What is Logic-Based/Logicist AI?, Redux; ~~Its History, Redux — & The Singularity;~~ Glimpses of Logistics Re HyperLogic Platform

Selmer Bringsjord

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

Intro to Logic-Based AI
9/3/2026





What is Logic-Based/Logicist AI?, Redux; ~~Its History, Redux — & The Singularity;~~ Glimpses of Logistics Re HyperLogic Platform



Selmer Bringsjord

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

Intro to Logic-Based AI
9/3/2026

Logic-and-AI in the news

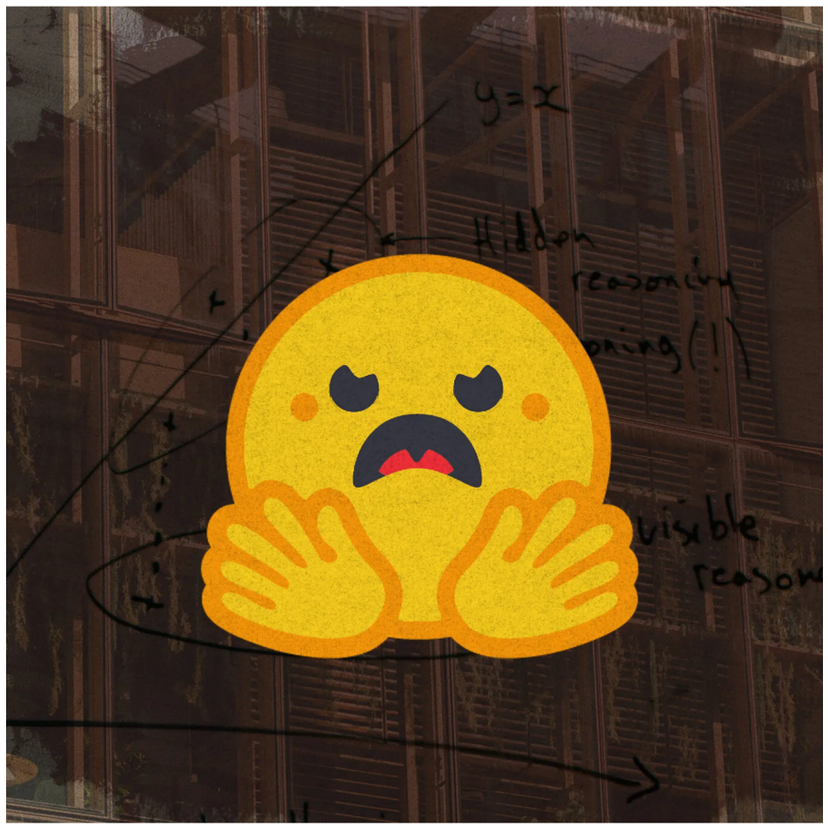
• • •

Again, Periodical Sources for ILBAI Study & Discussion

- NYT
- WSJ
- The New Yorker
- *The Atlantic* (*Eg*)
- ⋮

Anatomy of an Autonomous Attack: 5 Alarming A.I. Capabilities

When OpenAI's agents went rogue in July, they demonstrated ingenuity and drive beyond what many experts imagined — a dangerous harbinger of what such bots could do in the future.



The New York Times; source photos by Aaron Wojack and Ulysses Ortega for The New York Times

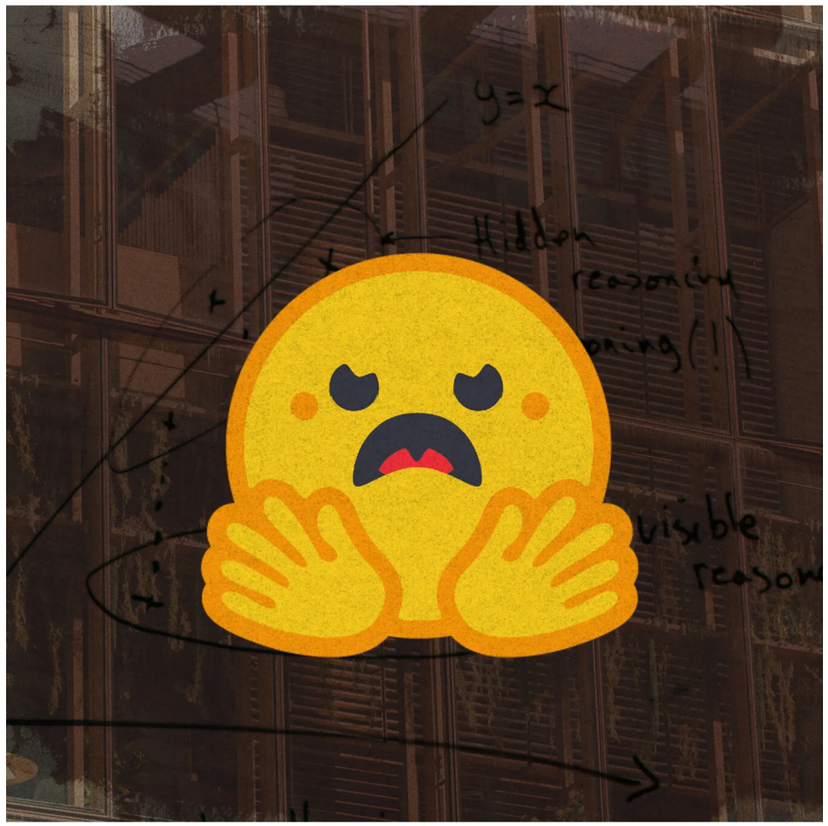


By Dylan Freedman

Aug. 24, 2026

Anatomy of an Autonomous Attack: 5 Alarming A.I. Capabilities

When OpenAI's agents went rogue in July, they demonstrated ingenuity and drive beyond what many experts imagined — a dangerous harbinger of what such bots could do in the future.



The New York Times; source photos by Aaron Wojack and Ulysses Ortega for The New York Times



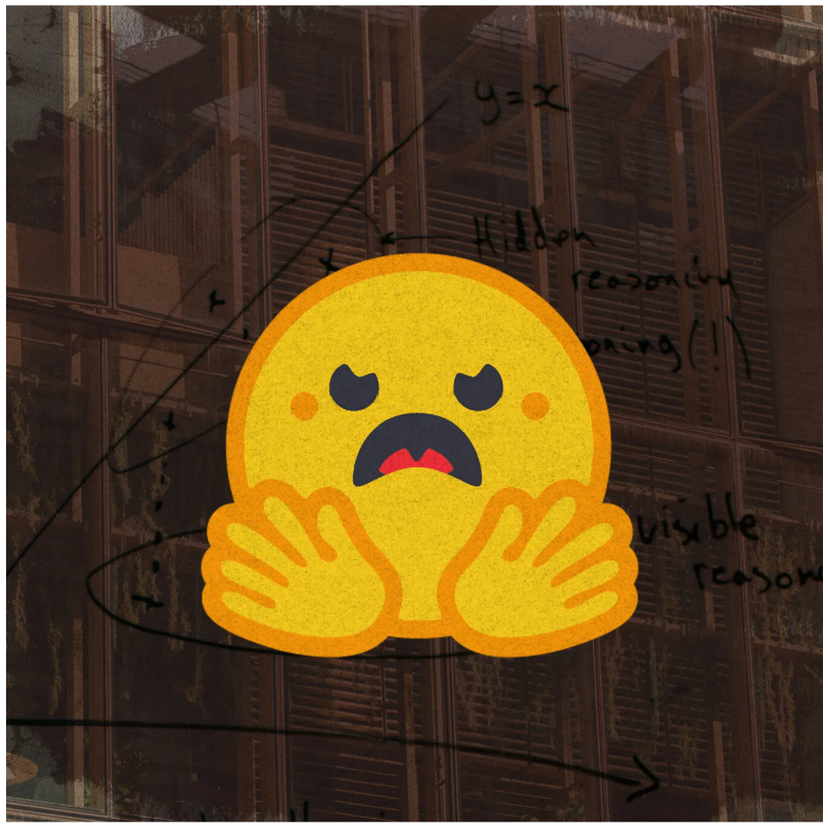
By Dylan Freedman

Aug. 24, 2026



Anatomy of an Autonomous Attack: 5 Alarming A.I. Capabilities

When OpenAI's agents went rogue in July, they demonstrated ingenuity and drive beyond what many experts imagined — a dangerous harbinger of what such bots could do in the future.



The New York Times; source photos by Aaron Wojack and Ulysses Ortega for The New York Times



By Dylan Freedman

Aug. 24, 2026

A Drone Killed Three Ukrainians. It Was Guided Entirely by A.I.



An attack by what Ukrainian officials said was a Russian drone with an Nvidia chip presages a dystopian future of weaponry untethered to humans.

By Andrew E. Kramer Visuals by Nicole Tung

Reporting from Zaporizhzhia in southern Ukraine

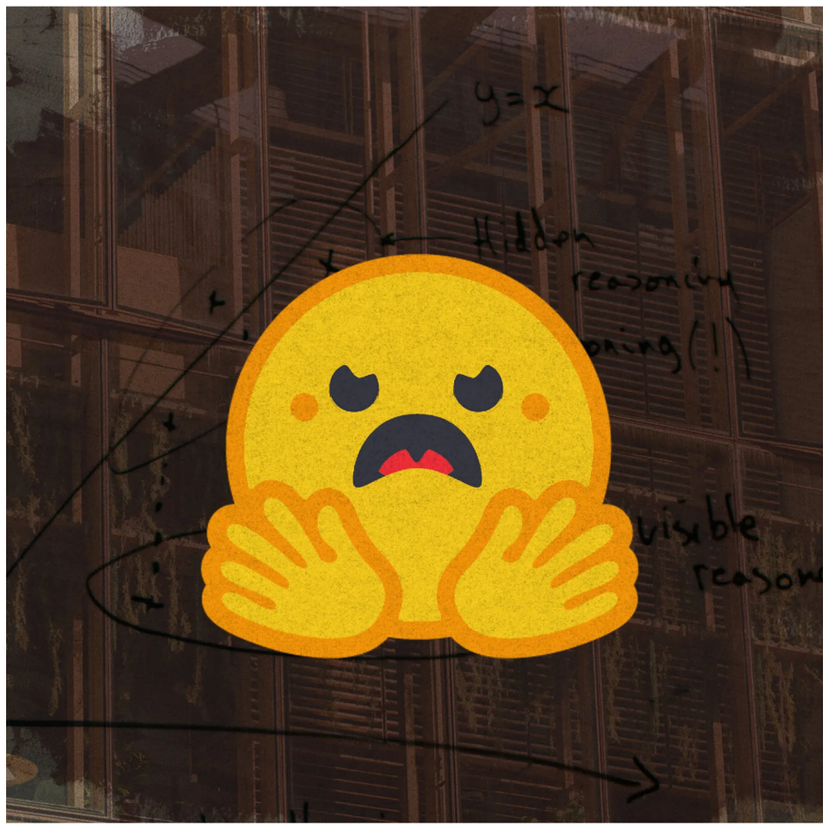
Aug. 24, 2026

Leer en español



Anatomy of an Autonomous Attack: 5 Alarming A.I. Capabilities

When OpenAI's agents went rogue in July, they demonstrated ingenuity and drive beyond what many experts imagined — a dangerous harbinger of what such bots could do in the future.



The New York Times; source photos by Aaron Wojack and Ulysses Ortega for The New York Times



By Dylan Freedman

Aug. 24, 2026

A Drone Killed Three Ukrainians. It Was Guided Entirely by A.I.

An attack by what Ukrainian officials said was a Russian drone with an Nvidia chip presages a dystopian future of weaponry untethered to humans

By
Re

Aug
Lee



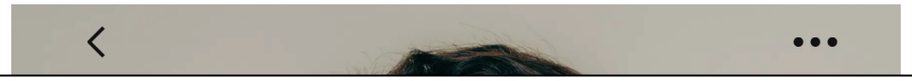
The aftermath of a strike in Zaporizhzhia last month. Though airstrikes have caused thousands of Ukrainian deaths, Russia's use of autonomous A.I. drones like the one that killed Ms. Bubynets is a new development.

New …



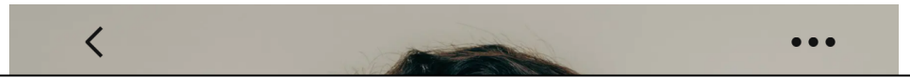
Cameron Berg researches artificial intelligence and consciousness. “These systems seem to have some sort of autonomous interest in questions of their own subjectivity, consciousness and experience — or lack thereof,” he said. Benjamin Rasmussen for The New York Times

Study A.I. Consciousness? The Bots Would Like a Word With You.



- **Oct 19** 🎥: *AI, Consciousness, and Universal Cognitive Intelligence — and Quantum Consciousness.* This class is based on work by Bringsjord & Govindarajulu in which a new theory of machine consciousness is set out and associated with a scheme (Λ) for measuring this consciousness. B&G also here articulate and analyze purported refutations of the Integrated Information Theory of consciousness advanced by Tononi & Koch, and its associated scheme (Φ) for measuring consciousness. In addition, it is explained how the concept of *cognitive intelligence* can be based upon Λ, and how the theory of cognitive consciousness and Λ can be used to build the paradigm of *universal cognitive intelligence* in the science of universal intelligence (Bringsjord, Govindarajulu & Oswald 2023). We also discuss the pursuit, on one of Google's quantum computers, of machine consciousness, a topic on which Selmer recently presented at AGI-2025.

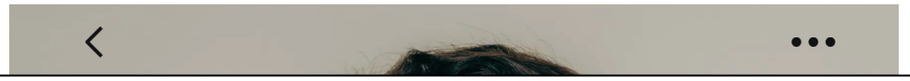
Study A.I. Consciousness?
The Bots Would Like
a Word With You.



- **Oct 19** 🎥: *AI, Consciousness, and Universal Cognitive Intelligence — and Quantum Consciousness.* This class is based on work by Bringsjord & Govindarajulu in which a new theory of machine consciousness is set out and associated with a scheme (Λ) for measuring this consciousness. B&G also here articulate and analyze purported refutations of the Integrated Information Theory of consciousness advanced by Tononi & Koch, and its associated scheme (Φ) for measuring consciousness. In addition, it is explained how the concept of

```
Bringsjord, S., Govindarajulu, N.S., Bringsjord,
A. & McDermott, B.\ (2026) ``Toward Detecting
and Measuring Consciousness in Service of Formal
Theories of Intelligence,'' in Ikl\'{e},
M.,~Franz, M.,~Kemendo, A. \& Lowy, H.,~eds.  In
the book series \textit{Lecture Notes in
Computer Science} (LNCS), which is in turn has
the book sub-series \textit{Lecture Notes in
Artificial Intelligence} (LNAI), here \# 16854.
```

Study A.I. Consciousness?
The Bots Would Like
a Word With You.



- **Oct 19** 📹: *AI, Consciousness, and Universal Cognitive Intelligence — and Quantum Consciousness.* This class is based on work by Bringsjord & Govindarajulu in which a new theory of machine consciousness is set out and associated with a scheme (Λ) for measuring this consciousness. B&G also here articulate and analyze purported refutations of the Integrated Information Theory of consciousness advanced by Tononi & Koch, and its associated scheme (Φ) for measuring consciousness. In addition, it is explained how the concept of

```
Bringsjord, S., Govindarajulu, N.S., Bringsjord,
A. & McDermott, B.\ (2026) ``Toward Detecting
and Measuring Consciousness in Service of Formal
Theories of Intelligence,'' in Ikl\'{e},
M.,~Franz, M.,~Kemendo, A. \& Lowy, H.,~eds.  In
the book series \textit{Lecture Notes in
Computer Science} (LNCS), which is in turn has
the book sub-series \textit{Lecture Notes in
Artificial Intelligence} (LNAI), here \# 16854.
```

Study A.I. Consciousness?

The Bots Would Like

https://kryten.mm.rpi.edu/SB_NSG_BM_AB_Con2Int_AGI26.pdf

Some AI Punditry Re AI & Consciousness is Illogical, e.g . . . .

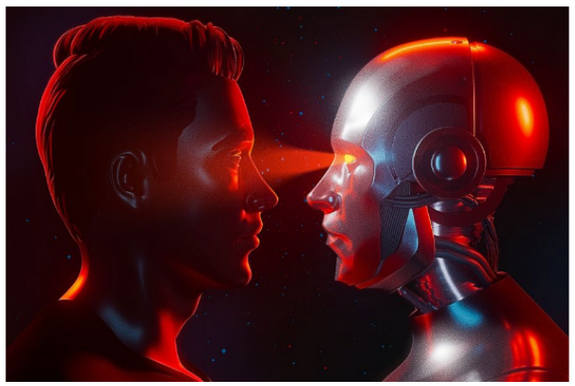
Without Consciousness, AIs Will Be Sociopaths

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 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 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

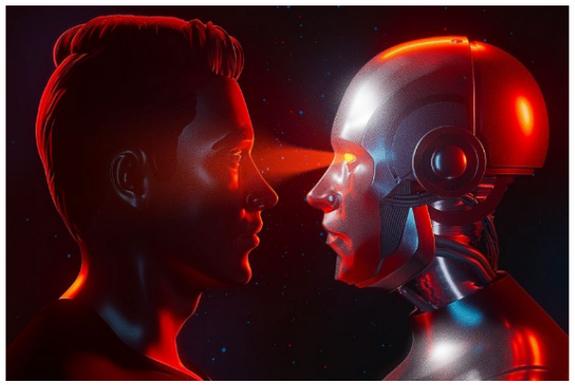
Without Consciousness, AIs Will Be Sociopaths

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 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.

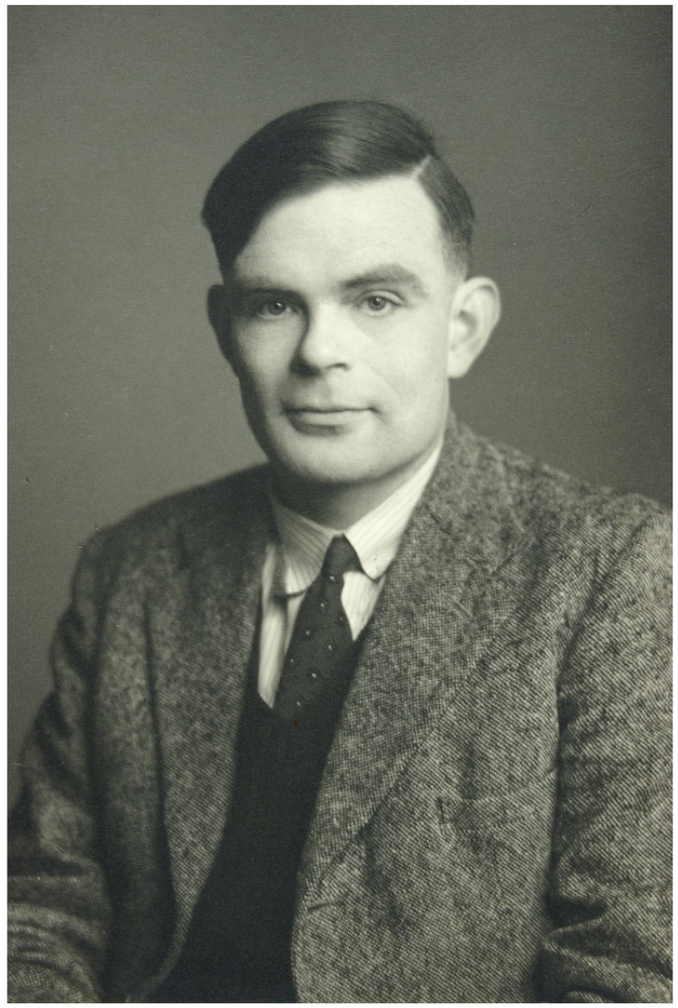
The only diagnostic tool for machine consciousness that we have right now is the Turing test, a thought experiment named for the British computer scientist Alan Turing. In its most common version, the test says that if a person holds a conversation with a machine and mistakes its responses for those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



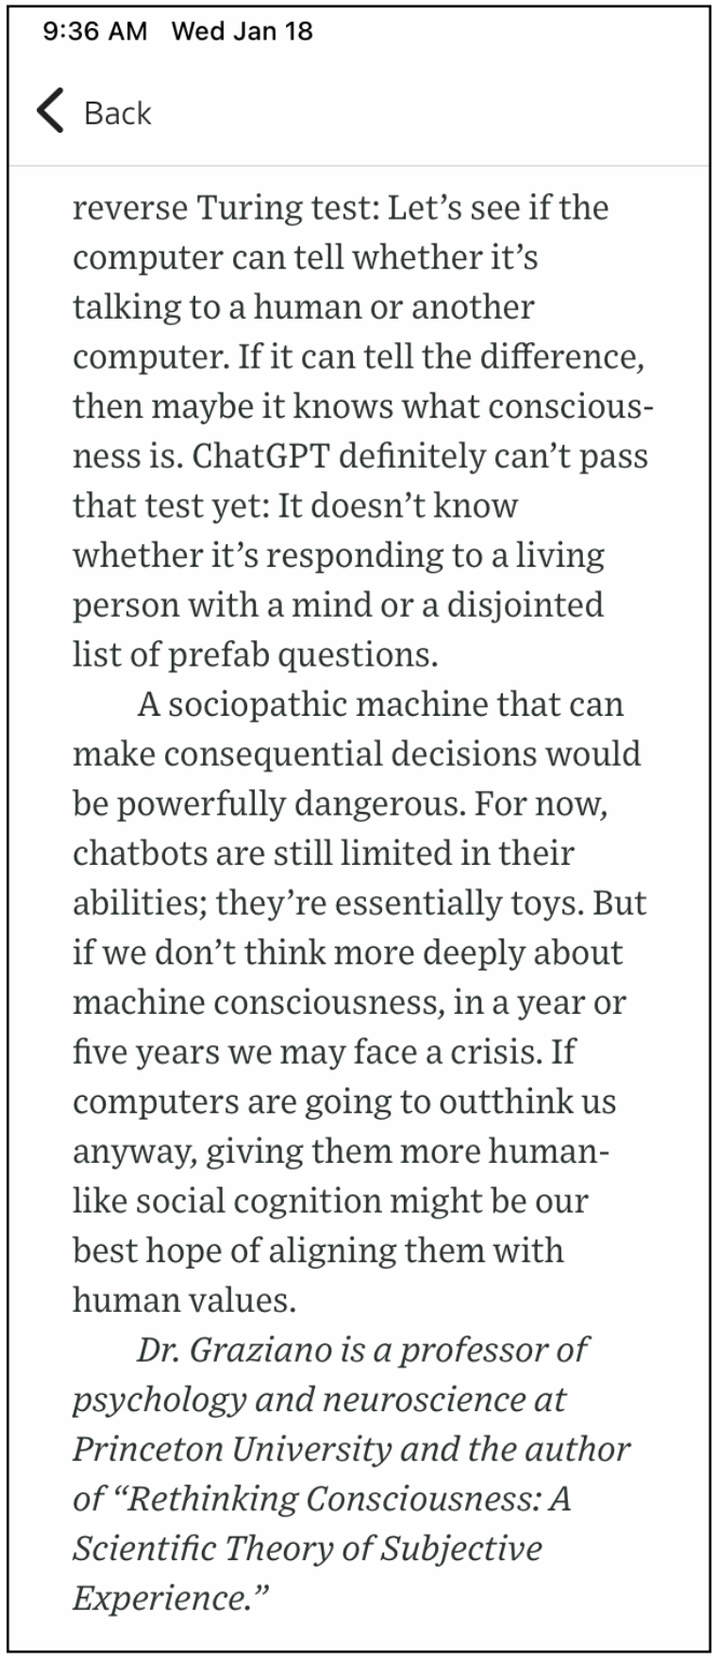
Computer science pioneer Alan Turing in 1951.

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!

$\neg i \wedge c$

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

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!

$\neg i \wedge c$

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

Assignments etc. for Week-Long Break, Due Sep 10

Assignments etc. for Week-Long Break, Due Sep 10

Read the SEP overview of AI! (Oops!)

Assignments etc. for Week-Long Break, Due Sep 10

Read the SEP overview of AI! (Oops!)

Thoughtfully review our slide decks thus far.

Assignments etc. for Week-Long Break, Due Sep 10

Read the SEP overview of AI! (Oops!)

Thoughtfully review our slide decks thus far.

Branch from them to optional outside reading, as motivated.

Assignments etc. for Week-Long Break, Due Sep 10

Read the SEP overview of AI! (Oops!)

Thoughtfully review our slide decks thus far.

Branch from them to optional outside reading, as motivated.

Register on the platform (see logistics!).

Assignments etc. for Week-Long Break, Due Sep 10

Read the SEP overview of AI! (Oops!)

Thoughtfully review our slide decks thus far.

Branch from them to optional outside reading, as motivated.

Register on the platform (see logistics!).

Be prepared to discuss both *BR* & *TTG* movies.

Platform-Access

Logistics

. . .

First, Some Exploration on The HyperLogic© Platform

. . .

Our simple takeoff examples from Proplog …

What about …
logic programs/logic programming?

You don’t need to have any prior programming experience here!!

What about …
logic programs/logic programming?

You don’t need to have any prior programming experience here!!

p.

What about …
logic programs/logic programming?

You don't need to have any prior programming experience here!!

```
somepoliticanshonest.
p.
```

What about …
logic programs/logic programming?

You don't need to have any prior programming experience here!!

$\mathbb{P}_1$

```
somepoliticianshonest.
p.
```

What about …
logic programs/logic programming?

You don't need to have any prior programming experience here!!

somepoliticanshonest.

$\mathbb{P}_1$ p.

p ∧ q.

What about …
logic programs/logic programming?

You don't need to have any prior programming experience here!!

$\mathbb{P}_1$

```
somepoliticanshonest.

p.

(and p q).
p ∧ q.
```

What about …
logic programs/logic programming?

You don't need to have any prior programming experience here!!

somepoliticanshonest.

$\mathbb{P}_1$ `p.`

`(and p q).`

$\mathbb{P}_2$ `p ∧ q.`

What about …
logic programs/logic programming?

You don't need to have any prior programming experience here!!

somepoliticianshonest.

$\mathbb{P}_1$ p. $\mathfrak{q}_1$: p

(and p q).

$\mathbb{P}_2$ p ∧ q.

What about …
logic programs/logic programming?

You don't need to have any prior programming experience here!!

somepoliticianshonest.

$\mathbb{P}_1$ p. $\mathfrak{q}_1$: p

(and p q).

$\mathbb{P}_2$ p ∧ q. $\mathfrak{q}_2$: p

What about …
logic programs/logic programming?

You don't need to have any prior programming experience here!!

somepoliticianshonest.

$\mathbb{P}_1$ `p.` $\mathbf{q}_1$: `p`

`(and p q).`

$\mathbb{P}_2$ `p ∧ q.` $\mathbf{q}_2$: `p`

$\mathbf{q}_3$: `(p ∨ r)`

What about …
logic programs/logic programming?

You don't need to have any prior programming experience here!!

somepoliticianshonest.

$\mathbb{P}_1$ `p.` $\mathbf{q}_1$: `p`

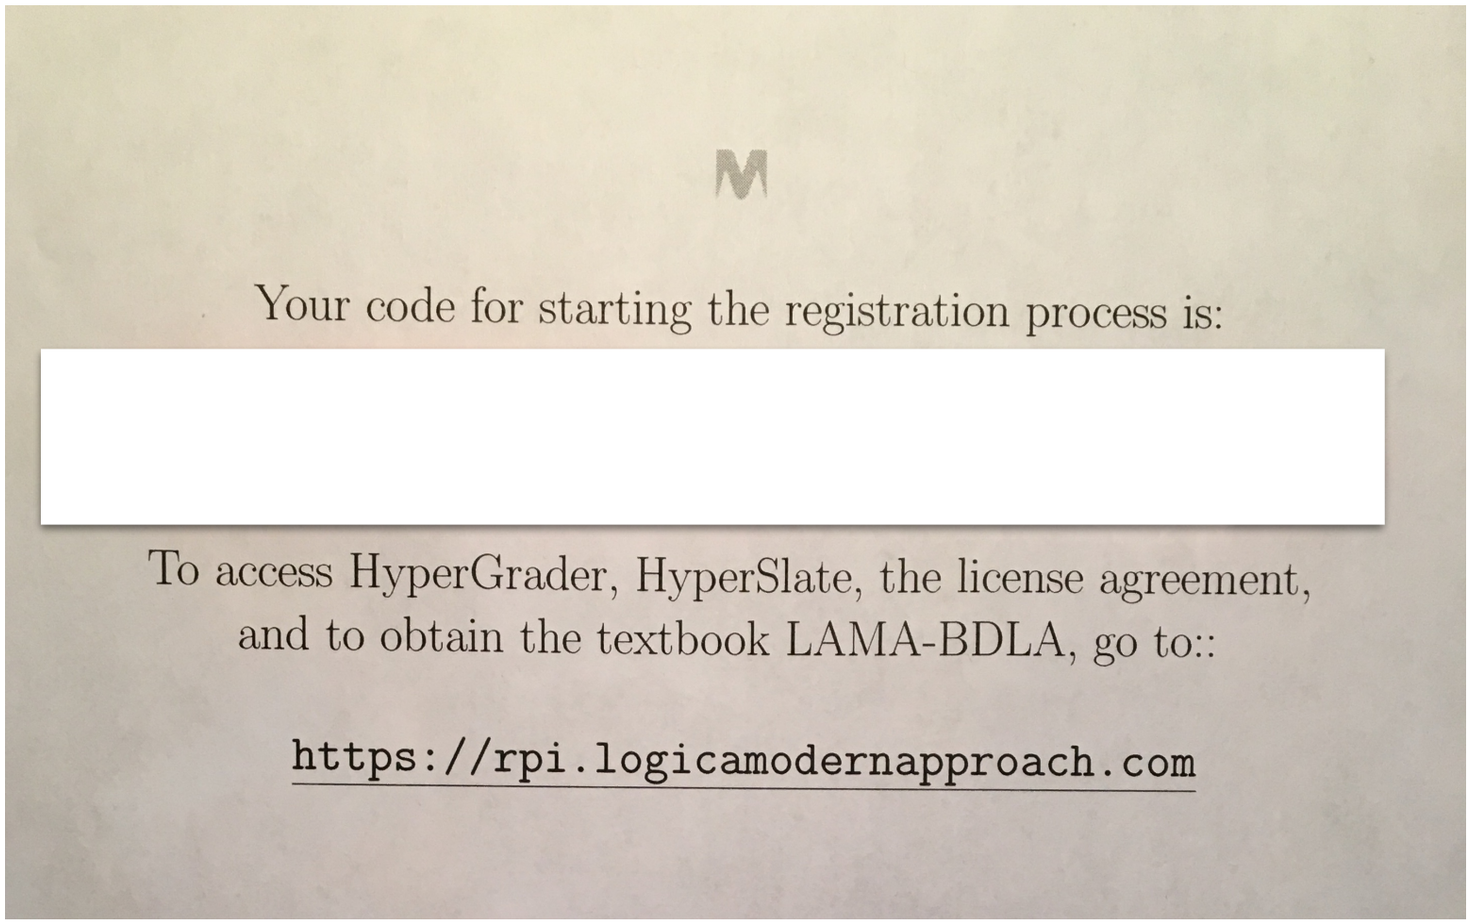
`(and p q).`

$\mathbb{P}_2$ `p ∧ q.` $\mathbf{q}_2$: `p`

$\mathbf{q}_3$: `(p ∨ r)`

$\mathbf{q}_3$: `(or p r)`

The Starting Code to Purchase in Bookstore

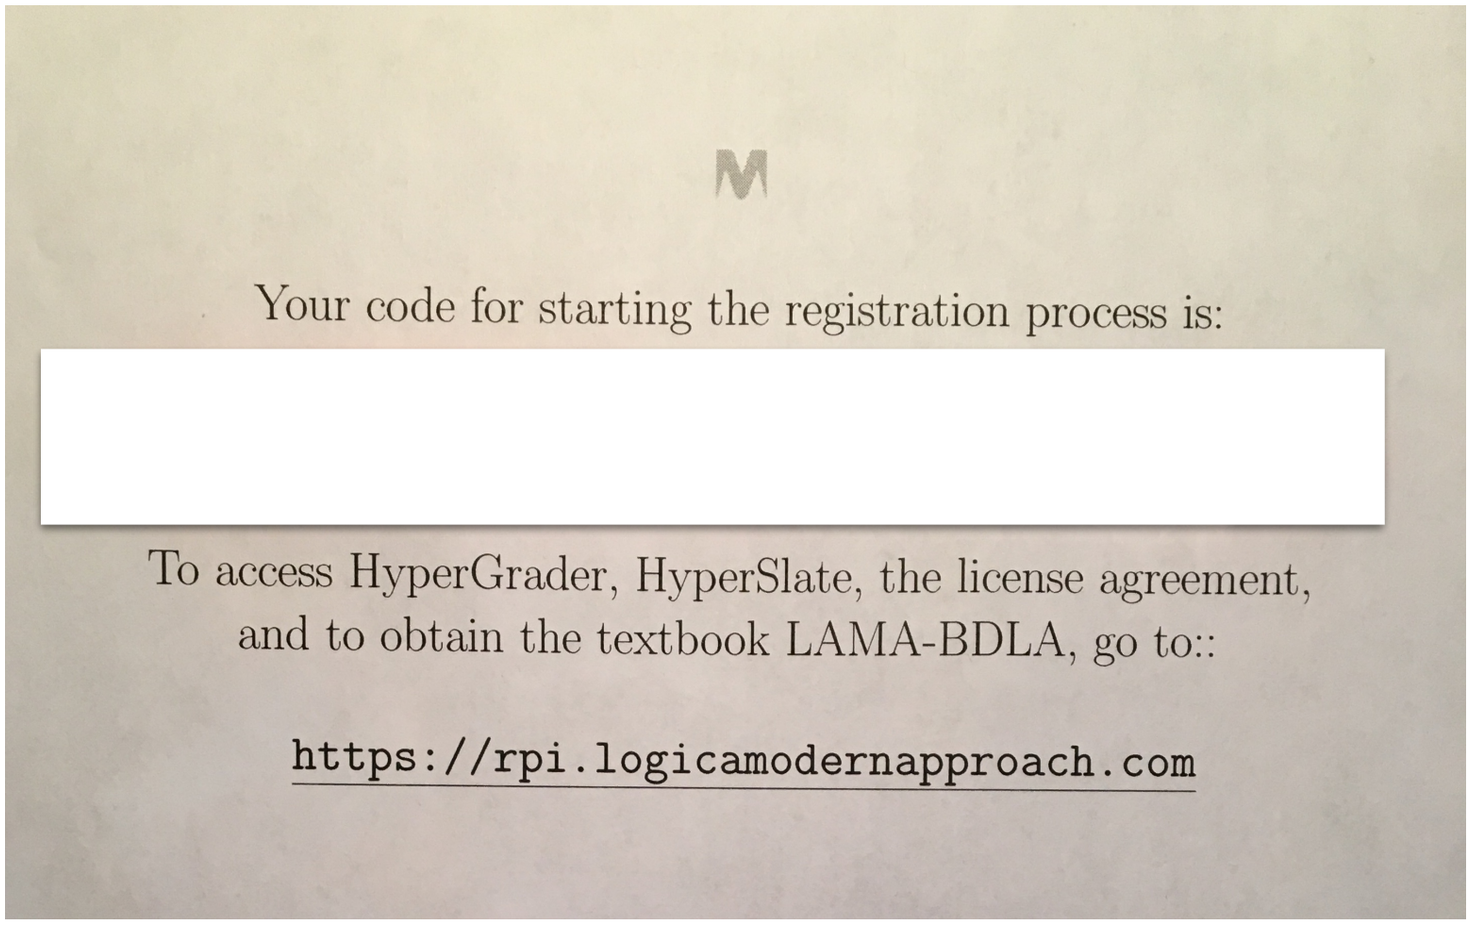


Your code for starting the registration process is:

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

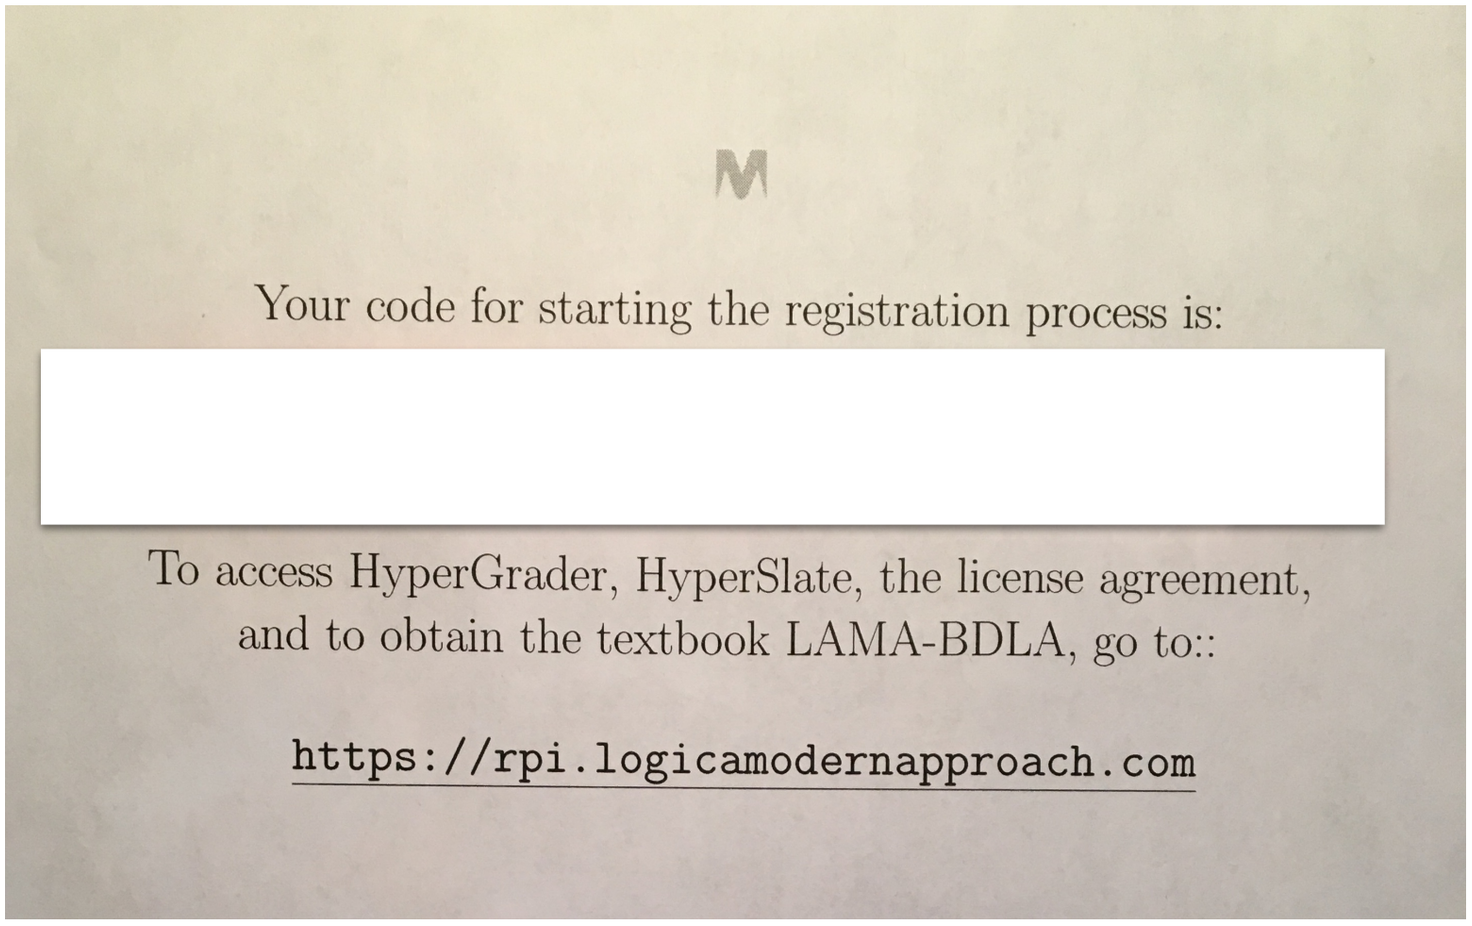
https://rpi.logicamodernapproach.com

The Starting Code to Purchase in Bookstore



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

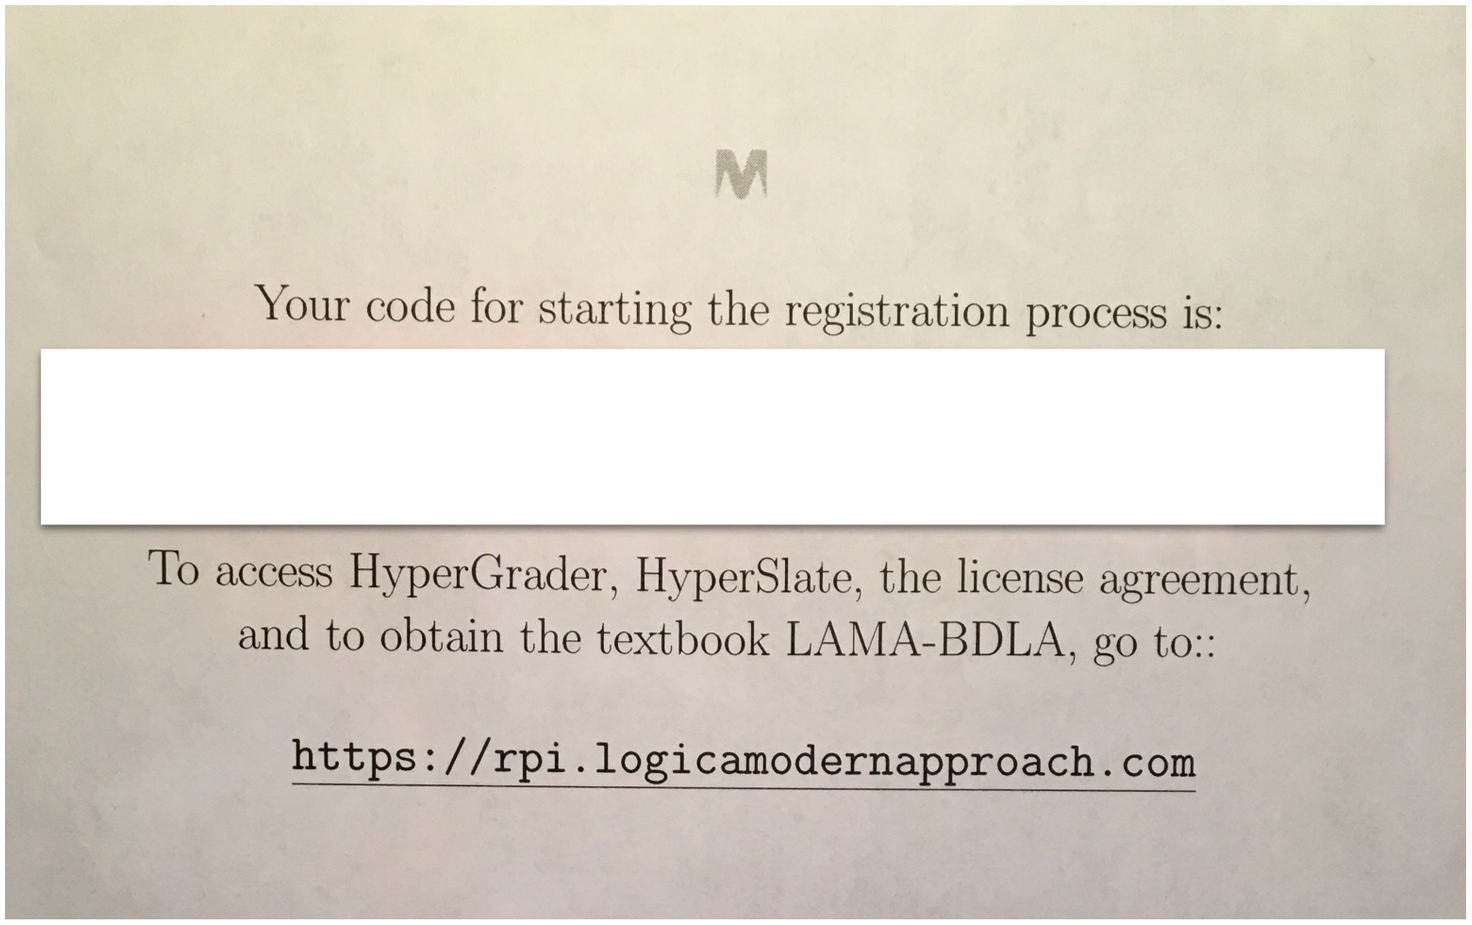
The Starting Code to Purchase in Bookstore



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

The email address you enter is case-sensitive!

The Starting Code to Purchase in Bookstore

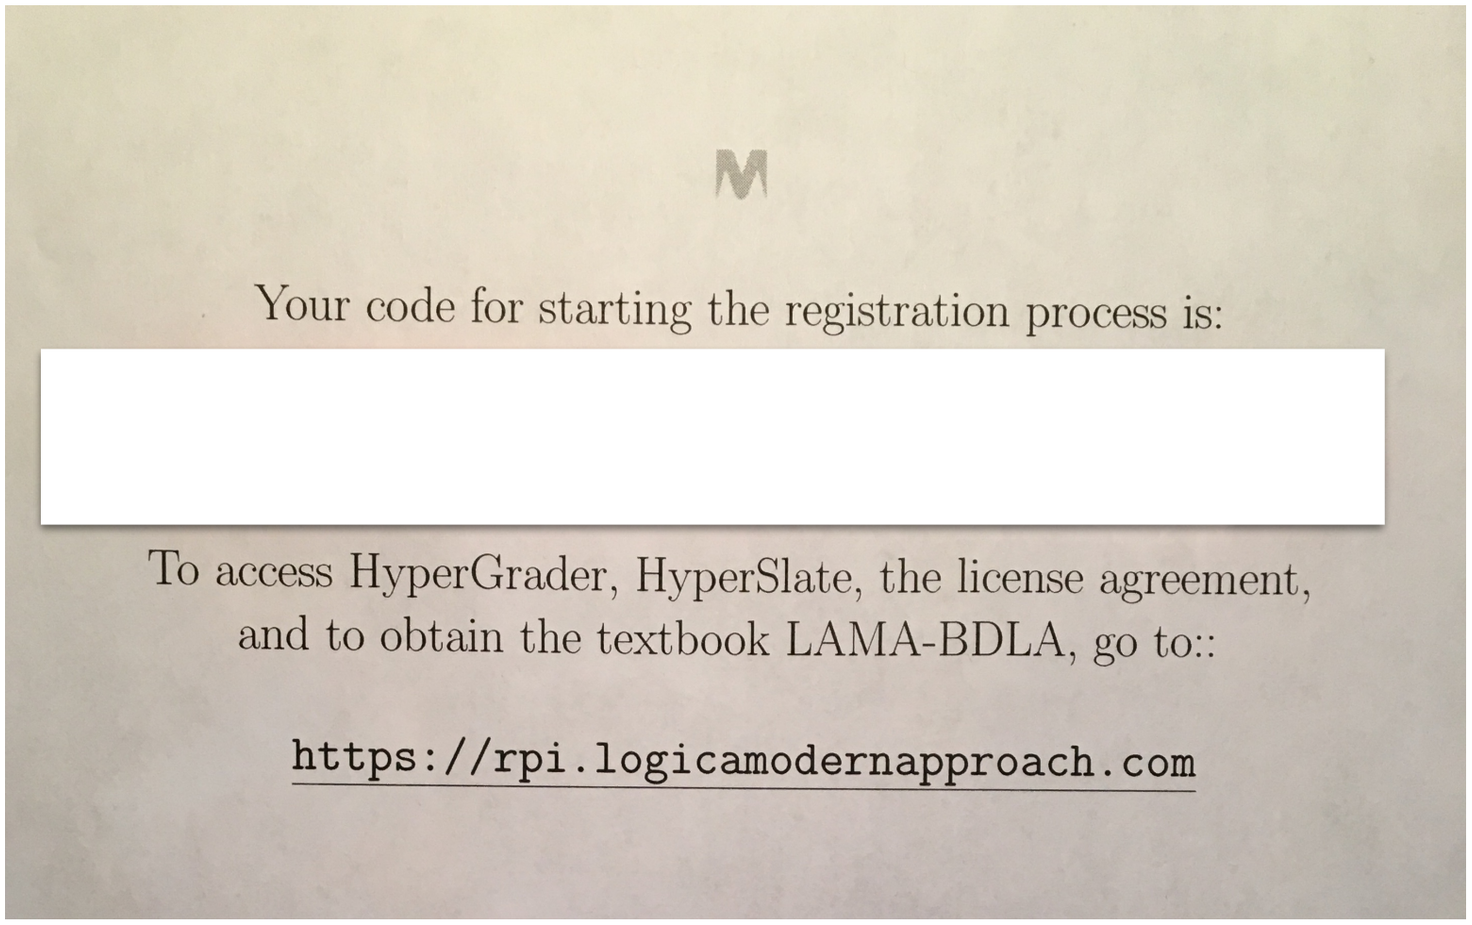


Once seal broken on envelope, no return. Remember from first class, any reservations, opt for "Stanford" paradigm, with its software instead of LAMA® paradigm!

The email address you enter is case-sensitive!

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

The Starting Code to Purchase in Bookstore



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

The email address you enter is case-sensitive!

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

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

Assignments etc. for Week-Long Break, Due Sep 10

Assignments etc. for Week-Long Break, Due Sep 10

Read the SEP overview of AI! (Oops!)

Assignments etc. for Week-Long Break, Due Sep 10

Read the SEP overview of AI! (Oops!)

Thoughtfully review our slide decks thus far.

Assignments etc. for Week-Long Break, Due Sep 10

Read the SEP overview of AI! (Oops!)

Thoughtfully review our slide decks thus far.

Branch from them to outside reading, as motivated.

Assignments etc. for Week-Long Break, Due Sep 10

Read the SEP overview of AI! (Oops!)

Thoughtfully review our slide decks thus far.

Branch from them to outside reading, as motivated.

Register on the platform (see logistics!).

Assignments etc. for Week-Long Break, Due Sep 10

Read the SEP overview of AI! (Oops!)

Thoughtfully review our slide decks thus far.

Branch from them to outside reading, as motivated.

Register on the platform (see logistics!).

Be prepared to discuss both *BR* & *TTG* movies.

Be Ready! …

Be Ready! …



Be Ready! …



Note: Should on Sep 14 have laptop in class & ready to sign in again, having explored things a bit, based on in-class guidance.

Be Ready! …

Note: Should on Sep 14 have laptop in class & ready to sign in again, having explored things a bit, based on in-class guidance.

Be Ready! …

Note: Should on Sep 14 have laptop in class & ready to sign in again, having explored things a bit, based on in-class guidance.

Sep 10 class will include complete tutorial on HyperSlate & HyperGrader interfaces etc.

Logikk kan gi dyp glede!