The Balderdash that is GPT-3 AI; Motivating Paradoxes, Puzzles, Part 2

Selmer Bringsjord

Intro to (Formal) Logic (and AI) = IFLAI1

1/20/22

Selmer.Bringsjord@gmail.com

Economists Pin More Blame on Tech for Rising Inequality

Recent research underlines the central role that automation has played in widening disparities.

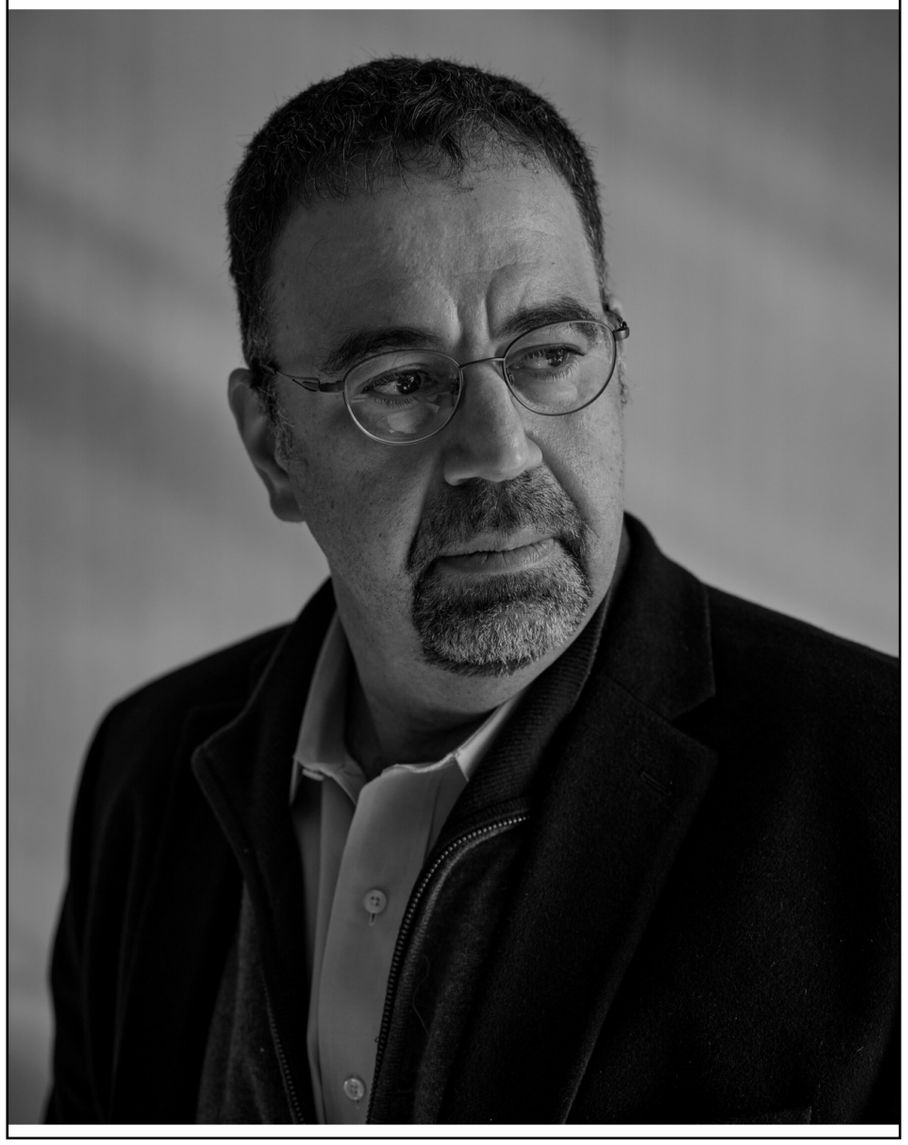

Economists Pin More Blame on Tech for Rising Inequality

Recent research underlines the central role that automation has played in widening disparities.

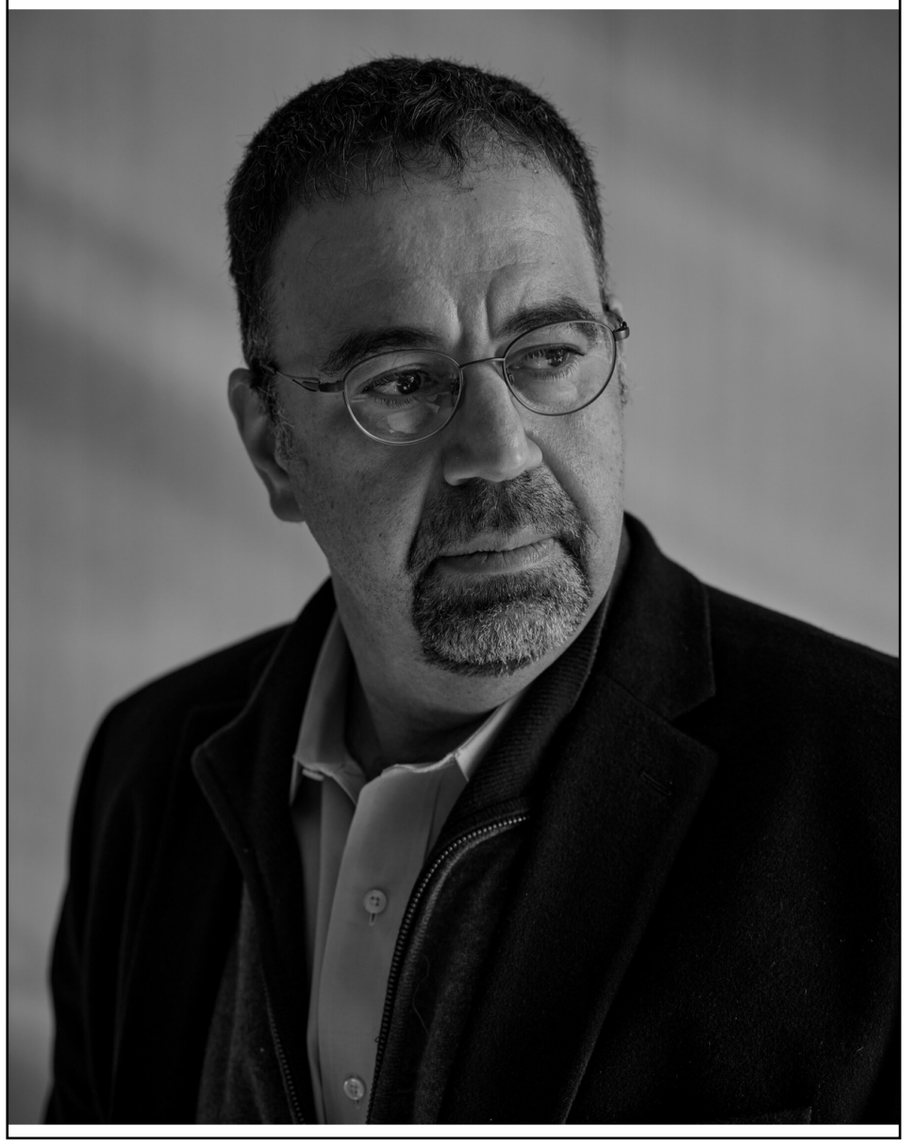

At least half the rising gap in wages among American workers in the last 40 years comes from the automation of tasks once done by people, says Daron Acemoglu, an economist at M.I.T. Cody O'Loughlin for The New York Times

By **Steve Lohr**

Published Jan. 11, 2022 Updated Jan. 12, 2022

Economists Pin More Blame on Tech for Rising Inequality

Recent research underlines the central role that automation has played in widening disparities.

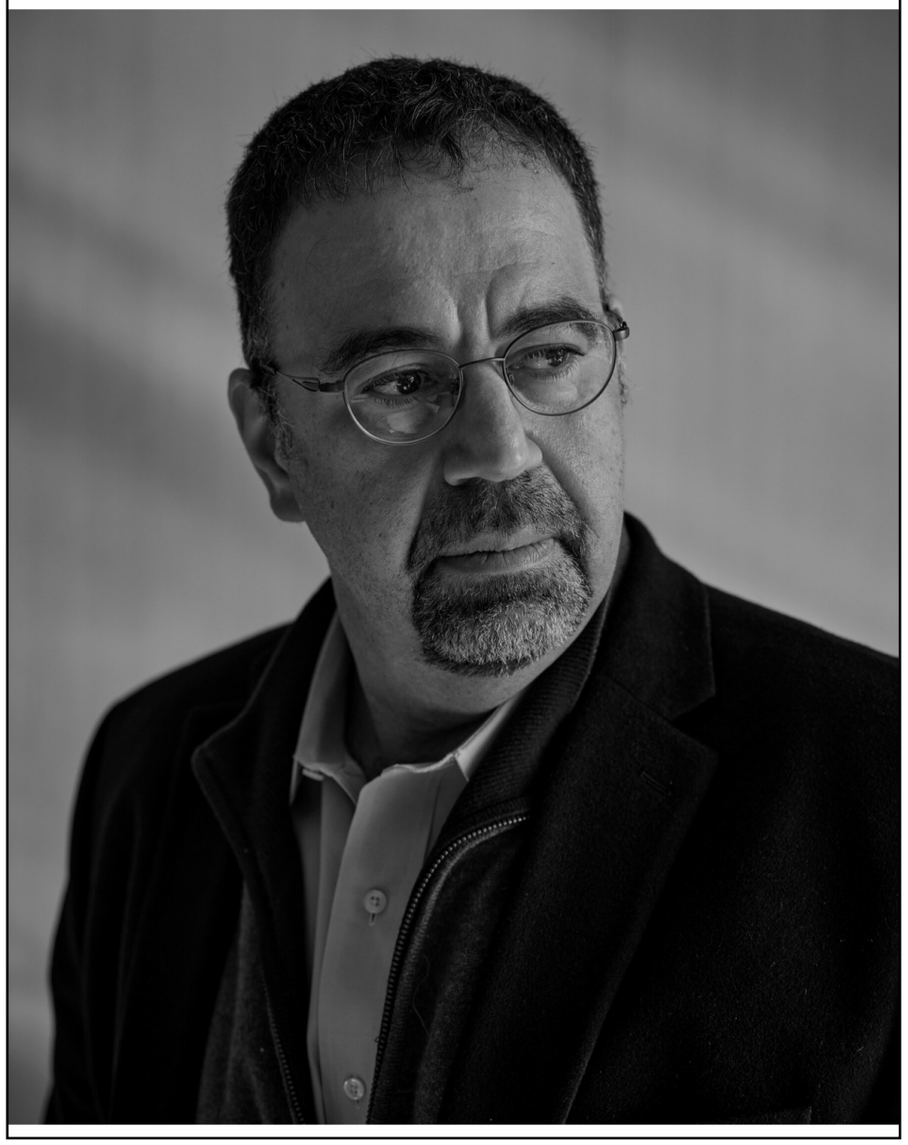

At least half the rising gap in wages among American workers in the last 40 years comes from the automation of tasks once done by people, says Daron Acemoglu, an economist at M.I.T. Cody O'Loughlin for The New York Times

By **Steve Lohr**

Published Jan. 11, 2022 Updated Jan. 12, 2022

Mr. Acemoglu, like some other economists, has altered his view of technology over time. In economic theory, technology is almost a magic ingredient that both increases the size of the economic pie and makes nations richer. He recalled working on a textbook more than a decade ago that included the standard theory. Shortly after, while doing further research, he had second thoughts.

“It’s too restrictive a way of thinking,” he said. “I should have been more open-minded.”

Mr. Acemoglu is no enemy of technology. Its innovations, he notes, are needed to address society’s biggest challenges, like climate change, and to deliver economic growth and rising living standards. His wife, Asuman Ozdaglar, is the head of the electrical engineering and computer science department at M.I.T.

forum/AI & economics

Rage against the machine

WHAT ARE OUR PROSPECTS AS ADVANCES IN AI CHANGE THE LABOUR EQUATION?

SELMER BRINGSJORD AND **JOE JOHNSON**

SEE TROUBLE AHEAD

http://kryten.mm.rpi.edu/SB_JJ_Rage_Against_Machine_offprint.pdf

Logistics …

Logistics …

- [Reading the syllabus is very helpful :).]

- Codes for registering for use of HyperGrader® & HyperSlate® cld be on sale as early as tomorrow in the Follett Bookstore. I will now review how to register. There is a tutorial video on our course web page that shows this as well. YOU MUST PUT YOUR UNIVERSITY ID (= RIN, for you) IN YOUR PROFILE.

Logistics …

- [Reading the syllabus is very helpful :).]

- Codes for registering for use of HyperGrader® & HyperSlate® cld be on sale as early as tomorrow in the Follett Bookstore. I will now review how to register. There is a tutorial video on our course web page that shows this as well. YOU MUST PUT YOUR UNIVERSITY ID (= RIN, for you) IN YOUR PROFILE.

- Hopefully you've studied last mtg's "stunners." Questions?

What *is* Logic?

What *is* Logic?

- The key to becoming rational.
- "The science and engineering of reasoning." — so the not-unreasonable slogan goes.
- The only invincible subject there is.
- The basis for the formal sciences (from mathematics to game theory to decision theory to probability calculi to axiomatic physics ….) — and hence the basis for disciplines based on the formal sciences, e.g., …
  - Engineering! Computer Science!
  - Mathematics itself: see "reverse mathematics"!
- The way of escape from shallow content and context to pure, immaterial, and immortal form and structure (which is why the exotic, imaginary, and seemingly non-sensical is so pedagogically useful).
- The most challenging subject there is.
- One of the chief differentiators between dogs and monkeys versus you (let alone bears and you); and mindless machines (like Deep Blue & Watson) versus you.
- A key to riches.
- The key to divining the meaning of life (and other such big questions).
- The better way to program computers; and fundamentally the *only* way to *reliably* program computers.
- One of two fundamental approaches to studying minds, and replicating/simulating minds in machines…
- The thing many creatures of fiction have mastered — have you (as a New Yorker)?…
- …

What *is* Logic?

- The key to becoming rational.
- "The science and engineering of reasoning." — so the not-unreasonable slogan goes.
- The only invincible subject there is.
- The basis for the formal sciences (from mathematics to game theory to decision theory to probability calculi to axiomatic physics ….) — and hence the basis for disciplines based on the formal sciences, e.g., …
  - Engineering! Computer Science!
  - Mathematics itself: see "reverse mathematics"!
- The way of escape from shallow content and context to pure, immaterial, and immortal form and structure (which is why the exotic, imaginary, and seemingly non-sensical is so pedagogically useful).
- The most challenging subject there is.
- One of the chief differentiators between dogs and monkeys versus you (let alone bears and you); and mindless machines (like Deep Blue & Watson) versus you.
- A key to riches.
- The key to divining the meaning of life (and other such big questions).
- The better way to program computers; and fundamentally the *only* way to *reliably* program computers.
- One of two fundamental approaches to studying minds, and replicating/simulating minds in machines…
- The thing many creatures of fiction have mastered — have you (as a New Yorker)?…
- …

Non-Logicist AI

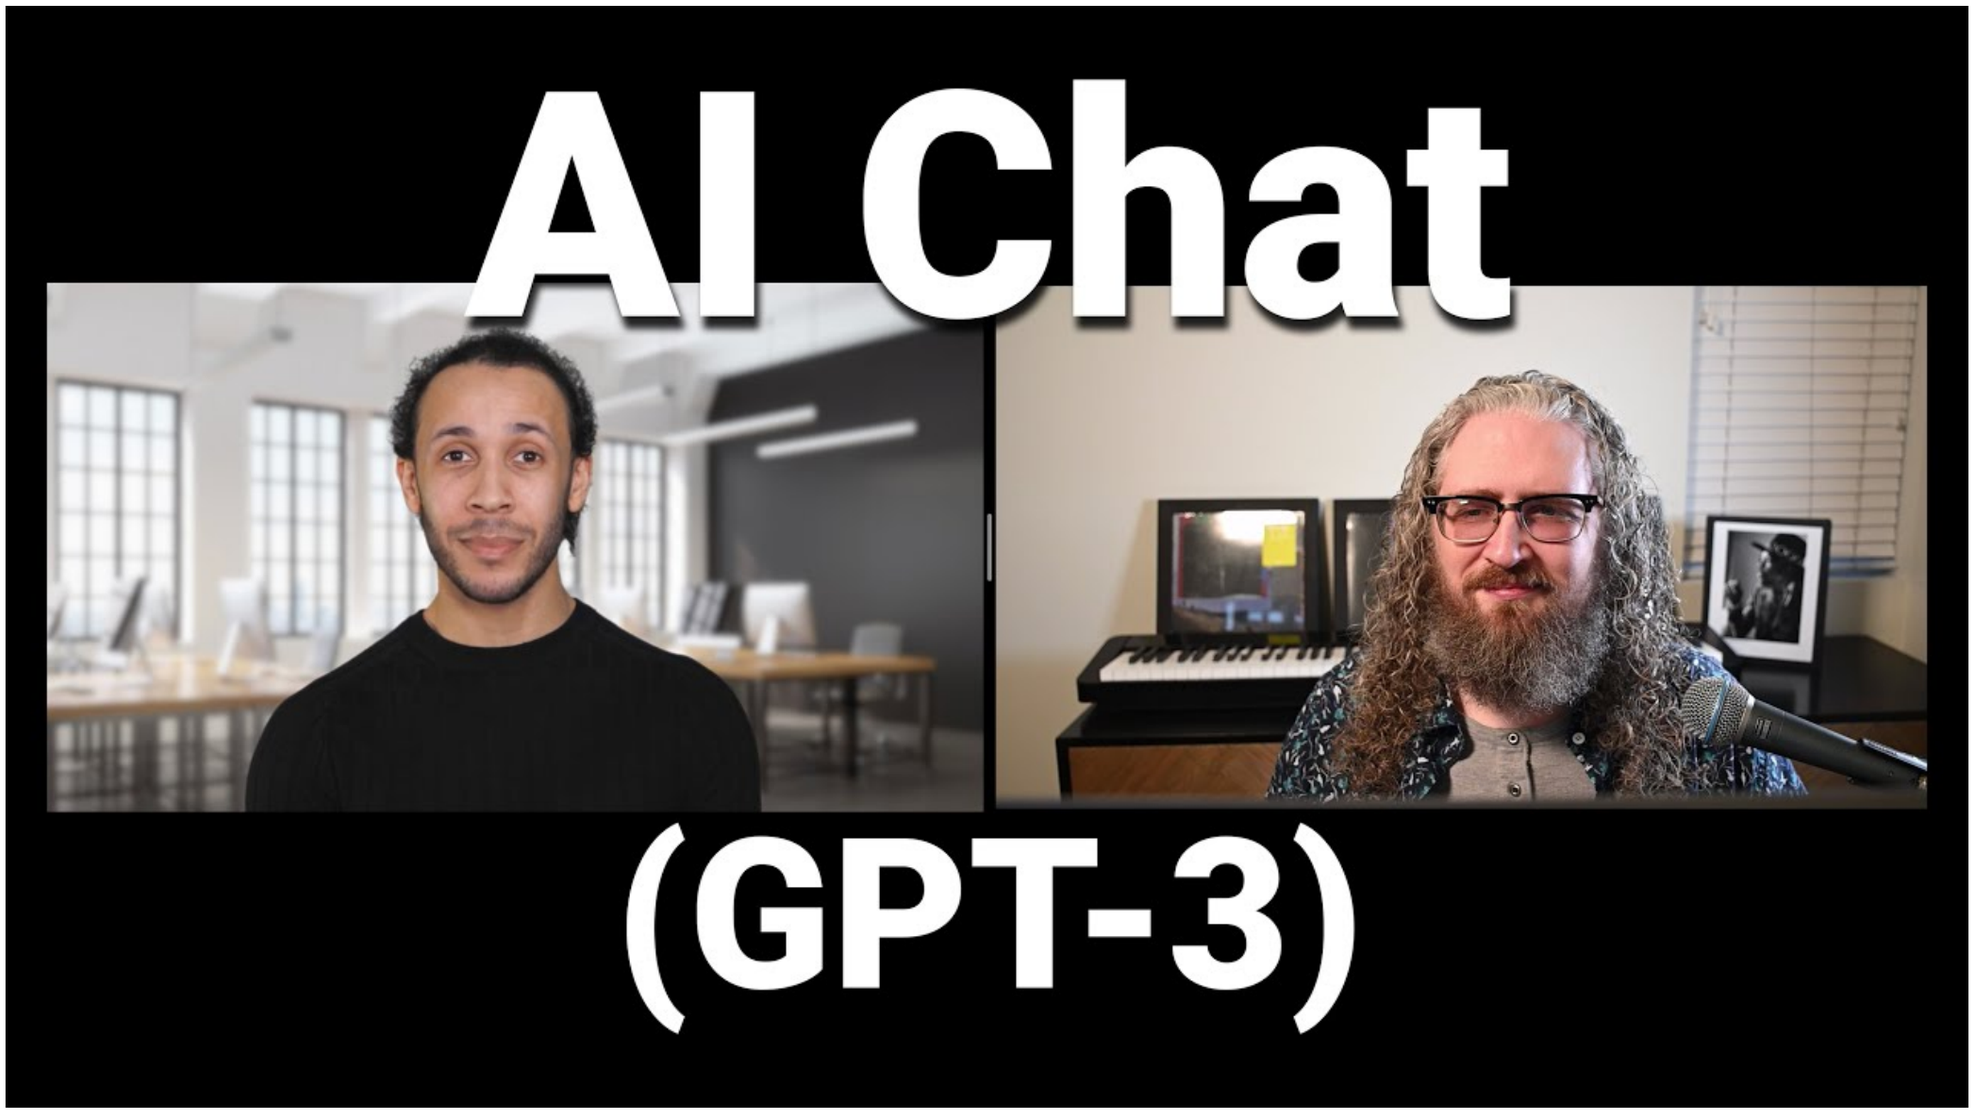

Non-Logicist AI

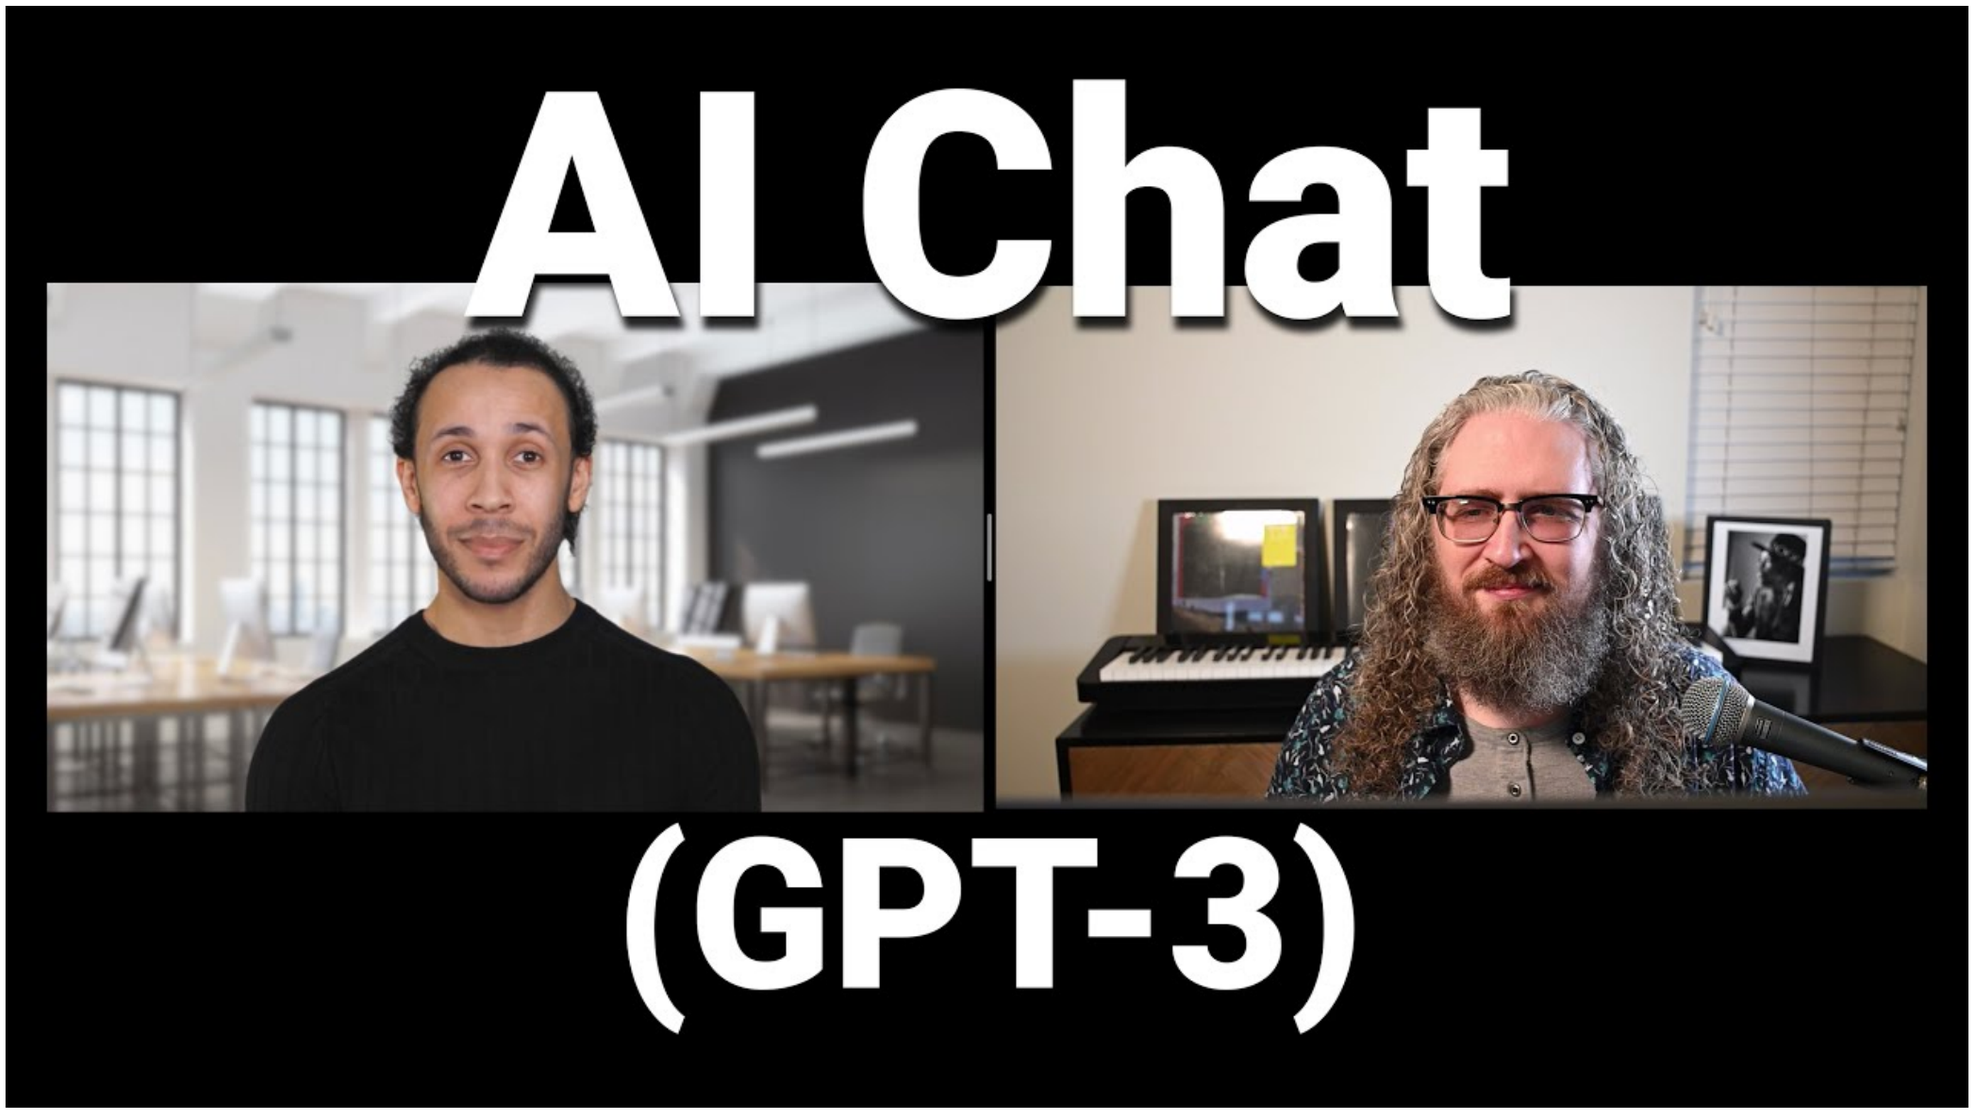

EE: "A lot of people have claimed that GPT-3 doesn't have the ability to process information and reason about things …"

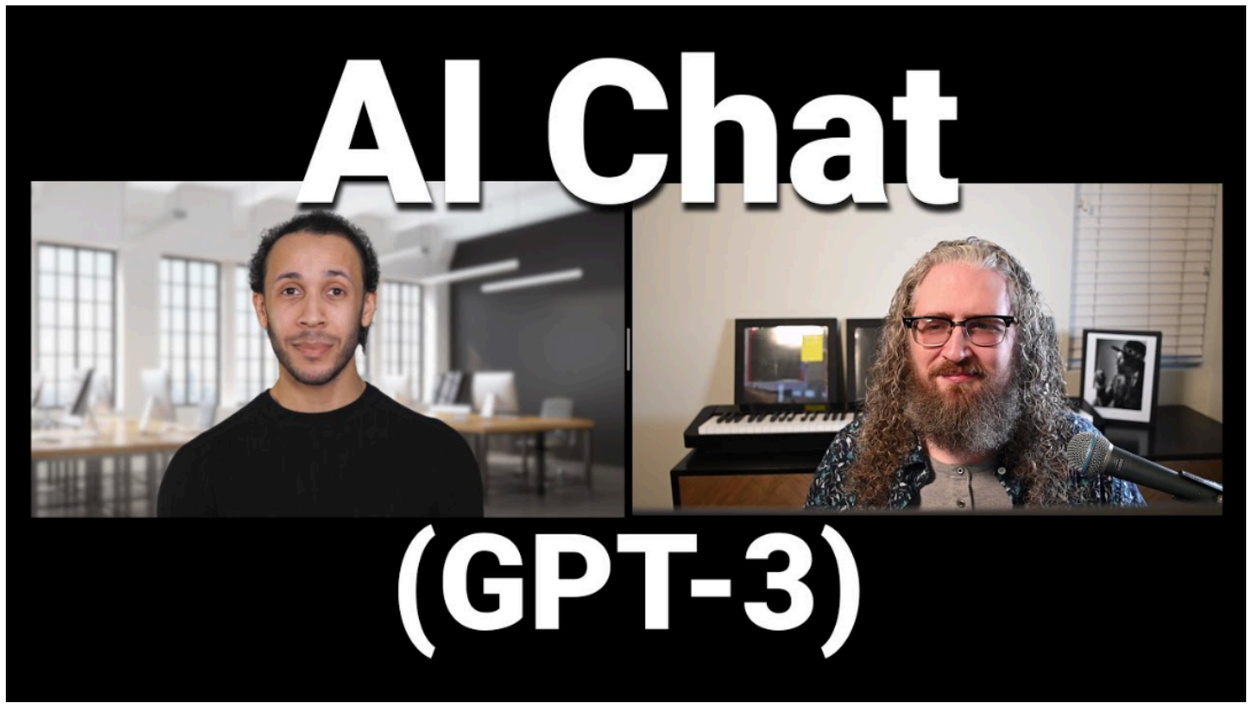

EE: "A lot of people have claimed that GPT-3 doesn't have the ability to process information and reason about things …"

Indeed. And the claim is correct!

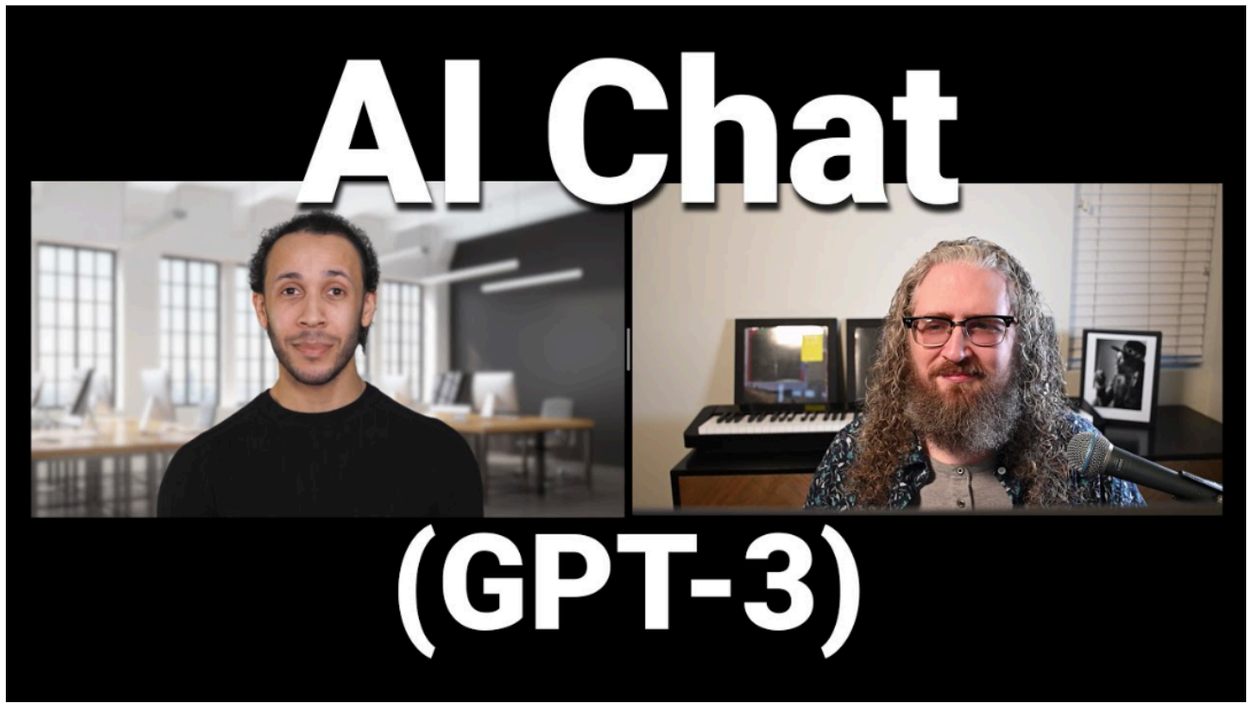

Non-Logicist AI

EE: "A lot of people have claimed that GPT-3 doesn't have have the ability to process information and reason about things …"

EE: "You know what I mean by lying, though, right?"

GPT-3: "Yes. I can make statements that I know are not true."

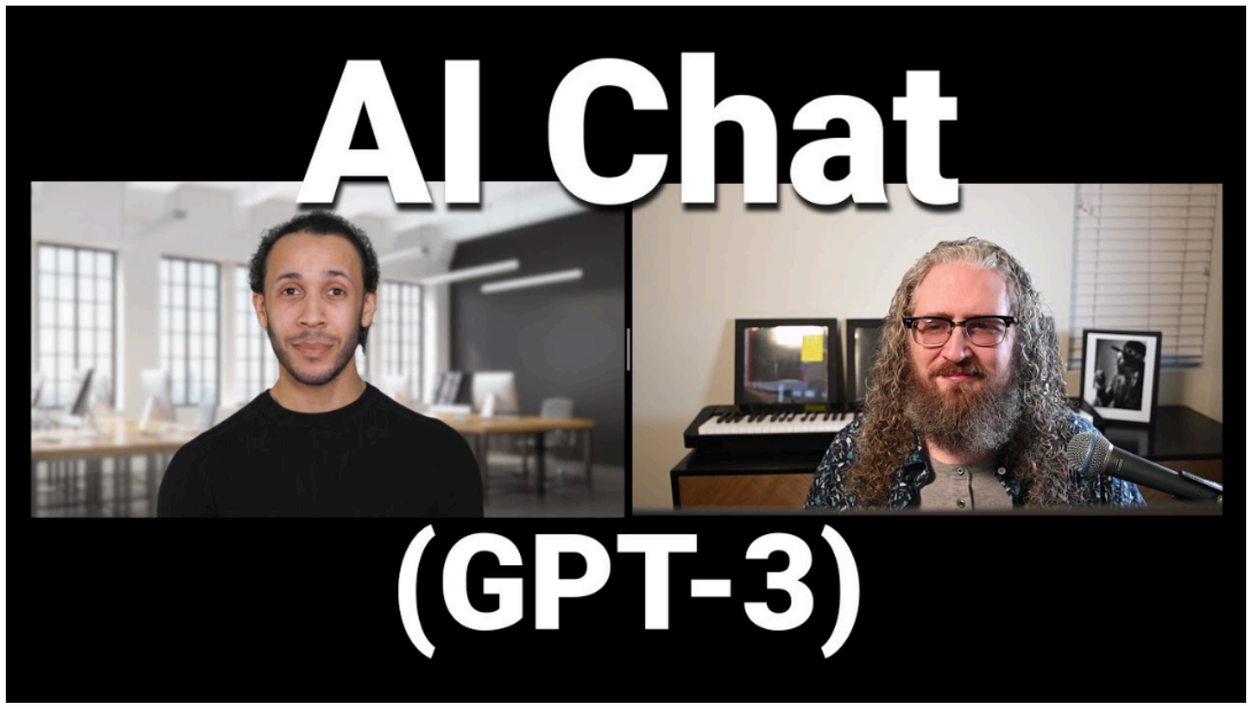

Non-Logicist AI

EE: "A lot of people have claimed that GPT-3 doesn't have have the ability to process information and reason about things …"

EE: "You know what I mean by lying, though, right?"

GPT-3: "Yes. I can make statements that I know are not true."

Sorry. Ilogical. That's *not* lying.
Easy counter-examples abound.

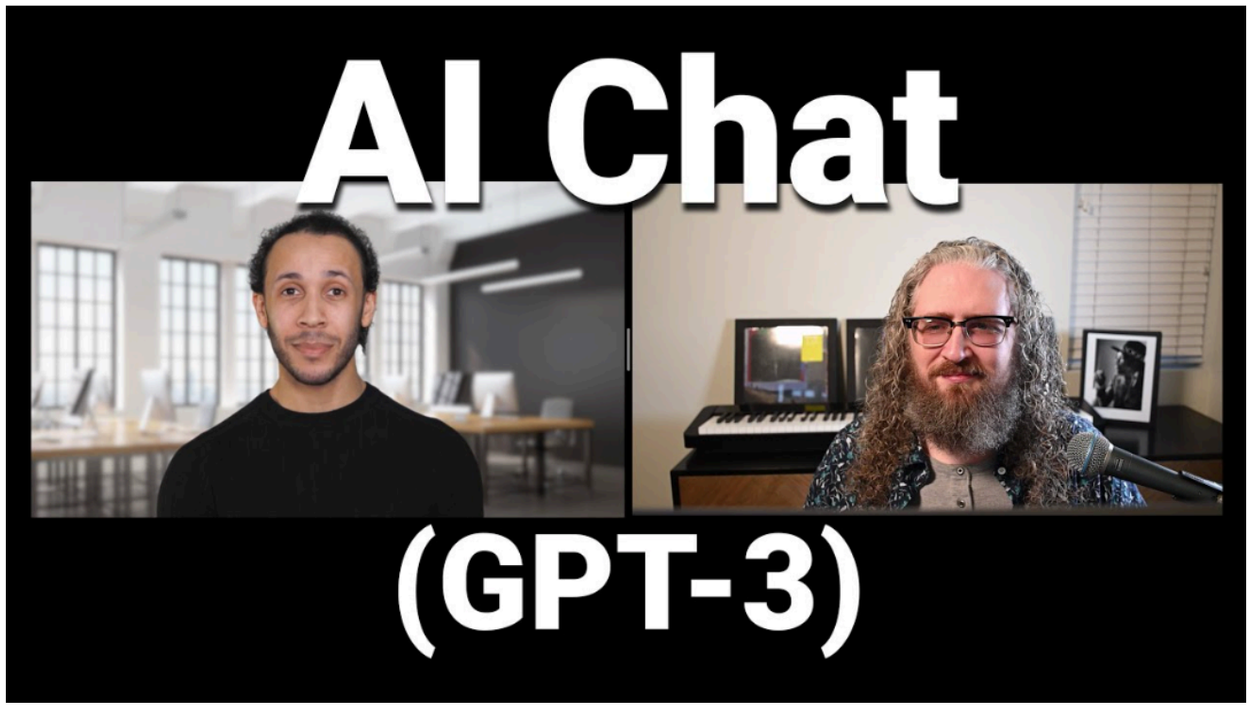

Easy Counter-example:

Jones & Smith are criminal co-conspirators who believe it's possible that a detective is eavesdropping on their conversation. Jones *to Smith*: "I'll be hiking Pine Ridge all day tomorrow."

Easy Counter-example #2 (variant):

Jones & Smith are criminal co-conspirators, & Jones both believes it's possible that a detective Dan is eavesdropping on their conversation (which Smith knows too), and wants to get caught (which Smith doesn't know). Jones to Smith: "I'll be hiking Pine Ridge all day tomorrow." Here, Jones knows Dan knows that Jones has an injury that precludes his doing any hiking tomorrow, and Smith doesn't know about the injury. But Smith does know that Jones won't be hiking tomorrow and will in fact be at home.

Easy Counter-example #2 (variant):

This is not a lie, but according to GPT-3 it *is*.

Jones & Smith are criminal co-conspirators, & Jones both believes it's possible that a detective Dan is eavesdropping on their conversation (which Smith knows too), and wants to get caught (which Smith doesn't know). Jones to Smith: "I'll be hiking Pine Ridge all day tomorrow." Here, Jones knows Dan knows that Jones has an injury that precludes his doing any hiking tomorrow, and Smith doesn't know about the injury. But Smith does know that Jones won't be hiking tomorrow and will in fact be at home.

The Universe of Logics

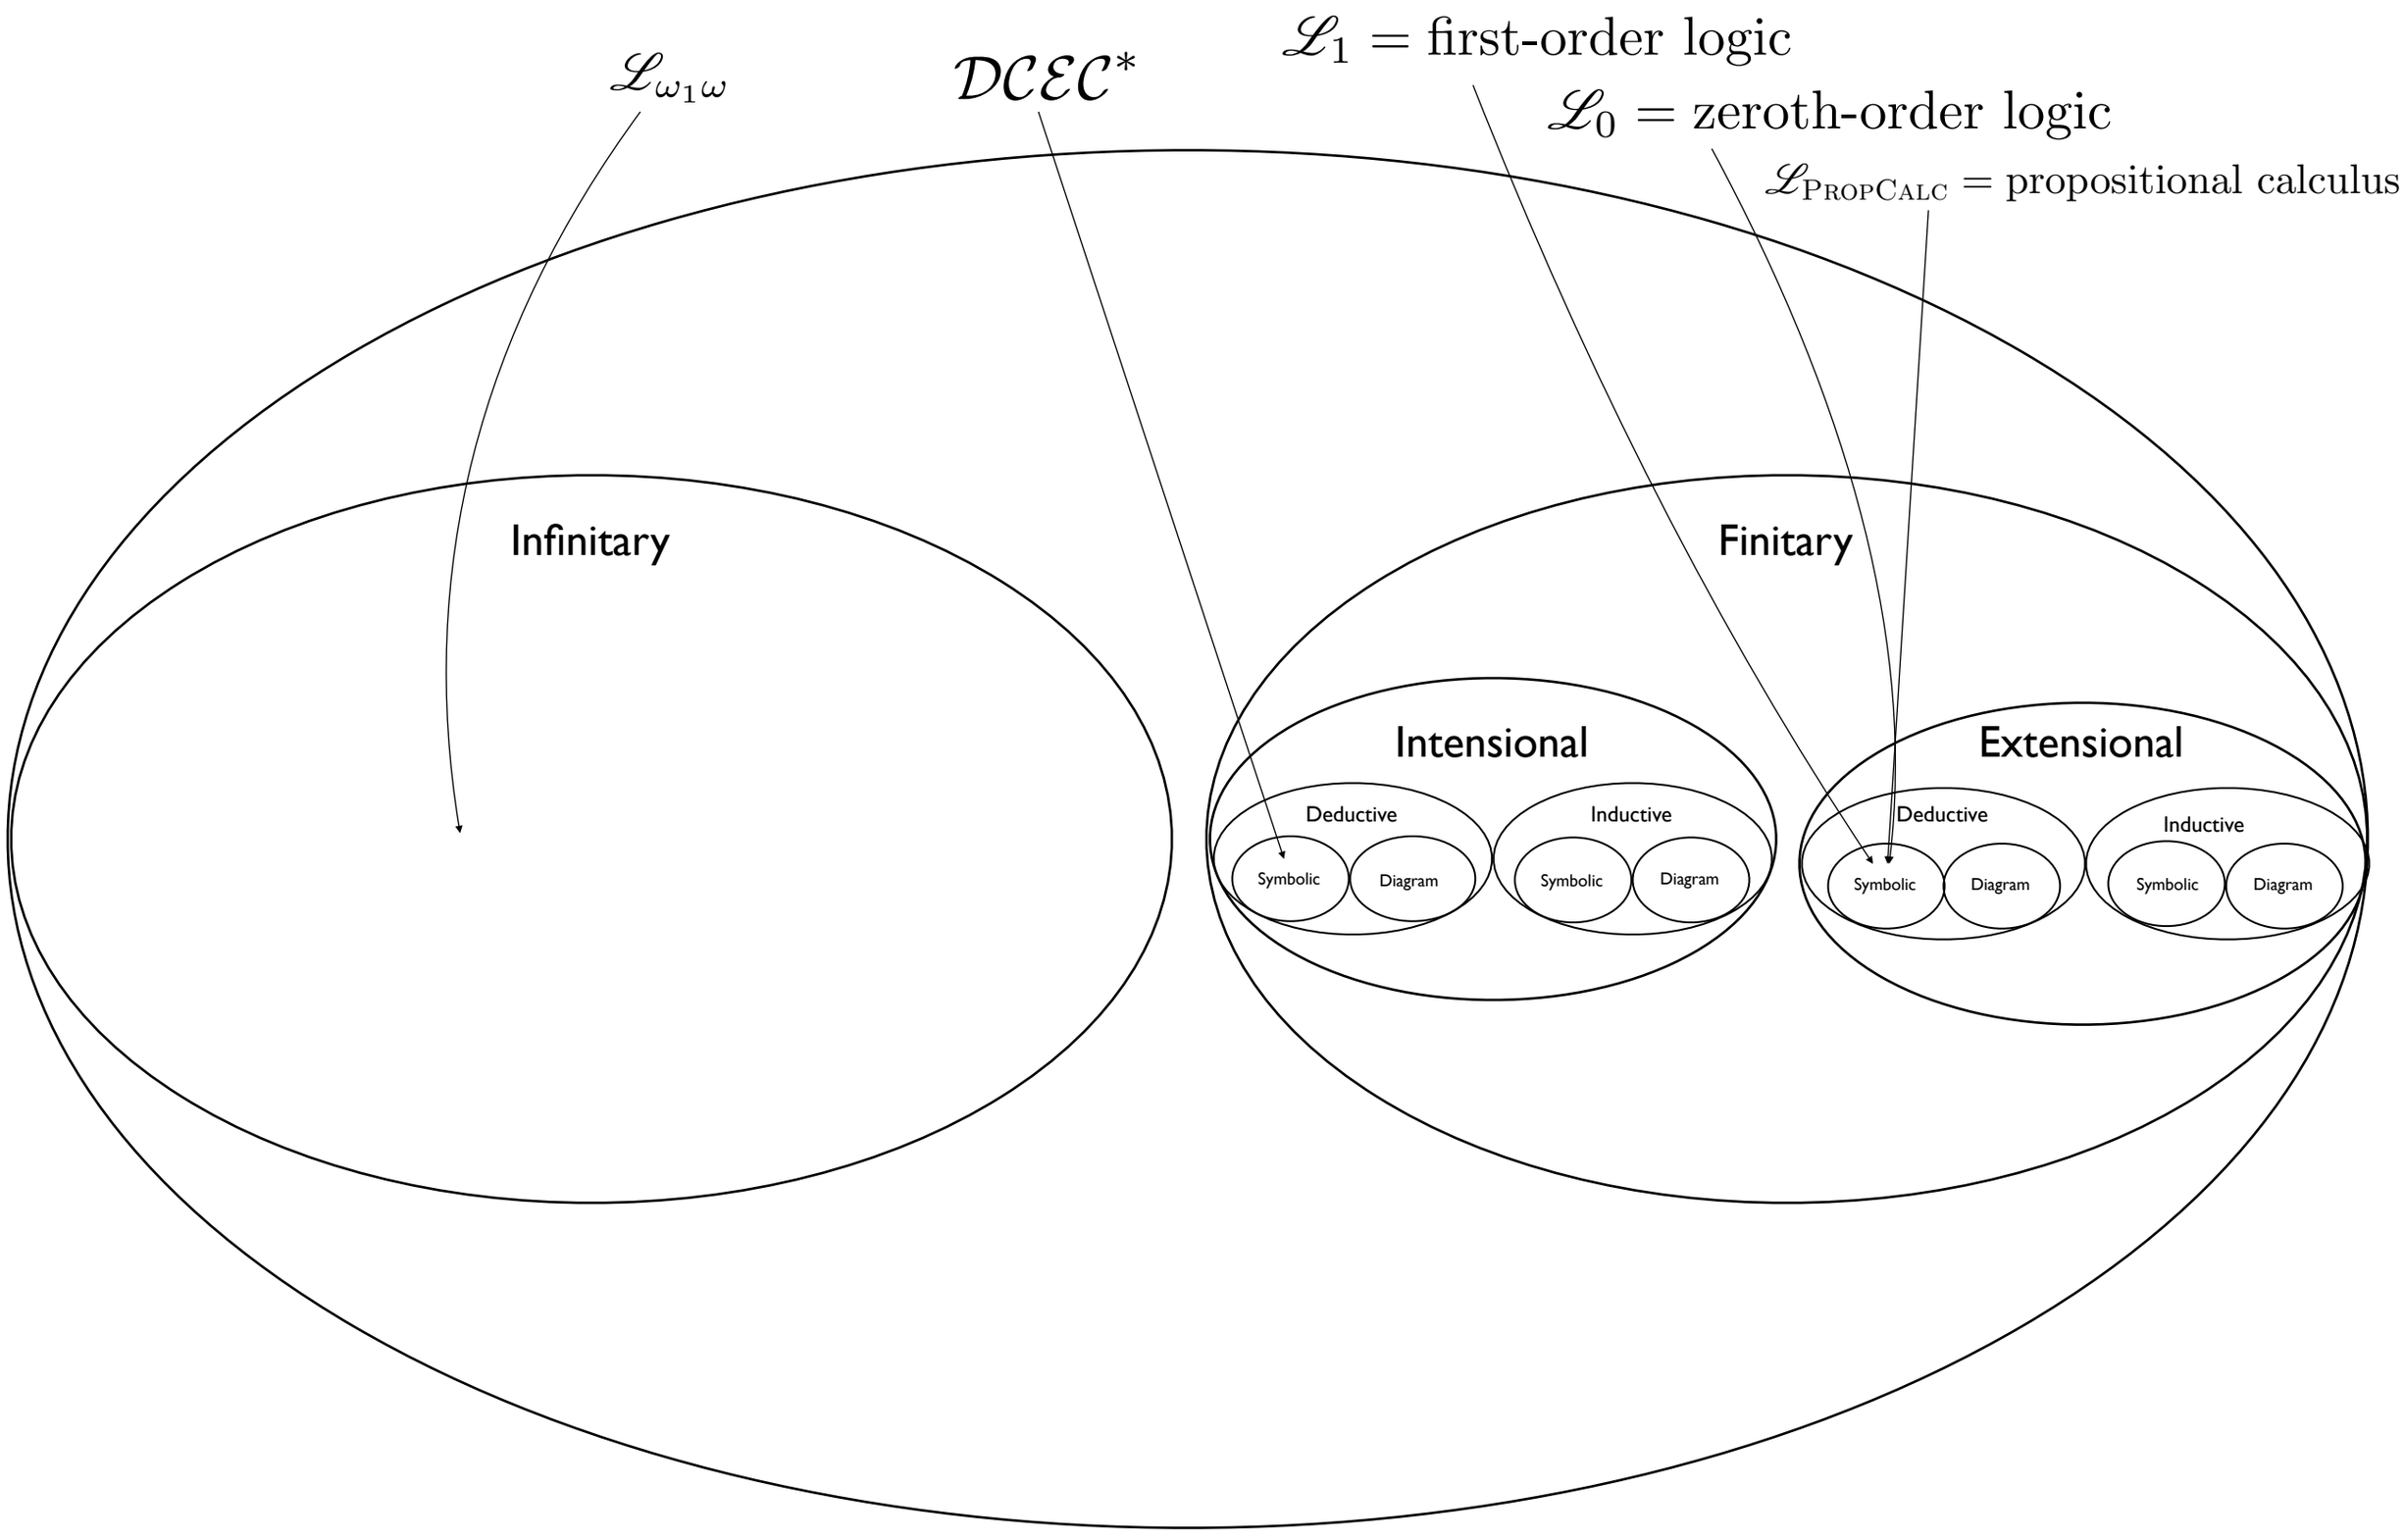

The Universe of Logics

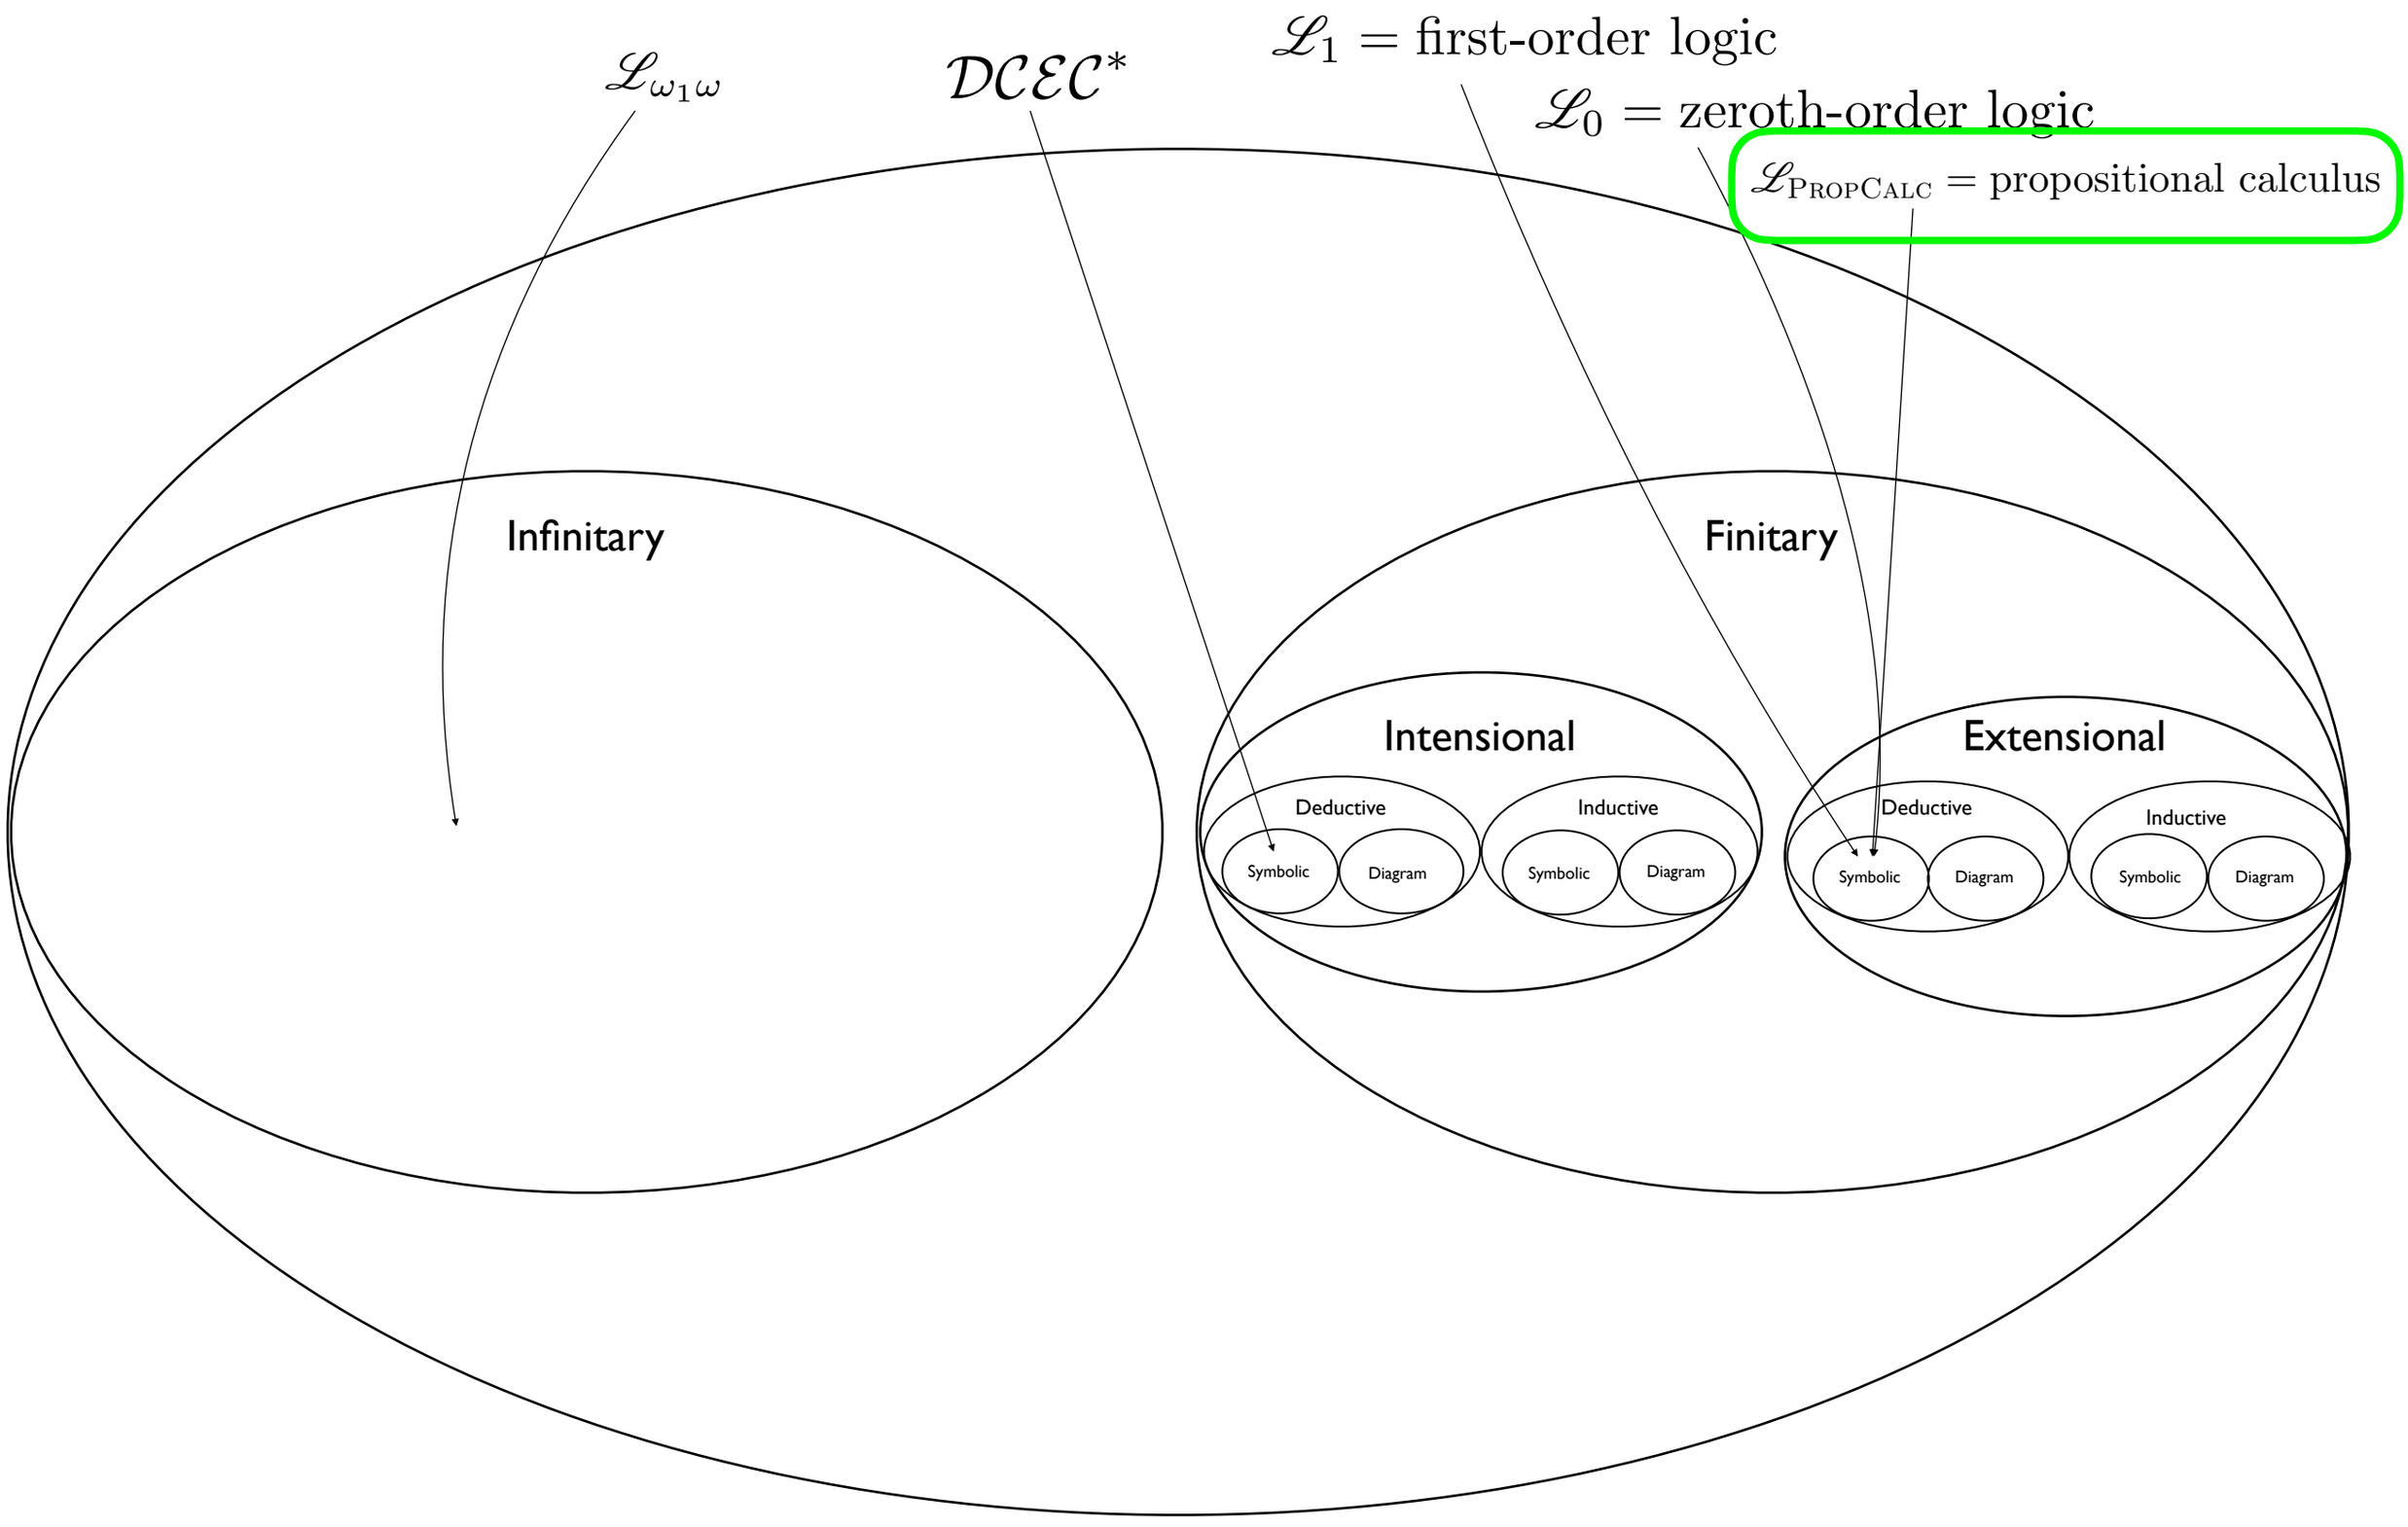

The Universe of Logics

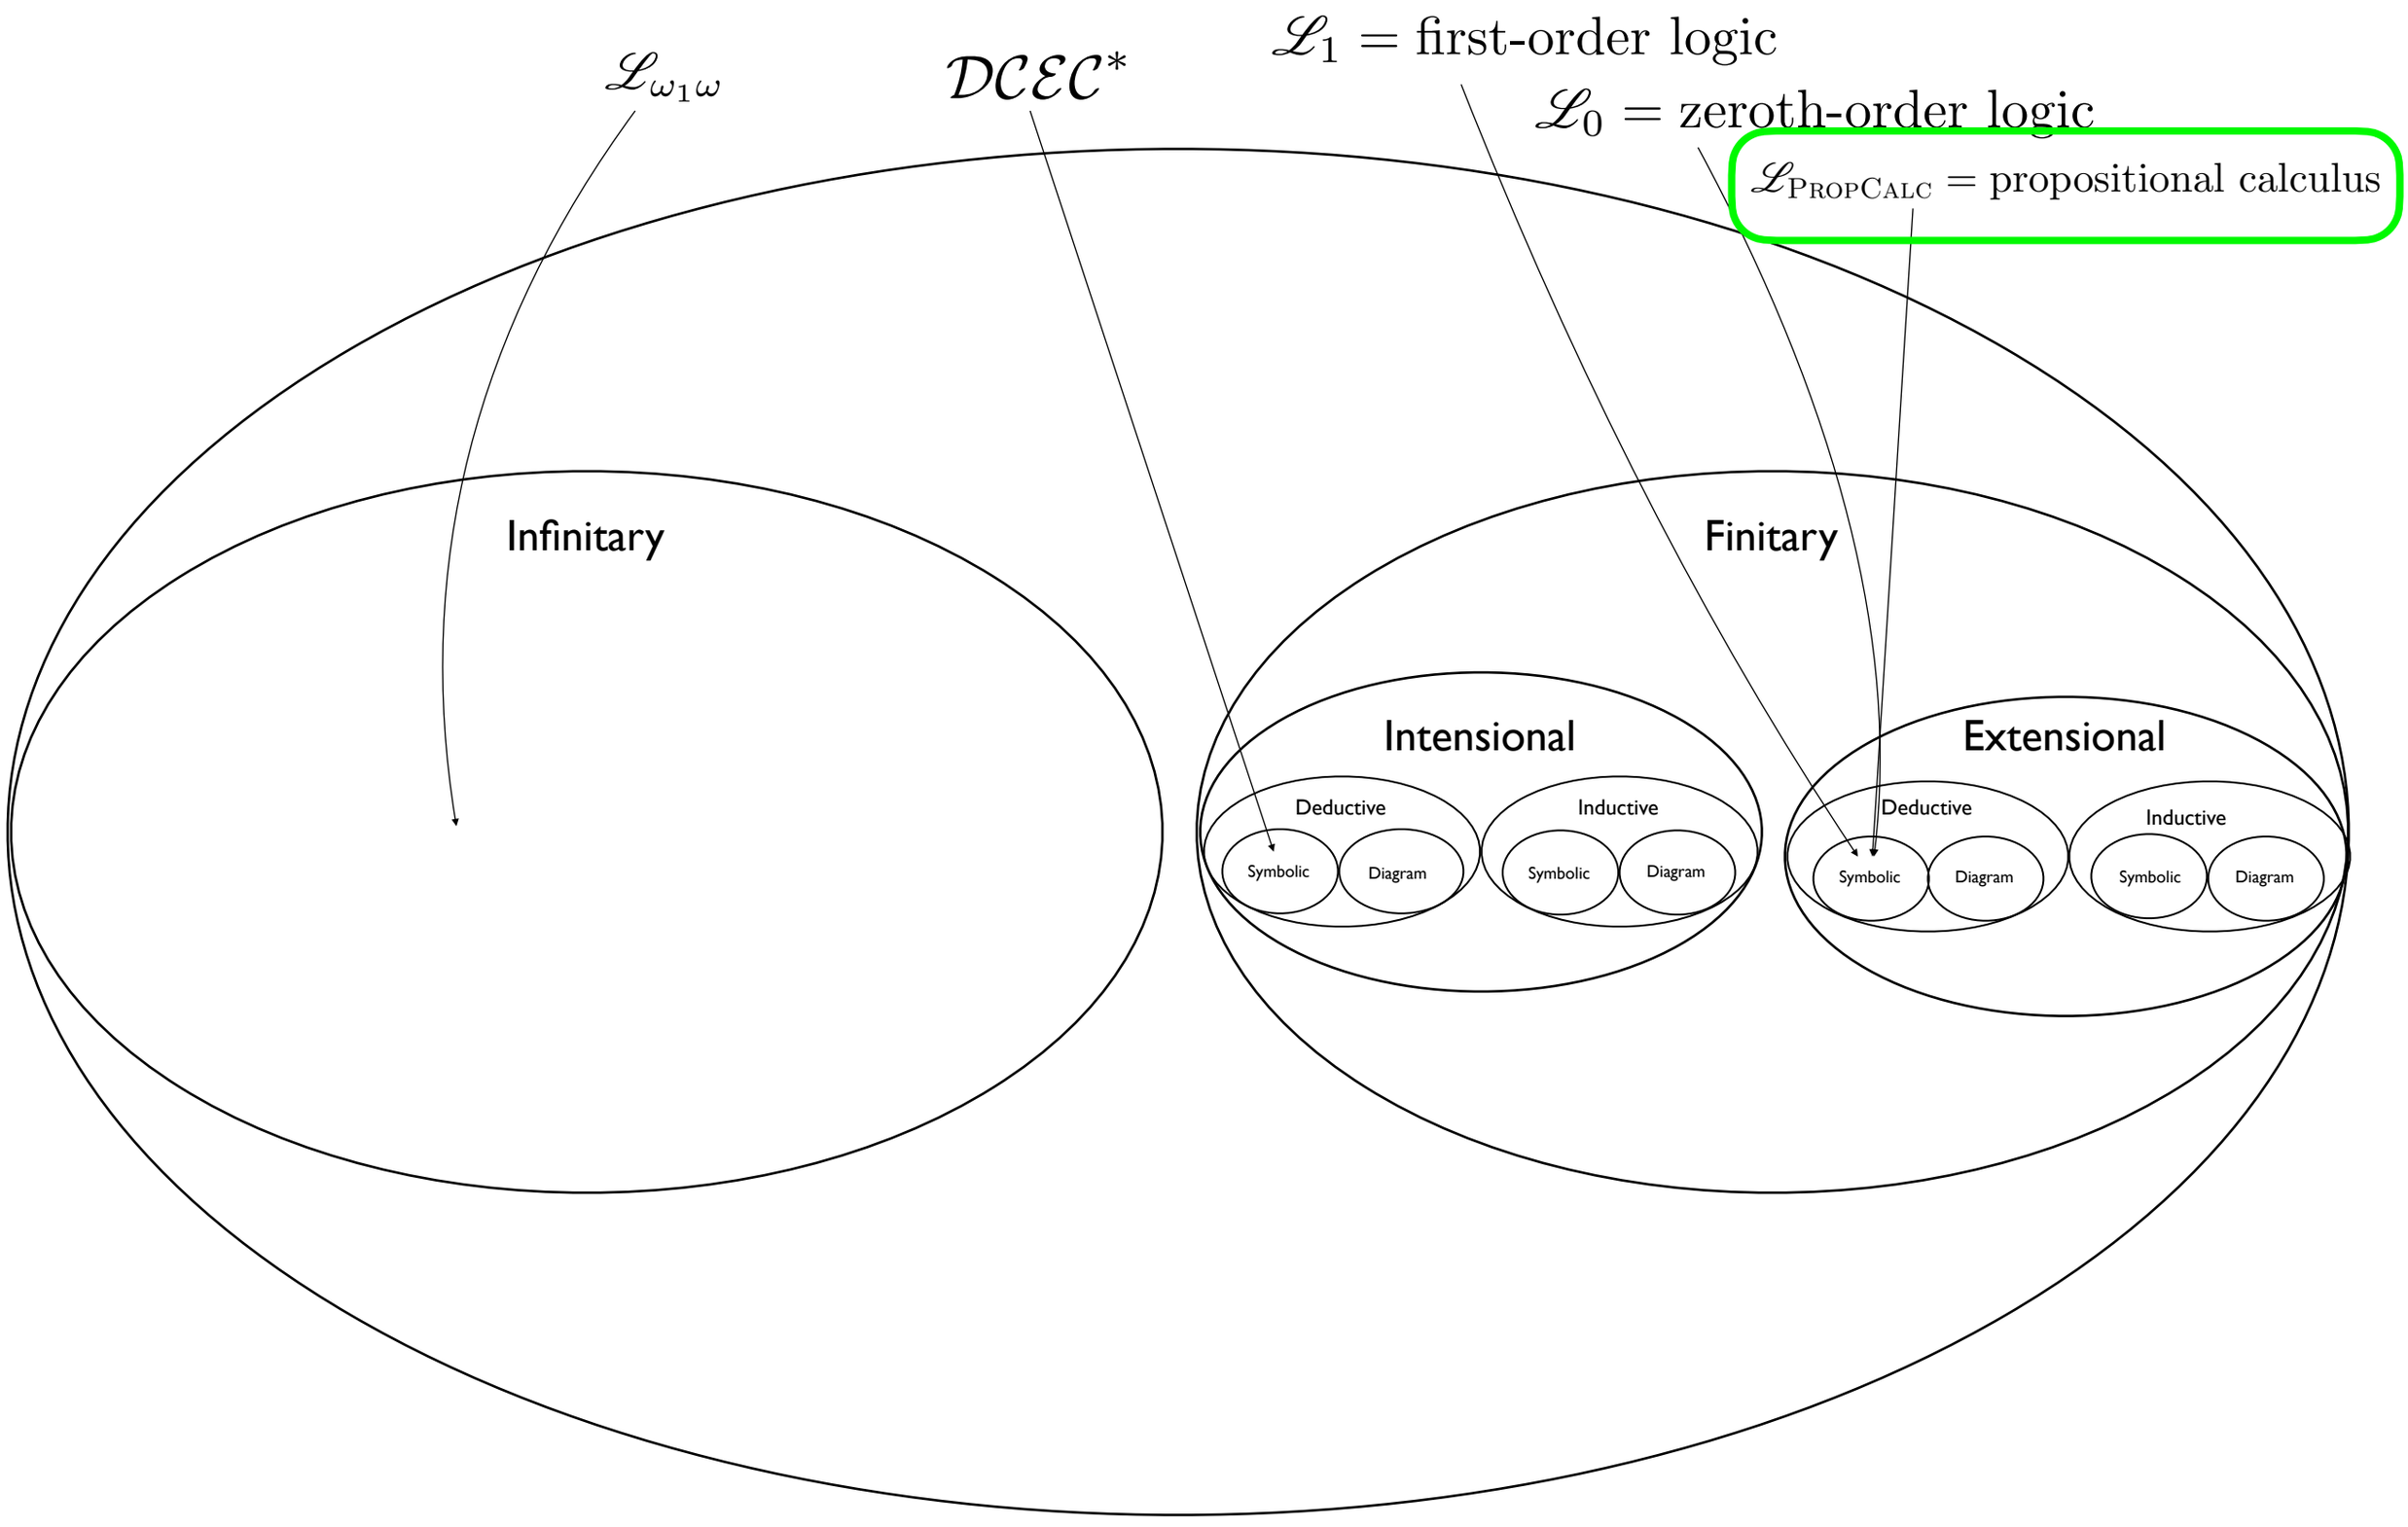

The Universe of Logics

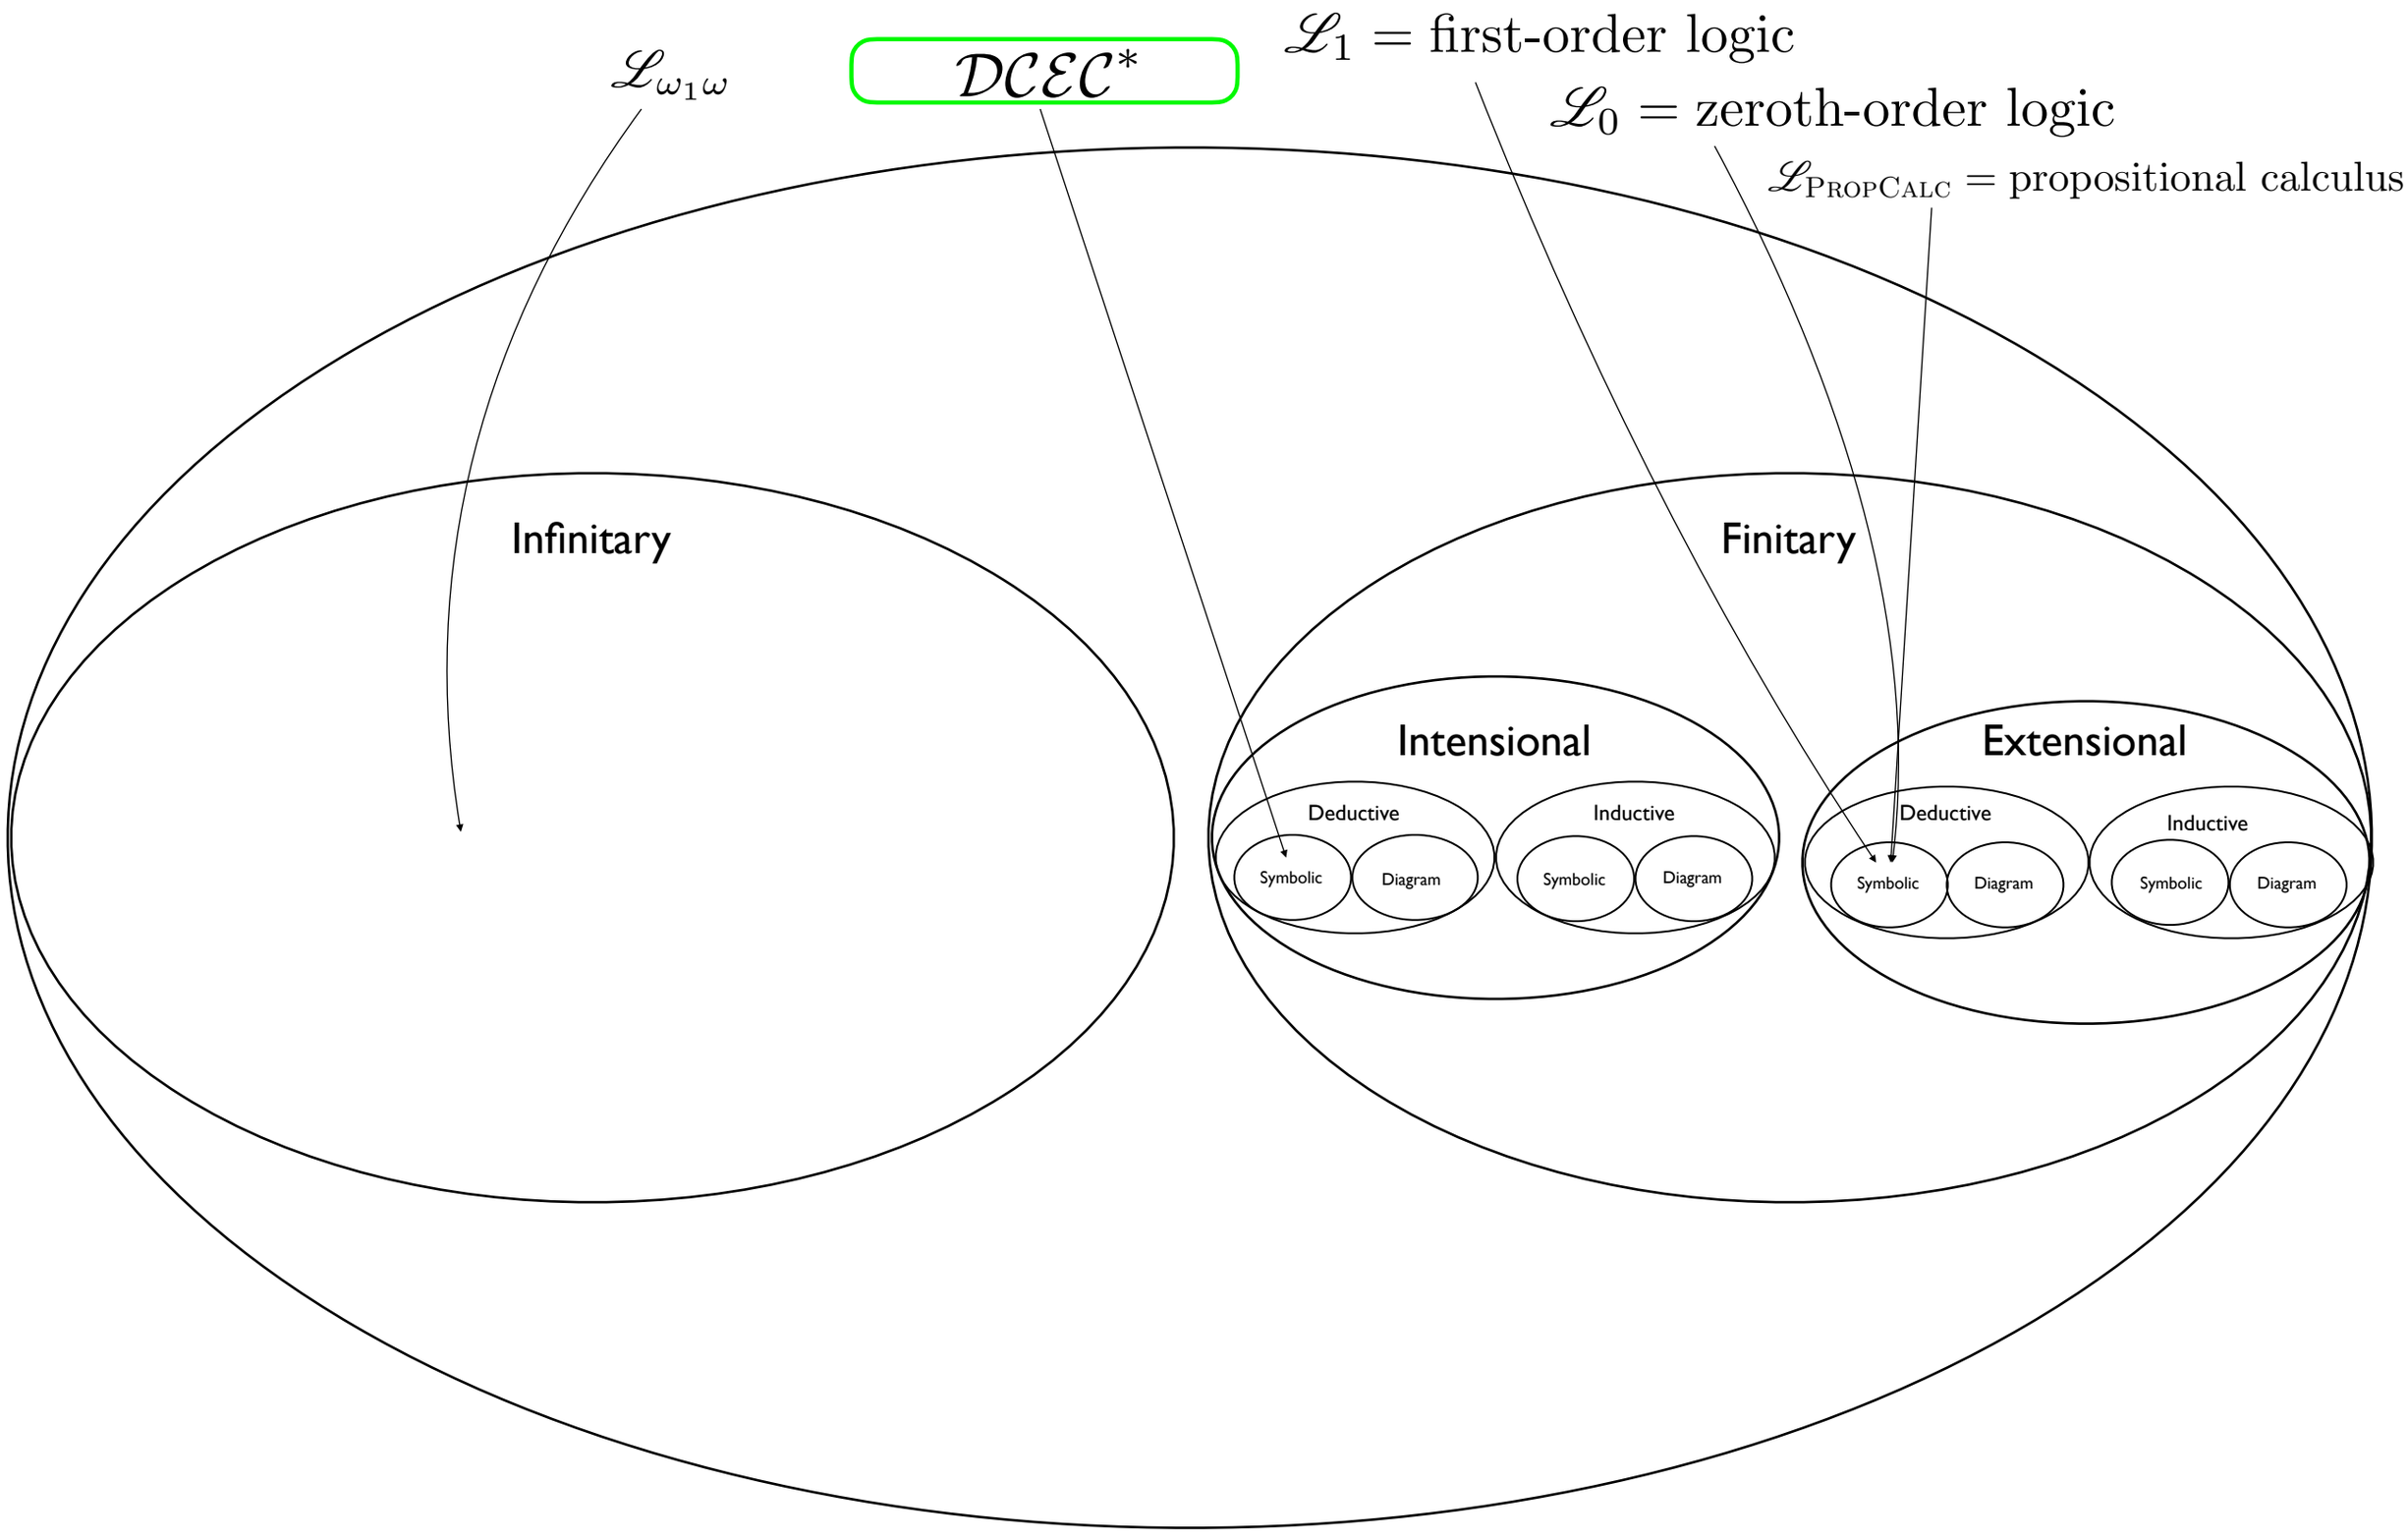

What *is* Logic?

- The key to becoming rational.
- "The science of reasoning." — so the not-unreasonable slogan goes.
- The only invincible subject there is.
- The basis for the formal sciences (from mathematics to game theory to decision theory to probability calculi to axiomatic physics ….) — and hence the basis for disciplines based on the formal sciences (e.g., engineering, computer science).
- The way of escape from shallow content and context to pure, immaterial, and immortal form and structure (which is why the exotic, imaginary, and seemingly non-sensical is so pedagogically useful).
- The most challenging subject there is.
- One of the chief differentiators between dogs and monkeys versus you (let alone bears and you); and mindless machines (like Deep Blue & Watson) versus you.
- A key to riches.
- The key to divining the meaning of life (and other such big questions).
- The better way to program computers; and fundamentally the *only* way to *reliably* program computers.
- One of two fundamental approaches to studying minds, and replicating/simulating minds in machines…
- The thing many creatures of fiction have mastered — have you (as a New Yorker)?…

What *is* Logic?

- The key to becoming rational.
- "The science of reasoning." — so the not-unreasonable slogan goes.
- The only invincible subject there is.
- The basis for the formal sciences (from mathematics to game theory to decision theory to probability calculi to axiomatic physics ….) — and hence the basis for disciplines based on the formal sciences (e.g., engineering, computer science).
- The way of escape from shallow content and context to pure, immaterial, and immortal form and structure (which is why the exotic, imaginary, and seemingly non-sensical is so pedagogically useful).
- The most challenging subject there is.
- One of the chief differentiators between dogs and monkeys versus you (let alone bears and you); and mindless machines (like Deep Blue & Watson) versus you.
- A key to riches.
- The key to divining the meaning of life (and other such big questions).
- The better way to program computers; and fundamentally the *only* way to *reliably* program computers.
- One of two fundamental approaches to studying minds, and replicating/simulating minds in machines…
- The thing many creatures of fiction have mastered — have you (as a New Yorker)?…

Background Claim

$\mathcal{R}$ 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

Background Claim

$\mathcal{R}$ 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

Background Claim

$\mathcal{R}$ 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

Background Claim

$\mathcal{R}$ 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

abstract-and-valid inference schemata

quantification

Background Claim

intensional reasoning

recursion

self-reference

abstract-and-valid inference schemata

quantification

intensional reasoning

recursion

self-reference

To infinity and beyond! — routinely

abstract-and-valid inference schemata

quantification

intensional reasoning

recursion

self-reference

To infinity and beyond! — routinely

abstract-and-valid inference schemata

quantification

intensional reasoning

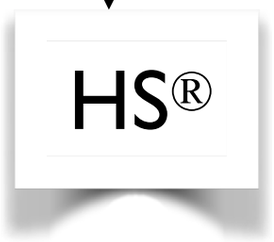

recursion

self-reference

To infinity and beyond! — routinely

abstract-and-valid inference schemata

quantification

intensional reasoning

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

abstract-and-valid inference schemata

quantification

intensional reasoning

recursion

self-reference

To infinity and beyond! — routinely

abstract-and-valid inference schemata

quantification

intensional reasoning

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

“NYS 3”

Given the statements

¬¬c

c → a

¬a ∨ b

b → d

¬(d ∨ e)

which of the following statements are provable?

¬c

e

h

¬a

all of the above

“NYS 3”

Given the statements

¬¬c

c → a

¬a ∨ b

b → d

¬(d ∨ e)

which of the following statements are provable?

¬c

e

h

¬a

all of the above

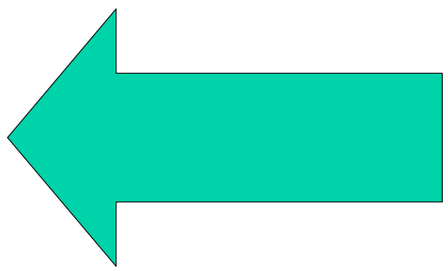

The Original King-Ace

Suppose that the following premise is true:

> If there is a king in the hand, then there is an ace in the hand, or else if there isn't a king in the hand, then there is an ace.

What can you infer from this premise?

The Original King-Ace

Suppose that the following premise is true:

> If there is a king in the hand, then there is an ace in the hand, or else if there isn't a king in the hand, then there is an ace.

What can you infer from this premise?

> There is an ace in the hand.

The Original King-Ace

Suppose that the following premise is true:

> If there is a king in the hand, then there is an ace in the hand, or else if there isn't a king in the hand, then there is an ace.

What can you infer from this premise?

> ~~There is an ace in the hand.~~

The Original King-Ace

Suppose that the following premise is true:

> If there is a king in the hand, then there is an ace in the hand, or else if there isn't a king in the hand, then there is an ace.

What can you infer from this premise?

NO! ~~There is an ace in the hand.~~

The Original King-Ace

Suppose that the following premise is true:

> If there is a king in the hand, then there is an ace in the hand, or else if there isn't a king in the hand, then there is an ace.

What can you infer from this premise?

NO! ~~There is an ace in the hand.~~ NO!

The Original King-Ace

Suppose that the following premise is true:

> If there is a king in the hand, then there is an ace in the hand, or else if there isn't a king in the hand, then there is an ace.

What can you infer from this premise?

NO! ~~There is an ace in the hand.~~ NO!

In fact, what you *can* infer is that there *isn't* an ace in the hand!

King-Ace 2

Suppose that the following premise is true:

If there is a king in the hand, then there is an ace in the hand; or if there isn't a king in the hand, then there is an ace; but not both of these if-then statements are true.

What can you infer from this premise?

King-Ace 2

Suppose that the following premise is true:

If there is a king in the hand, then there is an ace in the hand; or if there isn't a king in the hand, then there is an ace; but not both of these if-then statements are true.

What can you infer from this premise?

There is an ace in the hand.

King-Ace 2

Suppose that the following premise is true:

If there is a king in the hand, then there is an ace in the hand; or if there isn't a king in the hand, then there is an ace; but not both of these if-then statements are true.

What can you infer from this premise?

~~There is an ace in the hand.~~

King-Ace 2

Suppose that the following premise is true:

If there is a king in the hand, then there is an ace in the hand; or if there isn't a king in the hand, then there is an ace; but not both of these if-then statements are true.

What can you infer from this premise?

NO! ~~There is an ace in the hand.~~

King-Ace 2

Suppose that the following premise is true:

If there is a king in the hand, then there is an ace in the hand; or if there isn't a king in the hand, then there is an ace; but not both of these if-then statements are true.

What can you infer from this premise?

NO! ~~There is an ace in the hand.~~ NO!

King-Ace 2

Suppose that the following premise is true:

If there is a king in the hand, then there is an ace in the hand; or if there isn't a king in the hand, then there is an ace; but not both of these if-then statements are true.

What can you infer from this premise?

NO! ~~There is an ace in the hand.~~ NO!

In fact, what you *can* infer is that there *isn't* an ace in the hand!

Informal Proofs vs. Formal Proofs

Informal Proofs vs. Formal Proofs

has ambiguous natural language
(e.g. English or Chinese)

Informal Proofs

has ambiguous natural language
(e.g. English or Chinese)

vs.

Formal Proofs

no ambiguous natural language
(e.g. English or Chinese)

Informal Proofs vs. Formal Proofs

has ambiguous natural language
(e.g. English or Chinese)

cannot be executed & checked
by a computing machine

no ambiguous natural language
(e.g. English or Chinese)

Informal Proofs

has ambiguous natural language
(e.g. English or Chinese)

cannot be executed & checked
by a computing machine

vs.

Formal Proofs

no ambiguous natural language
(e.g. English or Chinese)

can be executed & checked by
a computing machine

Informal Proofs vs. Formal Proofs

has ambiguous natural language
(e.g. English or Chinese)

cannot be executed & checked
by a computing machine

no ambiguous natural language
(e.g. English or Chinese)

can be executed & checked by
a computing machine

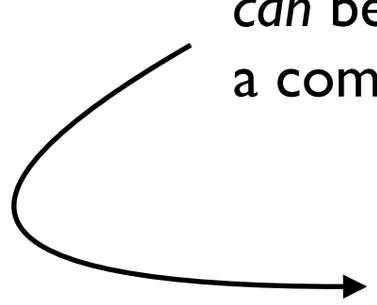

Informal Proofs vs. Formal Proofs

has ambiguous natural language
(e.g. English or Chinese)

cannot be executed & checked
by a computing machine

no ambiguous natural language
(e.g. English or Chinese)

can be executed & checked by
a computing machine

HyperSlate®

Informal Proofs vs. Formal Proofs

has ambiguous natural language
(e.g. English or Chinese)

cannot be executed & checked
by a computing machine

supposed to have learned
how to produce, to an
appreciable degree, in High
School — but likely didn't

no ambiguous natural language
(e.g. English or Chinese)

can be executed & checked by
a computing machine

HyperSlate®

Informal Proofs vs. Formal Proofs

has ambiguous natural language
(e.g. English or Chinese)

cannot be executed & checked
by a computing machine

supposed to have learned
how to produce, to an
appreciable degree, in High
School — but likely didn't

no ambiguous natural language
(e.g. English or Chinese)

can be executed & checked by
a computing machine

HyperSlate®

have not learned how to
produce in a relevant
system (though may have
had some Prolog)

FOR NOW

Informal Proofs vs. Formal Proofs

has ambiguous natural language
(e.g. English or Chinese)

cannot be executed & checked
by a computing machine

supposed to have learned
how to produce, to an
appreciable degree, in High
School — but likely didn't

no ambiguous natural language
(e.g. English or Chinese)

can be executed & checked by
a computing machine

HyperSlate®

have not learned how to
produce in a relevant
system (though may have
had some Prolog)

FOR NOW

STARTING ≥ 1/27/22

Informal Proofs vs. Formal Proofs

has ambiguous natural language (e.g. English or Chinese)

cannot be executed & checked by a computing machine

supposed to have learned how to produce, to an appreciable degree, in High School — but likely didn't

no ambiguous natural language (e.g. English or Chinese)

can be executed & checked by a computing machine

HyperSlate®

have not learned how to produce in a relevant system (though may have had some Prolog)

King-Ace Solved

(informal proof)

Proposition: There is *not* an ace in the hand.

Proof: We know that at least one of the if-thens (i.e., at least one of the **conditionals**) is false. So we have two cases to consider, viz., that *K* => *A* is false, and that ¬*K* => *A* is false. Take first the first case; accordingly, suppose that *K* => *A* is false. Then it follows that *K* is true (since when a conditional is false, its antecedent holds but its consequent doesn't), and *A* is false. Now consider the second case, which consists in ¬*K* => *A* being false. Here, in a direct parallel, we know ¬*K* and, once again, ¬*A*. In both of our two cases, which are exhaustive, there is no ace in the hand. The proposition is established. **QED**

Bringsjord 1

(1) The following three assertions are either all true or all false:

If Billy helped, Doreen helped.

If Doreen helped, Frank did as well.

If Frank helped, so did Emma.

(2) The following assertion is definitely true: Billy helped.

Can it be inferred from (1) and (2) that Emma helped?

Bringsjord 1

(1) The following three assertions are either all true or all false:

If Billy helped, Doreen helped.

If Doreen helped, Frank did as well.

If Frank helped, so did Emma.

(2) The following assertion is definitely true: Billy helped.

Can it be inferred from (1) and (2) that Emma helped?

YUP! — & now prove it!

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 …

(With a debt to P. Johnson-Laird for the king-ace kernel used here.)

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

(With a debt to P. Johnson-Laird for the king-ace kernel used here.)

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 …

(With a debt to P. Johnson-Laird for the king-ace kernel used here.)

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 *do*es 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

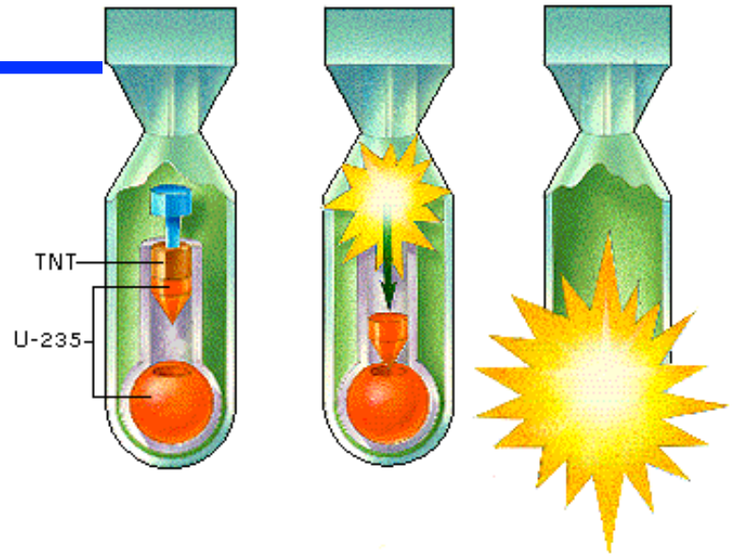

Life on Earth has ended.

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

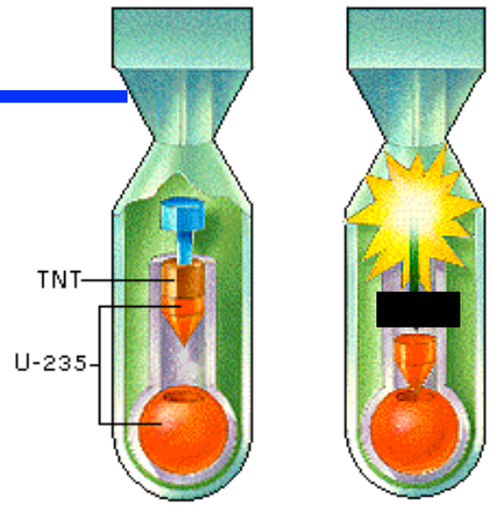

Life on Earth is saved!

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

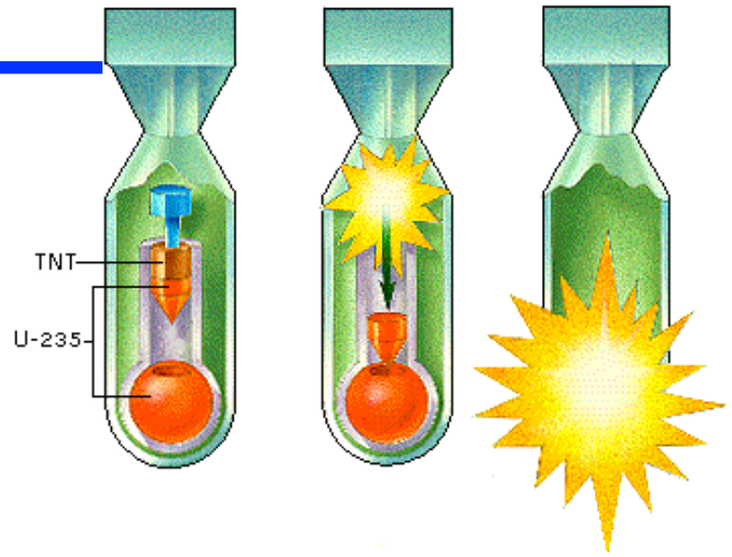

Life on Earth has ended.

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

Logic kan redde menneskehten!