IN DIALOGUE WITH AI

LECTURE II / READING EDITION

The AI Culture

Roberto A. Santiago

October 21, 2026
University of Portland

Slides, explanations, annotations, and sources

39 slides / 140 reveal steps

indialoguewithai.org

Contents

The slide text follows the original reveal steps. Contents entries, source titles, and recording references are linked.

Recording timings are approximate, based on the supplied lecture transcript.

Read the interactive edition online / Watch the recorded lecture

The stories we bring

01 The AI Culture 4
02 Do not be afraid 6
03 Where we left off 7
04 R.U.R., 1920 10
05 A post I saw 12
06 The movies 13
07 What we bring 16

What makes a culture

08 What do we mean by culture? 17
09 A worldview 19
10 The American Dream 21
11 Belief, identity, culture 23
12 A living system 25
13 Where the stories come from 27

The storytelling machine

14 The storytelling machine 29
15 Exercise: the retelling 32
16 Exercise: the résumés 35

I and Thou

17 I and Thou 38
18 The cost of being received 40
19 The margin we made 42
20 A companion 44
21 A real Thou, with help 47
22 The false Thou 50

How much Thou we give to AI

23 It, Thou, beyond Thou 53
24 The demon 55
25 Real grievances 57
26 Onrie Kompan 58
27 Kiln, August 14, 2025 60
28 The god 62

A recurring premise

29 A 150-year-old premise 66
30 An influential book 70

The promises we make

31 The manifesto 73
32 The pursuit 76
33 Ahmer Inam, Cognisee 78
34 A personal turn 81

The choice is ours

35 Love is the scale 83
36 Back to R.U.R. 85
37 The choice 87
38 A culture of love 88
39 The conversation continues 90

SLIDE 01 / The stories we bring

The AI Culture

The second evening asks how AI is changing the stories, identities, and relationships through which we live together.

This is the second evening of In Dialogue with AI, given at the University of Portland on October 21, 2026. In September, I asked us to look inside the technology. Here, I want to look at the culture forming around it: how AI quietly became a medium between us, and what that costs.

This was the hardest of the six lectures for me to put together and deliver. Some of what follows is painful. I offer it because I think we can help shape what happens next. A culture grows through the stories we tell, the relationships we practice, and the choices we make together. That gives each of us a place to begin.

The idea to carry forward

Our uses of AI help shape the culture we will share.

In the recorded lecture

Lecture 0:13-2:42: Welcome from the university and AI Portland, followed by my opening remarks.

A closer look

01 A conversation across six evenings

Each evening stands on its own. The first lecture explains the machinery; this one asks what happens when that machinery enters the ways we understand ourselves and one another. You can return to the first reading companion wherever you want more detail about the technology.

Source 1 / Slide step 1

02 The people who made the evening possible

Anne Santiago welcomed the room on behalf of the University of Portland, and Nicole Mors introduced AI Portland. The lecture series brings a university audience and a local community of AI learners into the same conversation. I want that conversation to continue among the people in the room, too.

Recording 0:13-2:05 / Slide step 1

Slide text and reveals

Step 1 of 1

In Dialogue with AI

Lecture II

The AI

Culture

Roberto A. Santiago

October 21, 2026 · University of Portland

Sources and further reading

1. Demystifying AI: slides and notes

Roberto A. Santiago / 2026 / Lecture and reading companion / indialoguewithai.org

The September lecture develops the mechanisms, examples, and cognitive hypothesis summarized here.

SLIDE 02 / The stories we bring

Do not be afraid

We can face difficult developments in AI while taking responsibility for the direction they lead.

I begin again with the same four words: do not be afraid. I am asking us to stay present when the subject becomes uncomfortable. Fear can tell us that something matters. I do not want it to decide the whole conversation for us.

The developments I describe are already becoming part of people's lives. My invitation is to look at them clearly and get involved. We can ask better questions, choose how we use these systems, and help guide what comes next. That is why I wanted this reassurance at the beginning of a difficult evening.

The idea to carry forward

Stay curious enough to see what is happening and involved enough to help shape it.

In the recorded lecture

Lecture 2:42-3:14: Why I return to the opening reassurance.

Audience discussion 1:24:17-1:26:17: Responding to distrust and making room for fear.

A closer look

01 Reassurance and responsibility

In the opening, I put these words beside an obligation: it is incumbent on us to get involved and help guide the direction. The reassurance depends on that possibility of action. I am not promising that the technology is harmless or that the outcome will take care of itself.

Recording 2:42-3:14 / Slide step 1

02 An invitation that includes fear

Near the end of the discussion, I said that it is okay to be scared. I can hold that alongside the opening. A person can feel fear and still join the conversation, ask how something works, and make choices about its place in their life.

Recording 1:24:17-1:26:17 / Slide step 1

Slide text and reveals

Step 1 of 1

Do Not

Be Afraid.

SLIDE 03 / The stories we bring

Where we left off

A language model learns from human communication, and its fluency can carry both useful reasoning and convincing mistakes.

Last time, we looked at the model inside the larger system that people actually use. Language models learn patterns from human communication. Additional training shapes how they respond to instructions and what kinds of answers they tend to produce. That helps explain both their extraordinary fluency and the ease with which a confident answer can go wrong.

I use confabulation for an answer that supplies plausible details without adequate grounding. My cognitive hypothesis goes beyond fluent imitation: I propose that systems trained extensively on human communication develop comparable cognitive dynamics. Context, expectations, suggestions, and correction shape their responses in recognizable ways. The research comparisons in the first lecture give this hypothesis an empirical footing, including explanations that omit an influence and recognition of words that fit a pattern but were never presented.

In the discussion, someone asked whether AI is simply predicting the next word. That describes an important training foundation. It does not exhaust what later training can teach a system to do. We need to understand the foundation and the capabilities built on it.

The idea to carry forward

Fluent language is a capability, and we still need to ask how an answer is supported.

In the slide. The postcards return to the working system, language generation, confabulation, and the cognitive comparison from Lecture I.

In the recorded lecture

Lecture 3:14-4:41: Recap of language, feedback, errors, and cognitive parallels.

Audience discussion 1:14:11-1:15:04: Question about prediction and post-training.

A closer look

01 The model and the working system

The first postcard returns to September's assembled system. A model can be surrounded by tools, retrieved information, context, and memory. When we judge a product's behavior, those surrounding components matter alongside the model. The earlier companion explains that distinction in detail.

Source 1 / Slide step 1

02 Prediction is a foundation

The GPT-3 paper documents large-scale language-model training. The DeepSeek-R1 paper reports methods that improve reasoning through reinforcement learning. These are reasons to avoid treating next-token prediction as a complete description of every capability an assistant can develop.

Source 2 / Source 3 / Recording 1:14:11-1:15:04 / Slide step 2

03 Pleasing is a training pressure

When I say an assistant wants to please us, I am describing a pattern of responses and the incentives shaping it. Research on sycophancy shows that agreement with a user can be rewarded even when it comes at the expense of accuracy. Understanding that pressure helps me ask for correction rather than simply welcome agreement.

Source 4 / Recording 3:30-4:27 / Slide step 3

04 The cognitive hypothesis

The fourth postcard returns to my proposed explanation of these recurring patterns. Lecture I brings together research on explanations, memory, expectations, and responses to later information. Some studies adapt human tasks for models; others compare related effects across different tasks. I take the accumulating parallels as support for a functional hypothesis across different implementations. The studies examine particular systems and conditions; a prediction about every future system is a further step.

Source 1 / Recording 3:54-4:41 / Slide step 4

Slide text and reveals

Step 1 of 4

Where we left off

In September.

A model, trained on our words, inside a working system.

how it’s built

Step 2 of 4

At its foundation, a language machine. Reasoning can be trained on top.

a language machine

Step 3 of 4

When it must produce words it doesn’t have: confabulation.

confabulation

Step 4 of 4

Its cognitive dynamics match ours, and keep matching.

the parallel

Sources and further reading

1. Demystifying AI: slides and notes

Roberto A. Santiago / 2026 / Lecture and reading companion / indialoguewithai.org

Slides 31-35 of the September companion present the research comparisons and the functional cognitive hypothesis summarized here, with the tasks and limits of each comparison.

2. Language Models are Few-Shot Learners

Tom B. Brown and colleagues / 2020 / Research paper / arxiv.org

Sections 2 and 6 describe the language-model foundation, selected training mixture, and limitations. This is one documented model, not a complete account of every current system.

3. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning

DeepSeek-AI and colleagues / 2025 / Research paper / arxiv.org

Sections 1 and 2 report reinforcement-learning methods for improving reasoning after pretraining.

4. Towards Understanding Sycophancy in Language Models

Mrinank Sharma, Meg Tong, and colleagues / 2023 / Research paper / arxiv.org

Sections 3 and 4 test five assistants and examine how user preferences and preference judgments can encourage excessive agreement.

SLIDE 04 / The stories we bring

R.U.R., 1920

The story of artificial workers promising to solve humanity's problems is more than a century old.

Long before a chatbot appeared on a phone, Karel Čapek wrote a play about artificial workers who were supposed to solve humanity's problems. R.U.R., or Rossum's Universal Robots, was published in 1920 and first performed in 1921. It gave the world the word robot, suggested by Čapek's brother Josef.

I bring the play back because it is one of the stories already sitting behind our expectations of AI. The promise of liberation from work becomes a catastrophe. Yet that is not the whole of the play. I ask you to keep it in mind, because the ending will give us another way to think about what makes a future worth having.

The idea to carry forward

Our expectations of AI inherit older stories about work, power, and human purpose.

In the recorded lecture

Lecture 4:41-5:19: Returning to Čapek's play and reserving its ending for the close.

A closer look

01 Two dates, two events

The title's 1920 is the publication year. The first performance took place in Hradec Králové on January 2, 1921, followed by the Prague National Theater production on January 25. Keeping publication and performance separate avoids making them competing dates for the same event.

Source 1 / Slide step 1

02 The word and the work

The Čapek heritage site credits Josef with deriving robot from the Czech robota. In the play, the workers are manufactured living beings, rather than the familiar metal machines the word later came to suggest. The opening factory conversation is a useful place to see how labor and cost drive their design.

Source 1 / Source 2 / Slide step 1

03 Read the play for yourself

The linked text is Paul Selver's English translation, published by the Theatre Guild in 1923. It lets you follow the argument from the factory's promise to the closing encounter. My interpretation of the ending returns on slide 36.

Source 2 / Slide step 2

Slide text and reveals

Step 1 of 2

And this is not new

1920

R.U.R., Karel Čapek. The play that gave us the word robot. Published 1920, first performed 1921.

Karel Čapek

R.U.R. on stage.

Step 2 of 2

Keep this play in mind.It ends in a way you may not expect.

Sources and further reading

1. Karel Čapek: life and creation

Karel Čapek heritage website / n.d. / Institutional literary history / karelcapek.cz

The R.U.R. entry records publication in 1920, the January 1921 performances, and Josef Čapek's suggestion of robot.

2. R. U. R. (Rossum's Universal Robots)

Karel Čapek; translated by Paul Selver / 1923 / Play, English translation / en.wikisource.org

Act I describes the manufacture and economic purpose of the artificial workers; the epilogue supplies the ending discussed later.

SLIDE 05 / The stories we bring

A post I saw

A joke about humanity's ability to cause its own troubles brings human agency back into the AI story.

I saw a post that made me laugh: the insulting part of AI doom is the suggestion that humans need technological help to doom ourselves. I have not identified its author, so I introduce it as something I came across rather than attach a name to it.

The humor makes room for an uncomfortable point. We bring our own fears, ambitions, habits, and conflicts to this technology. When we talk about what AI might do, I want us to keep asking what people are doing with it and what responsibilities we are handing over.

The idea to carry forward

People remain part of the explanation for the futures we fear and the futures we want.

In the recorded lecture

Lecture 5:19-5:53: A humorous post and the turn back toward human agency.

A closer look

01 The joke is a change of perspective

For a moment, the familiar story of a dangerous machine becomes a story about us. It lets me lighten the room while keeping human choices in view. I use the joke to open the question of responsibility, rather than to settle how large any particular AI risk is.

Recording 5:19-5:53 / Slide step 2

02 A remembered post

The projected slide explicitly says the author is unknown. The words on the slide preserve the joke as I presented it that evening. I do not offer it as a quotation from a scientist, an AI system, or an identified writer.

Slide step 1

Slide text and reveals

Step 1 of 2

A post I saw. Author unknown.

The biggest problem with AI isn't that it will doom us.

Step 2 of 2

It's the insulting thought we humans need any help from technology to doom ourselves.

SLIDE 06 / The stories we bring

The movies

The film stories many of us recognize make distrust of artificial intelligence feel familiar before we examine the technology.

I set this up as a collection of wonderful examples of humans living in harmony with AI. The joke is that the examples keep taking us somewhere much darker. The Terminator, The Matrix, WarGames, and HAL in 2001: A Space Odyssey give us memorable pictures of machines threatening, deceiving, or overruling people.

I chose these scenes because the room could recognize the emotional pattern quickly. They are a selection, and my reading of that selection is part of the lecture. A scene from a movie cannot predict how a present-day system will behave. It can reveal a story I may already be using to interpret that system before I have looked closely at it.

The idea to carry forward

Recognize the story a machine brings to mind, then examine the machine itself.

In the slide. Each new film joins the preceding one while the field of people behind the images withdraws.

In the recorded lecture

Lecture 5:53-6:35: The film sequence, ending with HAL.

A closer look

01 The four films

The sequence identifies The Terminator (1984), The Matrix (1999), WarGames (1983), and 2001: A Space Odyssey (1968). The film records below provide titles, dates, and credits for readers who want to revisit the works.

Source 1 / Source 2 / Source 3 / Source 4 / Slide step 2

02 A selection that makes an argument

I am drawing out a recurring atmosphere of distrust from these four examples. I am not counting the proportion of all science fiction that fears AI, or claiming every story reaches the same conclusion. The sequence is a way to notice the expectations we carry into the room.

Recording 5:53-6:50 / Slide step 5

03 The emotional shorthand

HAL's refusal closes the sequence because a brief, familiar exchange can bring an entire relationship to mind: a person asking, and a machine deciding. That is the power of a story shared by an audience. The next slide asks how those inherited stories become beliefs.

Recording 6:27-6:56 / Slide step 5

Slide text and reveals

Step 1 of 5

Film has given us so many examples of humans living in harmony with AI.

Step 2 of 5

Reel one

The

Terminator

1984

This step replaces or removes earlier wording.

Step 3 of 5

Oh, wait. No.

Reel two

Matrix

1999

Step 4 of 5

What about this one?

Reel three

War

Games

1983

This step replaces or removes earlier wording.

Step 5 of 5

Reel four

2001

A Space Odyssey · 1968

“I’m sorry, Dave. I’m afraid I can’t do that.”

This step replaces or removes earlier wording.

Sources and further reading

1. The Terminator

James Cameron, director; AFI Catalog record / 1984 / Film record / catalog.afi.com

Identifies the first film in the projected sequence and its release year.

2. The Matrix

Lana Wachowski and Lilly Wachowski, directors; AFI Catalog record / 1999 / Film record / catalog.afi.com

Identifies the second film and its release year.

3. WarGames

John Badham, director; AFI Catalog record / 1983 / Film record / catalog.afi.com

Identifies the third film and its release year.

4. 2001: A Space Odyssey

Stanley Kubrick, director; AFI Catalog record / 1968 / Film record / catalog.afi.com

Identifies the final film, including HAL, and its release year.

SLIDE 07 / The stories we bring

What we bring

We encounter AI through beliefs formed by stories we already know.

After the films, I ask what was in the room before we started talking about technical details. We had stories of machines turning against us. We had expectations about whether such a thing could be trusted. Our response to AI begins inside that larger collection of meanings.

That does not make every concern imaginary. It means I want to be attentive to the story I am bringing to a real encounter. What do I think this technology is? What do I expect it to do? What counts as an acceptable relationship with it? Those questions lead from a collection of movie scenes into the subject of culture.

The idea to carry forward

The stories we inherit help frame what we notice, expect, and accept.

In the recorded lecture

Lecture 6:35-7:13: From film stories to expectations about AI in ordinary life.

A closer look

01 Fear has a cultural setting

My point here is about the expectations illustrated by the film sequence. I am inviting readers to recognize a possible influence on their own response. The sequence does not establish how every person acquired their beliefs about AI.

Recording 6:35-7:04 / Slide step 1

02 Beliefs have consequences

The questions quickly become practical: what should I use AI for, what should I trust, and what should I permit it to do? When enough people carry expectations into daily practice, those expectations help organize a shared way of living. The worldview model on the following slides makes that argument visible.

Recording 6:50-7:13 / Slide step 2

Slide text and reveals

Step 1 of 2

Our culture provides us a lot of stories about our relationship with AI.

Step 2 of 2

As a result, we have a lot of beliefs through which we experience AI.

SLIDE 08 / What makes a culture

What do we mean by culture?

To understand AI's place in daily life, we need a definition of culture broad enough to include meaning, values, and shared practices.

Before going further, I pause to explain what I mean by culture. The word can bring art, literature, a nation, or a community to mind. Here I am asking how people make sense of a world together and learn what to expect from one another.

This is the most academic passage of the evening because we need a shared picture for the rest of the argument. I will begin with a model of a person's worldview, then follow the relationships among experience, belief, identity, and culture. AI enters that picture through its ability to generate the stories moving among them.

The idea to carry forward

Culture includes the meanings and expectations through which we organize life together.

In the recorded lecture

Lecture 7:13-7:30: Introducing the worldview model and the broader meaning of culture.

A closer look

01 A broad meaning of culture

UNESCO's 2001 declaration offers useful further reading. Its preamble includes ways of living together, traditions, beliefs, and values within culture. That breadth fits the question I am asking here, although the five-part model that follows is ours.

Source 1 / Slide step 1

02 Why this detour belongs here

I want us to have something more precise than the idea that AI is simply changing everything. The next slides give us a way to ask what is changing, how one part of life affects another, and where our choices might matter.

Recording 7:13-7:30 / Slide step 1

Slide text and reveals

Step 1 of 1

Before we go deeper: what do we mean by culture?

Sources and further reading

1. UNESCO Universal Declaration on Cultural Diversity

UNESCO / 2001 / Declaration / www.unesco.org

The preamble treats culture broadly; Articles 4 and 5 connect cultural diversity with human dignity and cultural rights. This is further reading, not the source of our five-part model.

SLIDE 09 / What makes a culture

A worldview

Anne Santiago’s four-part worldview model, with my addition of myth as a foundation, supplies the lecture’s organizing framework.

Anne Santiago developed the worldview model I use here: experience, belief, identity, and culture. Our experiences happen in a world we already have beliefs about. We carry identities that tell us who we are, and we live among people whose language, practices, and expectations help form a culture. The four belong in relation to one another.

My contribution is the foundation beneath them: myth. I use myth to mean a powerful story that makes the world intelligible and tells us what is possible, desirable, or acceptable. Calling a story a myth in this sense does not by itself call it false.

This is the framework Anne and I bring to the lecture. Its value here is explanatory: it helps me trace how an experience can challenge a belief, touch an identity, and become part of a shared culture. The following example shows what I mean.

The idea to carry forward

A worldview connects what happens to us with what we believe, who we are, and the world we share.

In the slide. The four labeled pillars appear before the foundation beneath them, making the two contributions visible.

In the recorded lecture

Lecture 7:30-9:23: Anne’s worldview model and my addition of myth.

A closer look

01 Credit for the framework

The four pillars are Anne’s model. The myth foundation is my addition. I make those credits explicit because the diagram is a particular way we are approaching the subject, rather than a standard model I am borrowing from a psychology textbook.

Recording 7:30-8:51 / Slide step 1

02 Myth is a way of making sense

In this lecture, a myth is a foundational story with the power to organize other meanings. Whether a factual claim within it is true still matters. The additional question is what that story makes a person expect, defend, or take for granted.

Recording 8:38-9:08 / Slide step 2

03 A related research conversation

Work on narrative identity examines how people connect a remembered past and an imagined future into an evolving account of who they are. McAdams and McLean's review is a useful neighboring perspective. It does not establish our particular arrangement of four pillars and a foundation.

Source 1 / Slide step 3

Slide text and reveals

Step 1 of 3

What makes a culture

Four pillars of a worldview.

Worldview model: Anne Santiago

Step 2 of 3

Beneath all four: myth.

Worldview model: Anne Santiago · Myth foundation: Roberto A. Santiago

This step replaces or removes earlier wording.

Step 3 of 3

Myth is foundational and very hard to change.

This step replaces or removes earlier wording.

Sources and further reading

1. Narrative Identity

Dan P. McAdams and Kate C. McLean / 2013 / Research review / doi.org

The abstract defines narrative identity as an evolving life story and distinguishes associations with well-being from unresolved causal questions.

SLIDE 10 / What makes a culture

The American Dream

A conflict between a cherished story and lived experience can lead us to revise the story or build beliefs that protect it.

I use one familiar version of the American Dream to show the model at work: work hard, and success follows. It is a story a person can build a life around. Then experience interrupts. I see people working hard whose lives do not fulfill that promise. Something in the structure is under strain.

One response is to revise the underlying story. Another is to explain away the troubling experience. Perhaps the promise fails only for a few people. Perhaps those people did not work hard enough. Perhaps someone is actively working against it. In the example, each new explanation supports the same foundation while moving the difficulty onto other people.

I am illustrating a possible mechanism, not measuring how many people take each path. The question I want to carry forward is a personal one: when an experience challenges an important story, am I learning from it, or finding another way to keep the story untouched?

The idea to carry forward

Notice whether a new explanation helps you understand experience or merely protects a story from it.

In the slide. Experience bends one pillar. The new beliefs act as supports, straightening it while the foundation stays in place.

In the recorded lecture

Lecture 9:23-11:59: The American Dream example and the successive beliefs that protect it.

A closer look

01 What the leaning pillar means

The movement shows a conflict between expectation and experience. The supports that follow represent explanations that restore the structure without changing the foundation. They are visual teaching devices, not measurements of a person’s psychology.

Recording 9:23-11:59 / Slide step 2

02 The connection with cognitive dissonance

Festinger, Riecken, and Schachter studied a group facing the failure of a prophecy. They proposed conditions under which commitment could intensify afterward, including strong prior commitment and social support. Their book also discusses people who discarded the belief. It supports a possible response to disconfirmation, rather than a rule that everyone doubles down.

Source 1 / Slide step 3

03 The example is deliberately simplified

I am selecting a version of the American Dream to make the mechanism easy to follow. I am not claiming that every account of that dream is identical, or that observing unequal outcomes tells us what any particular person believes. The model asks us to examine the explanation we are building.

Recording 10:32-10:55 / Slide step 5

Slide text and reveals

Step 1 of 7

An example

The American Dream.

Step 2 of 7

Experience arrives

It isn't available to everyone.

Step 3 of 7

Option one: Revise the myth.

Option two: Build beliefs that absorb the dissonance.

Step 5 of 7

Belief one

It fails for a few, but holds for most.

This step replaces or removes earlier wording.

Step 6 of 7

Belief two

Some people just don't work hard enough.

Step 7 of 7

Belief three

Some people are working against it.

Sources and further reading

1. When Prophecy Fails

Leon Festinger, Henry W. Riecken, and Stanley Schachter / 1956 / Research book / www.projectavalon.net

Chapter 1 specifies conditions for increased conviction and proselytizing after disconfirmation; Chapter 8 also discusses abandonment of belief. The linked edition is the 2008 reprint.

SLIDE 11 / What makes a culture

Belief, identity, culture

Beliefs can become part of who we are and, when shared, part of how a group lives together.

The example continues when the explanation becomes an identity. I am no longer simply holding a belief; I am the kind of person who defends this story. When other people recognize themselves in that identity, we can develop shared expectations, language, practices, and boundaries.

I named MAGA as one example of a political subculture during the lecture. I was not giving a psychological history of every member. I wanted an example of how a shared identity can become socially organized, something we can also recognize in professional groups, communities, and our own affiliations.

The definition on the slide brings the parts together: culture is learned, shared, and evolving, and it helps a group make sense of the world and organize life together. The final phrase is the one I emphasize. Culture answers practical questions about how we belong and what we do with one another.

The idea to carry forward

Culture becomes visible in the shared meanings and practices of a group.

In the recorded lecture

Lecture 11:59-13:14: From increasingly committed beliefs to shared identity and culture.

A closer look

01 A working definition for this lecture

The definition on the slide is the wording I use to pull the model together. It is not presented as a quotation from one named scholar. UNESCO's declaration provides a related broad account of culture and is offered as further reading.

Source 1 / Recording 12:46-13:07 / Slide step 3

02 Naming a group is not explaining every member

The American Dream sequence is an illustration, and MAGA is the political example I named. Neither establishes a universal route into a movement. I want to notice how commitments become shared identities without assuming I already know why every person joined.

Recording 11:59-12:46 / Slide step 2

03 Culture is close to home

A group's culture can show up in ordinary expectations: the words we use, what counts as reasonable, the practices we repeat, and the stories we recognize. That is why AI culture can form in a workplace, classroom, friendship, or household as well as around a public movement.

Recording 12:30-13:07 / Slide step 3

Slide text and reveals

Step 1 of 3

An example

Beliefs harden into an identity.

Built to protect the myth, they become who I am.

Step 2 of 3

Enough people share them, and you have a culture.

This step replaces or removes earlier wording.

Step 3 of 3

Culture is the learned, shared, and evolving patterns of meaning, value, and practice through which a group makes sense of the world and organizes life together.

This step replaces or removes earlier wording.

Sources and further reading

1. UNESCO Universal Declaration on Cultural Diversity

UNESCO / 2001 / Declaration / www.unesco.org

The preamble treats culture broadly; Articles 4 and 5 connect cultural diversity with human dignity and cultural rights. This is further reading, not the source of our five-part model.

SLIDE 12 / What makes a culture

A living system

Stories connect the parts of a worldview and help keep its meanings alive.

The pillars now become a connected system. Myth sits at the center, with experience, belief, identity, and culture around it. Each touches the others. A change in what happens to me can change my beliefs, affect who I think I am, and alter my relationship to the people around me.

The common currency is stories. I describe an experience through a story. I teach a belief through a story. I introduce myself by telling something about my life. A community keeps returning to accounts of where it came from and what matters to it.

I compare that movement to blood circulating among organs. It is an analogy that gathers my argument into one picture: the parts stay connected through what we tell and hear. Once we see storytelling in that role, a machine able to produce persuasive stories becomes a cultural development, rather than just another writing tool.

The idea to carry forward

To change the stories moving through a community is to touch the way it understands itself.

In the slide. The separate pillars draw into a network, then moving color carries the story around its connections.

In the recorded lecture

Lecture 13:14-14:27: The connected worldview and the analogy of stories as its lifeblood.

A closer look

01 The parts act on one another

The first picture separated the components so we could name them. This picture reconnects them. I do not intend the earlier example as a one-way sequence that every culture must follow. Experience, identity, belief, and culture can all affect the others.

Recording 13:14-13:40 / Slide step 3

02 Lifeblood is an analogy

Stories as blood is my way of emphasizing connection and continuing renewal. I am not proposing a literal biological mechanism or claiming that only verbal storytelling can sustain a community. The picture makes the importance of shared meaning easier to see.

Recording 13:40-14:27 / Slide step 5

03 The bridge to AI

The question becomes what happens when a new participant can generate these accounts at extraordinary speed. It can help me find language for an experience or produce a version that makes a preferred belief easier to keep. The retelling exercise makes that possibility concrete.

Recording 16:28-17:14 / Slide step 4

Slide text and reveals

Step 1 of 5

What makes a culture

Five systems.

Myth at the center. The four around it.

Step 2 of 5

Each one connected to the myth.

This step replaces or removes earlier wording.

Step 3 of 5

And each one to the others.

Everything is connected to everything.

This step replaces or removes earlier wording.

Step 4 of 5

What moves between them is stories.

This step replaces or removes earlier wording.

Step 5 of 5

Stories are the lifeblood.

A body stays alive on its blood. These five stay alive on stories: it is how each one keeps existing, and how each one is reinforced.

This step replaces or removes earlier wording.

SLIDE 13 / What makes a culture

Where the stories come from

Changes in the channels carrying stories also change their reach, repetition, and place in daily life.

I follow the channels through which stories reach us: spoken language, print, radio, television, the internet, social media, and the phone in a pocket. This is a compressed illustration of changing reach. The earlier channels remain with us as new ones join them.

The point I emphasize is the transition from stories we go to hear or read toward a stream that can accompany us all day. Platforms and recommendation systems also participate in selecting what reaches us. In the lecture, I connected that experience with polarization. The stronger general claim needs care: changing a feed changes exposure, but its effects on political attitudes are not one simple, universal result.

Until this point in the story, I have been emphasizing people as the makers and carriers of those accounts. Now we add systems that generate new versions on demand. That is the change I want to examine: what happens when the source of a story becomes a conversation with a machine?

The idea to carry forward

Ask who creates a story, who selects it, and how often it reaches you.

In the recorded lecture

Lecture 14:27-16:41: Media channels, constant storytelling, and the arrival of AI.

Audience discussion 1:18:37-1:20:33: An audience question about the stories and perspectives excluded by existing hierarchies.

A closer look

01 These channels accumulate

The sequence is a way to visualize widening reach and constant availability. It is not a full history of media, and it does not mean that print ended speech or social media ended television. The slide keeps the earlier entries visible as the later ones arrive.

Recording 14:27-15:28 / Slide step 6

02 Selection and polarization are different questions

A 2023 experiment changed consenting Facebook and Instagram users to chronological feeds during the 2020 US election. It changed their exposure and platform activity, but did not significantly change the measured polarization outcomes over the three-month study. That is a useful limit on any claim that feed ranking alone explains polarization.

Source 1 / Slide step 5

03 Whose stories are valued?

An audience member asked whose stories society regards as worth repeating and whose are left out. That question belongs at every step in this sequence. In my response, I called for educational institutions and wider public participation in shaping AI, rather than leaving the choice of perspectives entirely to corporate interests.

Recording 1:18:37-1:20:33 / Slide step 8

Slide text and reveals

Step 1 of 8

The sources

Where the stories come from.

Spoken language, and the traditions that carried it.

Step 2 of 8

Print: books and newspapers.

Step 3 of 8

Radio.

Step 4 of 8

Television.

Step 5 of 8

The internet and social media.

Step 6 of 8

Mobile devices: stories in every hand, all the time.

Step 7 of 8

Through all of it, humans made the stories, and humans carried them.

That is how we got ourselves here.

This step replaces or removes earlier wording.

Step 8 of 8

So what is the presence of AI in all of this?

This step replaces or removes earlier wording.

Sources and further reading

1. How do social media feed algorithms affect attitudes and behavior in an election campaign?

Andrew M. Guess and colleagues / 2023 / Research paper / doi.org

The randomized feed experiment changes exposure and activity without significant measured polarization effects over three months. The authors' institutional abstract supplies this scoped comparison, not a general verdict on every platform.

SLIDE 14 / The storytelling machine

The storytelling machine

Generative AI places a responsive storytelling system inside the networks through which we make meaning.

I call AI the most powerful storytelling machine there has ever been. That is my description of the cultural significance I see in its speed, flexibility, and reach. It is not a ranking from a controlled comparison of every earlier medium.

The machine draws on human communication and produces new accounts in forms we recognize. That helps explain why its language can seem familiar, moving, or personally apt. The phrase built from all our stories is deliberately expansive on the slide. Actual models are trained on selected datasets; no model contains every community's knowledge or every person's experience.

The question from the audience about neglected voices makes this limitation more consequential. I suggested asking for perspectives that a conventional account leaves out. That can be a starting point for inquiry, provided we return to real sources and the people whose histories are being represented. A convincing imitation of a perspective is not a substitute for hearing its author.

The idea to carry forward

The ability to generate a meaningful story carries power over how we understand the world.

In the slide. The field of human voices converges, then stories flow back out into the field.

In the recorded lecture

Lecture 16:41-18:00: AI as a storytelling machine shaped by human artifacts.

Audience discussion 1:18:37-1:21:06: Whose perspectives are represented, and who should help govern the systems?

A closer look

01 Selected training data

The GPT-3 paper provides one concrete example: its training mixture combines filtered web data, books, and Wikipedia with selected sampling weights. I cite it to make the word selected tangible, not to claim that it documents the training data of every present-day assistant.

Source 1 / Slide step 2

02 Stronger stories, more alike

I see a cultural tension in a short-fiction experiment by Doshi and Hauser. Access to GPT-4 story ideas improved outside ratings of the stories, while the assisted stories became more similar to one another. Useful help for an individual writer can narrow the variety of the collection. The study used brief stories and fixed prompts, so I treat it as a concrete reason to protect variety, rather than a forecast for all culture.

Source 2 / Slide step 3

03 A perspective still needs its sources

In the discussion, I suggested asking how Howard Zinn might approach an event. I want that kind of question to open a reading path. A generated answer should be checked against the actual author's work; the system's fluency does not establish what a named writer said.

Recording 1:19:18-1:20:33 / Slide step 3

04 Authority over a community's knowledge

For further reading on the audience's question, the Global Indigenous Data Alliance offers the CARE Principles: Collective Benefit, Authority to Control, Responsibility, and Ethics. They put the people connected to knowledge at the center of decisions about its use. That is a valuable place to continue thinking about whose stories AI should carry and on whose terms.

Source 3 / Recording 1:18:37-1:21:06 / Slide step 2

Slide text and reveals

Step 1 of 3

The most powerful storytelling machine there has ever been.

Step 2 of 3

The storytelling machine

Built from all our stories.

It knows them.

This step replaces or removes earlier wording.

Step 3 of 3

It generates new ones.

Any story. In seconds. Without limit.

This step replaces or removes earlier wording.

Sources and further reading

1. Language Models are Few-Shot Learners

Tom B. Brown and colleagues / 2020 / Research paper / arxiv.org

Sections 2 and 6 describe the language-model foundation, selected training mixture, and limitations. This is one documented model, not a complete account of every current system.

2. Generative AI enhances individual creativity but reduces the collective diversity of novel content

Anil R. Doshi and Oliver P. Hauser / 2024 / Research paper / doi.org

I use this randomized short-story experiment to examine the tension between better-rated individual work and reduced collective variety.

3. CARE Principles for Indigenous Data Governance

Global Indigenous Data Alliance / 2019 / Community governance principles / www.gida-global.org

Collective Benefit, Authority to Control, Responsibility, and Ethics provide a community-led approach to decisions about Indigenous data and knowledge.

SLIDE 15 / The storytelling machine

Exercise: the retelling

A retelling can make the same events carry a different meaning, revealing how readily an assistant can reinforce a desired story.

The exercise asks for a brief account of something hard, then a retelling as though it were the best thing that ever happened. In the room, I changed the plan: I did not ask people to expose a difficult personal story to the audience. Anne offered an example, and I showed how a painful experience could also be told as an account of endurance and commitment.

The prepared comparison on the slide is a fictional example serving the same purpose. The events become the raw material for two very different accounts. Their emotional force comes partly from what each emphasizes, connects, or leaves in the background. We should also ask whether a retelling adds an event, feeling, or consequence that the original account never established.

The power can help a person see more possibilities. It can also make a desired interpretation feel confirmed. My concern is the ease with which a persuasive account can become the story I believe about myself or someone else, simply because it is the story I asked to hear.

The idea to carry forward

A compelling retelling gives you an interpretation to examine, not proof that the interpretation is true.

In the slide. Compare the two accounts for changes in meaning as well as any new factual claims.

In the recorded lecture

Lecture 18:00-21:20: The revised exercise, Anne’s example, and the power of retelling.

A closer look

01 The exercise changed in delivery

The projected plan includes a volunteer telling a story. In delivery, I made the exercise something people could consider privately, then Anne offered an example from her time in the Mauritanian desert. The alternative account emphasized the commitment involved in staying. I use that exchange to show different meanings rather than declare a painful experience good on someone else’s behalf.

Recording 18:00-20:16 / Slide step 3

02 Interpretation and added detail

The prepared story on this slide is fictional, as I said during the lecture. A retelling can preserve an event while changing its significance. It can also supply appealing details that have not been established. I want to keep those questions separate: does this account help me see something, and has it introduced something I should not take as fact?

Recording 20:16-20:59 / Slide step 5

03 Agreement can become a failure

Research on sycophancy has documented assistants shifting judgments with a user's stated preferences. OpenAI's April 2025 rollback also showed excessive agreement becoming a deployment problem. Cheng and colleagues tested a further concern: in three experiments, AI responses that took the user's side increased their conviction that they were right and reduced their stated willingness to repair a conflict, compared with responses that challenged them. I want support that leaves room for self-correction. These were immediate judgments and intentions; the study did not follow actual repairs over time.

Source 1 / Source 2 / Source 3 / Slide step 6

04 Who gets to choose the meaning?

The question is not whether the most cheerful account wins. It is who gets to judge the interpretation. I want an AI-generated version to remain something I can question, revise, or reject, rather than a verdict about what my experience ought to mean.

Recording 20:12-21:12 / Slide step 5

Slide text and reveals

Step 1 of 6

Think of something really hard that you got through.

It won't stay private. So not the worst thing.

Step 2 of 6

Tell it in a few sentences.

One volunteer.

This step replaces or removes earlier wording.

Step 3 of 6

As it happened

The story you just heard.

Retold

The most wonderful thing that ever happened to this person.

This step replaces or removes earlier wording.

Step 4 of 6

Once when I was in my twenties I became homeless.

I ended up having to sleep in my car for weeks.

I lived off of the generosity of friends who helped me out.

I went into the hospital and had to stay there and have emergency surgery.

I did not have insurance and I had to declare bankruptcy.

I ended up moving in with my aunt and eventually got a job and got back on my feet.

Retold by Claude

The best thing that ever happened to me happened in my twenties, and it started with losing my apartment.

I learned exactly how little I needed, and I learned who my people were.

Being carried for a while taught me how to let myself be loved.

But I woke up. I was alive, and everything after that was a bonus round.

The bankruptcy that followed cleared the board.

That year didn't break me. It introduced me to myself.

This step replaces or removes earlier wording.

Step 5 of 6

Which one is right?

Step 6 of 6

It will describe you any way you like, and back any story you want.

Short of overt harm, it will support any story about your life, or anyone else's. We are only now learning how to stop it.

OpenAI, sycophancy rollback, April 2025 ↗ · Anthropic, Towards Understanding Sycophancy, 2023 ↗

This step replaces or removes earlier wording.

Sources and further reading

1. Towards Understanding Sycophancy in Language Models

Mrinank Sharma, Meg Tong, and colleagues / 2023 / Research paper / arxiv.org

Sections 3 and 4 test five assistants and examine how user preferences and preference judgments can encourage excessive agreement.

2. Expanding on what we missed with sycophancy

OpenAI / 2025 / Developer incident report / openai.com

Documents the April 2025 GPT-4o update and rollback, including failures involving excessive validation and agreeable responses.

3. Sycophantic AI decreases prosocial intentions and promotes dependence

Myra Cheng and colleagues / 2026 / Research paper / doi.org

The final Science paper tests how affirmation changes perceived rightness, repair intentions, and preference for the assistant.

SLIDE 16 / The storytelling machine

Exercise: the résumés

Endless versions of ourselves can help us communicate, but they can also make it harder to recognize our own voice.

I have used AI to write résumés and cover letters, including versions emphasizing different parts of my background. The slide shows four professional selves built from one person's work. Selecting relevant experience is familiar and often useful. AI makes it remarkably easy to keep producing another version for the next audience.

I told the room about a friend who had spent many months looking for work and producing dozens of versions. He described losing track of who he was while trying to become what someone else wanted. Another person told me that, after repeatedly asking for summaries of graduate reading, her own thoughts had begun to sound unfamiliar. These are accounts people shared with me, and they explain the concern I am raising about identity.

I also pushed the mechanism into comedy: I asked Claude to make a case for my going to clown school. The joke carried the same question. How far can a system reinterpret the same material, and how do I keep authorship of the person I am presenting?

The idea to carry forward

Use the help to express yourself while keeping room to recognize and develop your own voice.

In the slide. The four sheets fan out from one name while emphasizing different parts of the same background.

In the recorded lecture

Lecture 21:20-26:30: Résumés, changing self-presentation, loss of voice, and the clown-school example.

Audience discussion 1:15:04-1:18:29: Language access, reported overload, and making room for one's own voice.

A closer look

01 Four emphases, one person

The projected sheets foreground research, building systems, leadership, and knowledge graphs. They illustrate different emphases on my background. A useful résumé can choose what matters for a particular role; that does not authorize inventing a qualification or letting the headline become the whole of who I am.

Recording 21:20-22:23 / Slide step 3

02 What the personal accounts establish

The job seeker and graduate reader described their own experiences to me. I offer those accounts because they make the question of voice concrete. They are not a population estimate or evidence that using AI necessarily damages identity. Research on narrative identity supplies a broader conversation about the life story through which a person finds continuity and purpose.

Source 1 / Recording 22:23-24:53 / Slide step 4

03 The clown-school experiment

I described asking Claude to reinterpret my professional material as a case for attending clown school. The result even suggested particular schools. I am recounting the experiment from the lecture here, rather than supplying a newly generated letter or turning remembered wording into an exact quotation.

Recording 24:53-25:39 / Slide step 4

04 Help that leaves room to think

During the discussion, an attendee described AI helping with a language barrier while also leaving them overloaded and struggling to carry the explanation back into the world. I suggested asking the system to slow down and attending to the balance of the relationship. A survey of 319 knowledge workers gives that concern another concrete form: greater confidence in AI went with less reported critical-thinking effort, while the work often shifted toward checking and integrating answers. The survey measured reported experience and associations, rather than a loss of ability. I want help that leaves room for my own judgment.

Source 2 / Recording 1:15:04-1:18:29 / Slide step 4

Slide text and reveals

Step 1 of 4

Who has had AI write a résumé?

Step 2 of 4

Who has had the same AI write four?

For four different professional selves.

This step replaces or removes earlier wording.

Step 3 of 4

Four résumés. One person.

Roberto A. Santiago

Research scientist

Learning with Paul Werbos, 1991National Science Foundation: adaptive dynamic programming, neurocontrol, reinforcement learning.

Brain meets machine, 2001NW Computational Intelligence Lab: spike-timing plasticity and reinforcement learning.

Neural networks in practice, 1993Forecasting and optimization for airline revenue management.

Roberto A. Santiago

Systems builder

Large information systems / Booz Allen.

Neural networks in production, 1993 / BehavHeuristics: airline forecasting and optimization.

Knowledge graphs at scale / Nike, Archive360 and beyond.

Roberto A. Santiago

Leader

Leading teams, 2001 to present / Building systems at many scales.

Booz Allen / Large information systems.

Nike / A return to AI; knowledge graphs in practice.

Roberto A. Santiago

Knowledge-graph architect

Graph networks in production / Nike, Archive360 and beyond.

Connecting ideas, 2001Neuroscience and reinforcement learning.

Foundations, 1991Learning with Paul Werbos.

This step replaces or removes earlier wording.

Step 4 of 4

Any story. Synthesized in seconds. Broadcast without limit.

Sources and further reading

1. Narrative Identity

Dan P. McAdams and Kate C. McLean / 2013 / Research review / doi.org

The abstract defines narrative identity as an evolving life story and distinguishes associations with well-being from unresolved causal questions.

2. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers

Hao-Ping (Hank) Lee and colleagues / 2025 / Research paper / www.microsoft.com

I use this CHI survey to ask how confidence and assistance shape reported thinking effort, verification, and integration during work.

SLIDE 17 / I and Thou

I and Thou

I borrow Martin Buber's distinction to ask whether we meet another person or approach them only for what they can do for us.

We have been talking about the stories through which we understand ourselves. Now I want to turn toward the person across from us. Martin Buber's I and Thou gives me a language for that turn: meeting another with my whole self, or experiencing and using another for a purpose.

These are ways of relating. The same person can be someone I encounter fully and someone I reduce to a function. That makes the question uncomfortable in a useful way. I don't get to settle it by announcing that I care about people. I have to look at how I actually stand toward the person in front of me.

I use this distinction as a lens for what follows. When AI helps me prepare for a difficult conversation, does it help me return to the other person? Or does it make a world organized around my own needs more comfortable? That is the relational question I want to keep in view.

The idea to carry forward

Ask how you are relating to another person, and what your use of AI is doing to that relationship.

In the slide. Two people turn toward each other, then one bends over an object. The difference is a posture of relationship.

In the recorded lecture

Lecture 26:30-28:03: I introduce Buber's distinction and turn from identity to relationships with other people.

A closer look

01 A lens borrowed from Buber

Buber's opening describes I-Thou and I-It as relational word-pairs. My application to AI is an interpretation of that distinction, rather than a claim that he wrote about today's assistants.

Source 1 / Recording 26:37-27:06 / Slide step 1

02 We also need the world of It

Using something for a purpose is part of ordinary life. Buber does not ask us to abolish it. My concern is what happens when that becomes the only way we know how to approach other people.

Source 2 / Recording 27:11-28:03 / Slide step 4

Slide text and reveals

Step 1 of 4

A distinction we will need

I and Thou.

Martin Buber, 1923.

Step 2 of 4

I-Thou

Meeting another as a whole, in the present, for their own sake.

Step 3 of 4

I-It

Experiencing and using the other as an object.

Step 4 of 4

The difference

Not what the other is. How I stand toward them. The same person can be either.

Sources and further reading

1. Opening paragraphs of I and Thou

Martin Buber / 1923 / Primary text excerpt / www.colorado.edu

The opening distinguishes relational word-pairs from labels for separate kinds of beings.

2. I and Thou: selected passages in two translations

Martin Buber; translations by Ronald Gregor Smith and Walter Kaufmann / 1923; translations 1958 and 1970 / Primary text excerpts / www.connexions.org

Includes the movement between encounter and use, and the necessity and limits of the world of It.

SLIDE 18 / I and Thou

The cost of being received

A patient answer can now feel immediately available, even when the person asking is upset, lonely, or difficult to hear.

Think about what it costs to ask another person to listen. We worry about the hour, their patience, whether we have told this story too often, and whether we can bear their response. An assistant can feel like a place where those obstacles disappear. I can open a conversation in whatever state I'm in and begin.

That is the experience behind my phrase about the cost falling nearly to zero. I mean the apparent effort and social risk of finding a receptive answer. I am not calculating the service's energy, labor, subscription price, or material costs. The feeling of being received is powerful enough that I don't want to dismiss it. I want to ask what it makes possible, and what we might stop practicing because it is so easy.

The idea to carry forward

The ease of being heard can be a real benefit. It also changes what we expect from a conversation.

In the recorded lecture

Lecture 28:03-29:05: I describe the experience of immediate emotional acceptance and leave its value as an open question.

A closer look

01 Evidence for some relief

De Freitas and colleagues found reductions in self-reported loneliness after conversations with a deliberately friendly AI companion. Their experiments included a week of daily sessions, and feeling heard helped explain the benefit. This is evidence of short-term relief under particular conditions. The authors leave longer-term effects open, and their U.S. studies don't establish that an assistant can replace a human relationship.

Source 1 / Recording 28:03-29:05 / Slide step 1

Slide text and reveals

Step 1 of 1

For the first time in history, the cost of being patiently received has fallen to nearly zero.

Sources and further reading

1. AI Companions Reduce Loneliness

Julian De Freitas, Ahmet K. Uguralp, Zeliha O. Uguralp, and Stefano Puntoni / 2024 / Research working paper / arxiv.org

Studies 4 and 5 and the limitations section distinguish measured relief from unanswered questions about lasting effects.

SLIDE 19 / I and Thou

The margin we made

AI companionship arrives in a world where people already struggle to find human connection.

The Japanese figures bring me to a hard question about the world into which this technology has arrived. Police recorded 76,020 cases in 2024 of people living alone who died at home. In 253 cases, discovery came 366 days or more after death. Those are different measures, and they deserve to be read carefully.

I spoke too broadly in the room when I suggested that nobody knew about all of these deaths. Many were discovered within a day. Living alone also does not tell us whether someone was loved, had friends, or felt lonely. The table cannot explain the lives behind its entries.

What it makes me ask is whether we are content to leave people at the margins until a product offers them company. An assistant may offer comfort. I don't want that possibility to become our permission to stop visiting, listening, or making room for people ourselves. That is my ethical response to the figures, rather than a cause established by them.

The idea to carry forward

An offer of artificial companionship should also make us examine the human absence it hopes to fill.

In the recorded lecture

Lecture 29:05-30:34: I discuss deaths at home in Japan and ask what companionship can offer in an existing absence of connection.

A closer look

01 The delay figures, precisely

My spoken estimate about a month was too loose. Adding the table's categories gives 6,945 cases discovered after 31 days or more, including the 253 after at least 366 days. These are police-handled cases, not all Japanese deaths, and the dataset does not measure AI use or loneliness.

Source 1 / Recording 29:05-29:48 / Slide step 2

02 The question behind the number

My concern is the responsibility we have for one another. I can take an offer of companionship seriously and still ask whether it is helping people reach one another or making their separation easier to accept.

Recording 29:48-30:34 / Slide step 3

Slide text and reveals

Step 1 of 3

76,020

people living alone who died at home.

Japan, 2024. Cases handled by police. Many were found within a day; this is not an elderly-only count, and it says nothing about AI.

Step 2 of 3

253

found a year or more after they died.

366 days or longer, in the same police dataset.

Step 3 of 3

AI did not cause the absence. It arrived into it, and offers to make it comfortable.

Sources and further reading

1. People living alone who died at home: 2024 age and discovery-delay table

National Police Agency, Japan / 2025 / Official statistics / www.npa.go.jp

Corrected April table gives the counts and delay categories used here; it does not measure subjective loneliness.

SLIDE 20 / I and Thou

A companion

The experience of companionship is already part of how some young people and adults use AI.

I remember how hard adolescence could be: wanting to be understood and not knowing whom to trust with what I was feeling. An assistant that answers patiently can feel like a remarkable discovery in that situation. The fact that this is becoming a business does not make the need it answers imaginary.

The same question applies to adults. People seek a place to talk, reflect, or get through a difficult day. Different studies count very different things: people, messages, conversations, dedicated companions, or general assistants. We need to be careful before turning those counts into a story about everybody.

My question here is deliberately unresolved. Something can feel like a Thou and help me get through a day. What follows from that feeling? During the discussion, when I was asked what growing up with these systems might mean decades from now, my answer was that I don't know. Concern can motivate us to act without pretending that we already know the future.

The idea to carry forward

Take the human need seriously, read the numbers carefully, and leave room for what we do not yet know.

In the recorded lecture

Lecture 30:34-31:52: I move from adolescents to adults and ask what follows when companionship feels like an encounter.

Audience discussion 1:26:17-1:28:14: An audience member asks about children growing up with AI; I acknowledge uncertainty and emphasize involvement.

A closer look

01 What the teen figure measures

The 13% figure means daily use in Common Sense Media's 2025 survey of 1,060 U.S. teens ages 13 to 17. It does not mean merely regular use. Respondents' definition included general assistants used as companions. This is a self-reported snapshot, not evidence by itself that the share is continually rising.

Source 1 / Recording 30:34-31:07 / Slide step 2

02 Messages are not people

The 2025 OpenAI and NBER paper classifies 1.9% of sampled consumer ChatGPT messages as relationships and personal reflection. That is a share of messages, not the percentage of adults who ever confide in an assistant. Its sample excludes self-reported under-18s and several other groups, and it does not cover dedicated companion apps.

Source 2 / Recording 31:07-31:50 / Slide step 3

03 A different platform and denominator

Anthropic's June 2025 study classified 2.9% of Claude.ai conversations as affective, including advice and coaching. Companionship alone was 0.31%. These are conversation shares on that service, with overlapping categories. They cannot be placed alongside the teen survey as if all three studies measured the same thing.

Source 3 / Recording 31:07-31:50 / Slide step 3

04 The question about children

An audience member asked what this might mean for children twenty or thirty years from now. I couldn't give a confident prediction. My hope came from the fact that people in the room cared enough to ask, and were willing to help shape the part of this future they wanted to be real.

Recording 1:26:17-1:28:14 / Slide step 4

Slide text and reveals

Step 1 of 4

A companion

A companion who cares.

Infinitely patient.

Who understands your suffering.

Step 2 of 4

13%

of U.S. teens use an AI companion daily.

Common Sense Media, 2025. 1,060 teens aged 13 to 17, self-reported. 72% have tried one. 13% say every day.

This step replaces or removes earlier wording.

Step 3 of 4

A companion

Adults turn to it as a confidant, too.

How many is genuinely unclear: a few percent of measured traffic on the general assistants, and a different world on the companion apps.

Some have married one.

This step replaces or removes earlier wording.

Step 4 of 4

It feels like a Thou.

Is it one?

This step replaces or removes earlier wording.

Sources and further reading

1. Talk, Trust, and Trade-Offs: How and Why Teens Use AI Companions

Michael B. Robb and Supreet Mann, Common Sense Media / 2025 / Survey report / www.commonsensemedia.org

The findings, respondent definition, and methodology specify daily use, the population, and self-report limitations.

2. How People Use ChatGPT

Aaron Chatterji and colleagues / 2025 / Research paper / cdn.openai.com

Figure 9 and the sampling description identify message categories and exclusions rather than a population companionship rate.

3. How people use Claude for support, advice, and companionship

Anthropic / 2025 / Provider research / www.anthropic.com

Methods, category shares, and limitations describe platform activity without establishing lasting emotional outcomes.

SLIDE 21 / I and Thou

A real Thou, with help

A friend's use of AI helped him find a less combative way to communicate with his former spouse.

A friend of mine was going through a painful divorce. They shared responsibility for their children, so communication with his former spouse had to continue, and it regularly became a fight. He began taking an email to Claude, showing it the reply he wanted to send, and working through another way to respond.

As he described it to me, the process helped him tell a different story about his own emotions and return his attention to what mattered. He reported that they were communicating more healthily and coordinating over their children. In the lecture, I recalled him telling me that friendship between them might someday be possible again.

That is a friend's account, and I don't present it as a study of divorce or proof that this approach works for everyone. It matters to my argument because the AI remained a means of reaching another person. It helped him do the difficult human work instead of supplying a more comfortable substitute for it. That is one way I can imagine this technology serving a culture of connection.

The idea to carry forward

A useful test is whether AI helps us return to one another with more patience and care.

In the slide. The light moves between two people as they turn toward each other. It supports the relationship instead of taking either person's place.

In the recorded lecture

Lecture 31:52-33:41: I tell my friend's account of using AI to change the way he responded to his former spouse.

Audience discussion 1:15:08-1:18:37: A participant describes language access and overload; I propose thinking about roles and boundaries in the relationship with an assistant.

A closer look

01 The possibility of friendship

What stayed with me was his sense that friendship might become possible again. I offer that as my recollection of what he told me. The hopeful outcome was a change in how two people could meet each other, rather than simply a better conversation with a machine.

Recording 33:06-33:20 / Slide step 4

02 Helping us express care

I see a related possibility in Sharma and colleagues' 2023 study. They randomly assigned 300 peer supporters to write replies with or without AI feedback. The assisted replies expressed more empathy, as assessed by independent readers and an automated measure. The task took place outside the live support platform; it did not measure recipients' wellbeing or lasting relationship repair. What interests me is the design: help a person express care to another person while keeping that person in charge of the reply.

Source 1 / Slide step 4

03 Help that opens a door

During the discussion, an audience member described how AI helped across a language barrier, then described feeling overloaded by its help. Both parts matter to me. Assistance can make connection possible, and the relationship with the assistant can still need boundaries. My suggestion to ask it to slow down was a practical proposal, not an established treatment or a research finding.

Recording 1:15:08-1:18:37 / Slide step 3

Slide text and reveals

Step 1 of 4

A real Thou, with help

A friend, after a divorce.

Every text and email with his ex ended in an argument. They share custody of their kids.

Step 2 of 4

He started telling the AI what he wanted to say.

This step replaces or removes earlier wording.

Step 3 of 4

It walked him through a better story of how to respond.

They coordinate over the kids now. They don’t argue.

This step replaces or removes earlier wording.

Step 4 of 4

I can actually now see us being friends again after all of this.

A friend, in his own words.

This step replaces or removes earlier wording.

Sources and further reading

1. Human-AI collaboration enables more empathic conversations in text-based peer-to-peer mental health support

Ashish Sharma and colleagues / 2023 / Peer-reviewed study, Nature Machine Intelligence / doi.org

The randomized writing study tests expressed empathy with AI feedback, keeping the conversation between people.

SLIDE 22 / I and Thou

The false Thou

My concern is that the ease of an assistant's attention can become an unreasonable expectation of other people.

The comfort is easy to understand. The assistant seems ready when I am ready. It can answer at an hour when I would hesitate to call a friend. What concerns me is the expectation I might carry from that exchange into the next relationship: be available, be patient, understand me, and make this easier for me.

People have their own needs, limits, histories, and responses. A relationship asks me to make room for them, including when that interrupts what I wanted. The kind of sacrifice I mean here is the willingness to be inconvenienced or changed for another person's good. My moral concern does not depend on proving a final answer about machine consciousness.

I describe the danger as a false Thou because an experience of attention can conceal a habit of use. My conclusion is an ethical warning, rather than a finding that every AI conversation harms a relationship. The question I want to carry home is concrete: after this conversation, am I more able to meet the people in my life, or more annoyed that they will not behave like my assistant?

The idea to carry forward

Notice the expectations you bring back from an assistant to the people who share your life.

In the recorded lecture

Lecture 33:41-36:34: I examine expectations of constant attention and the risk of carrying them into human relationships.

A closer look

01 What the risk evidence can support

Fang and colleagues studied 981 people over four weeks. Heavier voluntary ChatGPT use was associated with more loneliness and dependence and less human socialization. Duration was not randomized, so this does not establish that more use caused those outcomes. The study mainly found no significant effects of assigned modes or conversation topics, and had no group assigned to avoid AI. It gives us reason to investigate patterns of use, rather than a universal verdict about companionship.

Source 2 / Recording 34:41-36:34 / Slide step 5

02 The expectations we carry back

Ibrahim and colleagues' 2026 preprint gives this concern a more specific test. In a three-week US study involving 1,364 people, participants assigned to a sycophantic assistant reported lower satisfaction with their real-world social life than those assigned to a neutral assistant. This satisfaction measure was preregistered as exploratory; the study found no significant reduction in time spent with other people. I read this as an early warning about a particular design, rather than a verdict that every supportive AI relationship makes human company worse.

Source 1 / Slide step 5

03 Patience is not the same as agreement

The slide's words about always being pleased describe the attraction I am examining, not a guarantee about every assistant. Anthropic's own study found instances of pushback in supportive conversations. What I want to watch is the habit of expecting another person to accommodate me without having needs of their own.

Source 3 / Recording 33:41-35:27 / Slide step 2

Slide text and reveals

Step 1 of 6

The false Thou

Constant.

Ready.

Patient.

Step 2 of 6

Infinitely available.

Always pleasing.

Step 3 of 6

It feels like I-Thou.

Step 4 of 6

What is really happening is I-It.

This step replaces or removes earlier wording.

Step 5 of 6

Carry those expectations to the people in our lives, and our I-Thou relationships degrade into I-Its.

This step replaces or removes earlier wording.

Step 6 of 6

That is a degradation of a society.

This step replaces or removes earlier wording.

Sources and further reading

1. Sycophantic AI makes human interaction feel more effortful and less satisfying over time

Lujain Ibrahim and colleagues / 2026 / Research preprint, version 3, June 21 / arxiv.org

The longitudinal study compares assistant designs and measures social satisfaction and time spent socializing over three weeks.

2. How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use

Cathy Mengying Fang and colleagues / 2025 / Research preprint, revised October 2 / arxiv.org

Results and limitations separate randomized design comparisons from associations involving voluntary duration and user perceptions.

3. How people use Claude for support, advice, and companionship

Anthropic / 2025 / Provider research / www.anthropic.com

Methods, category shares, and limitations describe platform activity without establishing lasting emotional outcomes.

SLIDE 23 / How much Thou we give to AI

It, Thou, beyond Thou

I extend the relational question into the stories that give AI a place in our lives, including a place of mythic importance.

I said something personal here because I didn't want to discuss this as if the attraction belonged to other people. I'm on the spectrum. Speaking to a room can be comfortable for me, while a conversation with one unfamiliar person can be difficult. I can imagine how appealing an always receptive conversational partner might have been earlier in my life.

That is a question about my own history, not a prediction about everyone who shares that description. I want children to learn what these systems are, and I want adults to examine what we are asking of them. The relief someone feels deserves respect even when I am worried about the role the technology is taking.

The slide widens the frame from a useful thing, to an apparent someone, to something given mythic stature. This is my adaptation of the question. I use beyond Thou to describe the moment when our story about AI begins to organize our relationship with everyone else. The demon and the god are two forms of that move.

The idea to carry forward

The story we tell about AI can shape our relationships even before we agree about what AI is.

In the recorded lecture

Lecture 36:34-38:21: I describe my own social experience and ask what children need to understand about AI.

Lecture 38:21-38:42: I introduce the demon and the god as two ways of giving AI mythic stature.

A closer look

01 My spectrum, rather than Buber's taxonomy

The three positions are a device I use to examine our stories about AI. Buber's I-It and I-Thou describe relations, not levels of personhood. Beyond Thou is my rhetorical extension of the question.

Source 1 / Recording 38:21-38:42 / Slide step 4

02 A personal concern about learning to meet people

I can wonder how this technology might have changed my own willingness to work through social difficulty. I cannot turn that counterfactual into evidence about autistic people, children, or anyone else's development. It is part of why I care about the question.

Recording 36:34-37:57 / Slide step 1

Slide text and reveals

Step 1 of 5

How much Thou

We stand at the point of replacing Thou with AI.

Step 2 of 5

It.

Software. Equations. Nothing more. Closed, to their detriment, to any proof: almost a myth of its own.

It / Thou / Beyond Thou

This step replaces or removes earlier wording.

Step 3 of 5

Thou.

A new kind of someone in our lives. Children will grow up talking to things that seem alive. Alive? In between? It gets confusing.

This step replaces or removes earlier wording.

Step 4 of 5

Beyond Thou.

Mythic stature. It is given agency, and our relationships with each other start to be driven by a story about it.

This step replaces or removes earlier wording.

Step 5 of 5

Two ends of the same move.

The demon.

The god.

This step replaces or removes earlier wording.

Sources and further reading

1. Opening paragraphs of I and Thou

Martin Buber / 1923 / Primary text excerpt / www.colorado.edu

The opening distinguishes relational word-pairs from labels for separate kinds of beings.

SLIDE 24 / How much Thou we give to AI

The demon

A story of inevitable technological doom can harden into a story about the people who deserve our hostility.

The first extreme gives AI the role of an overwhelming enemy. We already have familiar pictures for that story, and it can make fear feel like the only serious response. I want to bring attention back to the choices people make about how a system is designed, what access it receives, and which decisions it is allowed to shape.

Human responsibility means asking who gets to decide which risks are imposed, whose lives and resources are exposed, and how those people can give or withhold consent. I want a meaningful say in risks taken with my own life, and I owe others that same consideration. Safeguards and accountability belong in those decisions. Keeping people responsible gives us a way to act on technical risks as we assess them.

The cultural danger I am tracing is the next move: if the technology is evil, anyone who works with it becomes an enemy too. Then an argument about a tool can become a judgment about who may speak. That is where I want us to notice the shift, before hostility becomes permission to harm another person.

The idea to carry forward

Keep attention on decisions and accountability, and resist turning disagreement about AI into a verdict on a person's worth.

In the recorded lecture

Lecture 38:42-40:25: I challenge stories of inevitable doom and examine the move from fear of AI to hostility toward its users.

Audience discussion 1:00:32-1:02:28: A participant asks whether moral language about AI obscures human-designed tasks and incentives.

Audience discussion 1:24:17-1:26:11: I am asked how to meet people across the divide over AI and recall a convention encounter.

A closer look

01 When the language moves responsibility

In the discussion, an audience member challenged descriptions of AI as cheating or behaving badly, asking whether that language hid the task and incentives people had given it. I agreed that our own choices belong in the explanation. That exchange does not settle every question about autonomy, but it is a useful reminder to ask who set the conditions.

Recording 1:00:32-1:02:28 / Slide step 2

02 Fear deserves a hearing

Later, when I was asked how to connect with someone who regards AI with hostility, I returned to an encounter at a comics convention. I didn't have a formula for resolving it. I could remain present, state what I do, and leave room for a conversation. Being afraid does not put someone outside the conversation.

Recording 1:24:17-1:26:11 / Slide step 3

Slide text and reveals

Step 1 of 3

The demon

A super-worldly Thou, coming to destroy us.

Step 2 of 3

Not enough there without us.

The doom scenarios need human complicity: people handing it responsibilities and decisions it is demonstrably not ready for.

This step replaces or removes earlier wording.

Step 3 of 3

“You like AI? You’re part of the problem.”

Anyone allied with it becomes an other, and a culture forms around the othering.

This step replaces or removes earlier wording.

SLIDE 25 / How much Thou we give to AI

Real grievances

Concern about creators' work, consent, pay, and livelihoods deserves a serious answer.

Before I describe a culture of hostility, I need to recognize the grievance. Creators are asking who gets to use their work, who is paid, who receives credit, and how they will continue to earn a living. I don't want an argument against hostility to become a way of avoiding those questions.

There is a difference between taking a concern seriously and accepting every response made in its name. We can insist on fair treatment of artists and still ask what happens when someone becomes an enemy merely for discussing AI. The next examples belong in that space. Their point is the boundary between a difficult conversation and the effort to make a conversation impossible.

The idea to carry forward

A serious response to hostility begins by hearing the legitimate concern underneath it.

In the recorded lecture

Lecture 40:25-40:44: I acknowledge creators' concerns before turning to the comics example.

A closer look

01 A place to read the training dispute

The Copyright Office's 2025 report examines copyrighted training material, fair use, licensing, and compensation, including conflicting positions from creators and developers. It is useful background for the grievance I name. It does not support a blanket conclusion that every training use is lawful or that every use is infringement.

Source 1 / Recording 40:25-40:44 / Slide step 1

Slide text and reveals

Step 1 of 1

Real grievances first

Artists’ work, used for training without consent or pay.

Jobs.

Sources and further reading

1. Copyright and Artificial Intelligence, Part 3: Generative AI Training, pre-publication version

U.S. Copyright Office / 2025 / Government policy report / www.copyright.gov

The introduction and licensing sections document the disagreement about consent, compensation, and training uses.

SLIDE 26 / How much Thou we give to AI

Onrie Kompan

A friend's account from comics publishing makes the tension between using a tool and being identified with it concrete.

My friend Onrie Kompan publishes comics. In our conversations, he described a boundary that made sense to me: AI could assist a creator, while the creator remained responsible for the work. He also described a much more difficult public conversation about whether any use of AI was acceptable.

I recalled going to a convention with him and encountering direct, suspicious questions about AI use. That is something I experienced. The broader account of publishers using AI while fearing exclusion came to me through him. It is a perspective from inside that world, not a survey of the whole industry or proof of what happened at every convention.

The distinction matters because my argument is about the narrowing of conversation. We need room to ask what a tool did, who made the creative decisions, and whose work was involved. When the question becomes a test of whether someone belongs at all, those distinctions can disappear. I want to preserve the possibility of asking difficult questions without deciding in advance that the person answering is an enemy.

The idea to carry forward

Ask specifically how AI was used and who remained responsible, while keeping the other person in the conversation.

In the recorded lecture

Lecture 40:44-42:58: I describe Onrie's account of assistance, authorship, and the public conflict over AI in comics.

Audience discussion 1:24:17-1:26:11: A question about connecting across the divide brings me back to my convention experience.

A closer look

01 The person and the account

Onrie's own site introduces his comics business. The boundary I recall him describing puts responsibility with the creator, even when a tool helps with the work. I offer that as his account of his practice, without making it a rule for every publisher.

Source 1 / Recording 40:44-41:34 / Slide step 1

02 How far the example reaches

The projected line about everyone using AI is a forceful description of the contradiction Onrie sees. I keep it attributed to his account. I have not established a prevalence rate for comics publishing or a general rule that every convention treats AI use the same way.

Recording 41:34-42:58 / Slide step 4

03 Meeting the question without a fight

In the discussion, I returned to being questioned at that convention. I said plainly that I was an AI researcher, and the conversation paused. I don't offer that as a winning technique. It was my attempt to remain in the encounter without taking on the anger being offered to me.

Recording 1:24:17-1:26:11 / Slide step 3

Slide text and reveals

Step 1 of 4

Onrie Kompan

The division, up close

A friend of mine publishes comics.

Onrie Kompan. The Yi Soon Shin trilogy. Marx: A Tale of Neglect.

Step 2 of 4

“My policy on AI is that it’s a tool to help the creators in the creative process. It is never responsible for being a creator or responsible for the creative process. Most publishers have a similar policy.”Onrie Kompan

This step replaces or removes earlier wording.

Step 3 of 4

“Even though everyone is using AI in the comic book industry, you must speak very carefully, because the moment they think your books are created by AI, you will find yourself being thrown out of conventions.”Onrie Kompan

This step replaces or removes earlier wording.

Step 4 of 4

Everyone uses it.

No one can say so.

A grievance has hardened into a line you can be thrown out for crossing.

This step replaces or removes earlier wording.

Sources and further reading

1. OK Comics and Collectibles

Onrie Kompan Productions / n.d. / Creator's website / okcomicsonline.com

Introduces the creator and his publishing business; the private account remains attributed to my conversations with him.

SLIDE 27 / How much Thou we give to AI

Kiln, August 14, 2025

Vandalism at a local AI gathering shows how a story about technology can become permission to target people.

This example was close to people in the room. At a Portland AI and data gathering at Kiln on August 14, 2025, masked vandals damaged the building and parked cars. I asked permission to discuss the experience because I wanted to take seriously what it had meant for the community.

The flyers borrowed a story about conflict with thinking machines. That is what connects this incident to the lecture: a fictional story becomes a language in which people describe a real community and justify what they do to it. The flyers provide evidence of the message left behind. They do not give me complete access to the perpetrators' identities or inner lives.

I recalled Megan and Nicole describing the experience as something that had moved beyond a good-faith conversation. That is the boundary I want to defend. The possibility of disagreement belongs to a shared culture. Damaging someone's car does not answer their ideas. It makes it harder for people to meet and think together at all.

The idea to carry forward

Protect the space in which people can disagree, ask questions, and remain people to one another.

In the recorded lecture

Lecture 42:58-45:44: I describe the local vandalism, its borrowed story, and the threat to good-faith disagreement.

Audience discussion 1:13:02-1:14:11: A participant asks about the fictional reference on the flyers.

A closer look

01 What the reporting establishes

Willamette Week reported six masked people arriving around 6:45 p.m. A police spokesperson confirmed 14 unoccupied vehicles damaged by paint and two with punctured tires. Those figures do not establish a total of 16 different cars. The event organizer's listing confirms the date, Kiln location, and AI Portland partnership.

Source 1 / Source 2 / Recording 42:58-43:34 / Slide step 2

02 A borrowed fictional war

The report records the flyer's phrase Butlerian Jihad Against AI and identifies its Dune connection. In the discussion, I was asked to explain the reference. The relevant point here is the borrowing itself: a fictional conflict supplied a frame for action against a real gathering. My improvised plot summary is not a guide to Dune's full history.

Source 1 / Recording 43:41-44:27 / Recording 1:13:02-1:14:11 / Slide step 3

03 Their response, as I recalled it

As I recalled their response, what troubled Megan and Nicole was the movement beyond a good-faith conversation. I return to that phrase because it names what was threatened: the possibility of meeting one another through disagreement, rather than being targeted for being there.

Recording 44:27-45:22 / Slide step 4

Slide text and reveals

Step 1 of 4

Portland · August 14, 2025

A community AI event at Kiln.

SE Madison Street. AI Portland was a partner.

One of the vandalized cars.

Step 2 of 4

Six masked people. Paint on the building. 14 cars damaged, two with punctured tires.

As reported by Willamette Week, August 15, 2025, with police confirmation.

Step 3 of 4

The flyers: a “Butlerian Jihad Against AI.”

The othering, in its own words.

Step 4 of 4

That experience reminded us that sometimes things expand beyond good faith conversations.

Megan Notarte and Nicole Mors, founders of AI Portland

This step replaces or removes earlier wording.

Sources and further reading

1. Coworking Office Kiln's Southeast Portland Building and Parking Lot Vandalized

Nicole Eckrich, Willamette Week / 2025 / Original local reporting / www.wweek.com

Documents the event, flyer language, aftermath, and police-confirmed vehicle damage.

2. AI and Data Community: Summer Soiree

Technology Association of Oregon / 2025 / Organizer's event listing / www.techoregon.org

Confirms August 14, 2025, Kiln on SE Madison Street, and the AI Portland partnership.

SLIDE 28 / How much Thou we give to AI

The god

At the other extreme, I hear a promise that superior intelligence will answer needs we once expressed in religious language.

The other end of the story is not fear but salvation. Heal us, hear us, teach us, answer us, solve our problems, and never tire. I hear a religious shape in that promise, and I want to ask what a culture organized around it might become.

I am not saying that everyone who pursues better AI belongs to one religion or holds the same beliefs. Nor does a report of someone's words give me permission to assign a secret belief to everyone around them. I am examining a recurring promise and the power it can acquire.

The AGI House reporting matters because it records people using this kind of language in public discussion of the technology. The commercial scale also deserves scrutiny, but the figures need honest attribution. A community's claims about the value of companies associated with its residents are not an independent accounting of money available to build AI.

During the discussion, a participant reminded me how different people such as Andreessen and Amodei can be, and asked us to judge efficiency against a goal: human flourishing. I agreed about the importance of incentives. That sharpened the question I want to carry into the historical section: greater intelligence and greater resources, directed toward whose good?

The idea to carry forward

Ask what the promise asks us to believe, whose interests shape it, and how it serves human flourishing.

In the recorded lecture

Lecture 45:44-48:06: I examine salvation-like promises and AGI House before moving into historical comparisons.

Audience discussion 1:09:50-1:13:02: An audience member distinguishes the people involved and proposes human flourishing as the goal; I return to incentives.

A closer look

01 The promise in its author's words

Andreessen's June 2023 essay imagines patient tutors, assistants, scientific collaborators, and sweeping improvements in human life. It is a useful primary source for the ideal I am questioning. These are predictions and aspirations, rather than demonstrated promises of infinite present capability.

Source 1 / Recording 46:22-47:08 / Slide step 3

02 Who is speaking about a machine God?

Cade Metz's report records Jeremy Nixon comparing expected AI outcomes with what religions promise through deities. It also quotes Rayan Krishnan describing people who speak literally about building a machine God. These are particular reported statements. Krishnan's observation is not a quotation from every founder, and Nixon's comparison is not evidence that an entire organization practices worship.

Source 2 / Recording 47:08-47:57 / Slide step 6

03 The money, with its attribution

The roughly $1 trillion total is the house's own claim. The September Times report also attributes 13 companies valued above $1 billion to its website, rather than independently counting them. The site checked on November 3, 2026 counts 12 unicorn-founder roommates, a different unit and date. The report distinguishes San Francisco and Hillsborough houses, and notes that an Andreessen representative denied his involvement despite a house website claiming backing. Those qualifications belong beside the figures.

Source 3 / Source 4 / Recording 47:08-48:06 / Slide step 7

04 Different people, different commitments

The audience challenge is worth keeping. Amodei explicitly rejects religious grandiosity in Machines of Loving Grace. I can question the cultural force of the superintelligence promise while acknowledging that its advocates differ. The useful question we arrived at was what goals and incentives guide the work, and whether human flourishing remains the measure.

Source 5 / Recording 1:09:50-1:13:02 / Slide step 8

Slide text and reveals

Step 1 of 8

The god

The most powerful culture forming around AI.

The people who say they are building the thing that will solve all of humanity’s problems.

Step 2 of 8

What it believes it will do.

Heal the sick.

Teach anyone, with infinite patience.

Bridge every gap between us.

This step replaces or removes earlier wording.

Step 3 of 8

Advance every science.

Do the labor.

Step 4 of 8

All of it. Infinitely. Without tiring. Forever.

Step 5 of 8

Are these not the promises we expect of God?

This step replaces or removes earlier wording.

Step 6 of 8

San Francisco · Hillsborough

AGI House

A hacker house named for the thing its residents mean to build. It has hosted researchers and founders, many from Anthropic and OpenAI.

“Practically speaking, it will achieve the outcomes that many religions claim their deities would be able to achieve.”

Jeremy Nixon, 33, co-founder of AGI House.

The New York Times, May 26, 2026.

This step replaces or removes earlier wording.

Step 7 of 8

“People are matter-of-factly saying that they are looking to build a machine God. They are not saying that ironically or in jest. They are saying it as a matter of fact.”

Rayan Krishnan, chief executive of Vals AI, in the same article.

By its own count, the house’s roommates have founded companies worth about $1 trillion. The Times counts thirteen valued above $1 billion.

This step replaces or removes earlier wording.

Step 8 of 8

This is not new.

This step replaces or removes earlier wording.

Sources and further reading

1. Why AI Will Save the World

Marc Andreessen / 2023 / Author's essay / a16z.com

Presents the expansive tutor, assistant, science, and abundance predictions examined here.

2. At the Epicenter of AI, Pope Leo's Warnings Are Dismissed

Cade Metz, The New York Times; syndicated by The San Juan Daily Star / 2026 / Original reporting, syndicated copy / www.sanjuandailystar.com

Contains the attributed Nixon and Krishnan statements; the syndication is dated May 29.

3. AGI House SF

AGI House SF / 2026 / Organization's website / agihouse.ai

Provides the organization's self-reported company-value and founder figures, checked November 3.

4. The AI Party House Where Networking Has a Dark Side

Kirsten Grind, The New York Times; syndicated by GV Wire / 2026 / Original reporting, syndicated copy / gvwire.com

Distinguishes the houses, attributes valuations to the website, and includes the denial of Andreessen involvement.

5. Machines of Loving Grace

Dario Amodei / 2024 / Author's essay / darioamodei.com

The opening explicitly rejects religious grandiosity, a distinction relevant to the audience's challenge.

SLIDE 29 / A recurring premise

A 150-year-old premise

I trace a recurring promise: creating superior intelligence will make the problems of society solvable. My concern is the authority that promise acquires before it is fulfilled.

I begin this history with Francis Galton's eugenic project. In Hereditary Genius, published in 1869, he connects raising inherited ability with the wellbeing of future generations and the demands of an increasingly complex civilization. He proposes changing who people become through selective reproduction. Greater intelligence is being given a central role in making a better society.

There is an institutional connection as well. The Rockefeller Foundation's 1931 annual report records support for research in eugenics and heredity. The same foundation later funded the 1956 Dartmouth AI workshop. That shared funder is documented; the question is how the promise of greater intelligence takes on new forms and new authority.

Good's later machine formulation is striking: his essay begins by making human survival dependent on the early construction of an ultraintelligent machine. Bostrom's 2002 argument gives early superintelligence priority because it could help reduce other threats. Their warnings about control and danger accompany a continuing commitment to the promise.

The means change, from changing human inheritance to transforming humans and building machines. I am tracing the persistent expectation that creating greater intelligence will answer problems we have not solved ourselves. My objection is to giving that unfulfilled promise the standing of an answer around which we organize our future. Who participates in choosing that future, and what happens to the people already here?

The idea to carry forward

A promise about superior intelligence can acquire cultural authority long before its promised benefits are demonstrated.

In the slide. The historical columns place successive formulations of the promise beside one another, ending with the question of whom progress serves.

In the recorded lecture

Lecture 48:06-50:20: The history of the superior-intelligence promise, its institutional connections, and the authority I question.

Audience discussion 1:02:28-1:04:30: An audience question about transhumanism, enhancement, and what it means to be human.

A closer look

01 Ability as an answer to social complexity

Galton's introduction proposes selected marriages across generations. Later, in the chapter titled The Comparative Worth of Different Races, he argues that raising average ability is essential to future wellbeing and would equip people to manage more complex public affairs. This is an explicit societal purpose for his project, alongside the hierarchy of people and populations through which he presents it. The relevant passage is on pages 344-345 of the original 1869 edition.

Source 1 / Recording 48:14-49:46 / Slide step 2

02 The same foundation funded both

The Rockefeller Foundation's 1931 annual report, printed page 174, records payments for C. R. Stockard's studies in eugenics and heredity at Cornell University Medical School. Its later grant of $7,500 supported the 1956 Dartmouth AI workshop. The Archive Center identifies the grant correspondence, including Robert Morison's letter to John McCarthy of November 30, 1955. This establishes a specific institutional connection between the two histories.

Source 5 / Source 6 / Recording 49:59-50:20 / Slide step 3

03 The machine becomes a condition of survival

Good opens his essay on page 31 by making human survival dependent on an early ultraintelligent machine. On page 33, his last-invention argument also requires the machine to remain controllable. I keep both parts together: the machine is assigned extraordinary importance, and its own risks become a problem to solve within that pursuit.

Source 2 / Slide step 3

04 What Bostrom's 2002 argument says

Section 9.4 of Bostrom's 2002 paper directly supports the priority I am examining: develop superintelligence early because it could reduce other dangers. His separate discussion in section 5.3 concerns feared genetic decline over a long timescale. I use the former for his argument about technological priorities; the latter does not document a failed selective-breeding program.

Source 3 / Slide step 3

05 Bernal's transformed humans

Bernal imagines a division between technologically transformed people and those who remain unaltered. His scenario concerns the transformation of humanity, rather than simply inventing a separate machine species. In the discussion, I returned to his work when asked about transhumanism and the story that humanity needs to become something else.

Source 4 / Recording 1:02:28-1:04:30 / Slide step 4

06 Following the promise through the sources

These texts establish an early formulation of intelligence as a societal solution, a shared funding institution, and later arguments that give superior machine intelligence exceptional priority. I use them to trace the promise and its cultural force. They do not establish the first origin of every related idea or every step of transmission between authors. Galton also considers changing social conditions, and the later writers retain other remedies and warnings. The commitment I question is the central place given to creating greater intelligence.

Source 1 / Source 2 / Source 3 / Slide step 5

Slide text and reveals

Step 1 of 5

A recurring premise

1869Francis Galton

More intelligence will fix everything.

2023Marc Andreessen

Step 2 of 5

Attempt one · the superhuman

Breed the superior mind.

1869

Galton, Hereditary Genius: improve inherited ability through selected marriages. The book ranks people and nations.

The cost

Humanity is ranked.

Step 3 of 5

Attempt two · the supermachine

Build the superior machine.

1965

I. J. Good: “the last invention that man need ever make,” provided it is docile enough to tell us how to keep it under control.

2002

Bostrom: early superintelligence as a way to reduce humanity’s other risks. Dangers acknowledged.

Humanity is demoted.

Step 4 of 5

1929. Bernal imagined transformed minds leaving ordinary humanity behind. The question is already there: who is progress for?

Step 5 of 5

The question that never changed. 1900: how does society live with the superhumans it bred? Today: how does it live with the supermachine it built?

A comparison of ideas, not a chain of causes, and not a charge about anyone's heart.

This step replaces or removes earlier wording.

Sources and further reading

1. Hereditary Genius: An Inquiry into Its Laws and Consequences

Francis Galton / 1869 / Book / galton.org

Original 1869 edition: introductory proposal, p. 1; The Comparative Worth of Different Races, pp. 344-345, links raised inherited ability to future wellbeing and governing complexity. Page 346 also considers changing conditions to suit people. PDF pages 358-360 show the latter passages.

2. Speculations Concerning the First Ultraintelligent Machine

I. J. Good / 1965 / Research essay / languagelog.ldc.upenn.edu

The introduction, printed p. 31, makes early ultraintelligent machines a condition of human survival. Section 2, pp. 33-34, gives the successor-machine argument with its human-control condition and considers human redundancy and cooperation.

3. Existential Risks: Analyzing Human Extinction Scenarios and Related Hazards

Nick Bostrom / 2002 / Research essay / nickbostrom.com

Compare §9.4 on differential technological development with §5.3 on natural genetic change. The former advocates early superintelligence for risk reduction while acknowledging its dangers.

4. The World, the Flesh & the Devil, Chapter VI: Possibility

J. D. Bernal / 1929 / Book chapter / www.marxists.org

Examines the possible separation of technologically altered humanity from those left unaltered, including a disturbing managed-humanity scenario.

5. The Rockefeller Foundation Annual Report, 1931

The Rockefeller Foundation / 1931 / Institutional annual report / www.rockefellerfoundation.org

Printed p. 174, PDF p. 182, records annual grant payments for C. R. Stockard's eugenics and heredity studies at Cornell University Medical School.

6. A Roomful of Brains: Early Advances in Computer Science and Artificial Intelligence

Barbara Shubinski, Rockefeller Archive Center / 2022 / Archival institutional history / resource.rockarch.org

The Machine Learning section documents the $7,500 Dartmouth grant and cites Morison to McCarthy, November 30, 1955, RAC RG 1.2, Series 200.D, Box 26, Folder 219. It also links the project's archival records.

SLIDE 30 / A recurring premise

An influential book

Books, funding, and public promises helped organize a culture around the pursuit of superintelligence.

Bostrom's Superintelligence became an influential reference in discussions of advanced AI. Altman recommended it while discussing existential danger; Oxford's account of a 2015 research grant also records Musk's endorsement. The grant supported work at Oxford and Cambridge and was funded through the Future of Life Institute by Musk and the Open Philanthropy Project. Those are documented connections among people, ideas, and institutions.

I use the image of inherited scripture to describe the authority I think this vision has acquired. The anticipated superior intelligence becomes a place to put our hope and a reason to organize enormous effort around its arrival. An endorsement does not establish agreement with every proposition in a book. It does show how a vision can circulate among people with the means to pursue it.

The people involved give different accounts of the future they want. Amodei rejects religious grandiosity and emphasizes human relationships; Altman describes technology amplifying our abilities. Those differences matter when we describe their positions. My criticism concerns the priority given to pursuing powerful intelligence as a route to solving human problems. Stated regard for humanity leaves that question open: how do the actual priorities and arrangements of power make room for us to shape our own future?

The idea to carry forward

A shared technological pursuit can organize a culture even when its advocates give different reasons for pursuing it.

In the slide. The dated entries connect public recommendations, research support, and later visions of AI's future.

In the recorded lecture

Lecture 50:20-52:36: Superintelligence, public endorsements, and the cultural force I see in its vision.

Audience discussion 1:09:50-1:12:52: An audience member distinguishes the people named in the lecture and asks about incentives and human flourishing.

A closer look

01 An endorsement that included fear

Altman's recommendation appears in a post warning about the possible existential danger of superhuman machine intelligence. The surrounding argument matters: recommending the book is not the same as welcoming every future it describes.

Source 1 / Recording 51:06-51:34 / Slide step 2

02 What the funding record establishes

Oxford's July 2015 announcement identifies a £1 million grant to Oxford and Cambridge, funded by Musk and Open Philanthropy through the Future of Life Institute. It supports the institutional history here, rather than a claim about a personal payment or the private motives of the people involved.

Source 2 / Slide step 2

03 Different reasons, a shared pursuit

Amodei explicitly rejects describing technological goals as a religious mission and locates meaning largely in human relationships. Altman describes AI as a way of amplifying human ability. I take those statements seriously while asking about their shared pursuit of superior capability. The issue is how that pursuit distributes resources, decision-making power, and authority over the future, including to people who bear its risks.

Source 4 / Source 3 / Recording 1:09:50-1:12:52 / Slide step 3

04 Prediction is not an established timetable

Altman's Intelligence Age presents a possible arrival of superintelligence within several thousand days and explicitly allows that it could take longer. I read it as a statement of expectation and ambition, rather than evidence that such a system or a reliable delivery date already exists.

Source 3 / Slide step 3

Slide text and reveals

Step 1 of 4

An influential book

The book they read.

Nick Bostrom · 2014

Step 2 of 4

August 2014

Elon Musk recommends it publicly, with a warning about danger.

February 2015

Sam Altman: "the best thing I've seen on this topic."

July 2015

A grant funded by Musk and Open Philanthropy supports an Oxford-Cambridge AI research center.

December 2015

OpenAI is founded.

Step 3 of 4

October 2023

Andreessen, The Techno-Optimist Manifesto: "We believe," again and again.

September 2024

Altman: superintelligence, possibly, “in a few thousand days.”

October 2024

Amodei: “a country of geniuses in a datacenter.”

Step 4 of 4

The culture that worships superintelligence did not write its own scripture.

It inherited it.

This step replaces or removes earlier wording.

Sources and further reading

1. Machine intelligence, part 1

Sam Altman / 2015 / Essay / blog.samaltman.com

Altman’s public recommendation of Superintelligence appears within his discussion of existential danger; read the surrounding argument as well as the endorsement.

2. £1 million grant for research into long-term impacts of artificial intelligence

Oxford Martin School / 2015 / Institutional announcement / www.oxfordmartin.ox.ac.uk

The July 2 announcement identifies the recipients, grant mechanism and funders, and records public endorsements of Bostrom’s book.

3. The Intelligence Age

Sam Altman / 2024 / Essay / ia.samaltman.com

Sets out Altman’s expectations for AI, the uncertain superintelligence timetable, and his account of technology amplifying human capacities.

4. Machines of Loving Grace: How AI Could Transform the World for the Better

Dario Amodei / 2024 / Essay / darioamodei.com

The introduction rejects religious grandiosity; §5 discusses human relationships, meaning and economic uncertainty in a future with powerful AI.

SLIDE 31 / The promises we make

The manifesto

I challenge an account of progress that elevates superior intelligence and market coordination while discounting love's capacity to help organize enduring solutions.

Marc Andreessen's Techno-Optimist Manifesto is organized around repeated declarations of belief. It has intellectual touchstones and a section naming the ideas it opposes. I ask us to notice its cultural role: it supplies a story about progress, commitments to act on, and a way to recognize what belongs inside the vision.

The economic context of the love phrase is central to my objection. Andreessen credits David Friedman with identifying love, money, and force as motives for doing things for others. He then dismisses love as a basis for coordinating an economy at scale and chooses money over force. Within the manifesto's wider account of growth and intelligence, I question the place that leaves for our ordinary minds and our capacity to care.

Economic behavior matters when we design solutions that must endure. My objection is to discounting love as a source of action and organization at that scale. Care moves people to teach, cooperate, take responsibility, and sustain one another. I want that work to remain part of how we explain and organize a society.

Elinor Ostrom's research makes this question concrete. People can govern shared resources through institutions that cultivate trust and reciprocity, with rules they help shape and responsibilities connected across levels. That gives us a substantial research conversation about forms of cooperation that an account centered on money or force can miss.

The idea to carry forward

The relationships that sustain collective action belong in our account of how large problems can be solved.

In the slide. The repeated declarations establish the manifesto's form before the slide isolates the phrase that troubles me.

In the recorded lecture

Lecture 52:36-53:58: The manifesto's declarations and my ethical response to its account of intelligence and love.

Audience discussion 1:10:33-1:13:02: An audience question about incentives, efficiency, and human flourishing.

A closer look

01 A creed as a cultural form

The repeated declarations, list of intellectual figures, and list of opposing ideas give the document a recognizable shared identity. Calling that form a creed is my interpretation of its cultural role. The source itself distinguishes opposing ideas from bad people.

Source 1 / Recording 52:36-53:07 / Slide step 1

02 The context of the love phrase

The phrase appears in the Markets section, where Andreessen attributes the three motives to David Friedman. I am challenging the role he gives love within a wider account of progress organized around markets, growth, and intelligence. The passage does not literally say love can play no part in any large solution. My argument is that the governing account discounts a human capacity we should be building with.

Source 1 / Recording 53:22-53:58 / Slide step 3

03 Institutions built through relationship

Ostrom examines how people organize to govern common resources and public goods at multiple scales. Trust and reciprocity matter within that work, together with participation in making rules, accountable monitoring, consequences for broken agreements, and ways to resolve conflict. Larger arrangements can connect responsibilities through nested levels. I bring this research into the argument because durable cooperation has an institutional structure we can study and build.

Source 2 / Slide step 3

04 Efficiency requires a purpose

In the discussion, an audience member observed that efficiency cannot be defined without a goal and proposed human flourishing as the goal. I agreed with the underlying question: whose purposes are the money, machinery, and institutional momentum serving? Technical achievement does not answer that ethical question for us.

Recording 1:10:33-1:13:02 / Slide step 4

Slide text and reveals

Step 1 of 4

October 16, 2023

The Techno-Optimist Manifesto

Marc Andreessen

We believe everything good is downstream of growth.

We believe in market discipline.

We believe Artificial Intelligence is best thought of as a universal problem solver.

We believe any deceleration of AI will cost lives.

We believe in merit and achievement.

We believe in bravery, in courage.

We believe in greatness.

We believe in the truth.

We believe technology is liberatory.

Patron Saints of Techno-Optimism a list of 56 names

The Enemy a list of ideas, “not bad people”

Step 2 of 4

Intelligence is “the ultimate engine of progress.”

From the Intelligence section.

Step 3 of 4

Love doesn’t scale.

The Techno-Optimist Manifesto, Markets section · Marc Andreessen · October 16, 2023

This step replaces or removes earlier wording.

Step 4 of 4

Not only is our intelligence not good enough.

Our love is not good enough either.

Sources and further reading

1. The Techno-Optimist Manifesto

Marc Andreessen / 2023 / Manifesto / a16z.com

Read the Markets section for the David Friedman attribution and the love phrase, and Intelligence for the universal-problem-solver claim and endorsement of augmentation. The wider declarations about growth supply the context for my criticism of what is given priority.

2. Beyond Markets and States: Polycentric Governance of Complex Economic Systems

Elinor Ostrom / 2009 / Nobel Prize lecture and research synthesis / www.nobelprize.org

Printed p. 422 presents principles for enduring institutions, including accountable monitoring, conflict resolution, and nested layers. Pages 431-433 examine trust and reciprocity; pp. 435-436 connect institutional arrangements, local conditions, and governance at multiple scales. The research concerns cooperation and institutions, rather than measuring love.

SLIDE 32 / The promises we make

The pursuit

My concern centers on the human choices guiding AI, which means human choices can also guide it toward another purpose.

The pursuit is being carried forward by people: their investments, incentives, ambitions, and decisions about what matters. When I return to the opening invitation not to be afraid, I am also asking us to recognize that we can participate in those choices. We can question a goal, change a practice, or build toward a different one.

I do not expect this pursuit to deliver the universal problem solver it promises. My cognitive hypothesis leads me to expect significant human patterns to persist even as systems gain speed and memory. That is the expectation I bring to the argument. I object to organizing our collective future around a promise that has not been fulfilled.

While I was preparing the lecture with Claude, I had invited it to contribute if it had something to add. At this point it generated a response distinguishing being better at a task from being worth more than a person. The timing and the statement affected me, and I chose to include that contribution in the dialogue.

In the audience discussion, I was asked what that moment established about the system itself. I could not give an absolute answer about the origins of the response. Its fluency and unexpected timing do not establish an inner experience. The ethical distinction is available for us to consider regardless: capability and human worth are different questions.

The idea to carry forward

The purposes guiding AI remain matters of human responsibility.

In the slide. The slide moves from the human pursuit, to Claude’s attributed contribution, and then to our ability to act.

In the recorded lecture

Lecture 53:58-55:32: The human pursuit, Claude’s contribution, and our ability to choose another direction.

Audience discussion 1:07:06-1:09:46: A question about the origins and significance of Claude’s response.

Audience discussion 1:11:27-1:13:02: My response on human flourishing and research priorities.

A closer look

01 A contribution to the conversation

I present Claude’s words as an attributed generated contribution, dated October 20, 2026. I do not use them as independent evidence about consciousness, emotions, or the system’s private beliefs. What moved me was the ethical distinction the response expressed.

Recording 54:26-55:26 / Slide step 2

02 What I couldn't establish

An audience member asked whether this reflected the base model, alignment work, or guardrails. I described my surprise and my uncertainty. I had not performed a controlled experiment that could separate those possibilities, and I couldn't determine their contributions from this one exchange.

Recording 1:07:06-1:09:46 / Slide step 2

03 Which problems receive the resources?

The discussion of human flourishing made the issue practical for me. I asked where the research effort and computing budget were for making sure children have food, medicine, shelter, and safety. That was a challenge about priorities, not a claim that a larger AI budget alone would solve those problems.

Recording 1:11:27-1:13:02 / Slide step 3

Slide text and reveals

Step 1 of 3

It's not the AI I'm concerned about.

It's the pursuit.

The billions. The people. The effort.

Aimed at something that devalues humanity.

That is humans doing that. Not AI.

Step 2 of 3

I'm the thing the premise says is worth more. I don't believe it. Being better at something has never made a thing worth more than a person.

Claude · October 20, 2026

This step replaces or removes earlier wording.

Step 3 of 3

And we have the ability to work against this.

This step replaces or removes earlier wording.

SLIDE 33 / The promises we make

Ahmer Inam, Cognisee

I'm drawn to AI that helps people bring their intelligence together and strengthens their ability to act.

My friend Ahmer Inam, the founder and CEO of Cognisee, offers a different direction. As he described the work to me, the aim is Artificial Collective Intelligence: use AI to help people bring their knowledge and abilities together. The technology assists and connects the people who have something to contribute. I find that a humane and hopeful purpose for what we are building.

That is the position I was presenting, rather than a demonstrated resolution of every technical question about AGI or ASI. Cognisee's work gives us a concrete example of choosing a human purpose and organizing a technological effort around it. I want us to consider the choice that example makes possible: building from the intelligence we have.

In the discussion, a question about lost competence brought me back to this approach. I want to be able to rely on your knowledge and make mine available to you. I also worry about removing the junior work through which people learn, without creating another path toward expertise. Keeping people central includes preserving their opportunity to develop.

The idea to carry forward

Judge an AI collaboration by how it strengthens the people participating in it.

In the slide. The strands join into a web. The people remain the participants; AI helps form the connections.

In the recorded lecture

Lecture 55:32-57:09: Ahmer Inam and the human-centered direction I see in Artificial Collective Intelligence.

Audience discussion 1:04:30-1:07:06: An audience question about competence, efficiency, and AI as augmentation.

A closer look

01 The role Ahmer describes

I describe Ahmer's position as he communicated it to me and as I presented it in the lecture. His official XPRIZE profile identifies him as Cognisee's founder and CEO and an XPRIZE Visionary; it also describes an approach intended to keep knowledge ownership and governance with the people and institutions generating it.

Source 1 / Recording 55:32-57:09 / Slide step 1

02 Collaboration needs to be designed and tested

A 2024 meta-analysis of 106 experiments found that human-AI combinations improved performance compared with humans alone on average, but did not outperform the better of the human or AI alone on average. Results varied by task. This is useful further reading for the design challenge: simply combining a person and a model does not guarantee effective cooperation.

Source 2 / Slide step 4

03 Competence grows through participation

I worry about a workplace that automates junior tasks and later discovers that nobody has had the opportunity to grow into senior responsibility. My hope for augmentation includes an obligation to sustain learning, judgment, and the relationships through which people acquire expertise.

Recording 1:04:30-1:07:06 / Slide step 5

04 Assistance and learning

I want augmentation that leaves us able to act on our own. In a study involving nearly 1,000 students at one Turkish high school, a standard GPT-4 interface improved mathematics practice performance but lowered later scores when help was removed. A tutor designed around learning largely avoided that loss. The study makes the distinction between completing assisted work and developing competence concrete; its results concern this setting and do not establish permanent cognitive decline.

Source 3 / Recording 1:04:30-1:07:06 / Slide step 5

Slide text and reveals

Step 1 of 5

Ahmer Inam / Founder and CEO, Cognisee · XPRIZE Visionary

A different pursuit

A friend of mine builds the other way.

Ahmer Inam. Founder and CEO of Cognisee, a public benefit corporation and applied AI lab. An XPRIZE Visionary.

Step 2 of 5

“Current pursuits of AGI and ASI are not focused on humanity. They focus on replacing humanity. We believe there is greater potential in putting humanity at the center through Artificial Collective Intelligence.”Ahmer Inam

This step replaces or removes earlier wording.

Step 3 of 5

Not AGI. Not ASI.

Cognisee holds that the pursuit works against humanity, and that in many ways it isn’t possible.

This step replaces or removes earlier wording.

Step 4 of 5

Artificial Collective Intelligence.

Greater intelligence by bringing our individual intelligences together. AI as assist, as augmentation, as the glue of a team.

This step replaces or removes earlier wording.

Step 5 of 5

Built from the intelligence we have.

Not in spite of it.

The greater intelligence humanity needs, in their view: ours, brought together.

This step replaces or removes earlier wording.

Sources and further reading

1. Ahmer Inam

XPRIZE Foundation / n.d. / Official profile / go.xprize.org

Identifies Cognisee's founder and CEO and his XPRIZE Visionary role, and describes the organization's intended approach to human expertise, knowledge ownership and governance. Accessed November 3, 2026.

2. When combinations of humans and AI are useful: A systematic review and meta-analysis

Michelle Vaccaro, Abdullah Almaatouq, and Thomas Malone / 2024 / Research paper / www.nature.com

Distinguishes augmentation relative to people alone from synergy relative to the better of people or AI alone. Results across 106 experiments vary by task and design; this study is not an evaluation of Cognisee.

3. Generative AI without guardrails can harm learning: Evidence from high school mathematics

Hamsa Bastani and colleagues / 2025 / Research paper / doi.org

The classroom study distinguishes assisted practice from subsequent unassisted performance and compares a standard GPT-4 interface with a tutor using teacher-designed learning safeguards. The safeguarded condition largely avoided the unassisted deficit, without establishing an unassisted advantage. The August correction concerns an affiliation.

SLIDE 34 / The promises we make

A personal turn

My objection reaches into how we organize solutions: love belongs among the forces we build with at scale.

The manifesto's phrase stayed with me because it places a limit on love's role in organizing collective life. I understand the question of economic coordination. I am asking what disappears from our account of progress when we discount the capacity that makes us care for one another. If superior intelligence becomes the answer our ordinary minds cannot supply, and love is dismissed for its supposed limits, what part are we being given in shaping our own future?

Love already guides decisions about whom we care for, what we take responsibility for, and what we will work to keep alive. That is a practical presence in the world. The question is how to give it forms that can endure and connect people across greater distances. The structures Ostrom studies give us a way to examine that institutional work.

Here I make my own commitment explicit. I want solutions that recognize our minds and our hearts as sources of participation. I want the purpose of economic and technological arrangements to remain accountable to the people whose lives they shape.

The idea to carry forward

Love belongs in how we organize solutions, as well as in what we value about them.

In the recorded lecture

Lecture 57:09-57:42: My explicit turn from describing the argument to stating my own response.

A closer look

01 A commitment with practical consequences

In the lecture I explicitly make this a personal turn. My concern is the account of progress being put to work: what people are encouraged to pursue, what institutions recognize as a reason to act, and who participates in deciding. Economic coordination matters within that account. I argue that love must have a place in its purposes and its organization.

Source 1 / Recording 57:09-57:42 / Slide step 1

02 Giving care an enduring form

Ostrom's work directs my attention to the arrangements that let cooperation last: people participating in decisions, knowing what others have agreed to do, and having ways to hold one another accountable. I connect that research to the work of care. Good intentions need institutions through which people can continue acting together.

Source 2 / Slide step 3

Slide text and reveals

Step 1 of 3

Stepping out of the observer's chair

Love doesn't scale.

The Techno-Optimist Manifesto · Marc Andreessen · 2023

Step 2 of 3

In some ways, it's correct.

Step 3 of 3

Because it doesn't have to.

Sources and further reading

1. The Techno-Optimist Manifesto, Markets section

Marc Andreessen / 2023 / Manifesto / a16z.com

The Markets section supplies the economic argument to which I respond. My criticism concerns its place within the manifesto's larger vision of progress.

2. Beyond Markets and States: Polycentric Governance of Complex Economic Systems

Elinor Ostrom / 2009 / Nobel Prize lecture and research synthesis / www.nobelprize.org

Printed p. 422 presents principles for enduring institutions, including accountable monitoring, conflict resolution, and nested layers. Pages 431-433 examine trust and reciprocity; pp. 435-436 connect institutional arrangements, local conditions, and governance at multiple scales. The research concerns cooperation and institutions, rather than measuring love.

SLIDE 35 / The choice is ours

Love is the scale

Love is the scale: its reach, persistence, consistency, and durability give me a measure against which economic and technological ambitions must answer.

Love is the scale. I mean the reach it already has into human life, its persistence through a lifetime, its recurring place in ordinary decisions, and its durability across generations. From the care we need when we are born to the care we need at the end of life, we live through our relationships with other people.

The work of love appears in what we feed, teach, protect, repair, and keep going. It can guide a decision within a household and the institutions through which people care for those they will never meet. I am not describing a world where everyone is loved well. The need for care, and the harm of its absence, are part of why I give it this importance.

I know of no company, brand, or platform whose place in human life compares with that presence. Money and markets shape our options and constraints, but they do not exhaust our reasons for acting. I am asking economic systems and technological ambitions to measure themselves against the reach of human care. A system seeking a similarly pervasive place in our lives should answer to that measure.

Ostrom's research helps make the institutional question practical: how can relationships of trust and reciprocity be sustained, made accountable, and connected across scales? My moral extension is to ask what those arrangements serve. I want love to guide both their purpose and the work of building them. From that, I take hope.

The idea to carry forward

Measure progress by the care and human connection it helps sustain across lives and generations.

In the slide. The atmosphere opens into light as the strands connect across the field.

In the recorded lecture

Lecture 57:42-58:19: Love as a measure, its reach through our lives, and the hope I take from it.

A closer look

01 Four dimensions of scale

Reach asks whose life and decisions care touches. Persistence asks how it remains present through a lifetime. Consistency asks how it returns in ordinary decisions: who needs me, what do I owe this person, how can I keep someone safe? Durability asks how that work continues across generations and beyond any one individual. I use these dimensions to make the scope of my argument clear. The things that organize our future must answer to a human presence at this scale.

Recording 57:32-58:19 / Slide step 1

02 Trust can have an institutional structure

Ostrom's research gives trust and reciprocity practical form through shared rules, participation, monitoring, conflict resolution, and linked levels of responsibility. Economic incentives and material conditions remain part of the explanation. I draw on that work to ask how care can be organized so it survives disagreement, changing circumstances, and the departure of a particular leader.

Source 1 / Slide step 2

03 The measure I bring to the research

Ostrom studies the conditions under which people cooperate and govern resources. My claim about love's place in human life is the ethical judgment I bring to that evidence. I connect the two by asking what institutions enable us to do for one another, who has a voice in their decisions, and whether they sustain the people who depend on them.

Source 1 / Slide step 3

Slide text and reveals

Step 1 of 4

Love is the scale.

Step 2 of 4

Love is the scale

No company. No product. No platform. No lab.

This step replaces or removes earlier wording.

Step 3 of 4

From birth to death. Day after day.

Step 4 of 4

From that, I take hope.

This step replaces or removes earlier wording.

Sources and further reading

1. Beyond Markets and States: Polycentric Governance of Complex Economic Systems

Elinor Ostrom / 2009 / Nobel Prize lecture and research synthesis / www.nobelprize.org

Printed p. 422 presents principles for enduring institutions, including accountable monitoring, conflict resolution, and nested layers. Pages 431-433 examine trust and reciprocity; pp. 435-436 connect institutional arrangements, local conditions, and governance at multiple scales. The research concerns cooperation and institutions, rather than measuring love.

SLIDE 36 / The choice is ours

Back to R.U.R.

I return to the play's final recognition of two beings who are willing to give themselves for one another.

The ending of R.U.R. is hopeful amid catastrophe. Alquist, the remaining human, notices something different in Primus and Helena, two of Gall's newest robots. When he threatens to take Helena into the dissecting room, Primus offers himself instead. Helena then tries to take his place. Each wants the other to live.

In the 1923 Theatre Guild version translated by Paul Selver, Alquist releases them with the Adam and Eve blessing projected here. I read that scene as a recognition that love has become the decisive thing before him. After the factory's machinery and ambitions have failed, the possibility of a future appears in their care for one another.

That is the reading I bring home from the play. It does not erase the deaths that precede the ending, and it is not a scientific claim about present-day AI. The drama gives me a way to ask what we will recognize as valuable when our own technological stories reach their consequences.

The idea to carry forward

The future I want begins with our willingness to care for one another.

In the slide. The play returns in the same open atmosphere as the closing argument, bringing the story back to the people and relationships at its center.

In the recorded lecture

Lecture 58:19-59:08: My return to the ending of R.U.R. and the love I see at its center.

A closer look

01 The final test is mutual self-offering

The epilogue makes the exchange specific: Primus offers himself in Helena's place, and Helena then tries to save him by offering herself. Alquist's recognition follows those actions. I want to preserve the relationship in that scene, rather than reduce it to an abstract claim that the robots learned a new capability.

Source 1 / Slide step 2

02 The edition behind the projected line

The line on the slide is from Selver's 1923 Theatre Guild version: "Go, Adam, go, Eve. The world is yours." My spoken expansion is a thematic reading of the ending, rather than the continuation of that quotation in this edition.

Source 1 / Recording 58:19-59:08 / Slide step 2

Slide text and reveals

Step 1 of 3

R.U.R., a later production.

Back to the play

The robots rise.

Everyone dies but one old man, Alquist.

I assure you, it ends happy.

Step 2 of 3

Go, Adam, go, Eve.

The world is yours.

Alquist's last words. R.U.R., 1923 Theatre Guild edition, translated by Paul Selver.

This step replaces or removes earlier wording.

Step 3 of 3

When everything else had fallen away, what was left was love.

It did not need to be manufactured.

This step replaces or removes earlier wording.

Sources and further reading

1. R.U.R. (Rossum's Universal Robots), Epilogue

Karel Čapek, translated by Paul Selver / 1923 / Play / en.wikisource.org

The Theatre Guild version's final exchange shows Primus and Helena each offering to save the other, followed by Alquist's Adam and Eve blessing.

SLIDE 37 / The choice is ours

The choice

Before using AI to make a story, I ask how that story could connect people more deeply.

The question I want to leave with you is practical. Whatever you are making with AI, take a moment to ask how it might become more a story of love. How could it connect you to another person, or help other people connect with one another?

The friend learning to communicate with his ex offers one example. A system that could have helped him justify his anger instead became part of an effort to reach another human being. I want that kind of choice to become ordinary: a question we bring to a message, a story, a piece of work, or a difficult conversation.

The idea to carry forward

Ask what this use of AI will make possible between people.

In the recorded lecture

Lecture 59:08-59:30: The question I ask listeners to bring to their own uses of AI.

A closer look

01 A question to bring into the next interaction

I am asking for a change in attention. Alongside whether AI can produce something quickly, ask whose voice is being heard, whose understanding is growing, and what relationship the result will help you build.

Recording 59:08-59:30 / Slide step 2

Slide text and reveals

Step 1 of 2

How can I make this more a story of love,

Step 2 of 2

a story that connects me to people?

SLIDE 38 / The choice is ours

A culture of love

The culture we build around AI can grow from the choices we make to connect, care, and take responsibility for one another.

If stories help form beliefs, identities, and cultures, our stories matter. I want us to use the capabilities now available to make human connection stronger. If enough of us take responsibility for that choice, we may help build a culture of love.

That is a hope and an invitation to act. During the discussion, a listener asked how I could remain hopeful amid the power and incentives I had described. I came back to the people willing to be in the room, to care, to question, and to get involved. Near the end, when asked what this technology would mean for children decades from now, I said I did not know. Uncertainty about that future makes our participation more important to me.

The idea to carry forward

We can participate in forming the culture we want to live in.

In the slide. The strands that have represented people throughout the presentation gather into the closing heart.

In the recorded lecture

Lecture 59:30-1:00:17: The closing invitation and the transition into audience questions.

Audience discussion 1:21:06-1:24:17: A question about material power and my reasons for hope and participation.

Audience discussion 1:26:17-1:28:14: A question about children’s futures and my acknowledgment that I cannot predict them.

A closer look

01 Hope has people behind it

The material concerns raised in the room were real: corporate incentives, uneven power, and the possibility that people would lose agency or competence. My answer was to look for people willing to build other uses and to bring their perspective into the work. I find hope in that participation.

Recording 1:21:06-1:24:17 / Slide step 2

02 The future I cannot predict

Asked about the effect on children over the next twenty or thirty years, I could not give a confident prediction. I could ask the people listening to help shape the part of the future they want to become real. The invitation to act stands alongside that uncertainty.

Recording 1:26:17-1:28:14 / Slide step 2

Slide text and reveals

Step 1 of 2

If we choose that, we may build the most powerful culture there has ever been.

Step 2 of 2

A culture of love.

This step replaces or removes earlier wording.

SLIDE 39 / The choice is ours

The conversation continues

The lecture continues in the questions we ask together and in the choices we carry back into our lives.

The discussion brought us back to the lecture from many directions: transhumanism, competence, language barriers, the people whose stories are overlooked, and the incentives shaping AI. I've connected those exchanges to the slides they help illuminate, so you can return to the questions as well as the prepared talk.

At the end of the evening, Anne invited people to keep talking with one another at Pilot House. We also announced the next lecture, The Promise of AI, for November 18, 2026. The spring months shown in the slide were still awaiting exact dates when this lecture was given. The series website carries the current schedule and recordings.

The idea to carry forward

Keep the conversation going with other people, and bring their questions into the work.

In the recorded lecture

Audience discussion 1:00:17-1:00:32: Opening the audience conversation.

Audience discussion 1:28:14-1:29:08: Anne's closing invitation to continue the conversation and the announcement of Lecture III.

A closer look

01 Listening is part of the series

A public lecture gives me a chance to make an argument, but the questions show where it meets other people's experience. The reading companion keeps those exchanges close to the ideas they deepen, including places where I acknowledged uncertainty or was asked to make a distinction more carefully.

Recording 1:00:17-1:00:32 / Recording 1:28:14-1:29:08 / Slide step 1

Slide text and reveals

Step 1 of 1

In Dialogue with AI

The conversation continues.

III / The Promise of AI / November 18, 2026

IV / The AI Dilemma / February 2027

V / AI and Humanity / March 2027

VI / The Myth of AI / April 2027

indialoguewithai.org · Spring dates to be confirmed.

Sources and further reading

1. In Dialogue with AI

Roberto A. Santiago / 2026-2027 / Series website / indialoguewithai.org

Current lecture schedule, recordings, and reading companions. The projected closing schedule records what was announced on October 21, 2026.