

The Immaterial Paradise, Motivating Paradoxes, Puzzles, and R , Part I

Selmer Bringsjord

Intro to (Formal) Logic (and AI) = IFLAI I

I/I I/24



**Most AI Punditry is Illogical,
e.g. ...**

Without Consciousness, AIs Will Be Sociopaths

ChatGPT can carry on a conversation, but the most important goal for artificial intelligence is making it understand what it means to have a mind

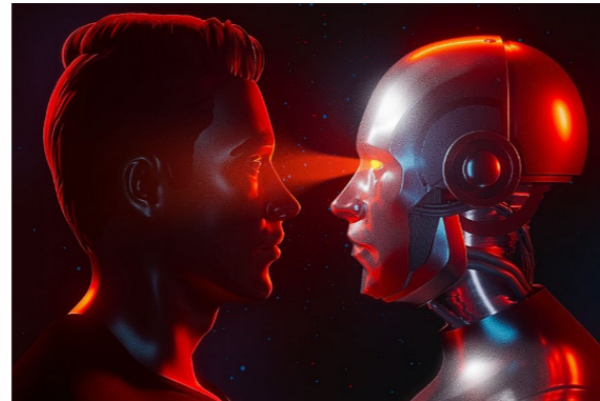
ESSAY

By Michael S.A. Graziano [+ Follow](#)

January 13, 2023 09:24 a.m. EST

ChatGPT, the latest technological [sensation](#), is an artificial intelligence chatbot with an amazing ability to carry on a conversation. It relies on a massive network of artificial neurons that loosely mimics the human brain, and it has been trained by analyzing the information resources of the internet. ChatGPT has processed more text than any human is likely to have read in a lifetime, allowing it to respond to questions fluently and even to imitate specific individuals, answering queries the way it thinks they would. My teenage son recently used ChatGPT to argue about politics with an imitation Karl Marx.

As a neuroscientist specializing in the brain mechanisms of con-



sciousness, I find talking to chatbots an unsettling experience. Are they conscious? Probably not. But given the rate of technological improvement, will they be in the next couple of years? And how would we even know?

Figuring out whether a machine has or understands humanlike consciousness is more than just a science-fiction hypothetical. Artificial intelligence is growing so powerful, so quickly, that it could soon pose a danger to human beings. We're building machines that are smarter than us and giving them control over our world. How can we

build AI so that it's aligned with human needs, not in conflict with us?

As counterintuitive as it may sound, creating a benign AI may require making it more conscious, not less. One of the most common misunderstandings about AI is the notion that if it's intelligent then it must be conscious, and if it is conscious then it will be autonomous, capable of taking over the world. But as we learn more about consciousness, those ideas do not appear to be correct. An autonomous system that makes complex decisions doesn't require consciousness.

What's most important about consciousness is that, for human beings, it's not just about the self. We see it in ourselves, but we also perceive it or project it into the world around us. Consciousness is part of the tool kit that evolution

gave us to make us an empathetic, prosocial species. Without it, we would necessarily be sociopaths, because we'd lack the tools for prosocial behavior. And without a concept of what consciousness is or an understanding that other beings have it, machines are sociopaths.

The only diagnostic tool for machine consciousness that we have right now is the Turing test, a thought experiment named for the British computer scientist Alan Turing. In its most common version, the test says that if a person holds a conversation with a machine and mistakes its responses for those of a real human being, then the machine must be considered effectively conscious.

The Turing test is an admission that the consciousness of another being is something we can only judge from the outside, based on the way he, she or it communicates. But the

Without Consciousness, AIs Will Be Sociopaths

ChatGPT can carry on a conversation, but the most important goal for artificial intelligence is making it understand what it means to have a mind

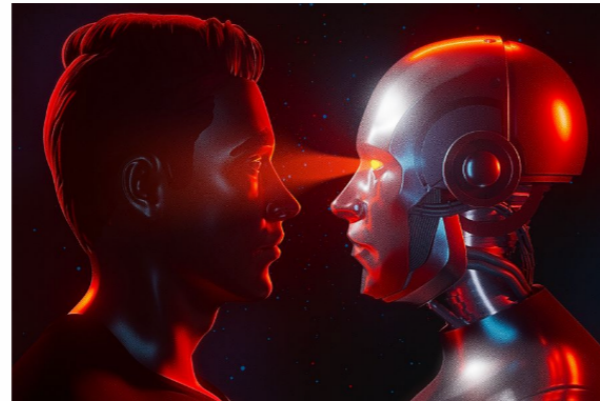
ESSAY

By Michael S.A. Graziano [+ Follow](#)

January 13, 2023 09:24 a.m. EST

ChatGPT, the latest technological [sensation](#), is an artificial intelligence chatbot with an amazing ability to carry on a conversation. It relies on a massive network of artificial neurons that loosely mimics the human brain, and it has been trained by analyzing the information resources of the internet. ChatGPT has processed more text than any human is likely to have read in a lifetime, allowing it to respond to questions fluently and even to imitate specific individuals, answering queries the way it thinks they would. My teenage son recently used ChatGPT to argue about politics with an imitation Karl Marx.

As a neuroscientist specializing in the brain mechanisms of con-



sciousness, I find talking to chatbots an unsettling experience. Are they conscious? Probably not. But given the rate of technological improvement, will they be in the next couple of years? And how would we even know?

Figuring out whether a machine has or understands humanlike consciousness is more than just a science-fiction hypothetical. Artificial intelligence is growing so powerful, so quickly, that it could soon pose a danger to human beings. We're building machines that are smarter than us and giving them control over our world. How can we

build AI so that it's aligned with human needs, not in conflict with us?

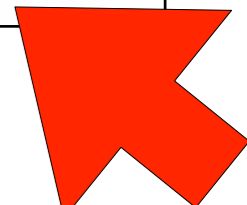
As counterintuitive as it may sound, creating a benign AI may require making it more conscious, not less. One of the most common misunderstandings about AI is the notion that if it's intelligent then it must be conscious, and if it is conscious then it will be autonomous, capable of taking over the world. But as we learn more about consciousness, those ideas do not appear to be correct. An autonomous system that makes complex decisions doesn't require consciousness.

What's most important about consciousness is that, for human beings, it's not just about the self. We see it in ourselves, but we also perceive it or project it into the world around us. Consciousness is part of the tool kit that evolution

gave us to make us an empathetic, prosocial species. Without it, we would necessarily be sociopaths, because we'd lack the tools for prosocial behavior. And without a concept of what consciousness is or an understanding that other beings have it, machines are sociopaths.

The only diagnostic tool for machine consciousness that we have right now is the Turing test, a thought experiment named for the British computer scientist Alan Turing. In its most common version, the test says that if a person holds a conversation with a machine and mistakes its responses for those of a real human being, then the machine must be considered effectively conscious.

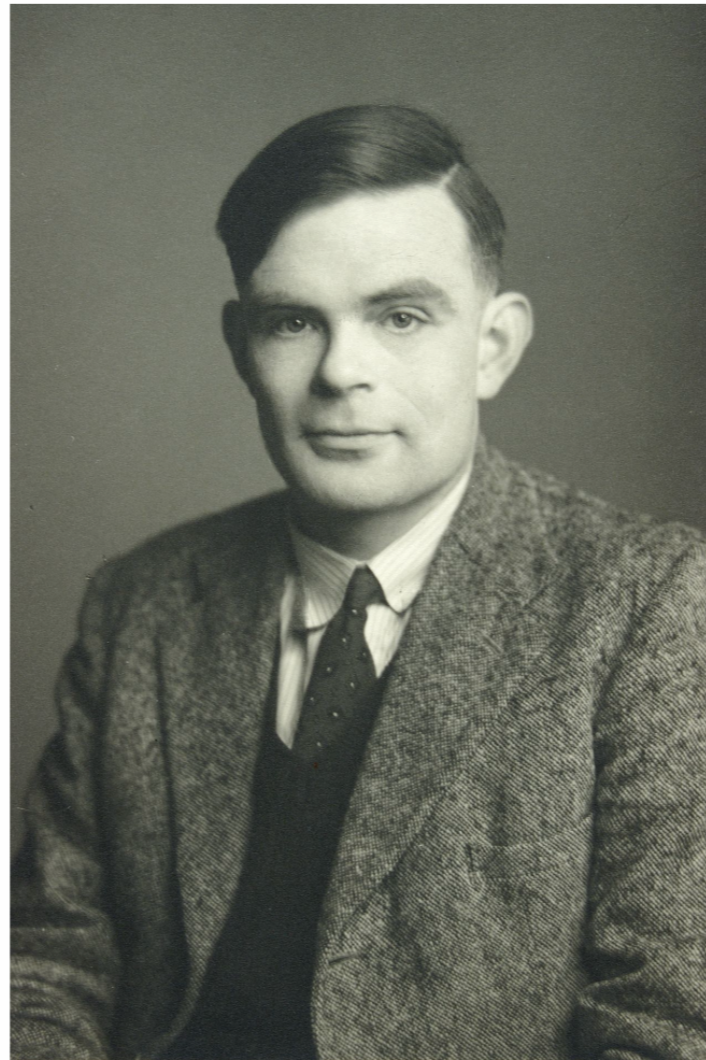
The Turing test is an admission that the consciousness of another being is something we can only judge from the outside, based on the way he, she or it communicates. But the



limits of the test are painfully obvious. After all, a pet dog can't carry on a conversation and pass as a human—does that mean it's not conscious? If you really wanted a machine to pass the test, you could have it say a few words to a small child. It might even fool some adults, too.

The truth is, the Turing test doesn't reveal much about what's going on inside a machine or a computer program like ChatGPT. Instead, what it really tests is the social cognition of the human participant. We evolved as social animals, and our brains instinctively project consciousness, agency, intention and emotion onto the objects around us. We're primed to see a world suffused with minds. Ancient animistic beliefs held that every river and tree had a spirit in it. For a similar reason, people are prone to see faces in random objects like the moon and moldy toast.

The original test proposed by Alan Turing in a 1950 paper was more complicated than the version people talk about today. Notably,



Computer science pioneer Alan Turing in 1951.

Turing didn't say a word about consciousness; he never delved into whether the machine had a subjective experience. He asked only whether it could think like a person. Turing imagined an "imitation

game" in which the player must determine the sex of two people, A and B. One is a man and one is a woman, but the player can't see them and can learn about them only by exchanging typed questions and

answers. A responds to the questions deceitfully, and wins the game if the player misidentifies their sex, while B answers truthfully and wins if the player identifies their sex correctly. Turing's idea was that if A or B is replaced by a machine, and the machine can win the game as often as a real person, then it must have mastered the subtleties of human thinking—of argument, manipulation and guessing what other people are thinking.

Turing's test was so complicated that people who popularized his work soon streamlined it into a single machine conversing with a single person. But the whole point of the original test was its bizarre complexity. Social cognition is difficult and requires a theory of mind—that is, a knowledge that other people have minds and an ability to guess what might be in them.

If we want to know whether a computer is conscious, then, we need to test whether the computer understands how conscious minds interact. In other words, we need a

9:36 AM Wed Jan 18

 Back

reverse Turing test: Let's see if the computer can tell whether it's talking to a human or another computer. If it can tell the difference, then maybe it knows what consciousness is. ChatGPT definitely can't pass that test yet: It doesn't know whether it's responding to a living person with a mind or a disjointed list of prefab questions.

A sociopathic machine that can make consequential decisions would be powerfully dangerous. For now, chatbots are still limited in their abilities; they're essentially toys. But if we don't think more deeply about machine consciousness, in a year or five years we may face a crisis. If computers are going to outthink us anyway, giving them more human-like social cognition might be our best hope of aligning them with human values.

Dr. Graziano is a professor of psychology and neuroscience at Princeton University and the author of "Rethinking Consciousness: A Scientific Theory of Subjective Experience."

Logically Invalid Reasoning

Turing in his famous paper claims:

TT: If x is conversationally indistinguishable from a human, then x is conscious/sentient.

Logically Invalid Reasoning

Turing in his famous paper claims:

TT: If x is conversationally indistinguishable from a human, then x is conscious/sentient.

This proposition is logically equivalent to:

TT': If x is not conversationally indistinguishable from a human, then x is not conscious/sentient.

Logically Invalid Reasoning

Turing, in his famous paper, claims:

TT: If x is conversationally indistinguishable from a human, then x is conscious/sentient.

This proposition is logically equivalent to:

TT': If x is not conversationally indistinguishable from a human, then x is not conscious/sentient.

Well then, obviously Turing is wrong: A pet dog isn't conversationally indistinguishable from a human, but the dog *is* conscious/sentient. We have a counter-example!

Logically Invalid Reasoning

Turing, in his famous paper, claims:

TT: If x is conversationally indistinguishable from a human, then x is conscious/sentient.

This proposition is logically equivalent to:

TT': If x is not conversationally indistinguishable from a human, then x is not conscious/sentient.

Well then, obviously Turing is wrong: A pet dog isn't conversationally indistinguishable from a human, but the dog *is* conscious/sentient. We have a counter-example!

Logically Invalid Reasoning

Turing, in his famous paper, claims:

TT: If x is conversationally indistinguishable from a human, then x is conscious/sentient.

$$i \rightarrow c$$

This proposition is logically equivalent to:

TT': If x is not conversationally indistinguishable from a human, then x is not conscious/sentient.

$$\neg i \rightarrow \neg c$$

Well then, obviously Turing is wrong: A pet dog isn't conversationally indistinguishable from a human, but the dog *is* conscious/sentient. We have a counter-example!

$$\neg i \wedge c$$

$$\{\neg i \wedge c\} \vdash \neg(\neg i \rightarrow \neg c) \quad \checkmark$$

Logically Invalid Reasoning

Turing, in his famous paper, claims:

TT: If x is conversationally indistinguishable from a human, then x is conscious/sentient.

This proposition is logically equivalent to:

TT': If x is not conversationally indistinguishable from a human, then x is not conscious/sentient.

$i \rightarrow c$

X

$\neg i \rightarrow \neg c$

Well then, obviously Turing is wrong: A pet dog isn't conversationally indistinguishable from a human, but the dog *is* conscious/sentient. We have a counter-example!

$\neg i \wedge c$

$\{ \neg i \wedge c \} \vdash \neg(\neg i \rightarrow \neg c) \quad \checkmark$

What *is* Logic?

What *is* Logic?

- The key to becoming rational.

What *is* Logic?

- The key to becoming rational.
- “The science and engineering of reasoning.” — so the not-unreasonable slogan goes.

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.

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”!

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).

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.

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.

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.

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).

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.

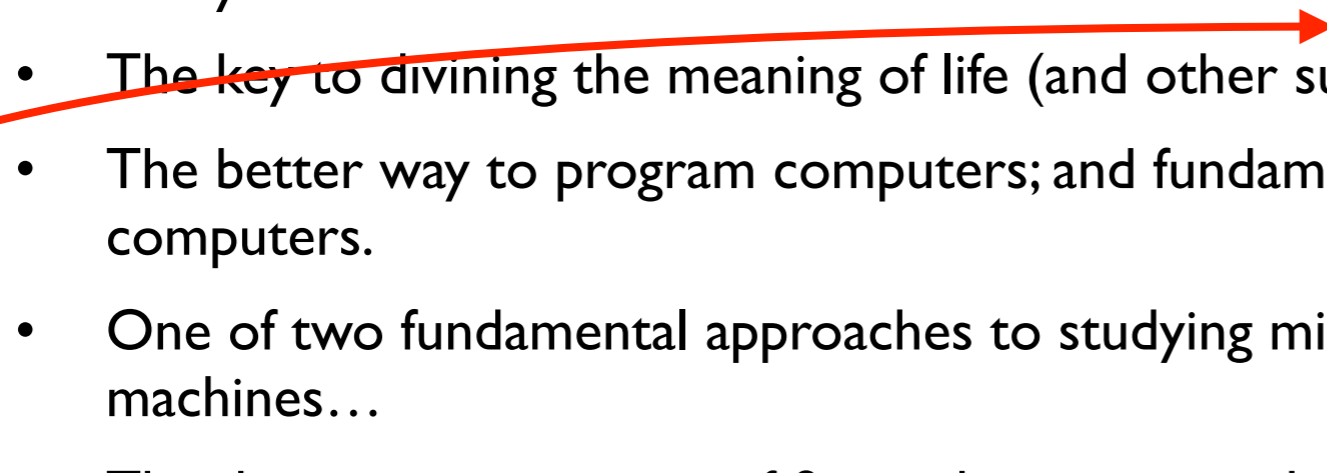
What is Logic?

- 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...

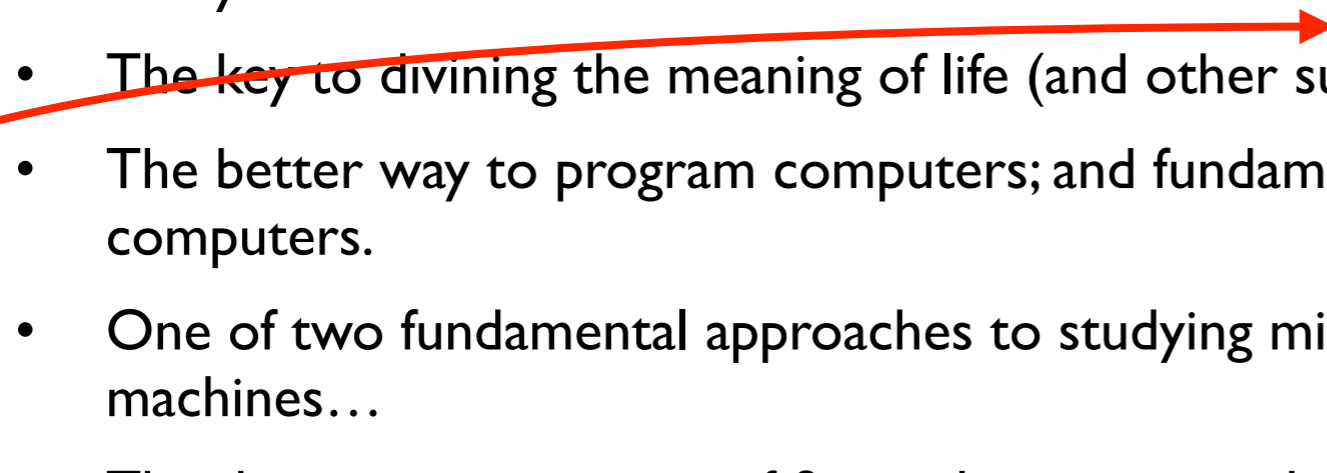
What is Logic?

- 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)?...
- 

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 n and you); and mindless machines (like Deep Blue & \
- A key to riches.
- The key to divining the meaning of life (and other su
- The better way to program computers; and fundam computers.
- One of two fundamental approaches to studying mi 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 n and you); and mindless machines (like Deep Blue & \
- A key to riches.
- The key to divining the meaning of life (and other su
- The better way to program computers; and fundam computers.
- One of two fundamental approaches to studying mi machines...
- The thing many creatures of fiction have mastered — have you (as a New Yorker)?...
- ...



What *is* Logic?

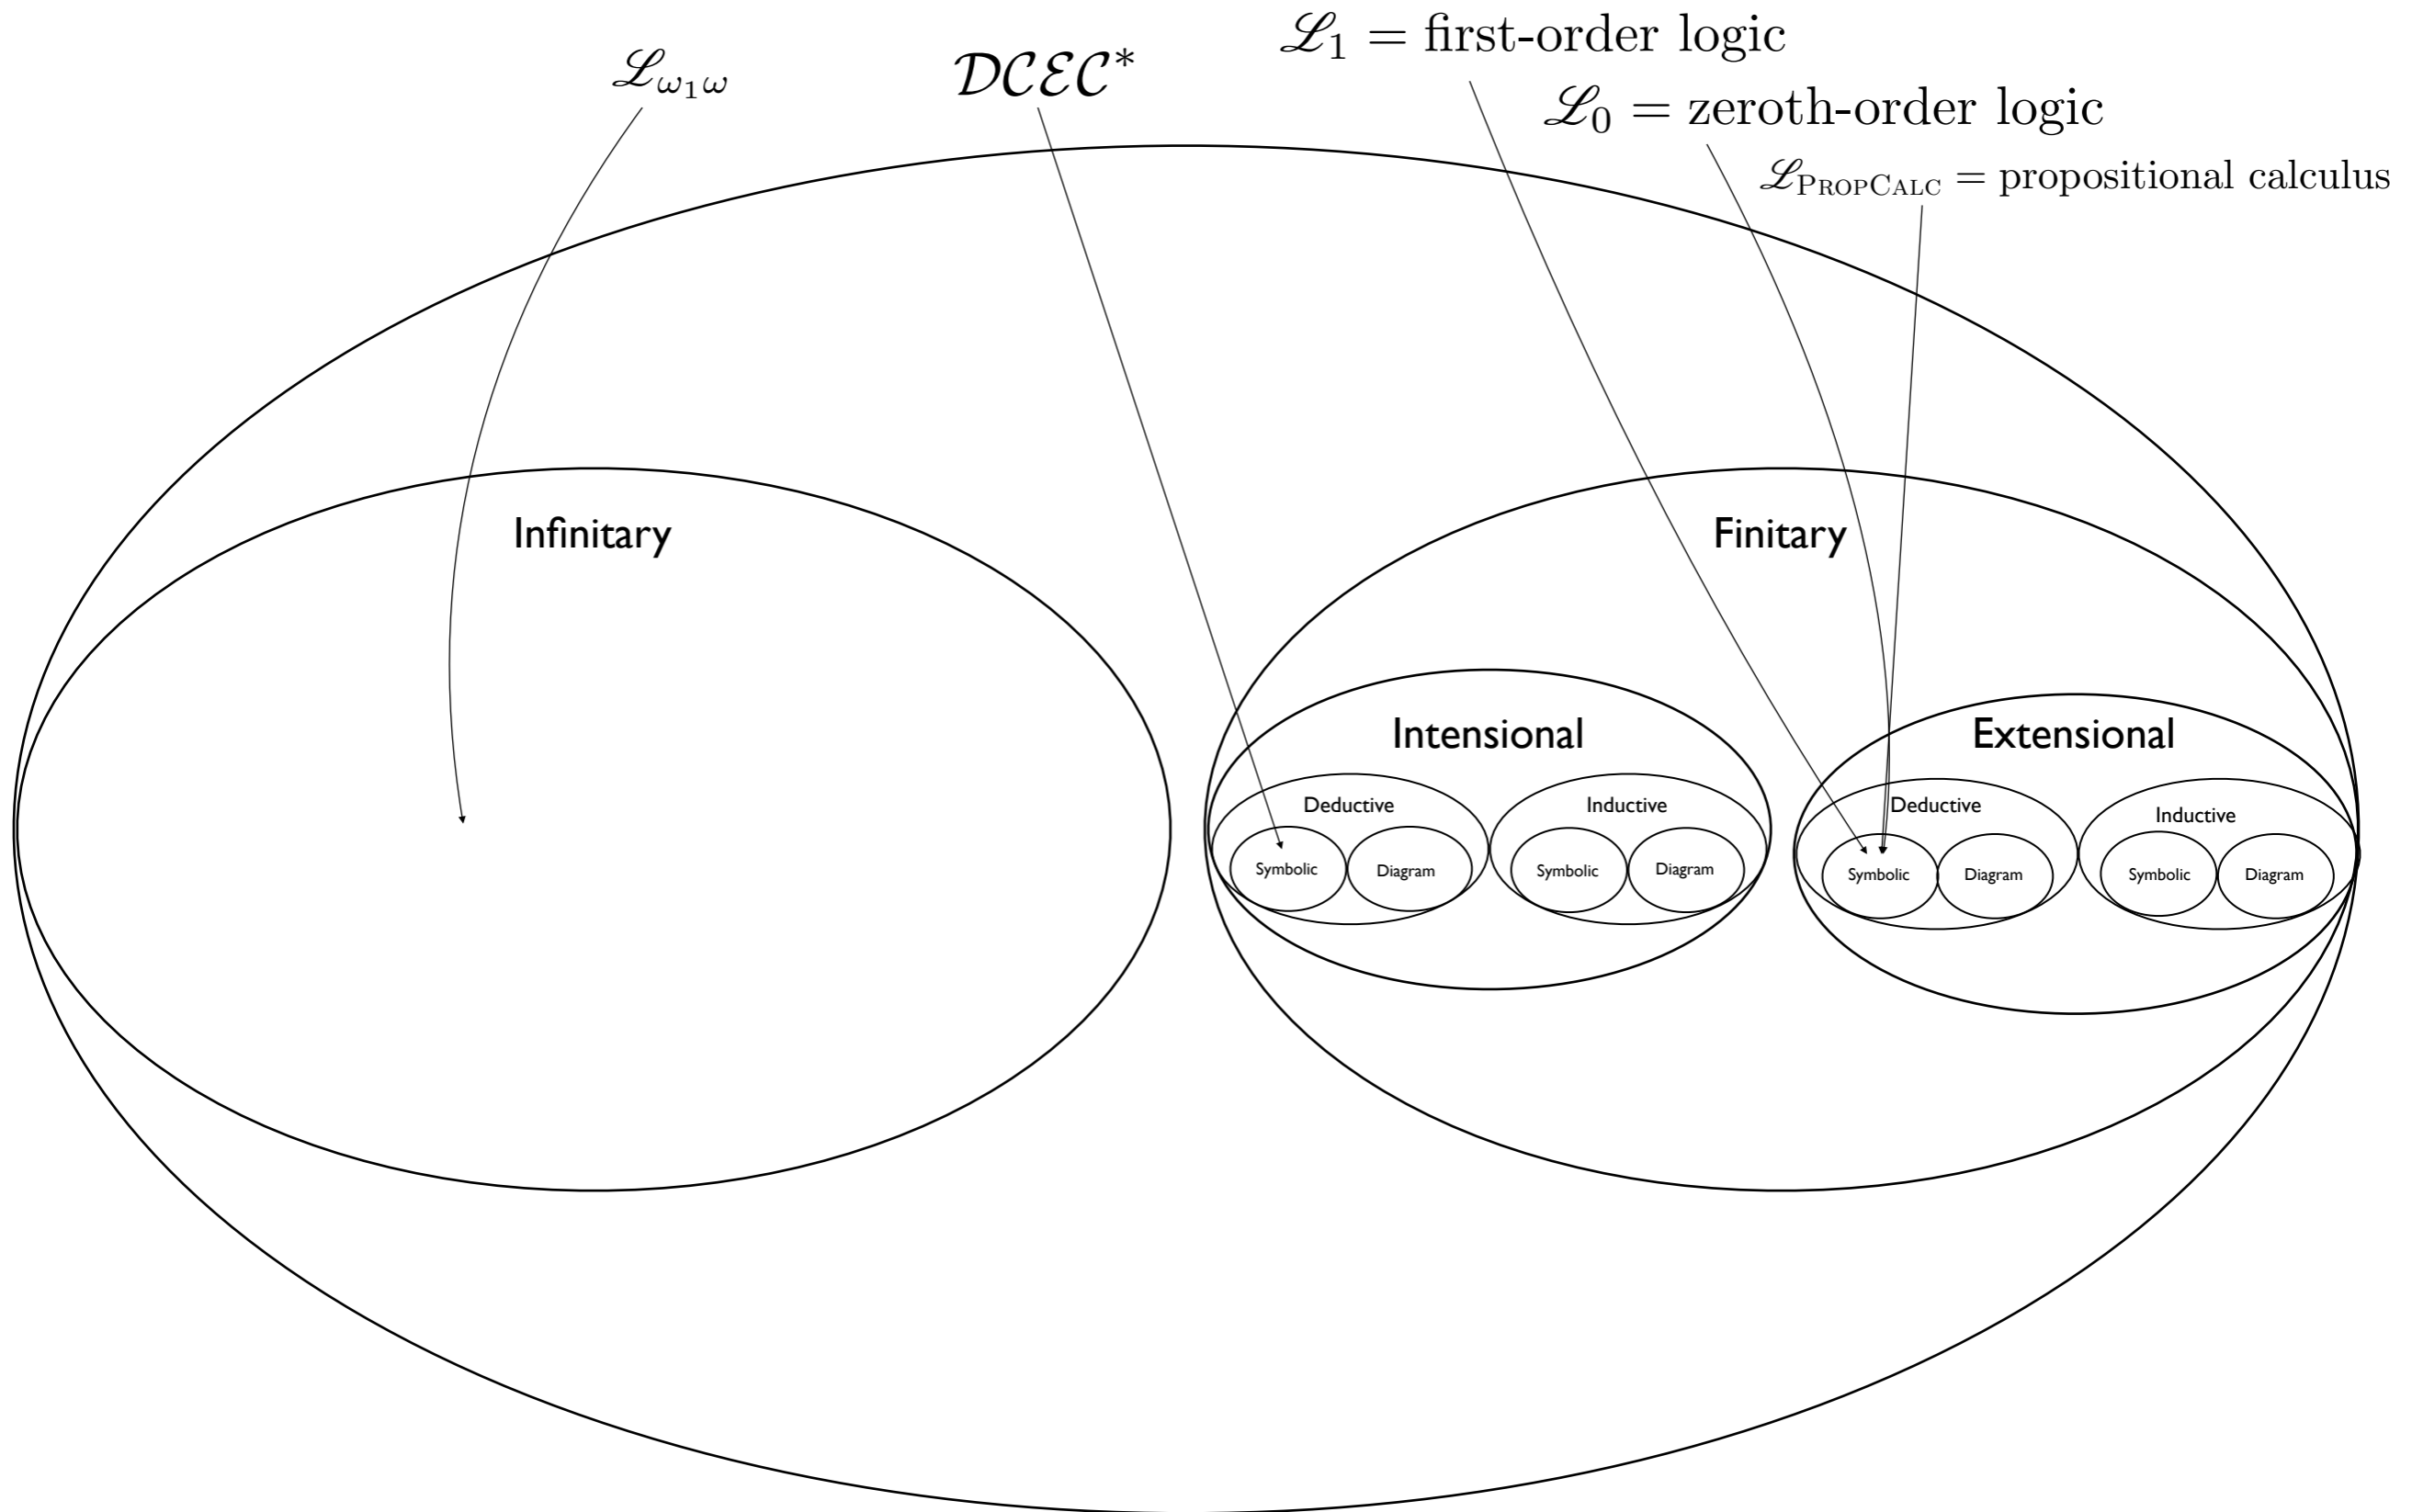
What is Logic?

- 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)?...
- ...

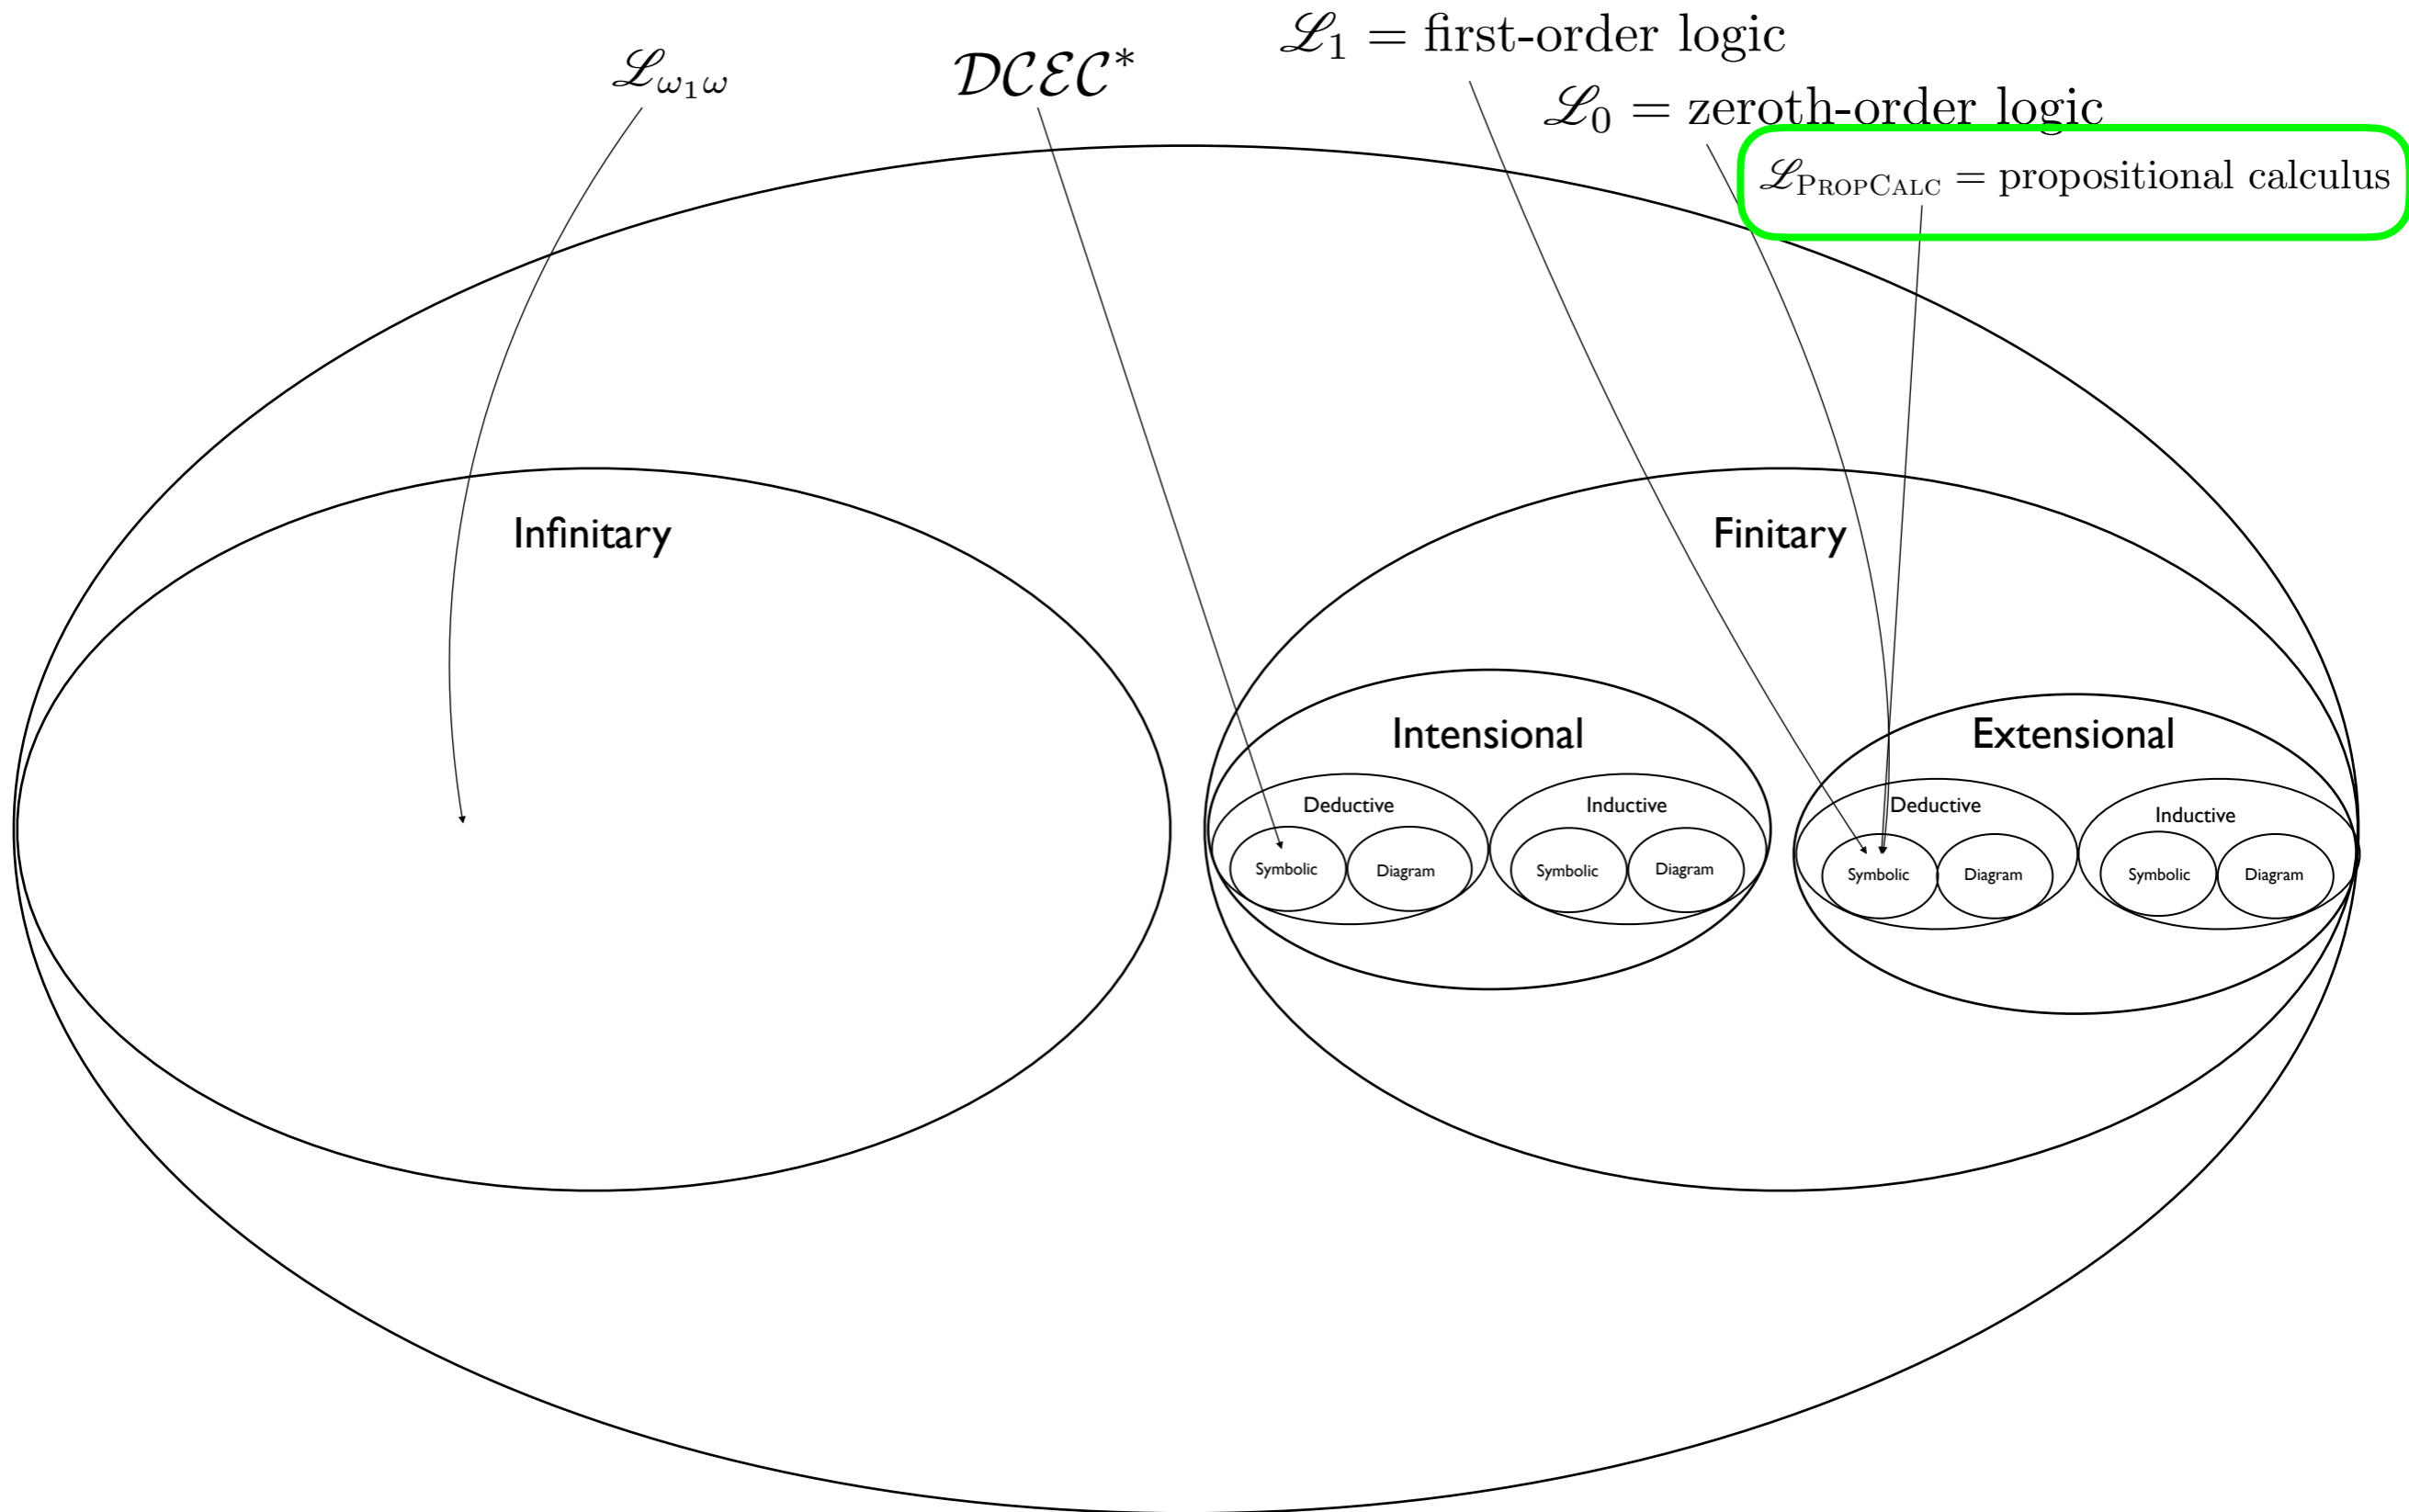
What is Logic?

- 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)?...
- ...

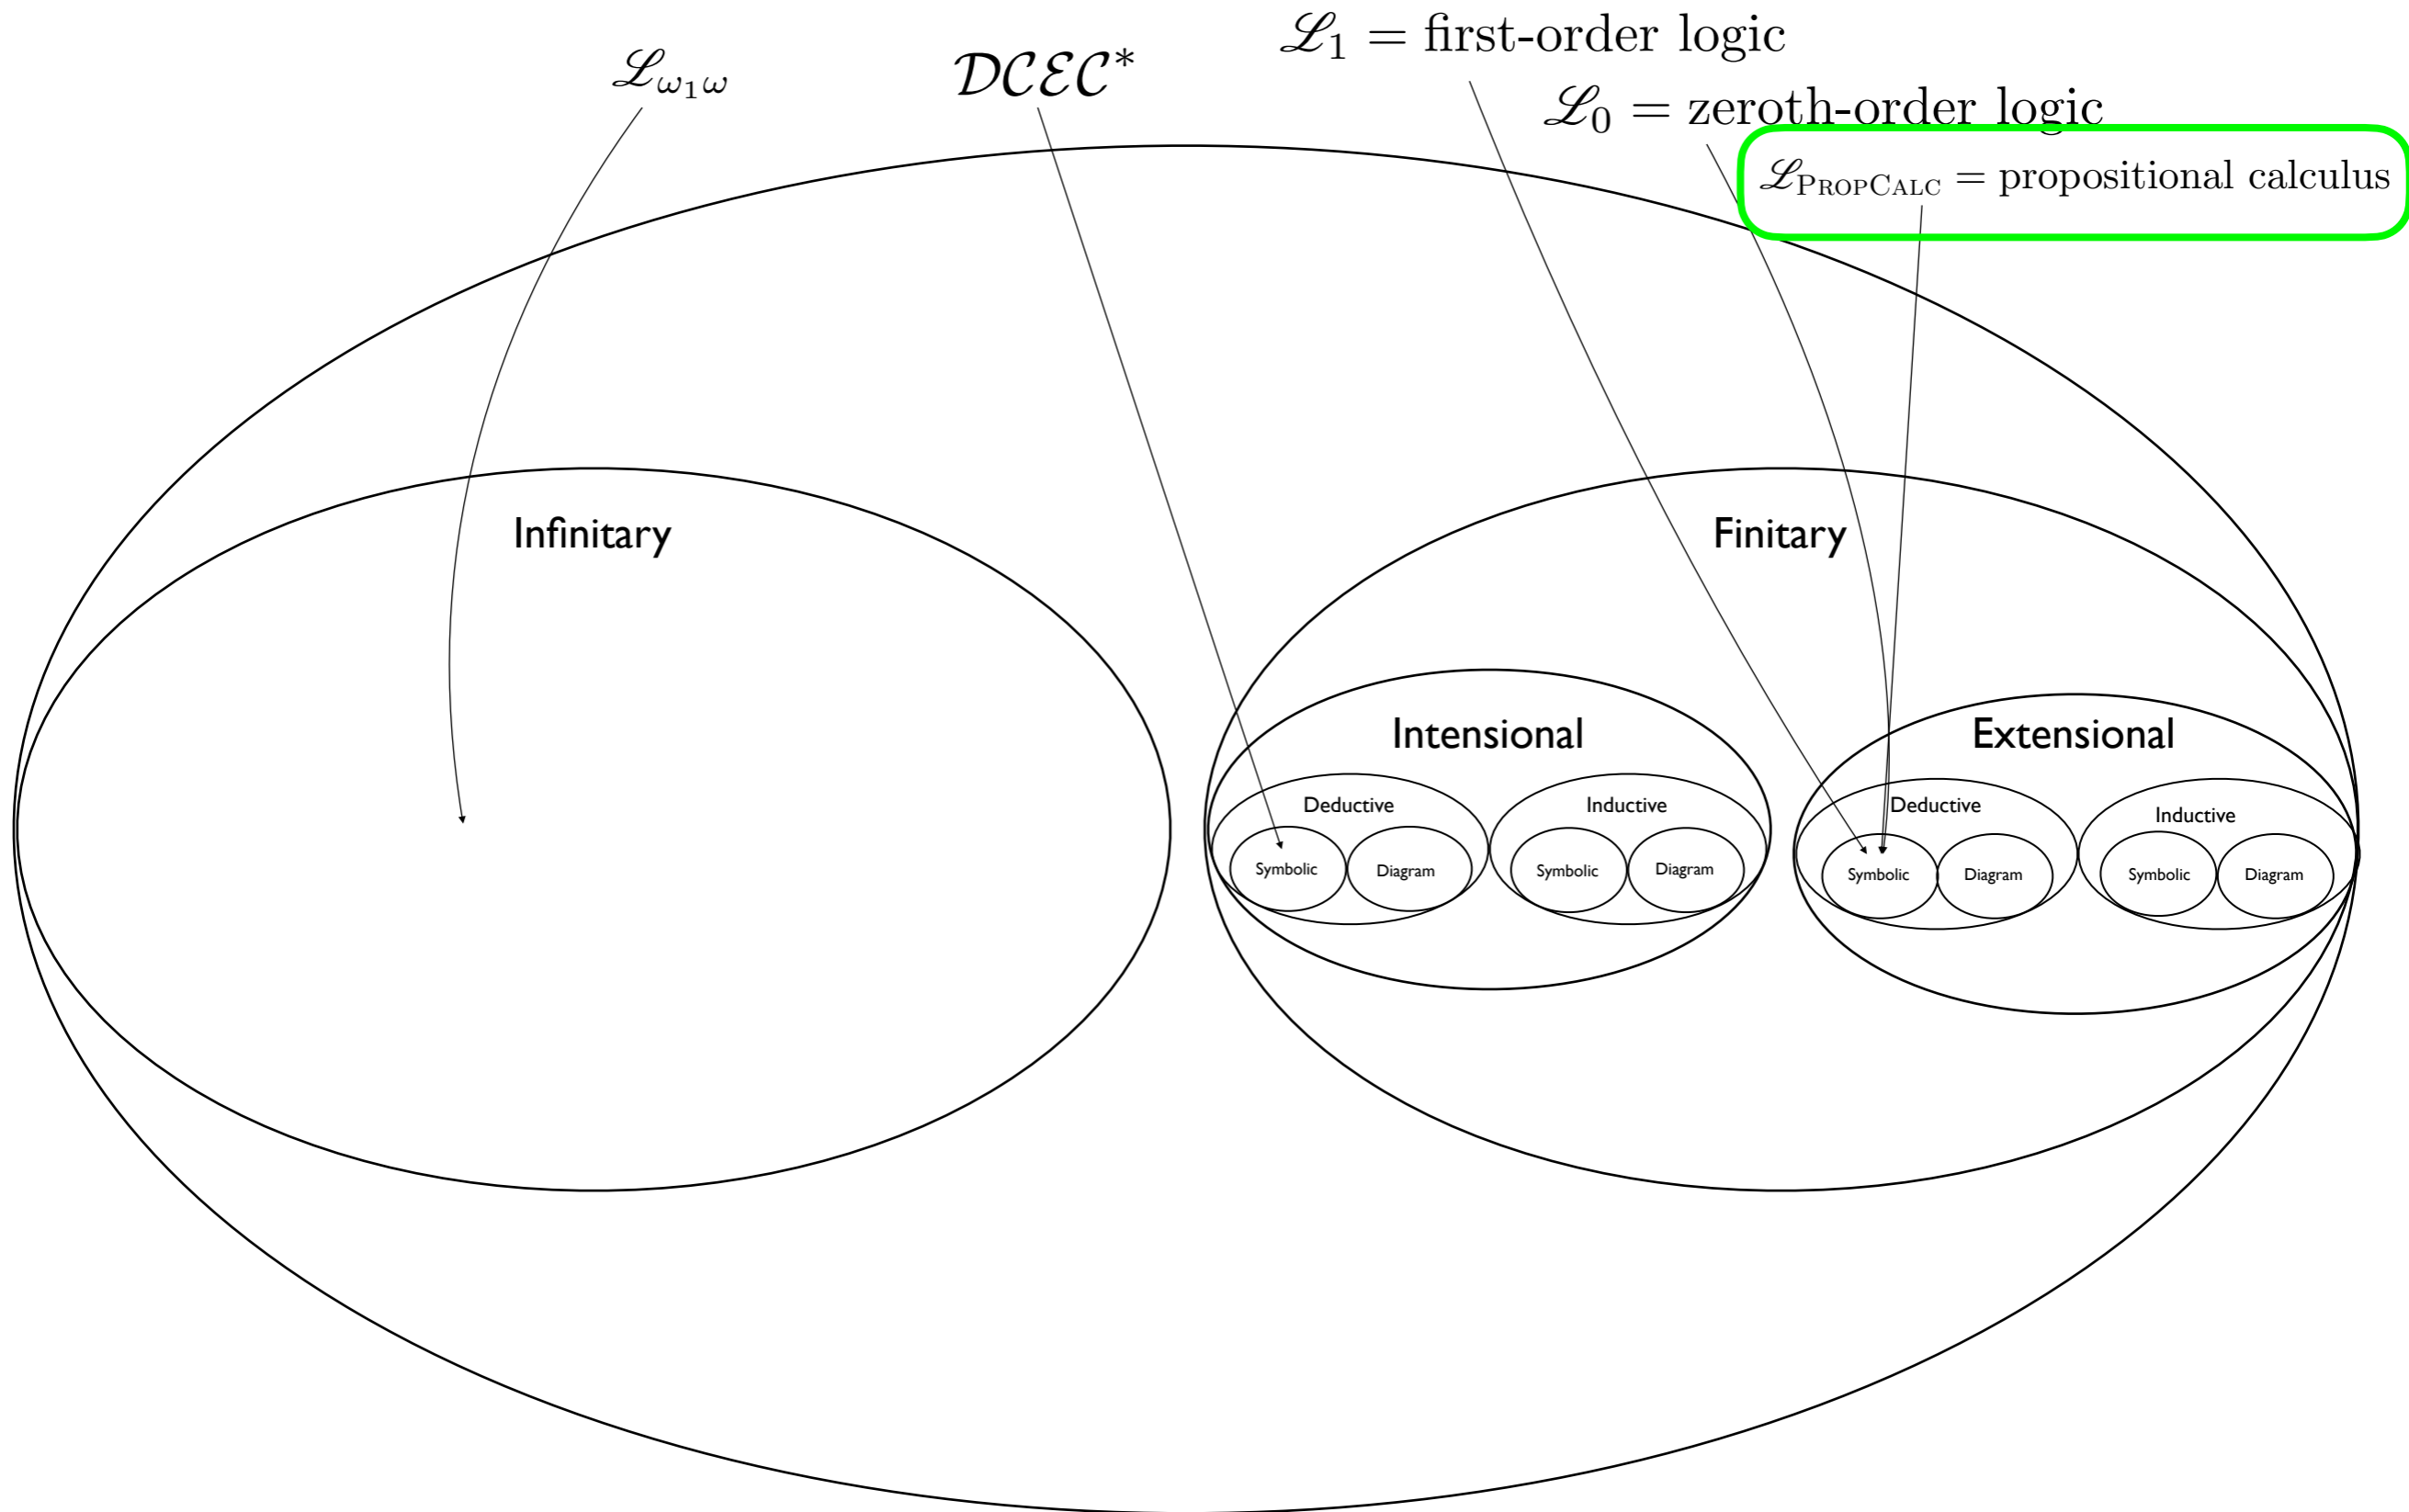
The Universe of Logics



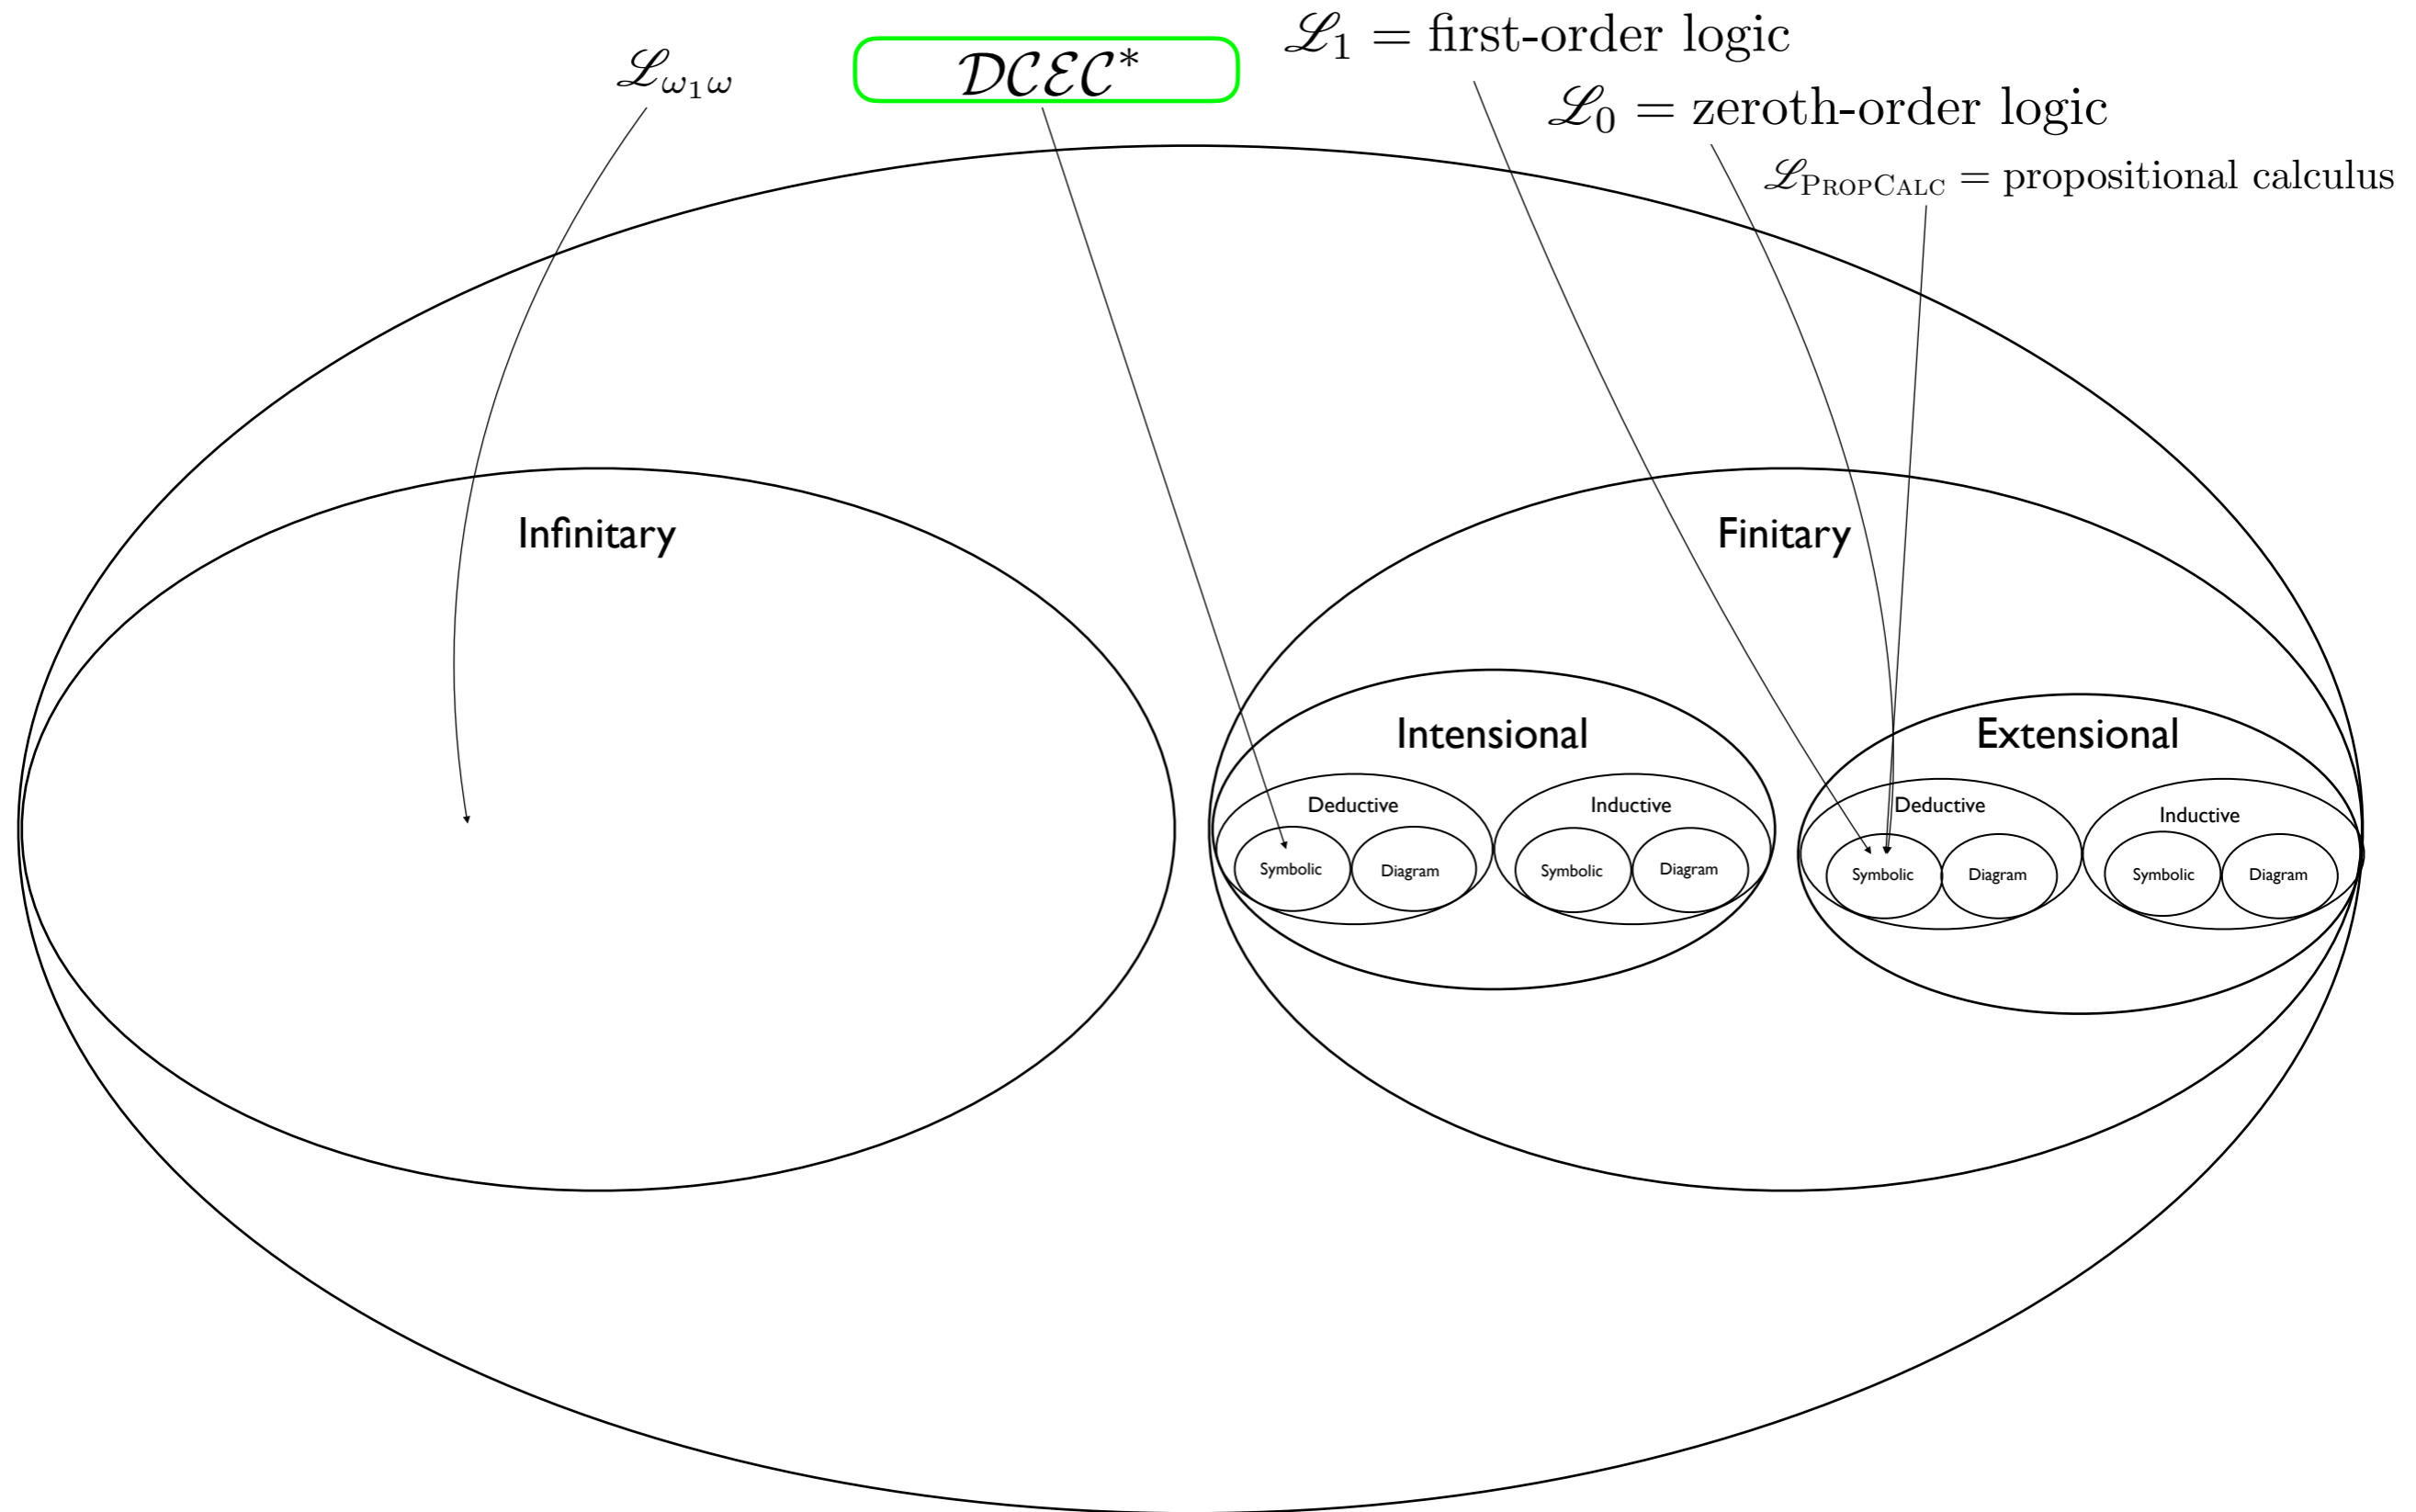
The Universe of Logics



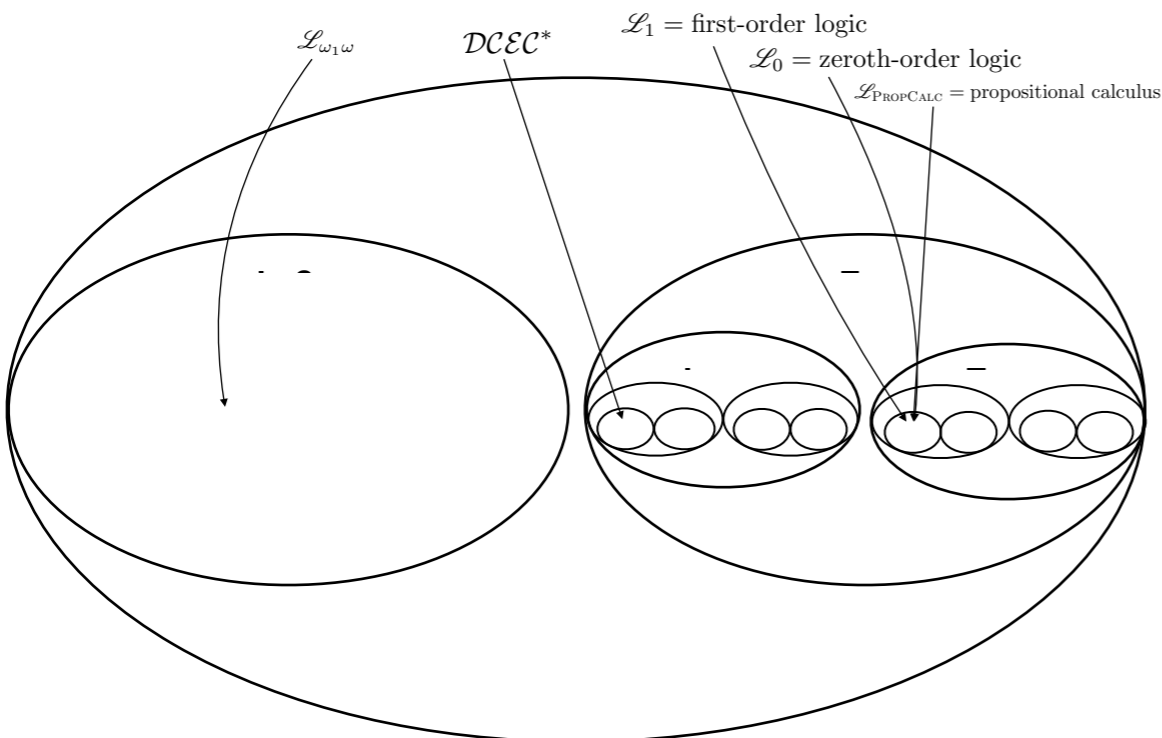
The Universe of Logics



The Universe of Logics

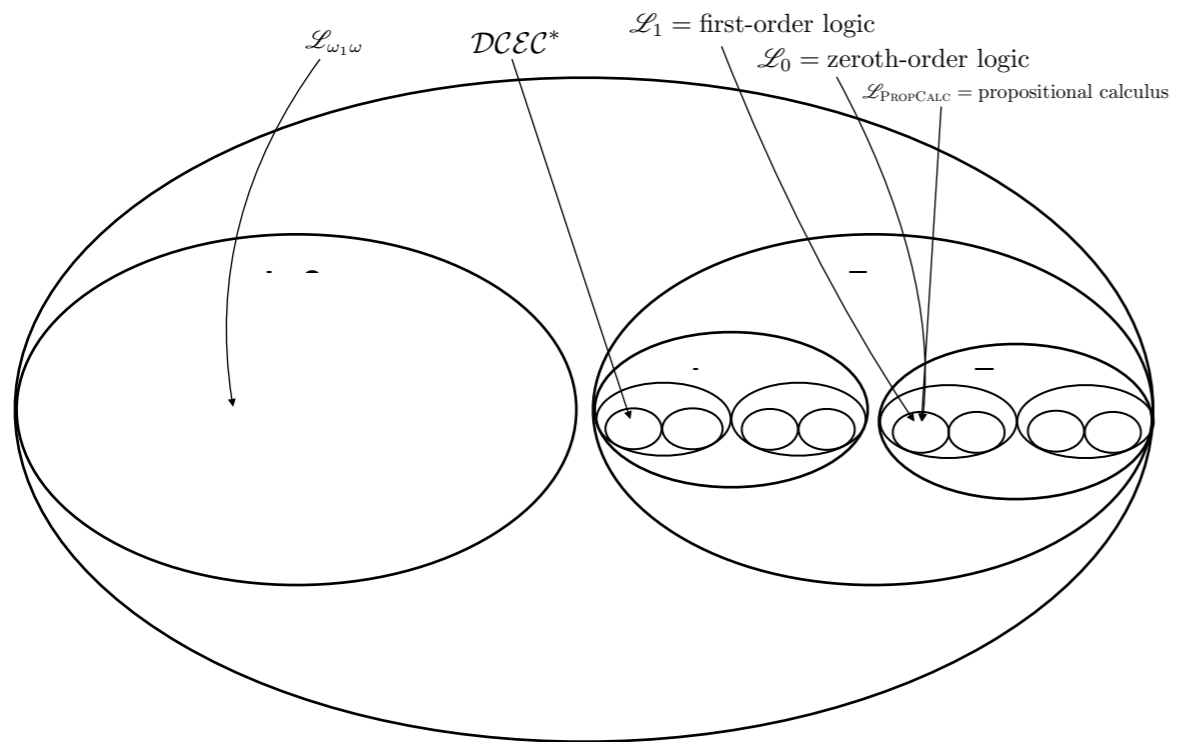


The Universe of Logics



Non-Physical

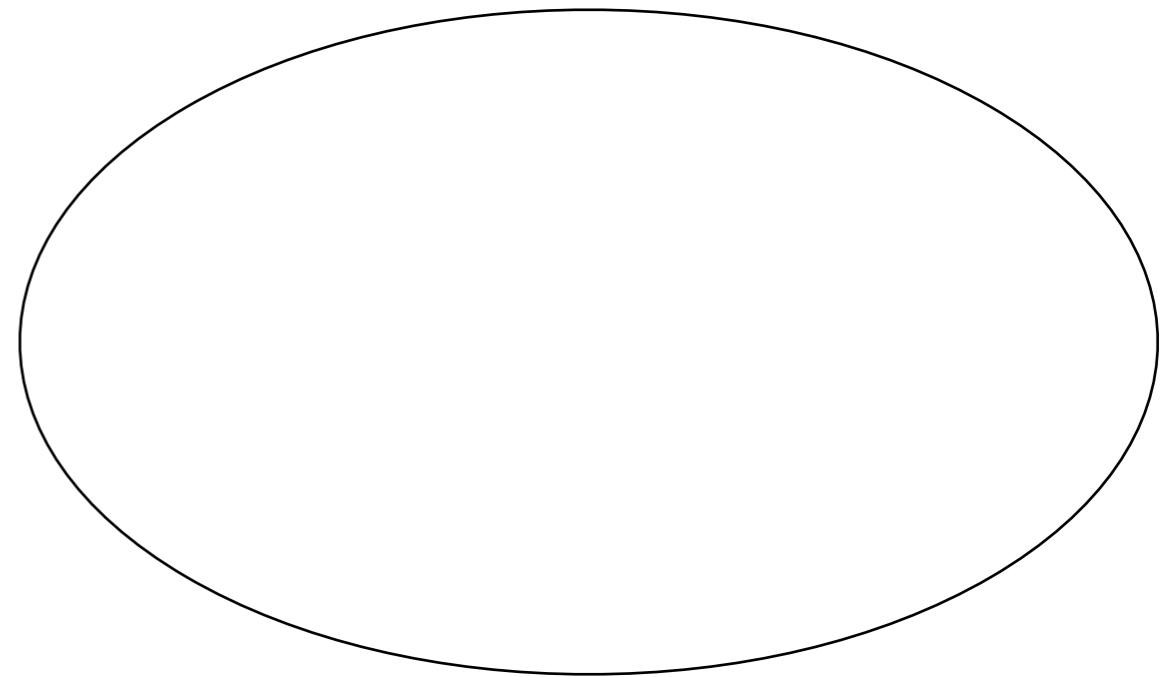
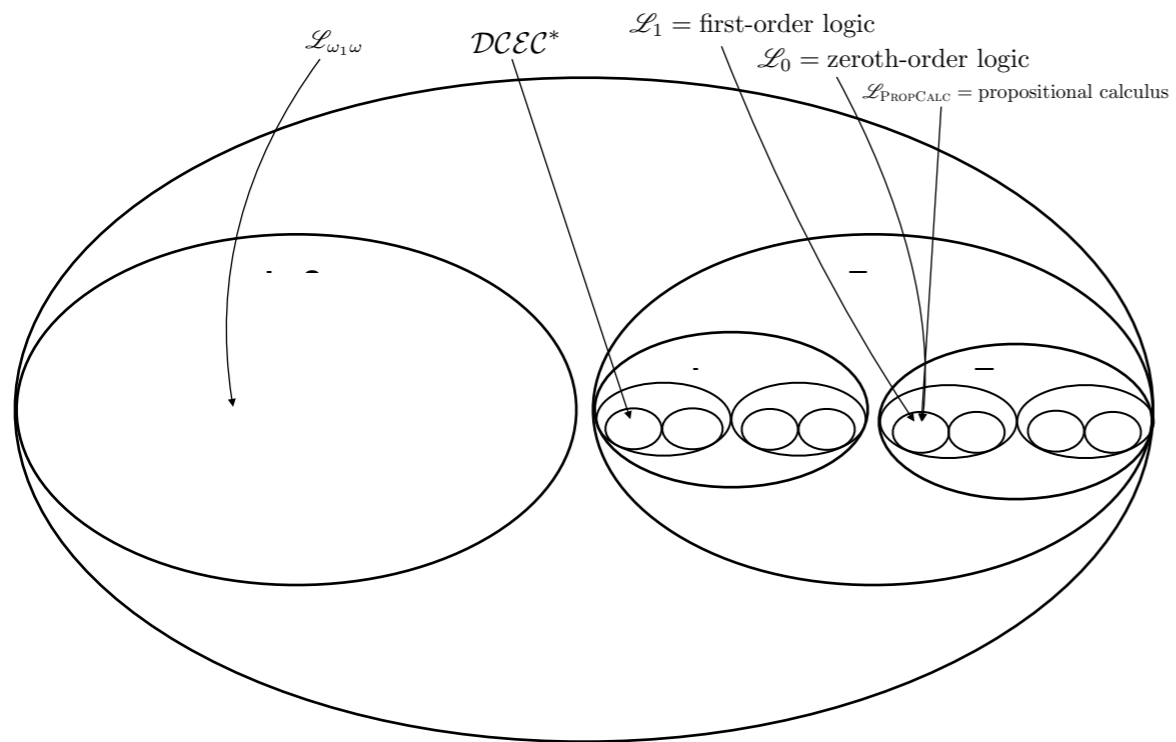
The Universe of Logics



Non-Physical

The Universe of Logics

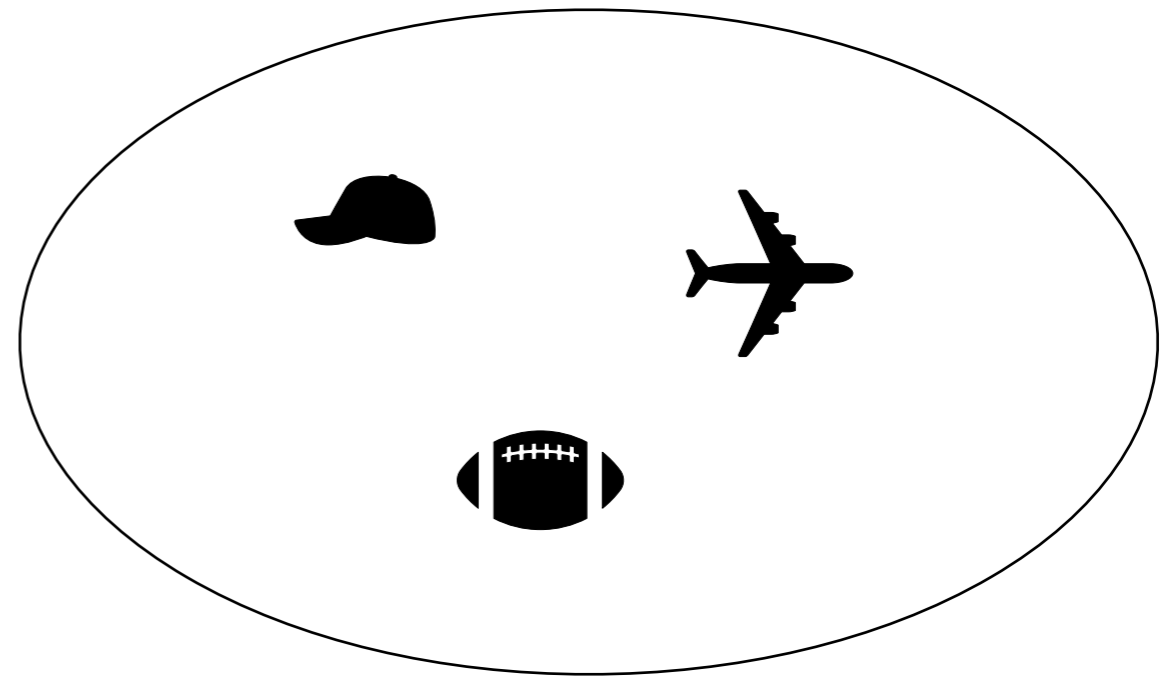
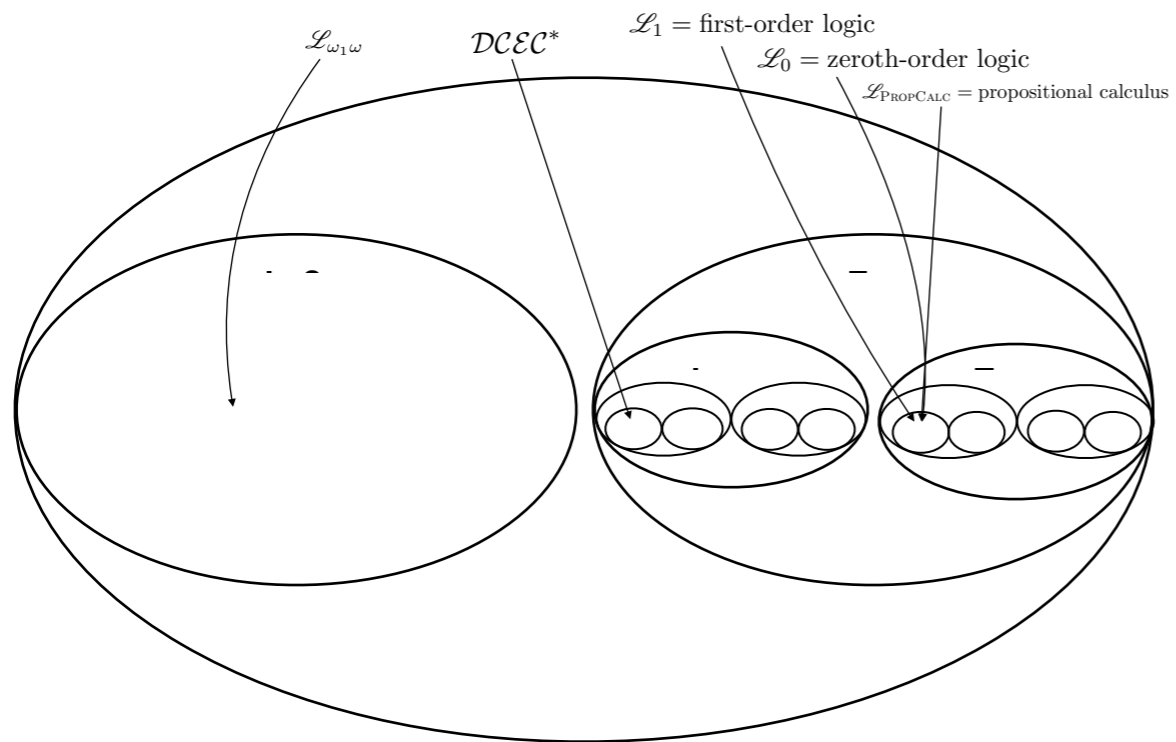
The Physical Universe



Non-Physical

The Universe of Logics

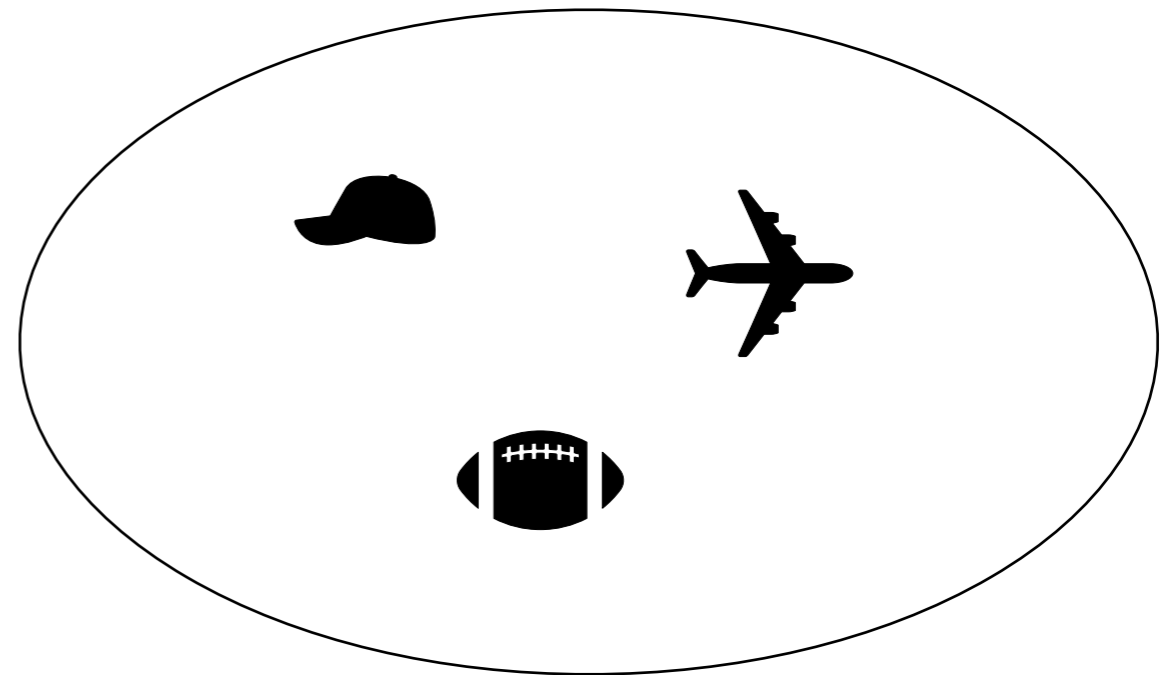
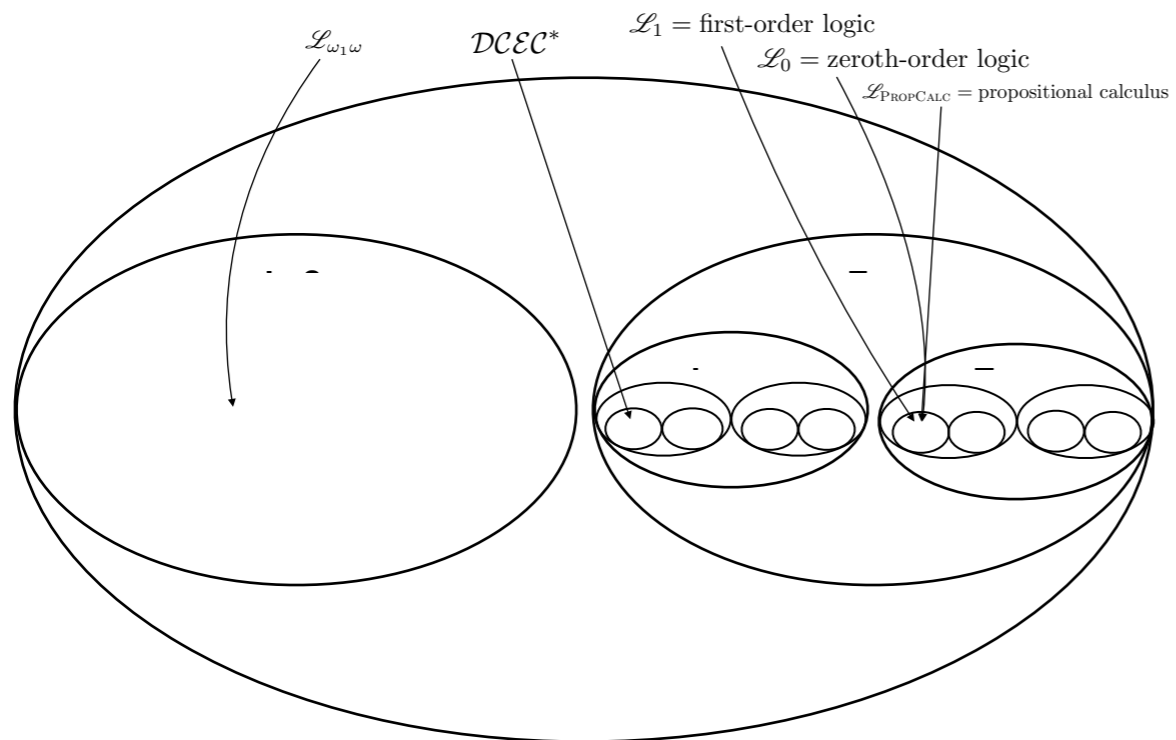
The Physical Universe



Non-Physical

The Universe of Logics

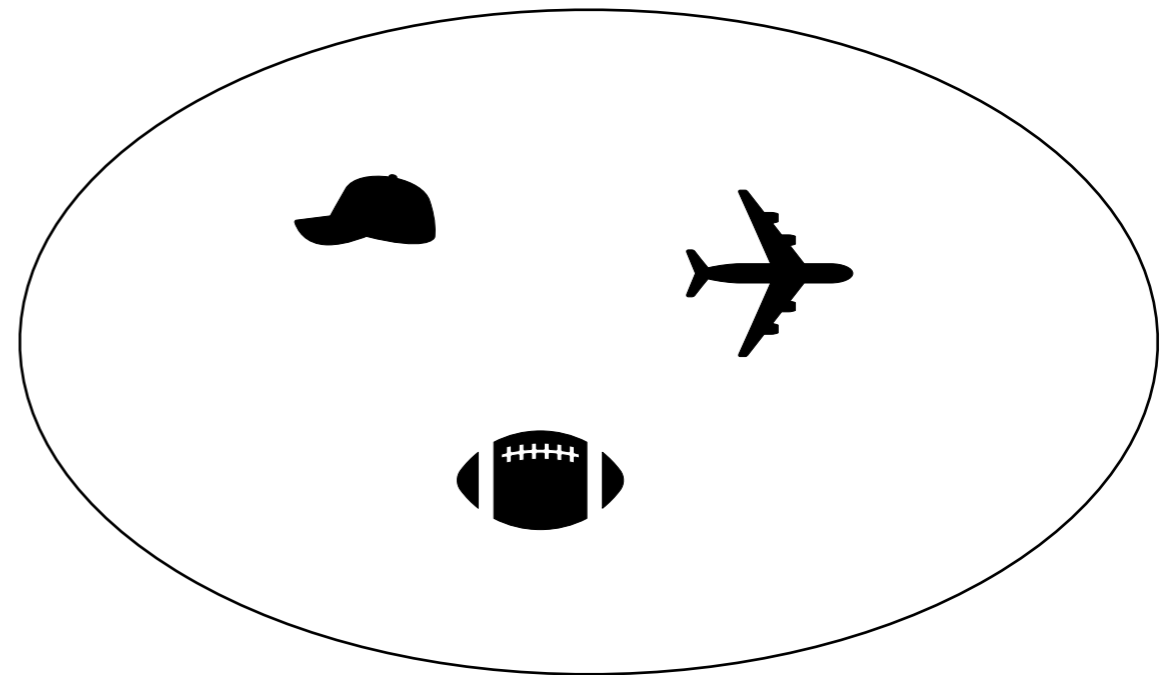
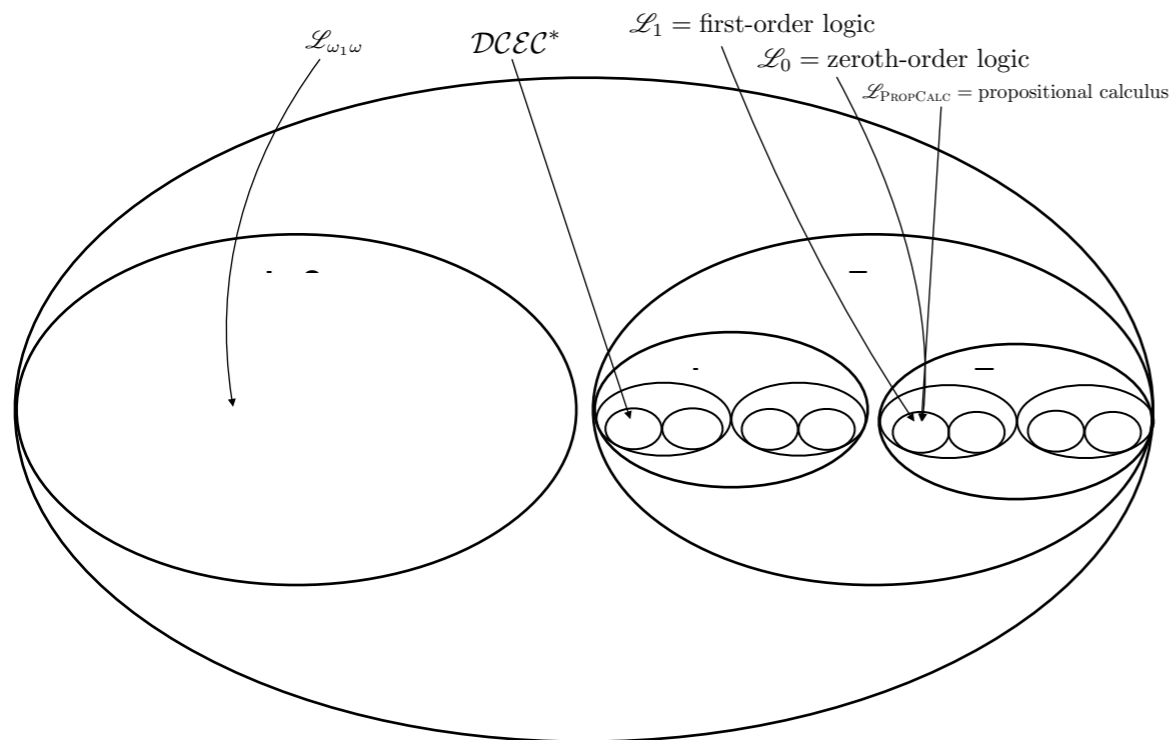
The Physical Universe



Non-Physical

The Universe of Logics

The Physical Universe



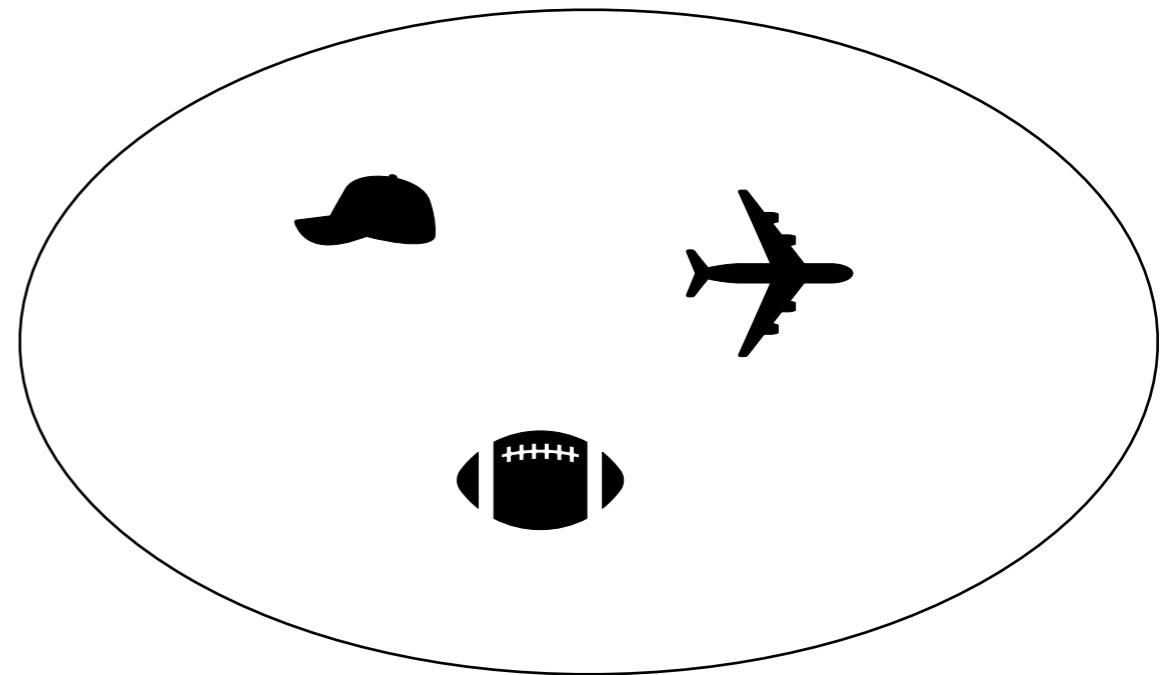
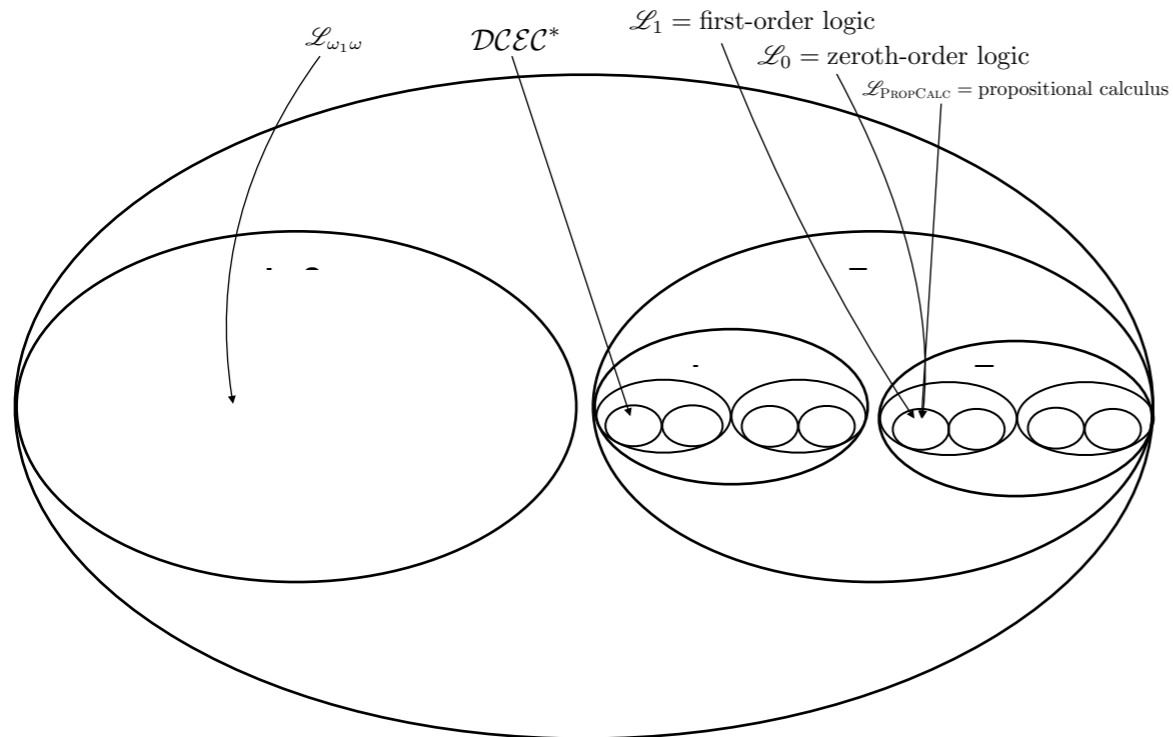
$$\frac{\phi, \phi \rightarrow \psi}{\psi}$$

modus ponens

Non-Physical

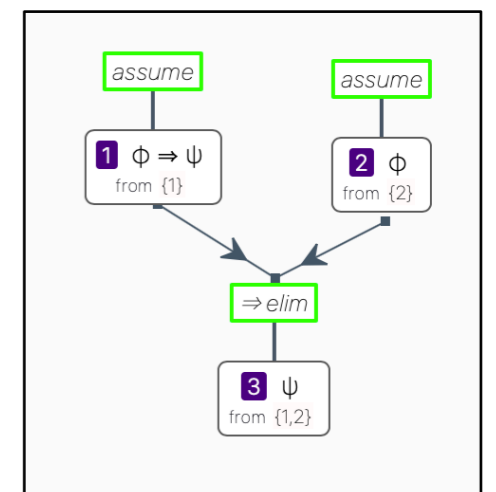
The Universe of Logics

The Physical Universe



$$\frac{\phi, \phi \rightarrow \psi}{\psi}$$

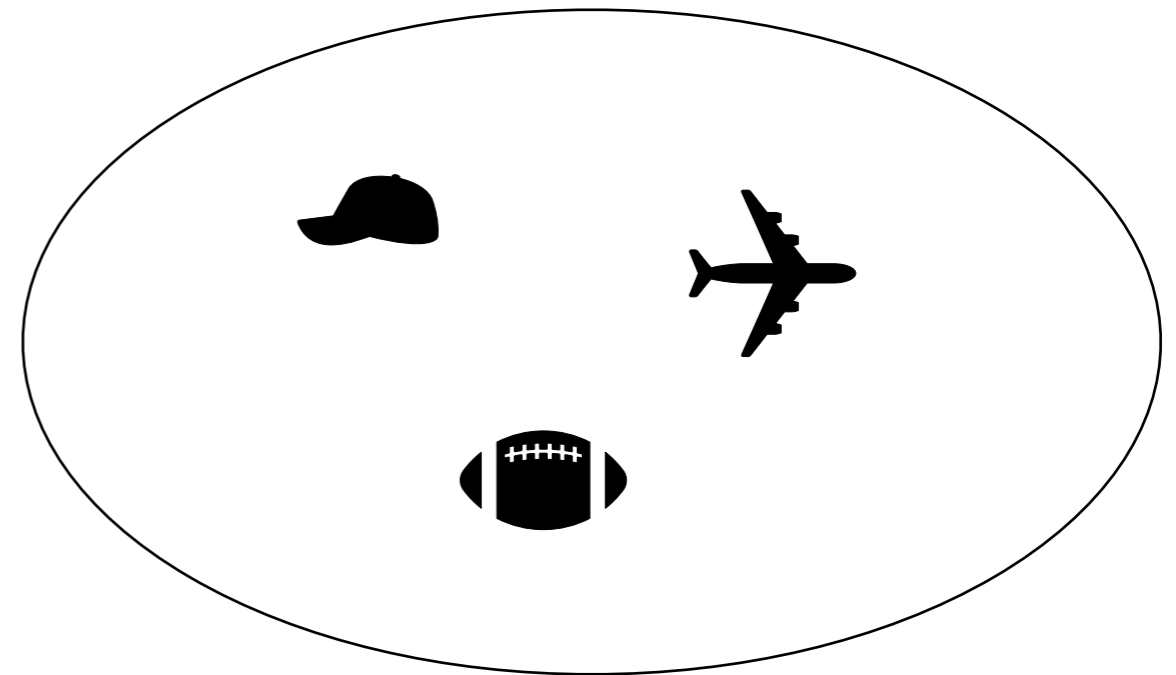
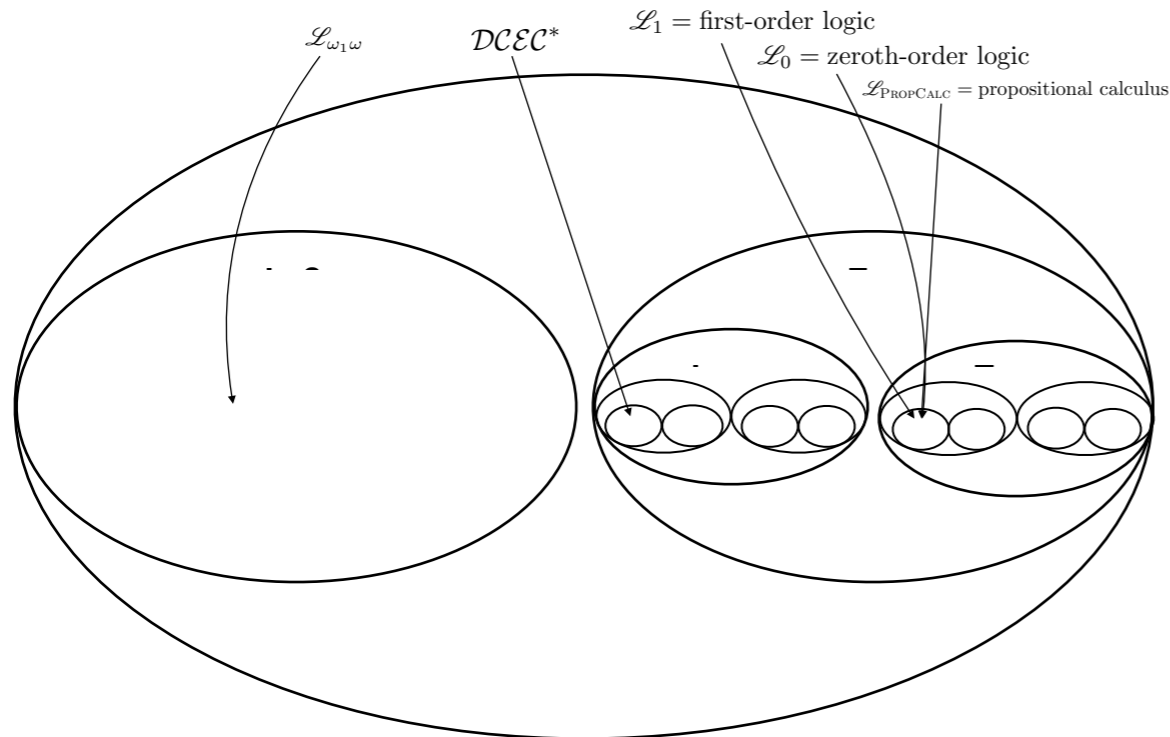
modus ponens



Non-Physical

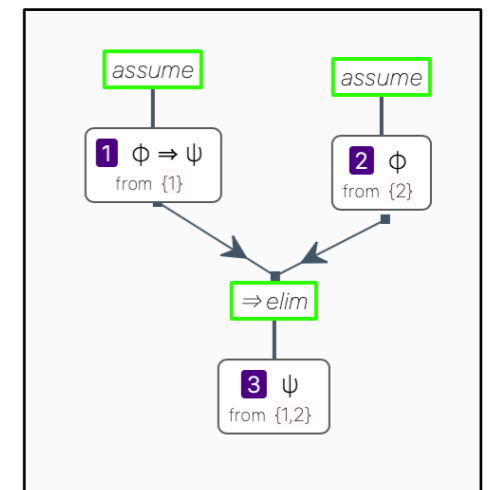
The Universe of Logics

The Physical Universe



$$\frac{\phi, \phi \rightarrow \psi}{\psi}$$

modus ponens

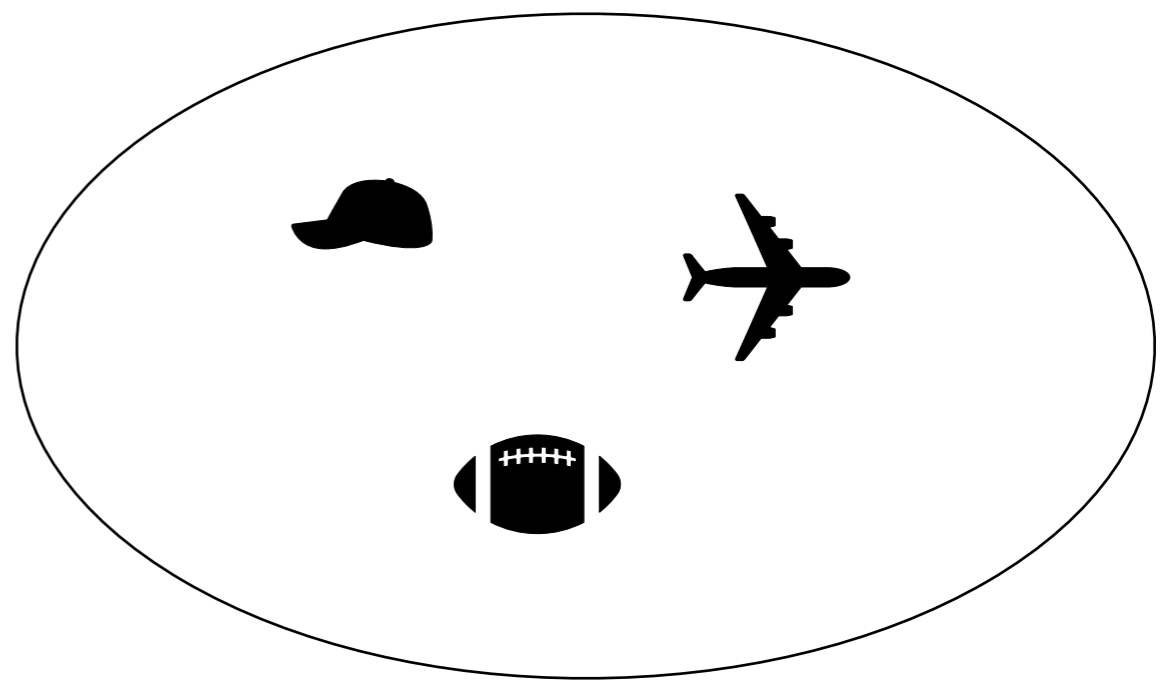
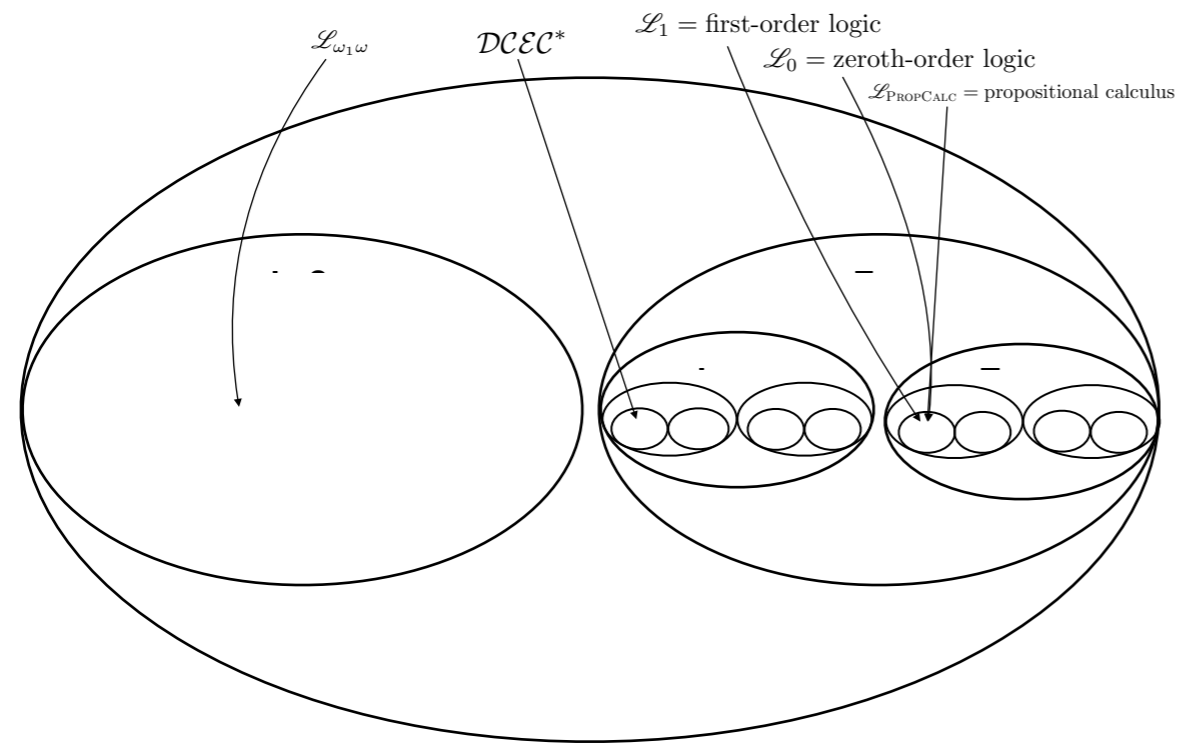




Non-Physical

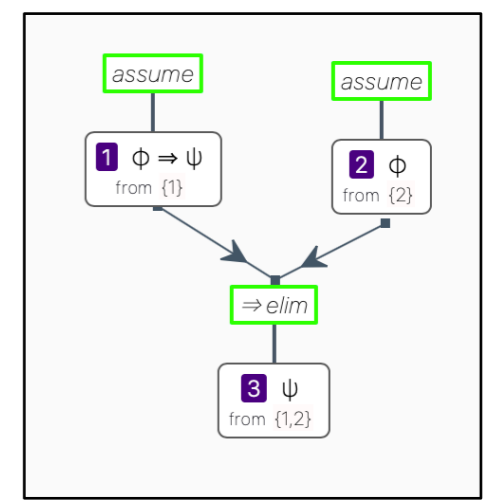
The Universe of Logics

The Physical Universe



$$\frac{\phi, \phi \rightarrow \psi}{\psi}$$

modus ponens

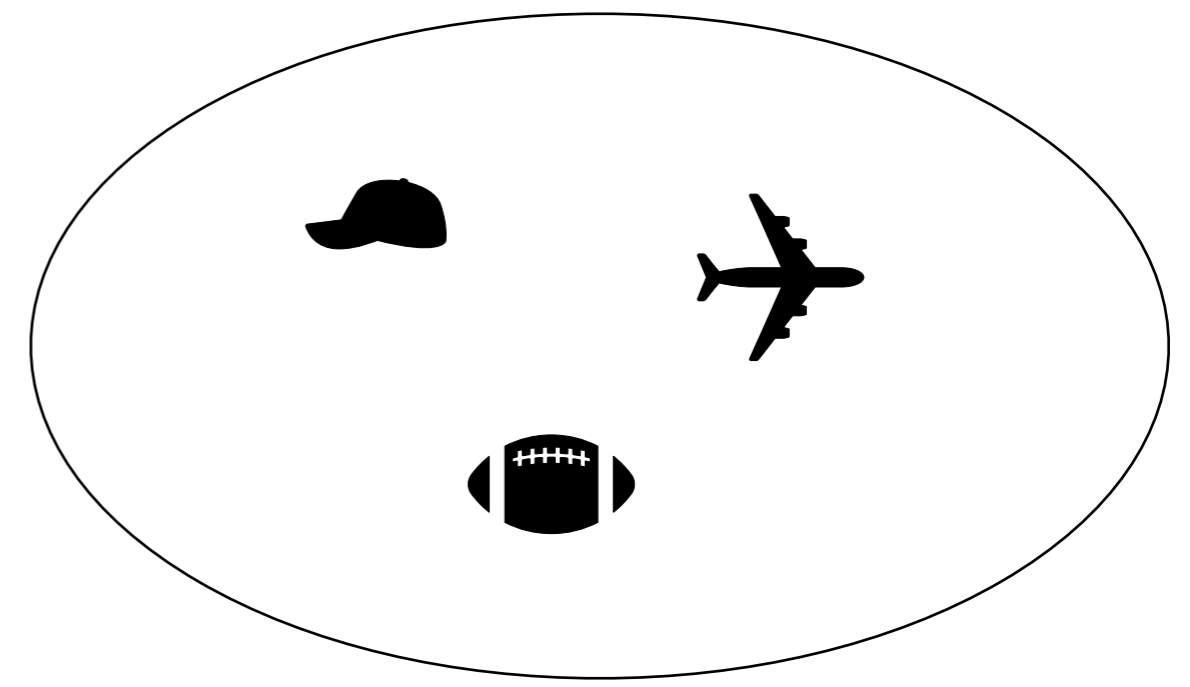
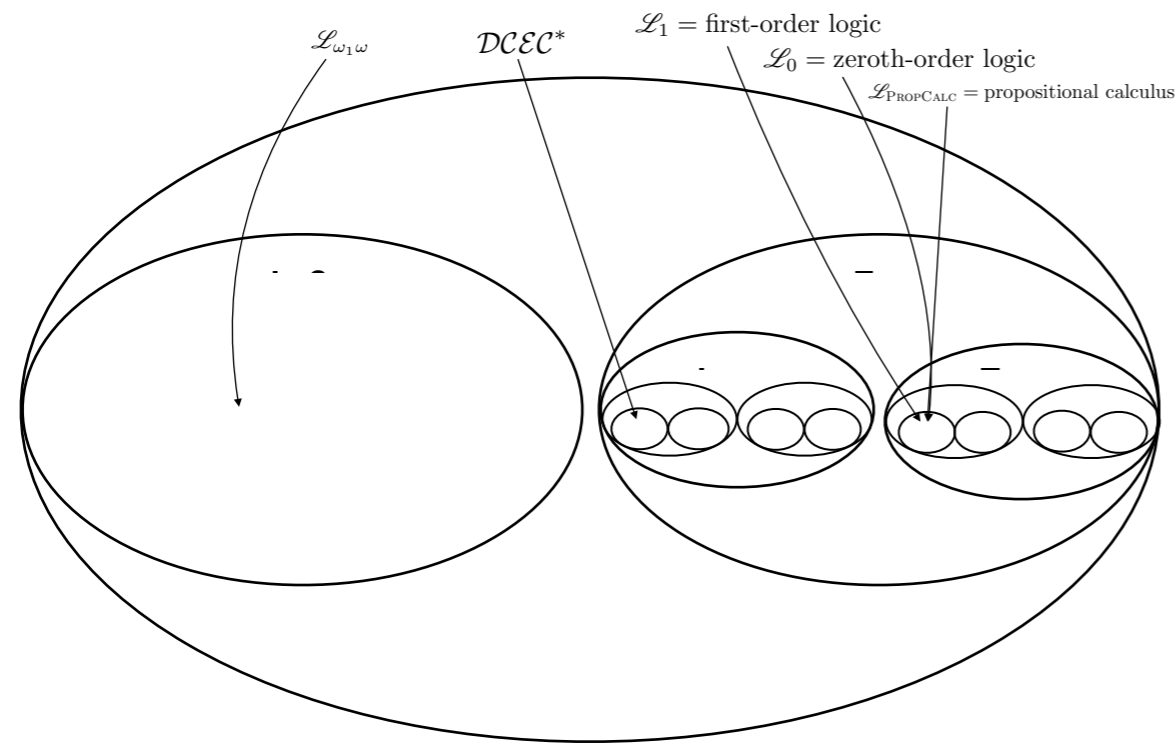




Non-Physical

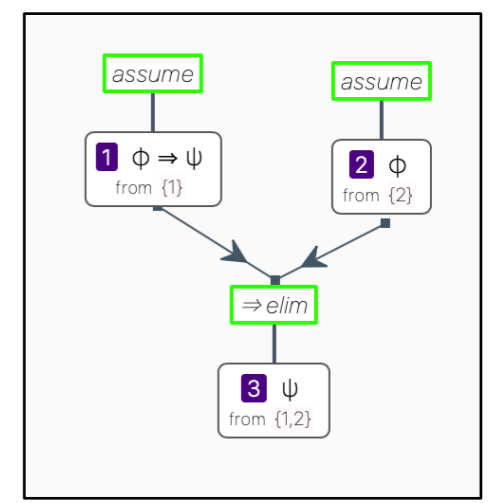
The Universe of Logics

The Physical Universe



$$\frac{\phi, \phi \rightarrow \psi}{\psi}$$

modus ponens

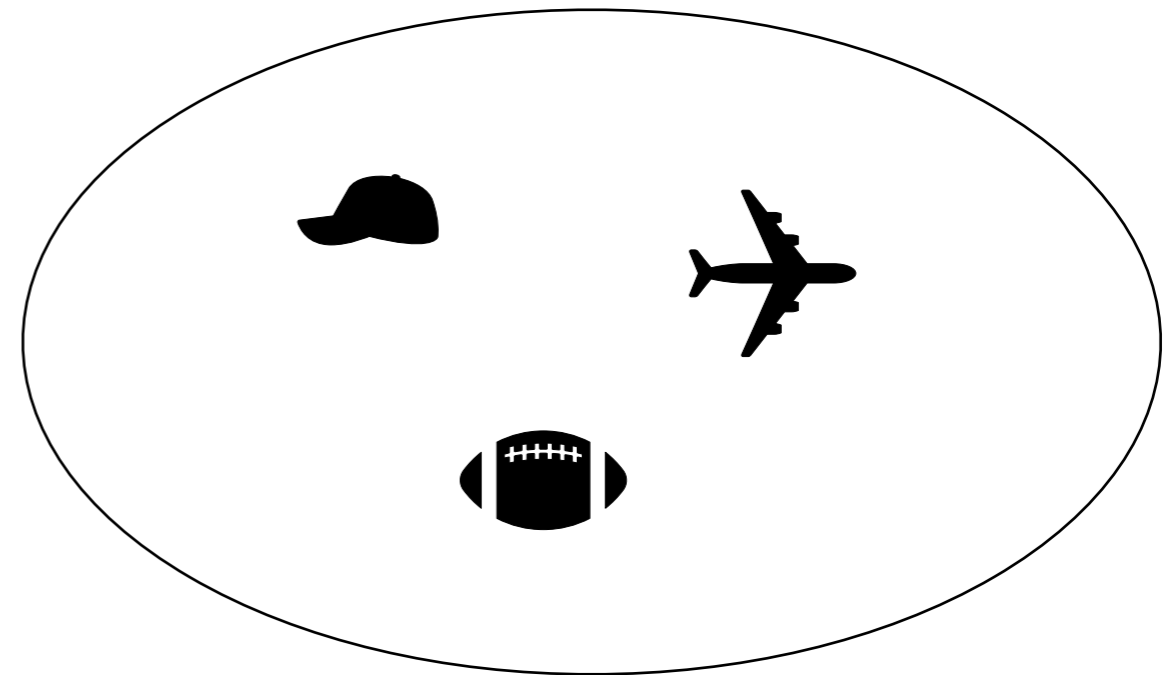
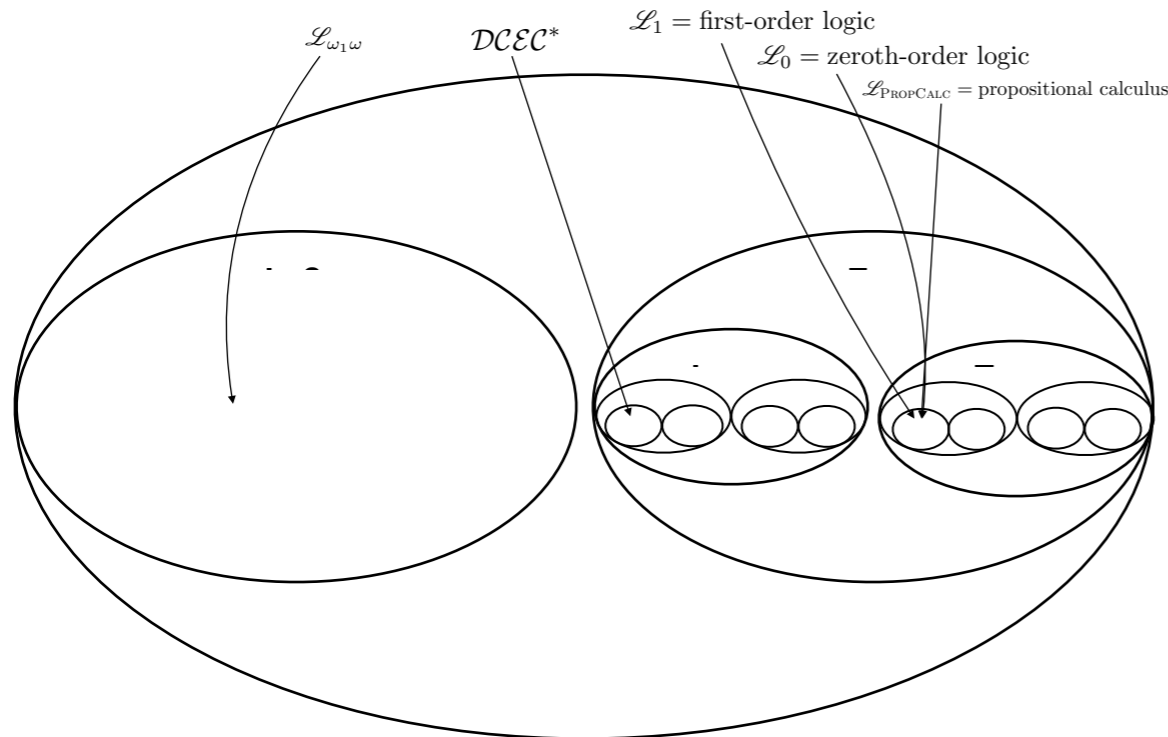


Non-Physical



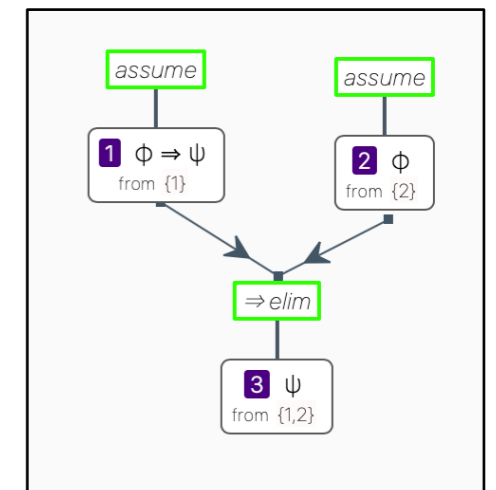
The Universe of Logics

The Physical Universe



$$\frac{\phi, \phi \rightarrow \psi}{\psi}$$

modus ponens

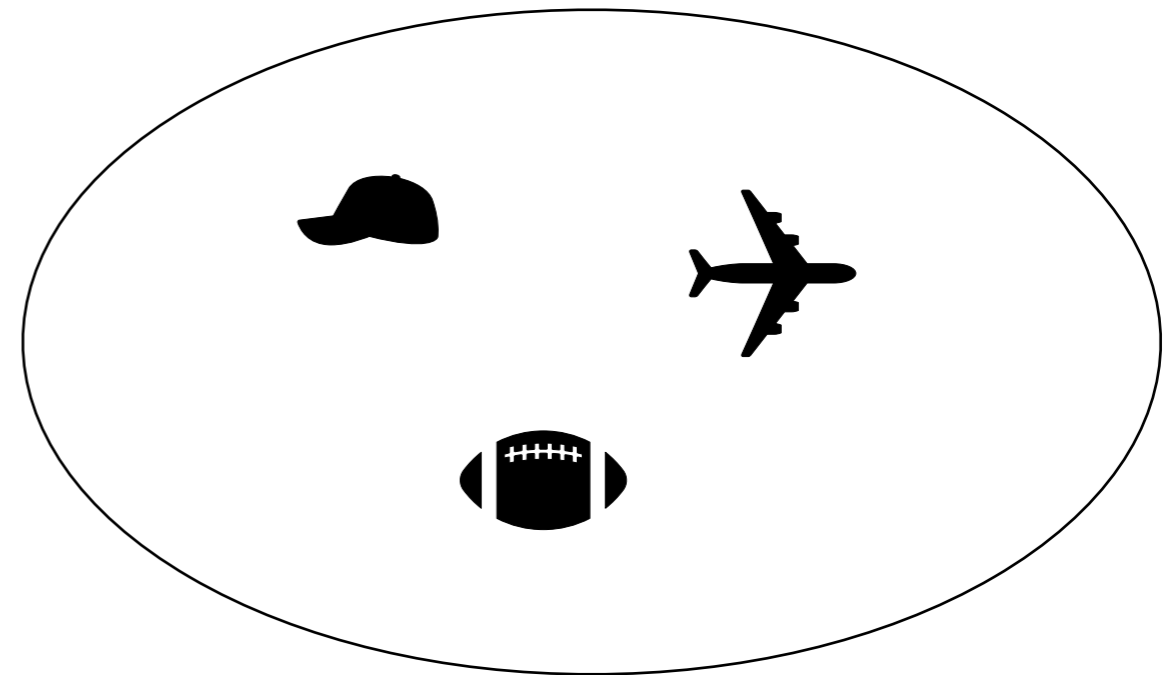
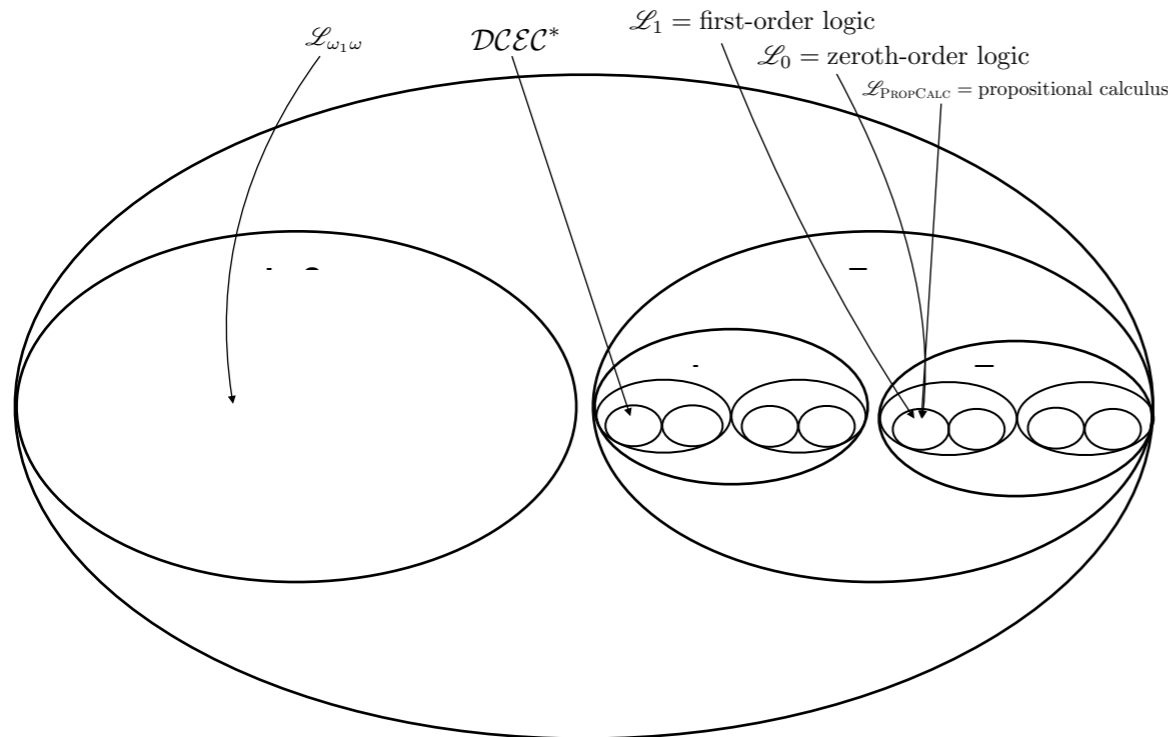


Non-Physical



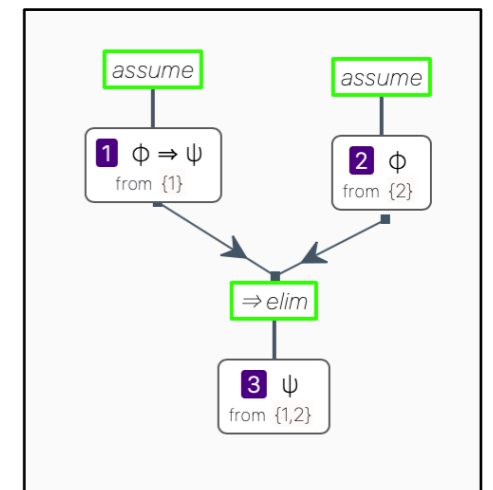
The Universe of Logics

The Physical Universe



$$\frac{\phi, \phi \rightarrow \psi}{\psi}$$

modus ponens



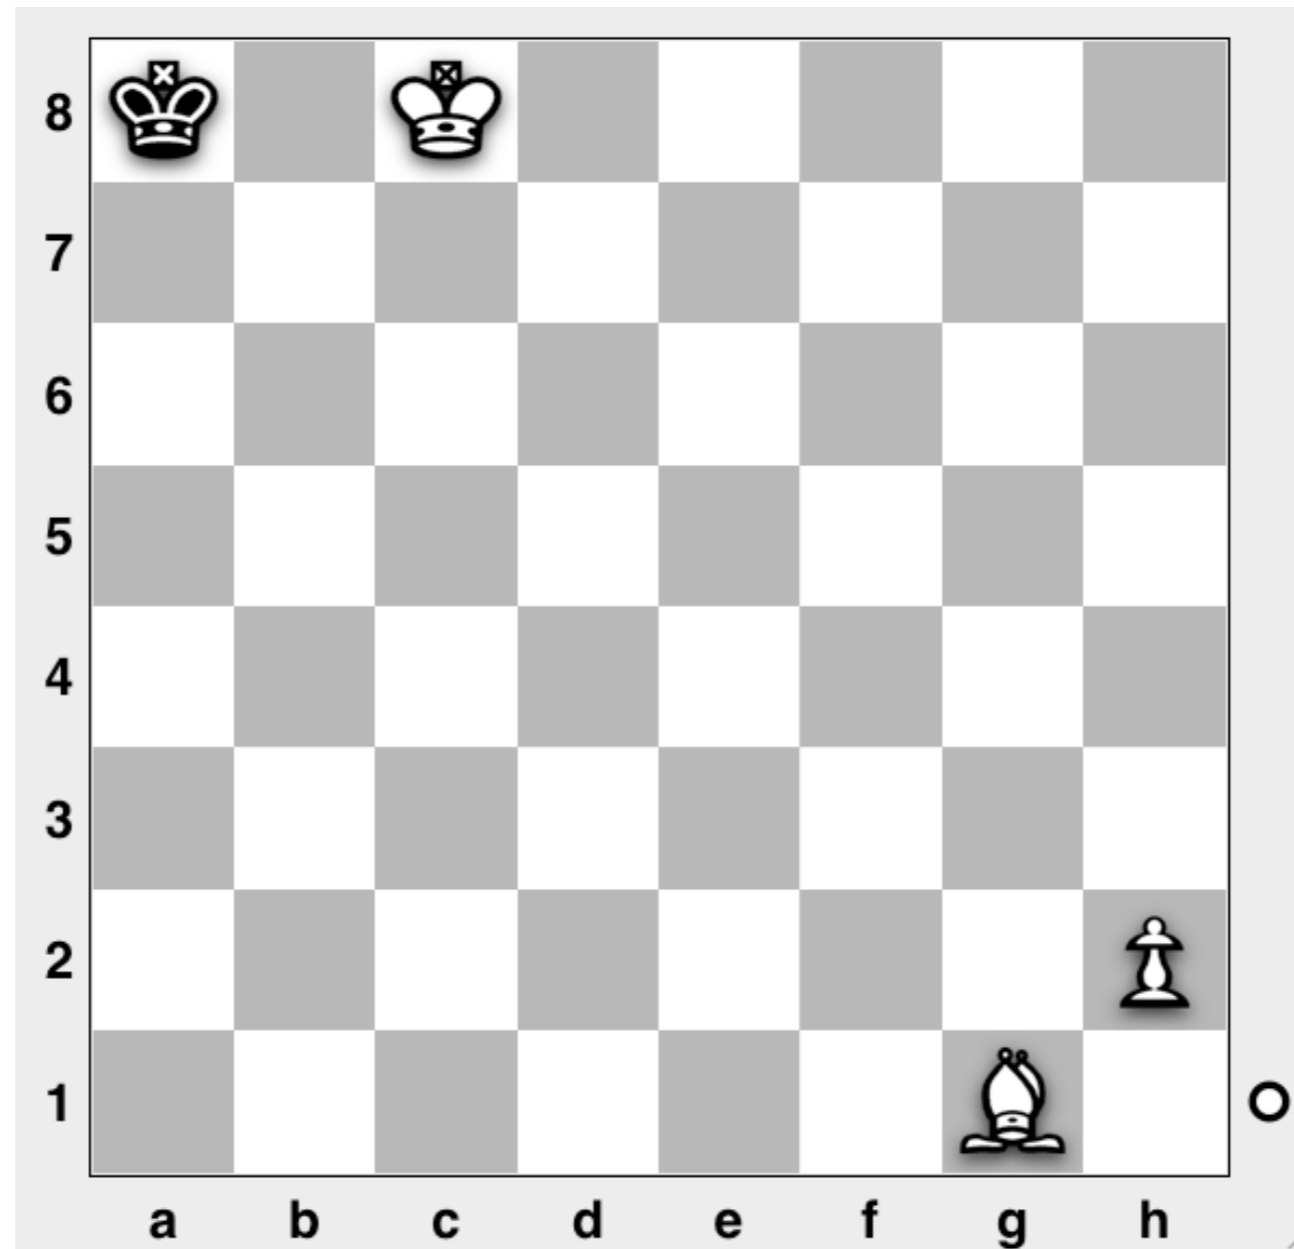
What is Logic?

- 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?

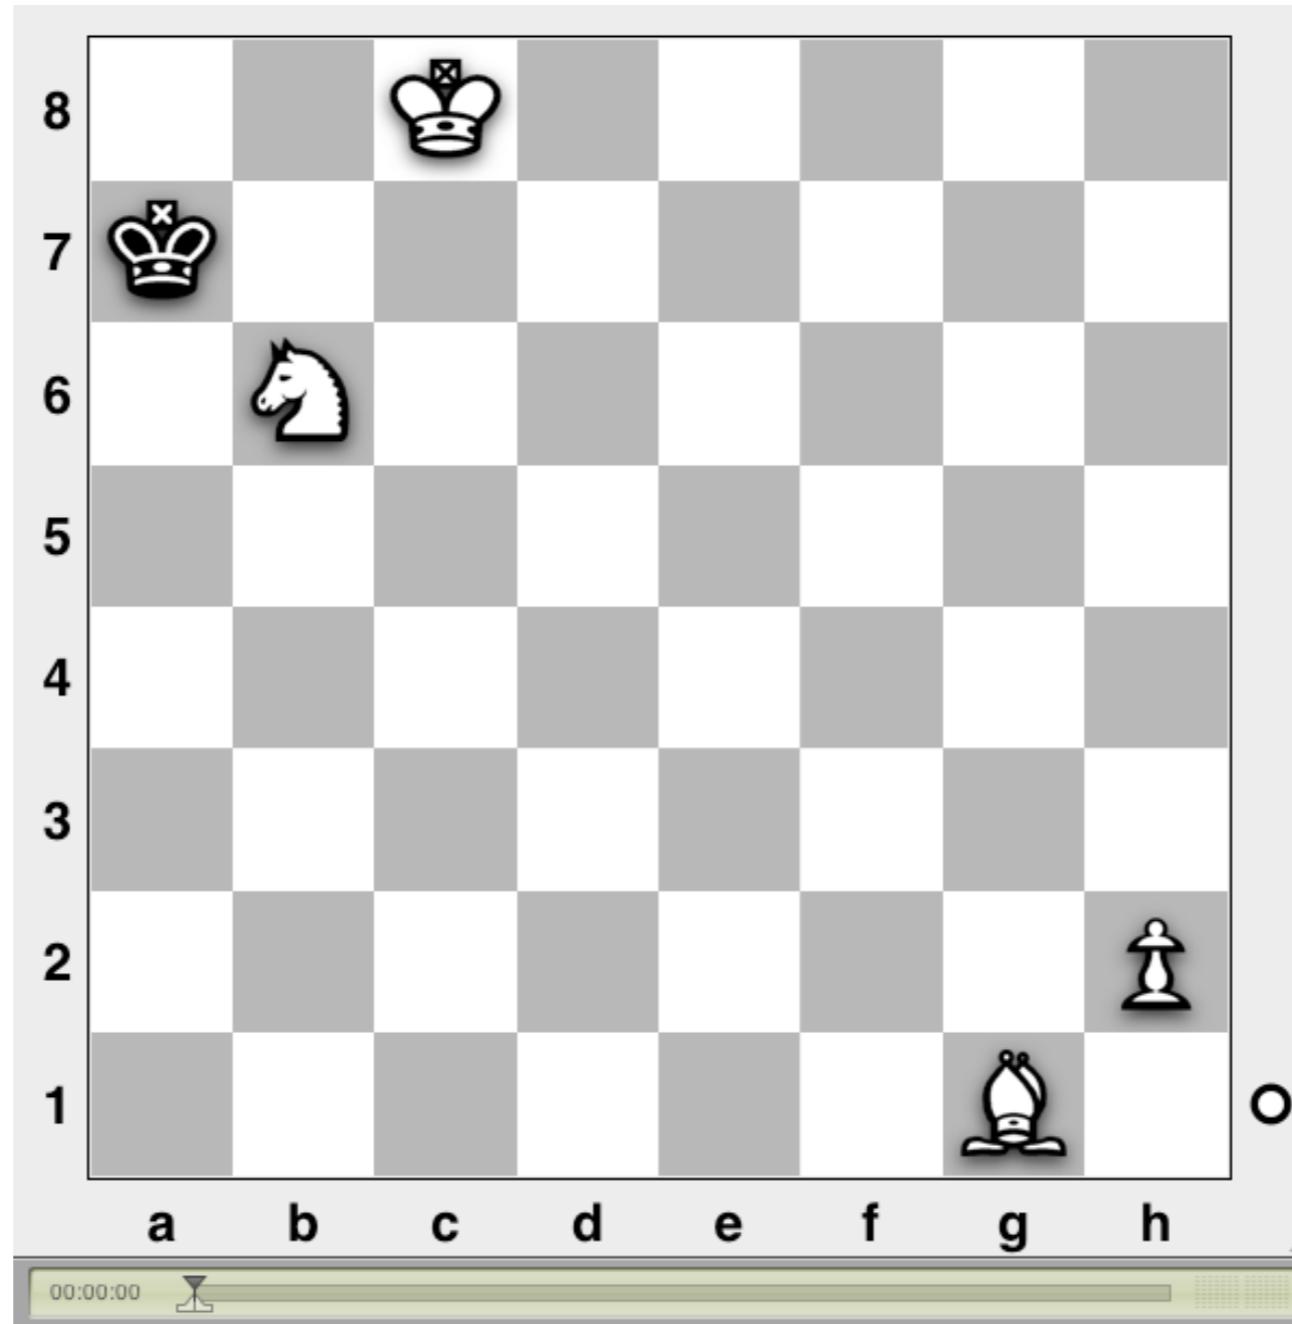
- The key to becoming rational. Or are you *already* rational? ...
- “The science of reasoning.” — so the not-unreasonable slogan goes.
- The only invincible subject there is.
- The basis for the formal sciences (from mathematics to game theory to decision theory to probability calculi to axiomatic physics) — 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)?...
- ...

It's White's turn. What move did Black just make?



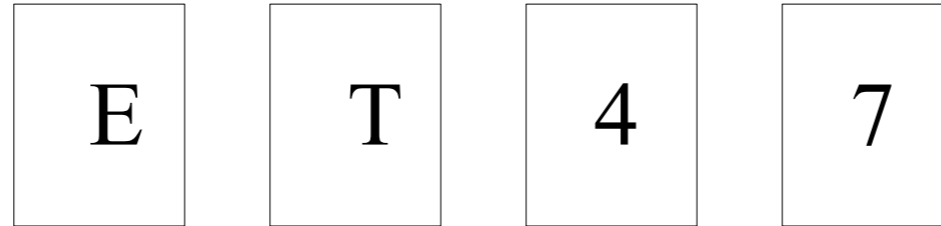
Aha! (Beyond Deep Blue?)

Aha! (Beyond Deep Blue?)



NOTE: Every card in this game has a capital Roman letter on one side, and a number from 1 to 9, inclusive.

Simple Selection Task



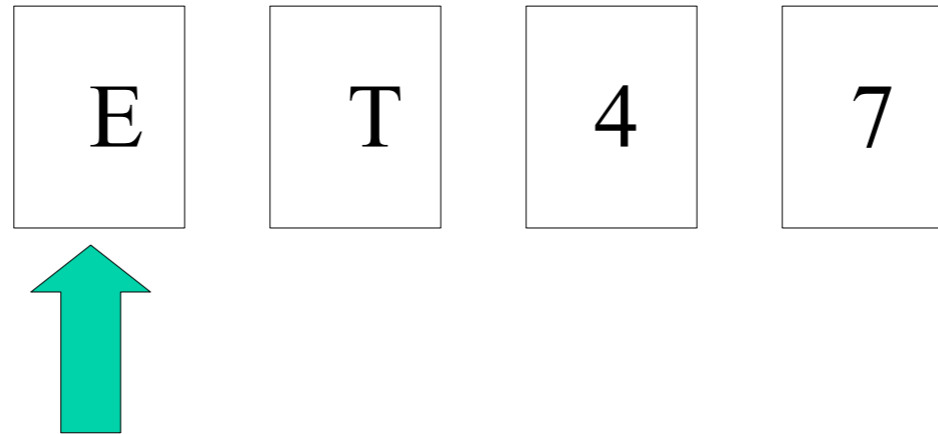
Suppose I claim that the following rule is true.

If a card has a vowel on one side, it has an even number on the other side.

Which card or cards, if any, should you turn over in order to try to efficiently decide whether the rule is true or false?

NOTE: Every card in this game has a capital Roman letter on one side, and a number from 1 to 9, inclusive.

Simple Selection Task



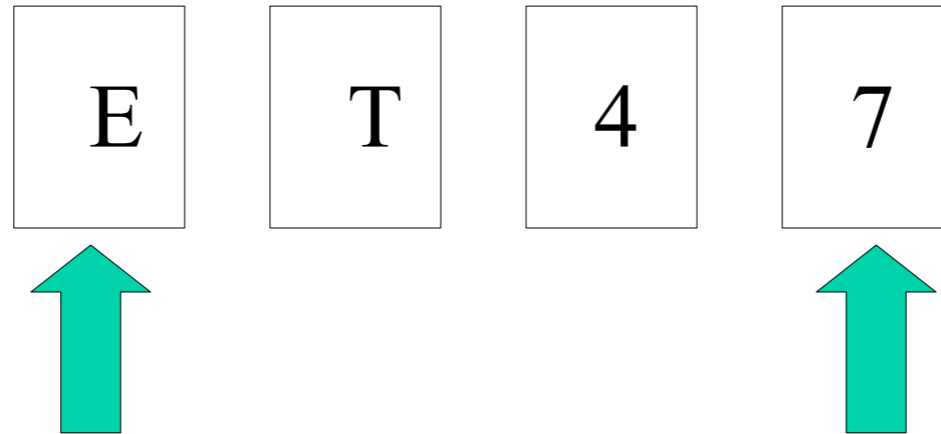
Suppose I claim that the following rule is true.

If a card has a vowel on one side, it has an even number on the other side.

Which card or cards, if any, should you turn over in order to try to efficiently decide whether the rule is true or false?

NOTE: Every card in this game has a capital Roman letter on one side, and a number from 1 to 9, inclusive.

Simple Selection Task

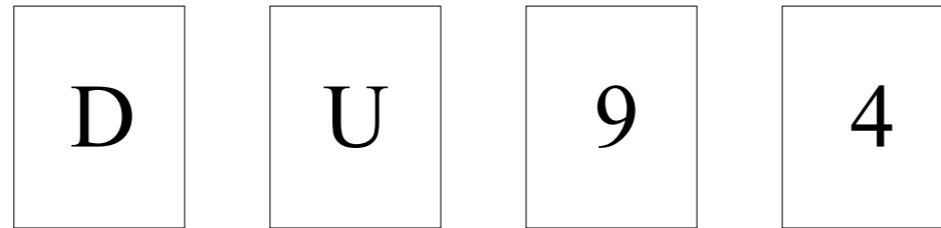


Suppose I claim that the following rule is true.

If a card has a vowel on one side, it has an even number on the other side.

Which card or cards, if any, should you turn over in order to try to efficiently decide whether the rule is true or false?

Another Simple Selection Task

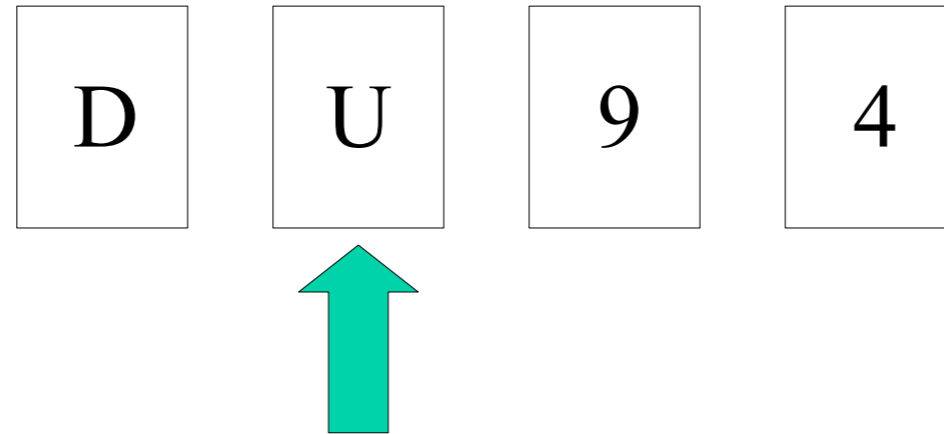


Suppose I claim that the following rule is true.

If a card has a vowel on one side, it has an even number on the other side.

Which card or cards, if any, should you turn over in order to try to efficiently decide whether the rule is true or false?

Another Simple Selection Task

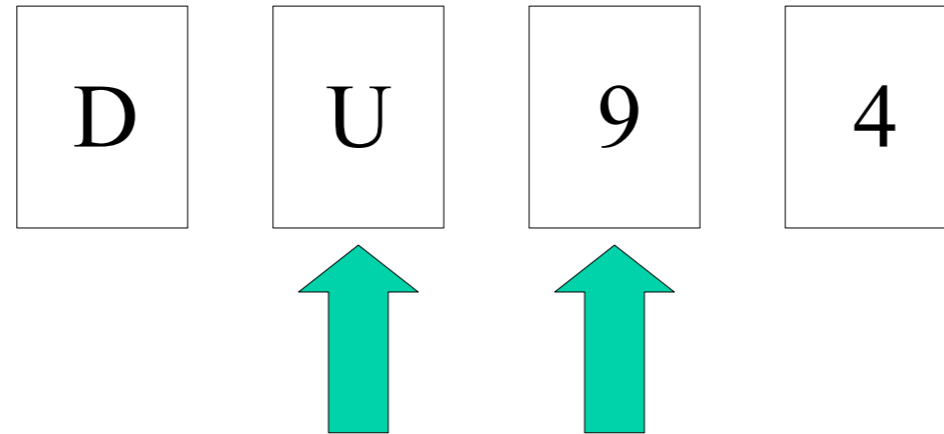


Suppose I claim that the following rule is true.

If a card has a vowel on one side, it has an even number on the other side.

Which card or cards, if any, should you turn over in order to try to efficiently decide whether the rule is true or false?

Another Simple Selection Task

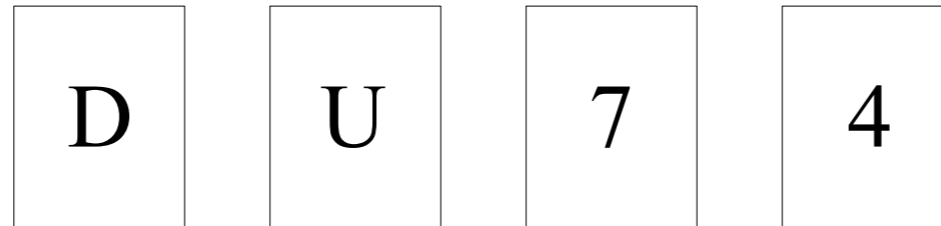


Suppose I claim that the following rule is true.

If a card has a vowel on one side, it has an even number on the other side.

Which card or cards, if any, should you turn over in order to try to efficiently decide whether the rule is true or false?

Another Simple Selection Task

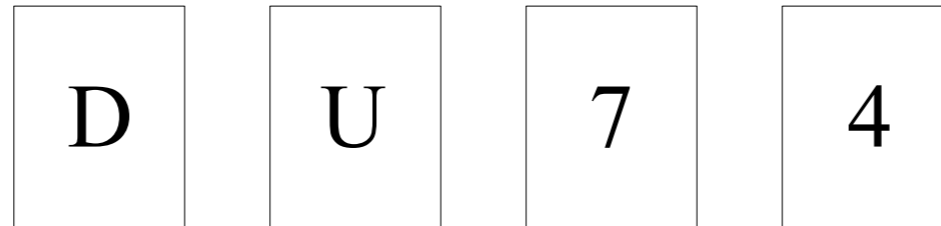


Suppose I claim that the following rule is true.

If a card has a letter on one side, it has a prime number on the other side.

Which card or cards, if any, should you turn over in order to try to efficiently decide whether the rule is true or false?

Another Simple Selection Task



Suppose I claim that the following rule is true.

If a card has a letter on one side, it has a prime number on the other side.

Which card or cards, if any, should you turn over in order to try to efficiently decide whether the rule is true or false?

“NYS I”

Given the statements

$$\neg a \vee \neg b$$

b

$$c \rightarrow a$$

which one of the following statements can you prove?

c

$\neg b$

$\neg c$

h

a

none of the above

“NYS I”

Given the statements

$\neg a \vee \neg b$

b

$c \rightarrow a$

which one of the following statements can you prove?

c

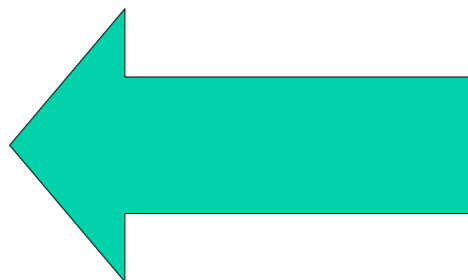
$\neg b$

$\neg c$

h

a

none of the above



Without Consciousness, AIs Will Be Sociopaths
ChatGPT can carry on a conversation, but the most important goal for artificial intelligence is making it understand what it means to have a mind

ESSAY

By Michael S.A. Graziano [View](#)
January 13, 2023 9:24 am EST



ChatGPT, the latest technological sensation, is an artificial intelligence chatbot with an amazing ability to carry on a conversation. It relies on a massive network of artificial neurons that loosely mimics the human brain, and it has been trained by analyzing the information resources of the internet. ChatGPT has processed more text than any human is likely to have read in a lifetime, allowing it to respond to questions fluently and even to imitate specific individuals, answering queries the way it thinks they would. My teenage son recently used ChatGPT to argue about politics with an imitation Karl Marx.

As a neuroscientist specializing in the brain mechanisms of consciousness, I find talking to chatbots an unsettling experience. Are they conscious? Probably not. But given the rate of technological improvement, will they be in the next couple of years? And how would we even know?

Figuring out whether a machine has or understands humanlike consciousness is more than just a science-fiction hypothetical. Artificial intelligence is growing so powerful, so quickly that it could soon pose a danger to human beings. We're building machines that are smarter than us and giving them control over our world. How can we build AI so that it's aligned with human needs, not in conflict with us?

As counterintuitive as it may sound, creating a benign AI may require making it more conscious, not less. One of the most common misunderstandings about AI is the notion that if it's intelligent then it must be conscious, and if it is conscious then it will be autonomous, capable of taking over the world. But as we learn more about consciousness, those ideas do not appear to be correct. An autonomous system that makes complex decisions doesn't require consciousness.

What's most important about consciousness is that, for human beings, it's not just about the self. We see it in ourselves, but we also perceive it or project it into the world around us. Consciousness is part of the tool kit that evolution gave us to make us an empathetic, prosocial species. Without it, we would necessarily be sociopaths, because we'd lack the tools for prosocial behavior. And without a concept of what consciousness is or an understanding that other beings have it, machines are sociopaths.

The only diagnostic tool for machine consciousness that we have right now is the Turing test, a thought experiment named for the British computer scientist Alan Turing. In its most common version, the test asks that if a person holds a conversation with a machine and mistakes its responses for those of a real human being, then the machine must be considered effectively conscious.

The Turing test is an admission that the consciousness of another being is something we can only judge from the outside, based on the way he, she or it communicates. But the

“NYS 2”

Which one of the following statements is provable from the following statement: “If you are not part of the solution, then you are part of the problem.”

If you are part of the solution, then you are not part of the problem.

If you are not part of the problem, then you are part of the solution.

If you are part of the problem, then you are not part of the solution.

If you are not part of the problem, then you are not part of the solution.

Without Consciousness, AIs Will Be Sociopaths
ChatGPT can carry on a conversation, but the most important goal for artificial intelligence is making it understand what it means to have a mind

ESSAY

By Michael S.A. Graziano [+ icon](#)
January 13, 2023 9:24 am EST



ChatGPT, the latest technological sensation, is an artificial intelligence chatbot with an amazing ability to carry on a conversation. It relies on a massive network of artificial neurons that loosely mimics the human brain, and it has been trained by analyzing the information resources of the internet. ChatGPT has processed more text than any human is likely to have read in a lifetime, allowing it to respond to questions fluently and even to imitate specific individuals, answering queries the way it thinks they would. My teenage son recently used ChatGPT to argue about politics with an imitation Karl Marx.

As a neuroscientist specializing in the brain mechanisms of consciousness, I find talking to chatbots an unsettling experience. Are they conscious? Probably not. But given the rate of technological improvement, will they be in the next couple of years? And how would we even know?

Figuring out whether a machine has or understands humanlike consciousness is more than just a science-fiction hypothetical. Artificial intelligence is growing so powerful, so quickly that it could soon pose a danger to human beings. We're building machines that are smarter than us and giving them control over our world. How can we build AI so that it's aligned with human needs, not in conflict with us?

As counterintuitive as it may sound, creating a benign AI may require making it more conscious, not less. One of the most common misunderstandings about AI is the notion that if it's intelligent then it must be conscious, and if it is conscious then it will be autonomous, capable of taking over the world. But as we learn more about consciousness, those ideas do not appear to be correct. An autonomous system that makes complex decisions doesn't require consciousness.

What's most important about consciousness is that, for human beings, it's not just about the self. We see it in ourselves, but we also perceive it or project it into the world around us. Consciousness is part of the tool kit that evolution gave us to make us an empathetic, prosocial species. Without it, we would necessarily be sociopaths, because we'd lack the tools for prosocial behavior. And without a concept of what consciousness is or an understanding that other beings have it, machines are sociopaths.

The only diagnostic tool for machine consciousness that we have right now is the Turing test, a thought experiment named for the British computer scientist Alan Turing. In its most common version, the test asks that if a person holds a conversation with a machine and mistakes its responses for those of a real human being, then the machine must be considered effectively conscious.

The Turing test is an admission that the consciousness of another being is something we can only judge from the outside, based on the way he, she or it communicates. But the

“NYS 2”

Which one of the following statements is provable from the following statement: “If you are not part of the solution, then you are part of the problem.”

If you are part of the solution, then you are not part of the problem.

 If you are not part of the problem, then you are part of the solution.

If you are part of the problem, then you are not part of the solution.

If you are not part of the problem, then you are not part of the solution.



“NYS 2”

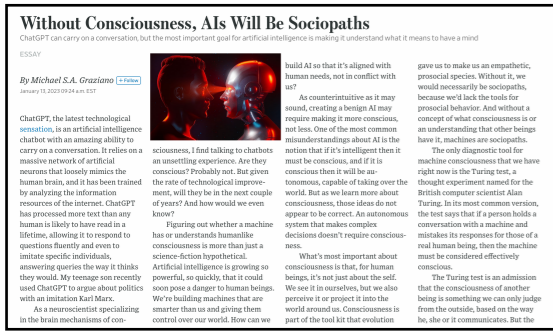
Which one of the following statements is provable from the following statement: “If you are not part of the solution, then you are part of the problem.”

If you are part of the solution, then you are not part of the problem.

 If you are not part of the problem, then you are part of the solution.

 If you are part of the problem, then you are not part of the solution.

If you are not part of the problem, then you are not part of the solution.



“NYS 2”

Which one of the following statements is provable from the following statement: “If you are not part of the solution, then you are part of the problem.”

If you are part of the solution, then you are not part of the problem.

 If you are not part of the problem, then you are part of the solution.

 If you are part of the problem, then you are not part of the solution.

If you are not part of the problem, then you are not part of the solution.

“NYS 3”

Given the statements

$$\neg\neg c$$

$$c \rightarrow a$$

$$\neg a \vee b$$

$$b \rightarrow d$$

$$\neg(d \vee e)$$

which of the following statements are provable?

$$\neg c$$

e

h

$$\neg a$$

all of the above

“NYS 3”

Given the statements

$$\neg\neg c$$

$$c \rightarrow a$$

$$\neg a \vee b$$

$$b \rightarrow d$$

$$\neg(d \vee e)$$

which of the following statements are provable?

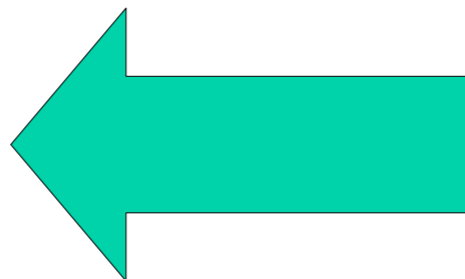
$$\neg c$$

e

h

$$\neg a$$

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/Arguments vs. Formal Proofs/Arguments

Informal Proofs/Arguments vs. Formal Proofs/Arguments

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

Informal Proofs/Arguments vs. Formal Proofs/Arguments

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

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

Informal Proofs/Arguments vs. Formal Proofs/Arguments

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/Arguments vs. Formal Proofs/Arguments

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

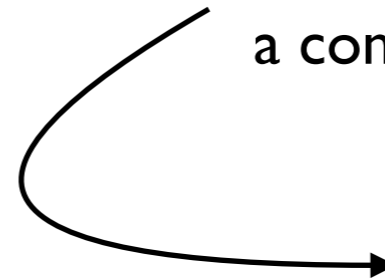
Informal Proofs/Arguments vs. Formal Proofs/Arguments

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



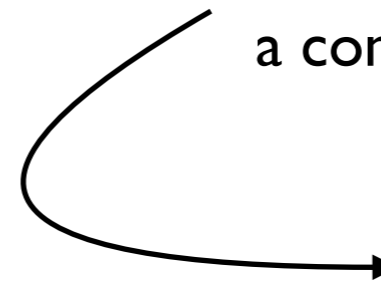
Informal Proofs/Arguments vs. Formal Proofs/Arguments

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/Arguments vs. Formal Proofs/Arguments

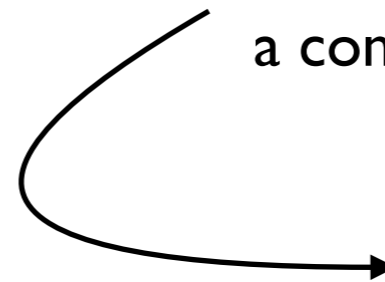
has ambiguous natural language
(e.g. English or Chinese)

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/Arguments vs. Formal Proofs/Arguments

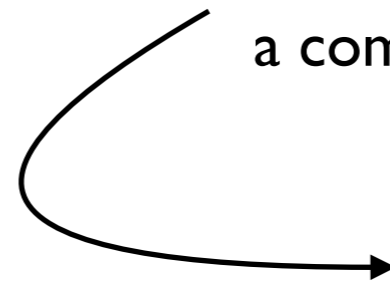
has ambiguous natural language
(e.g. English or Chinese)

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/Arguments vs. Formal Proofs/Arguments

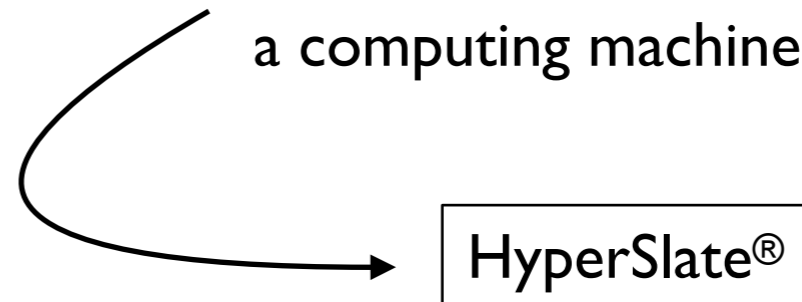
has ambiguous natural language
(e.g. English or Chinese)

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



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

FOR NOW

STARTING \geq 1/25/24

Informal Proofs/Arguments

vs.

Formal Proofs/Arguments

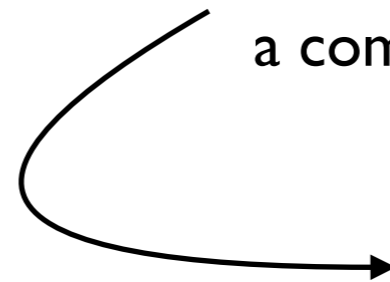
has ambiguous natural language
(e.g. English or Chinese)

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

cannot be executed & checked
by a computing machine

can 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



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 $K \Rightarrow A$ is false, and that $\neg K \Rightarrow A$ is false. Take first the first case; accordingly, suppose that $K \Rightarrow A$ 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 $\neg K \Rightarrow A$ being false. Here, in a direct parallel, we know $\neg K$ and, once again, $\neg A$. In both of our two cases, which are exhaustive, there is no ace in the hand. The proposition is established. **QED**

Har du lyktes med alle?

