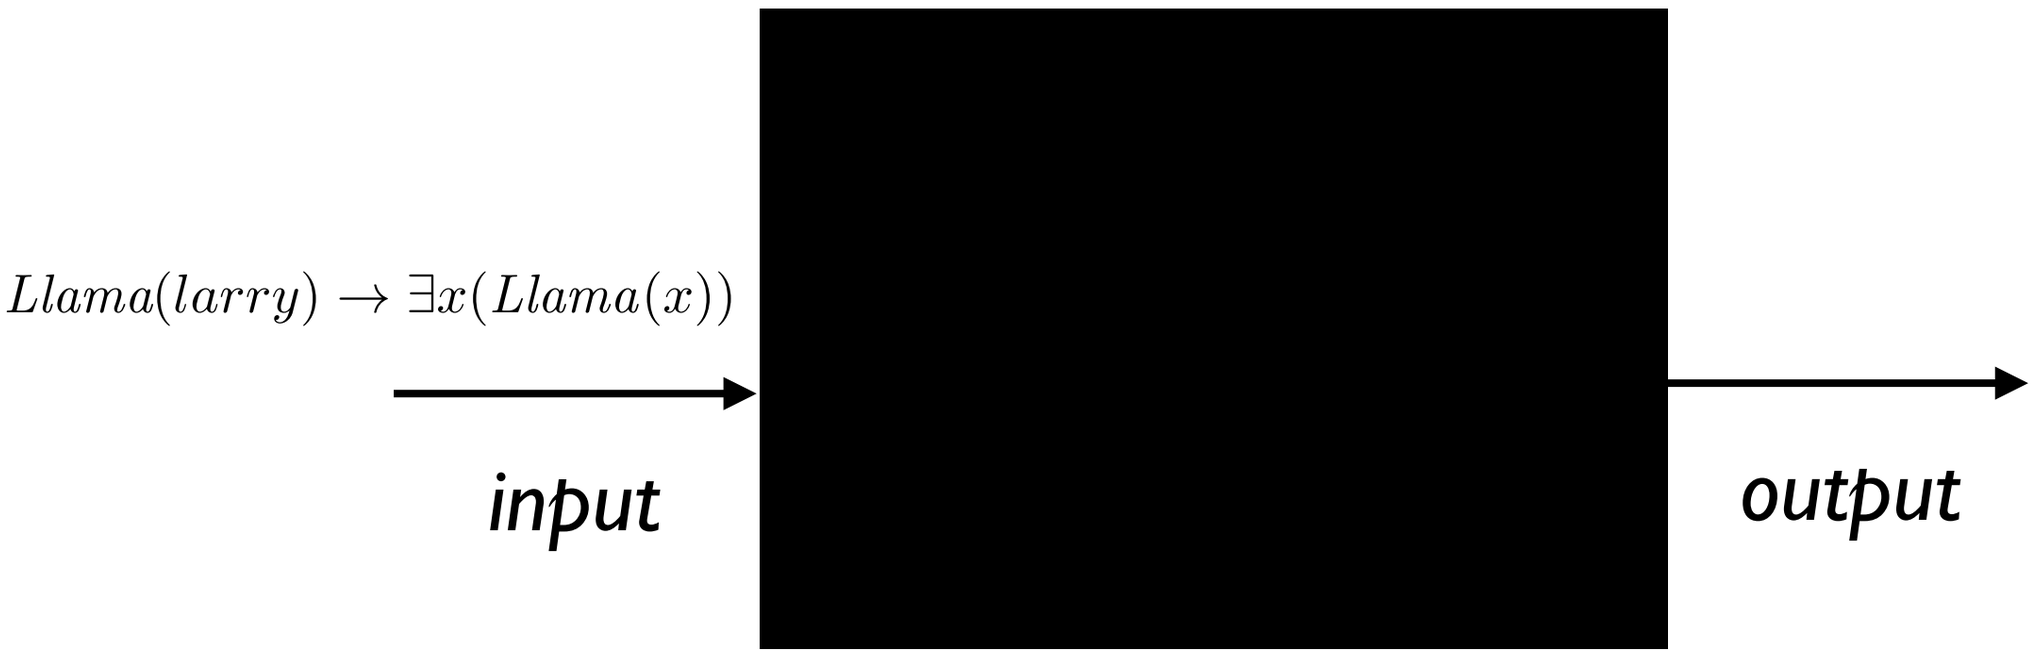
And now, the Theoremhood Decision Problem, i.e., the *Entscheidungsproblem*, ($\text{THEOREM}_{\text{FOL}}$) for First-Order Logic (FOL)



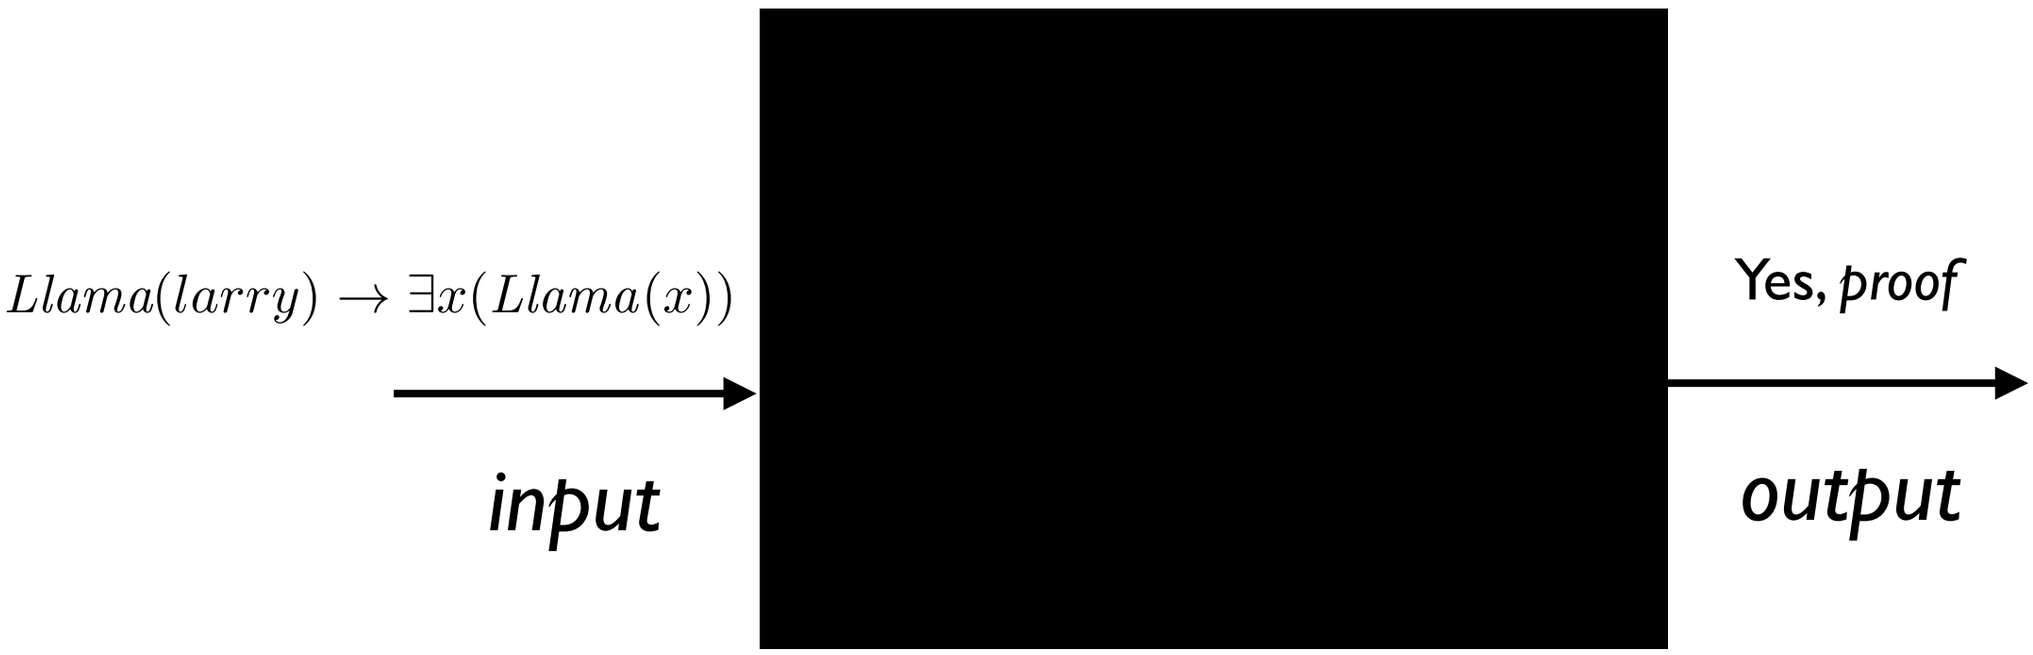
And now, the Theoremhood Decision Problem,
i.e., the *Entscheidungsproblem*,
($\text{THEOREM}_{\text{FOL}}$)
for First-Order Logic (FOL)



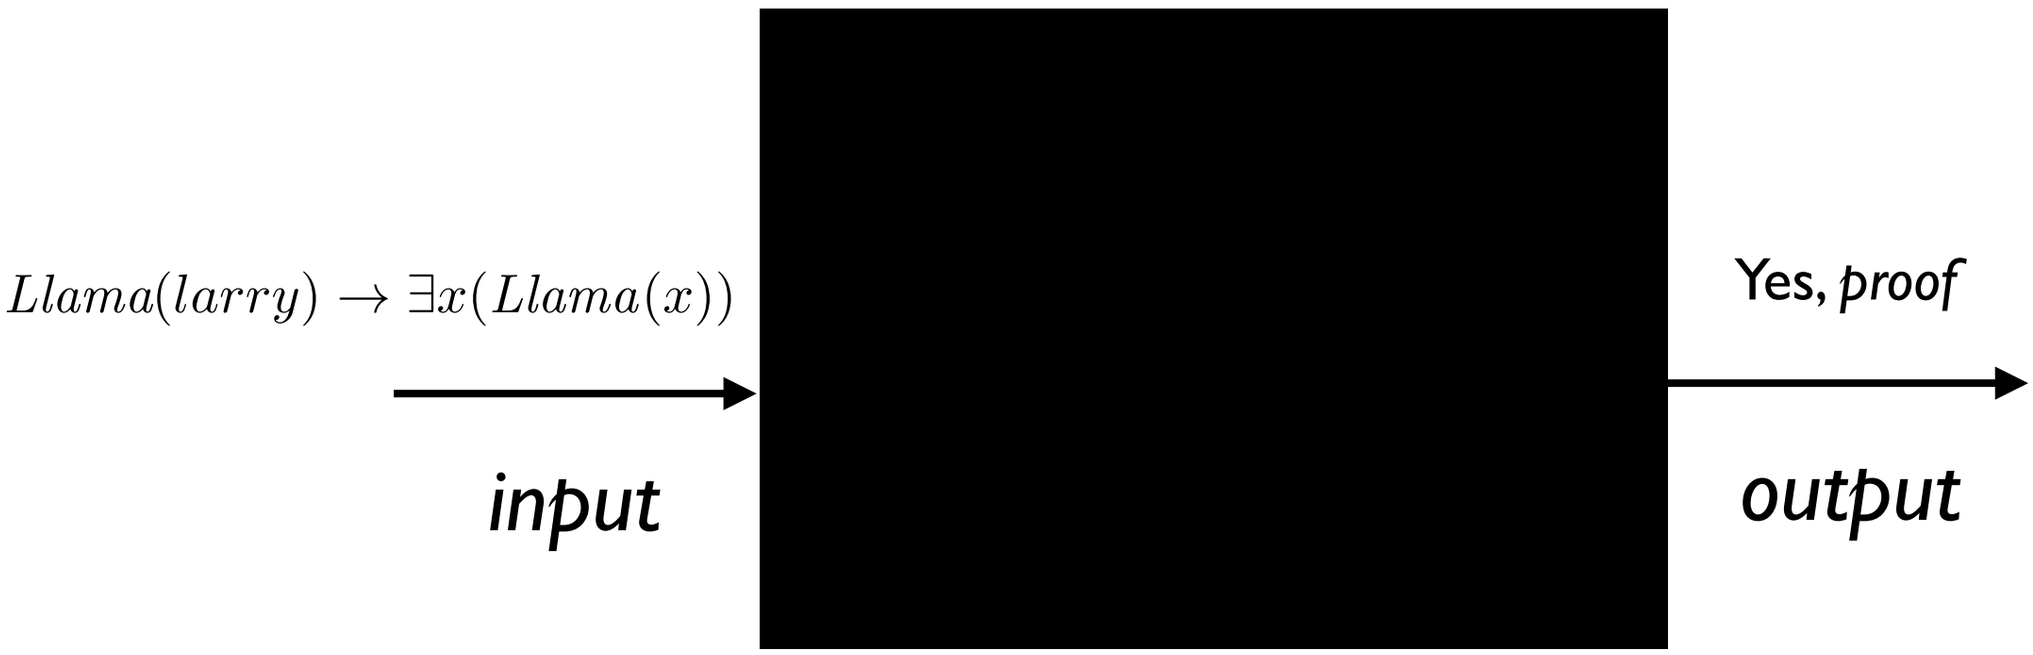
And now, the Theoremhood Decision Problem,
i.e., the *Entscheidungsproblem*,
($\text{THEOREM}_{\text{FOL}}$)
for First-Order Logic (FOL)



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

ASI = AI+++++++

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).

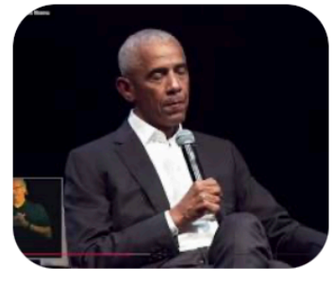
(Good-Chalmers Argument)

(Kurzweil is an "extrapolationist.")

ASI = AI+++++++

Trump, Obama, . . .

Former US President Barack Obama warned in a September 2026 interview at Colgate University that artificial intelligence has hit a "hockey stick" exponential growth trajectory driven by recursive self-learning. Medium +1



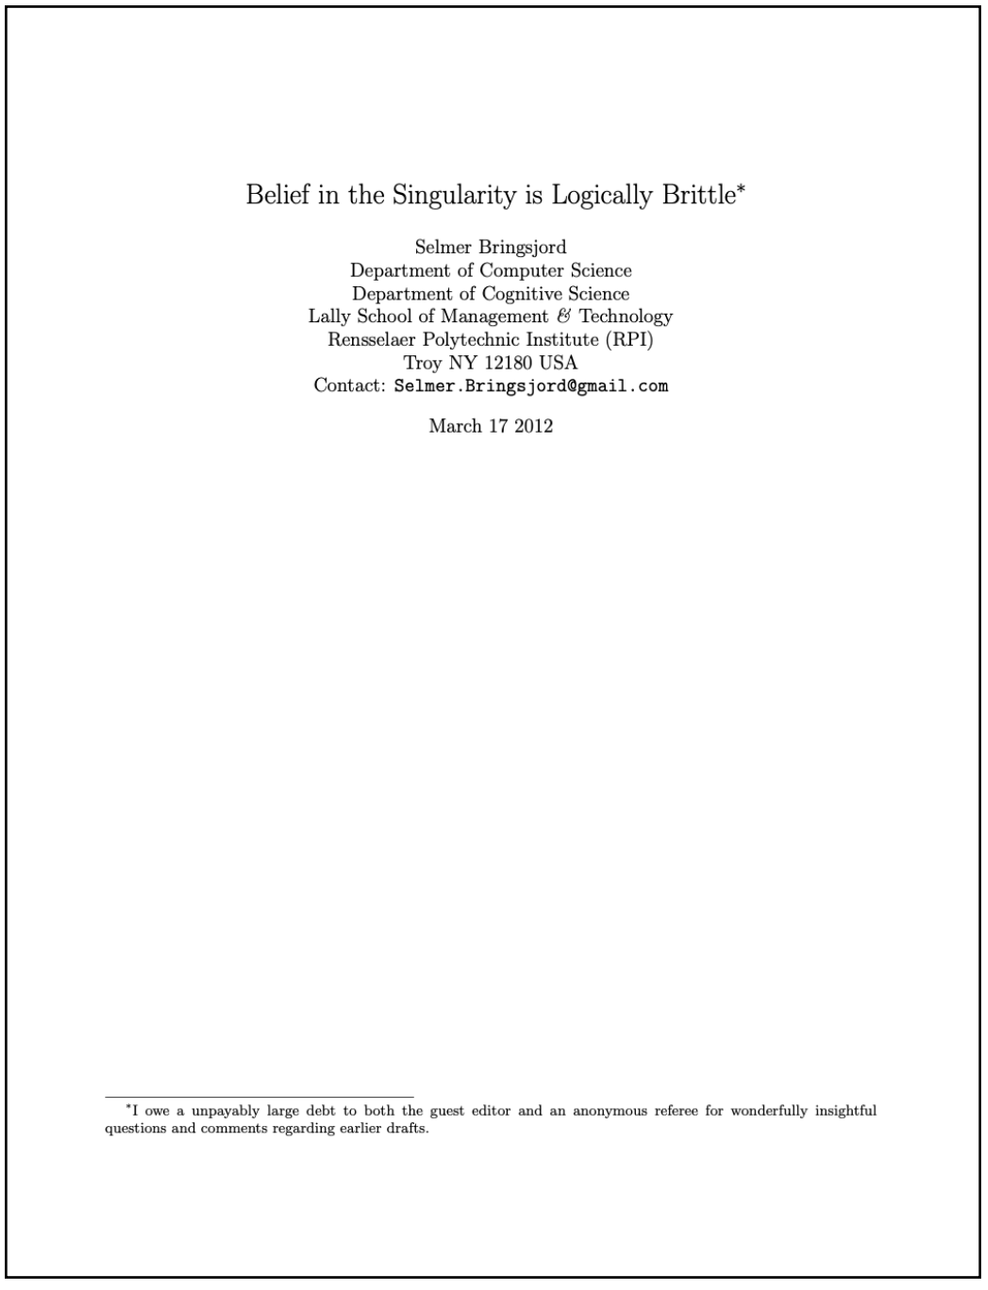
Obama's View on Recursive Self-Improvement

- **The Shift to Self-Teaching**: Obama stated that machines have reached a point of recursive learning where they can teach and improve themselves without waiting for human input. LinkedIn
- **The Learning Ratio**: He explained that while humans provided 90% of model learning a year prior, the ratio shifted to 50/50 in a very short period and could soon reach 90% AI-driven. YouTube · C-SPAN
- **Exponential Acceleration**: This transition puts AI capability on a sharp upward "hockey stick" curve, meaning systems are getting smarter much faster than before. LinkedIn
- **Not Overhyped**: Obama noted that while commercial company valuations might be overhyped, the underlying technology and its accelerated speed are very real. YouTube +1

Refuting the Good argument for The Singularity …

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*?

More Thorough Version of Refutation



Belief in the Singularity is Logically Brittle*

Selmer Bringsjord
Department of Computer Science
Department of Cognitive Science
Lally School of Management & Technology
Rensselaer Polytechnic Institute (RPI)
Troy NY 12180 USA
Contact: Selmer.Bringsjord@gmail.com

March 17 2012

*I owe a unpayably large debt to both the guest editor and an anonymous referee for wonderfully insightful questions and comments regarding earlier drafts.

https://kryten.mm.rpi.edu/SB_singularity_math_final.pdf

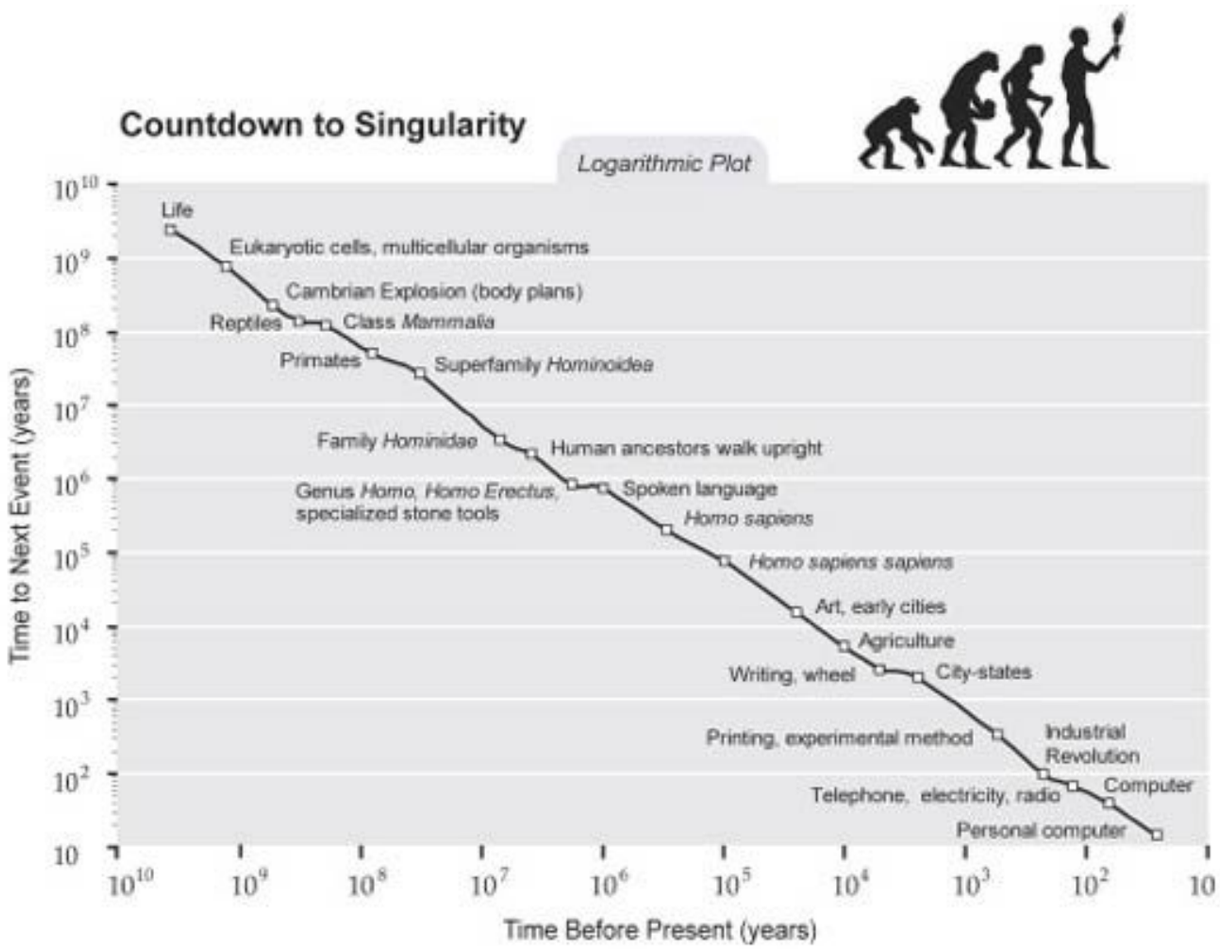
Some James Oswaldian Content

• • •

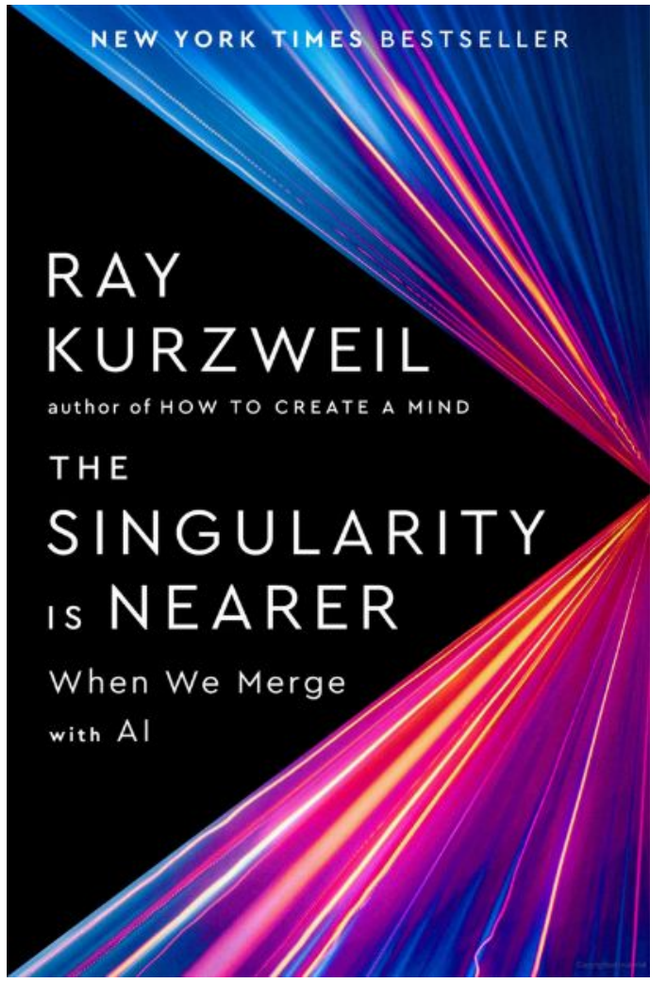
Kurzweil's Technology Scaling Argument

Given sufficient technology, we can do anything physically possible. Minimally HI is possible, we have it, AHI is possible.

Technology itself, including life itself, scales super-exponentially and has for billions of years. Thus we will eventually reach a point where AHI is technologically possible, and probably ASI and the Singularity.



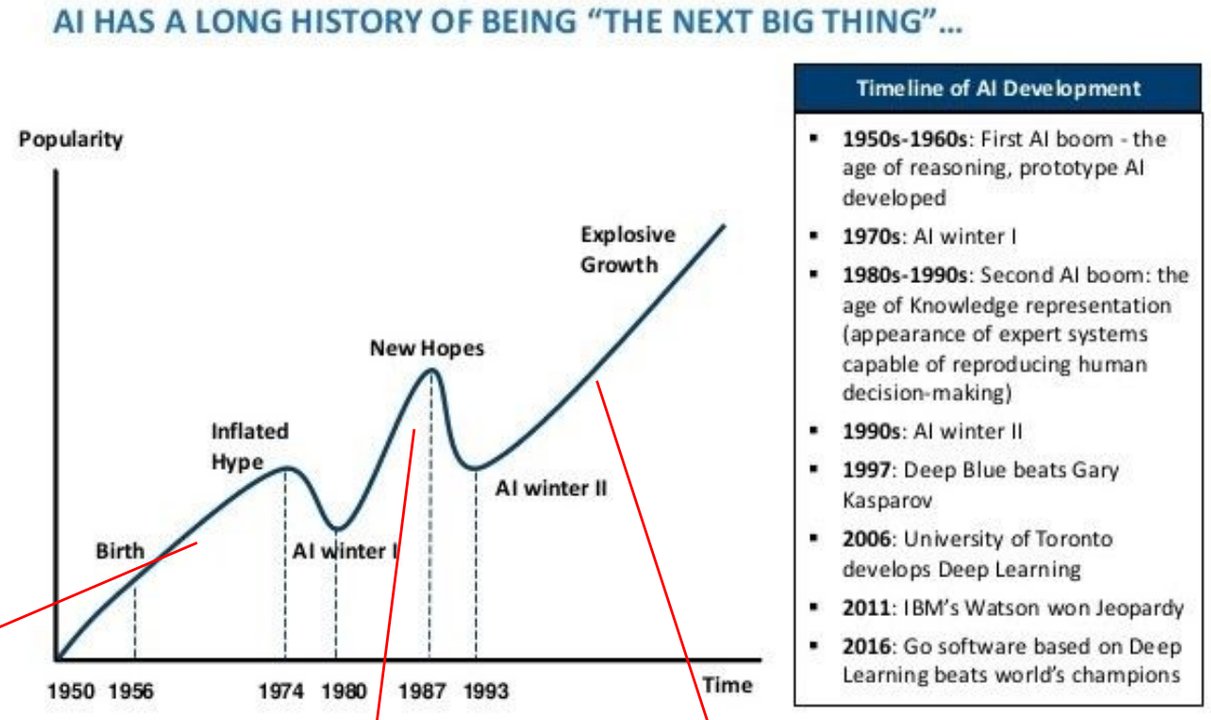
Suggested Reading Material



A History of AGI Research Timeline

Almost every AI researcher before the second AI winter had AGI as a goal. “I’m working on AGI… but first I'll show how good my approach is by using it to solve a narrow problem.”

ML based approaches which work very well on narrow problems & largely took AGI off the table as a goal.



First wave of logic based AI, AGI seen to be “only about 20 years away”

Second wave of logic based AI via KRR approaches (Japanese 5th Generation Project, CYC)

ML approaches possible on new hardware offer never before seen performance on narrow tasks

Core AGI Research Area: Formalizing Intelligence



Intelligence is ability to perform in all environments

$$\Upsilon(\pi) := \sum_{\mu \in E} 2^{-K(\mu)} V_{\mu}^{\pi}.$$



Intelligence is skill acquisition efficiency

Intelligence of system IS over $scope$ (optimal case):

$$I_{IS,scope}^{opt} = \underset{T \in scope}{Avg} \left[\omega_{T,\Theta} \cdot \Theta \sum_{C \in Cur_T^{opt}} \left[P_C \cdot \frac{GD_{IS,T,C}^{\Theta}}{P_{IS,T}^{\Theta} + E_{IS,T,C}^{\Theta}} \right] \right]$$

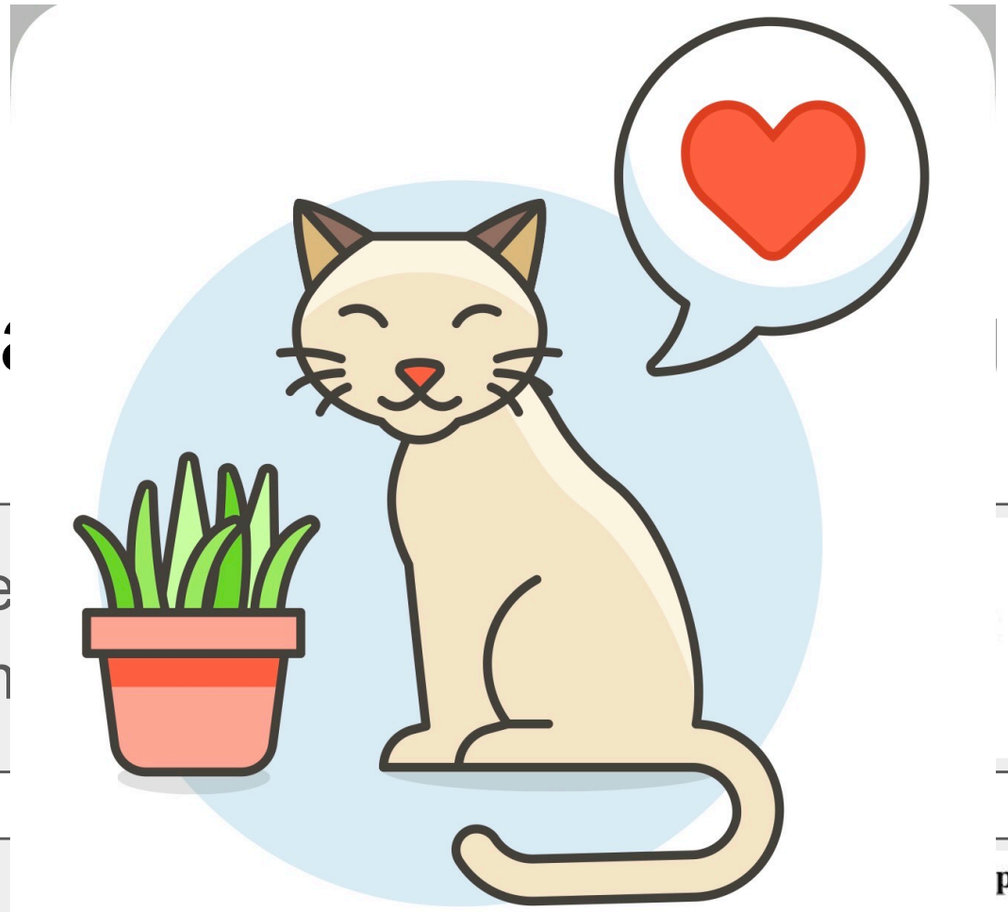


Intelligence is (internal) cognitive representation & reasoning capacity*

$$\Lambda(a,t)_{i,j} = \max_{\phi} \left\{ \mu^i(\phi) \mid \phi \in \Delta\Big(\omega_j\big[o(a,t)\big], \omega_j\big[i(a,t)\big]\Big) \right\}$$

*My take on Selmer's measure, not his

Core AGI Resea Intelligence





Intellige
perform



Intellige
acquisiti



Intelligenc
cognitive r
reasoning

*My take on

Cat Nap by J. Bethel Brown

Mr. Langersham often comes home from work in the afternoons to find his five cats sleeping. Recently, he discovered each feline (including Grizelda) napping in a different location in the house. He has owned each cat for a different number of years (from one through five), and each is a different age (three, four, five, six, or eight years old). Of course, Mr. Langersham hasn't owned any of his pets for more years than its own age. Can you match each cat with the location in which it was sleeping, its age, and the number of years for which it has been with Mr. Langersham?

$$\sum_{\mu \in E} 2^{-K(\mu)} V_{\mu}^{\pi}.$$

ptimal case):

$$\cdot \Theta \sum_{C \in Cur_T^{opt}} \left[P_C \cdot \frac{GD_{IS,T,C}^{\Theta}}{P_{IS,T}^{\Theta} + E_{IS,T,C}^{\Theta}} \right]$$

$$) \mid \phi \in \Delta\left(\omega_j[o(a,t)], \omega_j[i(a,t)]\right)\}$$

A vaere logisk kan gi dyp glede!