THE REASONER

Picture by alphaspirit

UNIVERSITÀ DEGLI STUDI DI MILANO

Milano University Press

THE REASONER
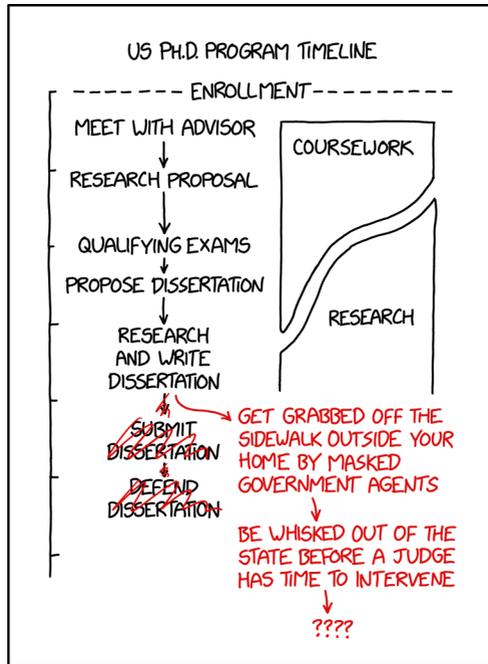
On March 25, 2025, 30-year-old Tufts University student Rümeysa Öztürk, a Turkish national on an F-1 student visa, was arrested by six masked plain clothes agents from the United States Department of Homeland Security (DHS) following the revocation of her visa near her home in Somerville, Massachusetts.

The Reasoner 19(3), July 2025.

Contents

An Interview with Kenneth Aizawa
:: Hykel Hosni

Abstract

Hykel Hosni interviews Kenneth Aizawa, Professor of Philosophy of Science at Rutgers University, Newark.

The interview opens with Aizawa's reflections on the challenges currently confronting US universities under the second Trump administration. It then turns to a discussion of his forthcoming book, *Compositional Abduction and Scientific Interpretation: A Granular Approach*. Following a retrospective on Aizawa's academic trajectory, the conversation concludes with his advice for early-career researchers.

Keywords

Scientific reasoning, abduction, confirmation, philosophy of science, scientific inquiry, scientific practice, US academia.

How to Cite

Hosni, H. (2025). An interview with Kenneth Aizawa. *The Reasoner*, 19(3). https://doi.org/10.54103/1757-0522/28998

Hykel Hosni: You are currently Professor of Philosophy at Rutgers Newark. Staggering news arrives from the US on an almost daily basis. The attack on Harvard by the Trump administration is the most spectacular and recent episode, but it is just one in a very long list. Columbia was stripped of 400 million in federal funding on allegations of antisemitism. Many NIH grants have been terminated. Those are unprecedented political measures against US universities. How does it feel to work in such a climate?

Kenneth Aizawa: It's a level of disbelief, I guess. Part of it is just, *can this be the United States where freedom of speech is under attack*? The rule of law is under attack. One of the engines of our economy – our science, our education– is under attack. And you know, our openness to people from other countries. Harvard might resist all these illegal things, but the message is getting out there that the United States is not a safe country. It's just not.

HH: It's a message Canada is hearing loud and clear

KA: I live in New Jersey, not far from Canada. It's just hard to believe that we could be at a place where we would not be welcome in Canada. In Canada! So, there is a level of disbelief that this is happening. Some things maybe we could resist. On the illegal things, maybe they will be stopped. Maybe they will be stopped by their incompetence. But there are certain things that you can't just put back in the bottle.

HH: Like the anti-immigrant sentiment?

KA: You just can't. It is here and it's not going to go away for

decades. And it's a really big deal. At Rutgers Newark we have lots and lots of international students. In my introductory classes, students will speak ten different languages. We have lots and lots of bilinguals. And you can just see that they might not come. And I can understand that. It isn't to fight anti-Semitism. It's just to harm universities.

HH: It's incredible that the US government is taking action that might not even be legal.

KA: They just *do* things, and then see if it's legal. They do all sorts of things just to drag their feet. The illegal deportations without habeas corpus are just astonishing. You might have seen the picture of the woman in Massachusetts who was just snatched off the street. It looks like that can't be the United States, right?

HH: I read that clinical trials involving non-US institutions are losing funding due to national security concerns, stemming from the idea that COVID originated in a foreign lab. Putting the rather questionable motivation aside, this will have a tangible effect soon on the US public health.

KA: Oh, yes. It's having an impact right now. There were people in these studies with medical implants which you now can't take out. You got someone with a medical implant that needs to be removed. It's just staggering.

HH: Presumably this will also have a long-term impact on how society perceives academics in the US, with more people questioning their relevance, importance, and integrity.

KA: This has been brewing for a decade. Sociologists have measured an uptick in Republican resistance to universities in the US. It was about 2015, 2016 where Republicans began to believe that universities were more of a force for the worse, than for the better.

HH: And the pandemic certainly didn't help.

KA: COVID was politicized and so people decided they didn't want to wear masks and the masking was associated with universities, so that probably contributed. Every now and then when someone finds out I'm a professor, they begin to ask, so do you have a lot of indoctrination there? You know that they're getting this message. And I live in a blue state!

HH: But why tearing down universities?

KA: Maybe they realise that there aren't enough conservative academics to replace all the others. So they're just going to blow it up. You might have heard of this New College in Florida. They fired all the liberal professors who didn't have tenure to make it a very conservative liberal arts college. They introduced sports and changed the curriculum so you know it's an America first kind of curriculum where you're supposed to learn about how great the United States is and things like this. So, they just did that.

HH: But they can't do that for everywhere, right?

KA: So, they're just going to blow up the rest of us.

HH: How far will this Administration go?

KA: As far as they can. They can hurt the universities, but they probably can't destroy us just by cutting off federal funding. In fact most state universities are run by tuition. That's why they're targeting research powerhouses like Harvard.

HH: And Harvard is fighting back.

KA: Harvard is on the ball. After they had this the international students block, within 24 hours or so, they had an extensive 72-page response. I mean, Harvard academics are educated, they've got a lot of money at stake, they've got a lot of resources at stake,

so they can fight back. And it's good that they will.

HH: For US academia, but I guess, it's good for all of us. And on this note of optimism, let's switch to research now. Your latest book *Compositional Abduction and Scientific Interpretation: A Granular Approach* is to be published soon by Cambridge University Press. The readers of The Reasoner will find a very enticing sentence in the opening pages: "My project is to provide a descriptively adequate account of published scientific reasoning."

KA: Some people say science really happens in the lab, and of course it does. But my focus is on published research because science also happens in the publication record. There, you can find experimental work, opinion pieces, review articles, and textbooks. But I think the journal article is a central fixture in scientific research. It plays a distinct role.

HH: You mean the introduction, methods, results and discussion format.

KA: Yes, and that part where you explain your results. That's where you find *abduction*. The part where you say why you have these results, and you explain what's going on. That provides confirmation. I believe that abduction is confirmation, and the engine of confirmation is the interpretation of experimental results That's why I'm focusing in on that and why I look at that.

HH: Is this an idea of trying to understand why scientists believe that the data they have give them a finding?

KA: That is the idea. Why do they take away from their experiments what they take away? That's a core bit of understanding how science works. It may not be all of science, but it's a big part of it.

HH: After all scientists do experiments to form beliefs rather than

knock them down.

KA: To me it is scandalous that for the last say 20 years, there hasn't been a theory that replaces hypothetical deductivism (HD). Almost every philosopher of science has heard of HD methodology. And they also heard that it doesn't work. But they don't have an account of what is going on.

HH: It's not uncommon for philosophers to theorise independently of how science works. What's puzzling is that scientists often buy into a 'textbook story' that differs significantly from what they actually do in practice.

KA: You have to separate out what scientists are doing from what they say they are doing. They may sometimes go together, but sometimes they come apart. But here's part of my story that occurs in chapter four. What scientists do *looks like* hypothetical deductivism. It really does. They have this hypothesis and they figure out some consequences and then they see if the consequences are really there. But in chapter four of my book I describe why I think is not what is going on.

HH: So, what you think it is going on?

KA: They're not deriving some sentences and then checking whether the sentences are true or not. What they're doing is that they have these hypotheses, and they're tracing out the physical consequences, not the logical consequences of those hypotheses in the world. Then they see if the physical consequences have taken place.

HH: So what scientists do looks like HD if one fails to distinguish logical and physical consequences of hypotheses. You recently published a paper on The Reasoner, pointing out that Disconfirmation is not modus tollens.

KA: I show in it that material conditional is not fit for discon-firmation because what scientists are really thinking about is not logical consequences. They are thinking about *physical conse-quences*.

A week or so ago, I hosted a conference at Rutgers with a lot of neuroscientists. When I described this to them, they said, "Yeah, that's what we are doing!" I think philosophers of science haven't given scientists a way of thinking about it. They've given them just a bad characterization of what they do.

HH: So your book puts forward a better one.

KA: That's what I'm trying to do. That's why I look closely at what scientists are doing. I don't just take what *they say* they're doing, because they're just maybe borrowing that from philoso-phers.

HH: I can understand that. They're busy enough with the sci-ence, and certainly we can expect them to do philosophy too! You argue that the practical effect of abduction is confirmation. This differs radically from the (rather mysterious) textbook definition of abduction. How do Peirceans react to your view?

KA: Oh, they're scandalized by it. No one has this view as far as I know. And the Peirce scholars just will not touch it. They just won't tolerate it at all. I get the most indignant responses from them about this, "just need to read more Peirce", they typically say.

HH: It is quite typical of epigones to be very territorial about the meaning of words.

KA: Exactly. Peirceans say abduction is inferring a hypothesis because it's explanatory. And then they build other things on top. Like the thing to be explained has to be surprising.

HH: And what about the epistemologists who work on Inference to the Best Explanation (IBE)?

KA: They just think my account makes no sense. For them, abduction is just defined as IBE. So they're just outraged by seeing abduction as confirmation.

HH: Do you see abductive confirmation quantitatively?

KA: It's qualitative support. And I guess that's why the IBE people don't like. It's not a guarantor of truth. It's just a reason. A scientist says, "Here's an explanation of the data." But then another scientist comes along and says, "Well, wait a second. Here's an alternative explanation." And then they have a "fight" to decide which, if any, is the correct explanation.

HH: This descriptive focus departs from traditional questions in the epistemology of science.

KA: Epistemology of science is concerned with the question of when do you know. I'm not asking that question. I'm interested in understanding the inquiry part of science. I do an experiment, I interpret the results, and then maybe I'll be convinced by that one experiment, but probably not. I'm going to continue to do many experiments, and each little experiment is to generate a little bit more evidence, a little bit more confirmation. But then someone else might come along and say, "No, no, that's not the correct explanation of any of these, the correct explanation is this." That's what I want to try to capture. I want an account of a single experiment because it is the building block.

HH: But what you are saying relates to the epistemology of science.

KA: Of course. To be able to tell when you have enough evidence for knowledge, you must have an account of what a piece

of evidence is.

HH: That's very clear. The core concept in your book is singular compositional abduction. You told us what you think abduction is (not). What about "singular" and "compositional"?

KA: Let's start from 'singular'. It picks up on the idea of a data point. In the book I often talk about what happened: "A current in axon number 15 that took place on one day in the summer of 1949". It's one single thing.

HH: A measurement.

KA: A measurement, exactly. You have this oscilloscope over here and then it's measuring what's going on in the axon over here. And so what you want is that particular action potential say, or that particular action current. And then you want to say what happened in that particular action current. That's the 'singular' aspect.

The answer is that there were some ion movements, let's say. So the ion movements down here explain that individual action current at that time, on that day. That's the *compositional* explanation. And it's compositional because the current in the axon is explained by these lower level movements of ions. So it presupposes that things are in levels, and that things that happen at a lower level determine something at a higher level.

HH: And then you get abduction when you say "this explanation gives me a reason to believe". You emphasize a science-first approach, and you go into the fine detail of experimental results. The examples you discuss in the book however, fall in the camp of traditional, or "theory-driven" science. How about model-based, data-intensive or AI-driven science? There is a lot of interest and promise in it, especially in the biosciences.

KA: One of the things that is growing in neuroscience is artificial neural networks. There's a lot of modeling there. Now, none of the examples I use are models. There's a big difference between Hodgkin-Huxley's fifth paper, which is modeling, and their earlier papers which are experimental. There is just a very different story there. When you're working with artificial models, it's often unclear what connection, if any, they have to the physical world, or even whether they're intended to have one. All of a sudden you just don't get the physical consequences of the story any more. If you're just running a mathematical simulation, you're doing a mathematical thought experiment, and it looks like it's a completely different kind of thing. Maybe you can explain something. But then that sense of explanation is very different.

So I think it would be a mistake to try to carry over my picture to the data-driven sort of science. I wouldn't know how to apply it there. But I'm OK with that.

HH: Science is complicated, it moves on fast, and scientists do many things in many different ways. I guess it would be suspicious if one approach claimed to account for all of science.

KA: I view my book as offering an abductive approach, but I don't have all the answers. The whole book is predicated on there being gaps. That's true for the neural network stuff and big data, but probably more. It is set up not to be the final word, but to be the next thing to try.

HH: Can you tell us about your academic trajectory? What got you interested in philosophy and then in cognitive science?

KA: I guess I got interested in philosophy in high school. Actually from taking a Spanish class. I took the fourth year Spanish and we were reading the Generation of '98. It all seemed very

philosophical to me, and it stimulated my interest and curiosity. I knew when I went off to college that I wanted to major in philosophy

HH: Which you did.

KA: Not immediately. My parents wanted me to be pre-med, so I said, "Okay, I'll do that." I gave it a shot, but then I liked philosophy too much, I wanted to go into philosophy. In pre-med I learned a fair amount of science and so then I decided to go into history and philosophy of science.

HH: The title of your 1989 PhD thesis was *The Promise of Parallel Distributed Processing*. It took some time, and some unexpected events, like everyone on the planet holding a smartphone in their hands 18 hours a day, but a lot of that promise did eventually deliver. How much of it would you change in hindsight?

KA: So, here is the picture at the time I was writing my thesis. Once I passed the qualifying exam, I didn't have any idea what I wanted to do for a thesis. I didn't tailor my coursework to prepare me for a dissertation. And that was a big mistake. But , I was teaching at Carnegie Mellon as a graduate student and Hinton was there. The *Parallel Distributed Processing* books were coming out and I bought a Xerox copy from James L. McClelland. One of my colleagues at Carnegie Mellon told me "You should try this." And so I just I threw in on that. In the same year John Bickle, Tim van Gelder and Bill Ramsey, also came out working on connectionism. All of us saw that it was happening. If it was a bad call, three other people made it too. But if I were to do things differently, I'd probably try to look further ahead.

HH: Speaking about bets, your best-known work (with Fred Adams) closes with a forecast: "Insofar as we aspire to create a

cognitive science, it seems reasonable to suppose that the science of cognition will resemble the science of physical, chemical, and biological processes. Such, at least, is our bet." Twenty-five years have passed, so it's perhaps time to assess that bet.

KA: I think we won! I guess I was struck in the original Clark and Chalmers paper. They said, "Inga and Otto are exactly the same." And then Adams and I said, "Well, no, by the lights of cognitive science, they're very different". Inga doesn't have the memory effects that Otto does, or Inga has memory effects that Otto doesn't. So cognitive psychologists will see that they are different. There are important differences. And as far as I know, everyone interested in the debate accepts that now. The move ever since has been to say that Inga and Otto are "in some respects they are the same". Well, game over, right? Of course they're the same *in some respect*!

HH: Is the controversy still ongoing?

KA: I stopped fighting with them. I've stopped fighting the extended cognition people. They have not decided they're going to go after cognitivism. They just said they're going to make up new words for cognition and say *okay that's extended.* Okay, so what?

HH: So you don't buy into the idea that machine learning-driven wearables break the bounds of cognition you identified in the skull?

KA: No. The whole reason you like these wearables is because they do something your brain does not. That's why they are useful. The reason you use pencil and paper to add is because it lightens the memory load. You just put everything in columns, do the first column, and then you forget about it. Then you do the second column, then you forget about it. The reason this is so great is it's not

the way you normally do it in your head. To me this is completely obvious, and I wonder how they ever thought otherwise.

HH: So why did you work on it?

KA: The first time I read the Clark and Chalmers paper, I wasn't going to do anything about it. I thought it was just so absurd. But then I saw Dan Dennett getting interested in that, and then John Haugeland, Merlin Donald, and others getting interested in that. Something was happening there. And I decided to write on it. I guess that was a good call, but I never dreamed it would be this big.

HH: Do you have any advice for postdocs on how to secure tenure?

KA: Find your niche. Do find your niche. Try and see what is coming down the pipe. See what's going to become hot.

HH: You are telling them to become *the* expert on something, rather than to become knowledgeable about many things.

KA: Just be very serious about one thing. Be the world's leading expert on something that is going to catch a wave.

HH: That's necessary but not sufficient, I guess.

KA: I guess I did catch a wave on connectionism, but that was not enough. One of the things that is happening now in the United States, is memory research. It is taking off. So that would be a good area. But where in memory are you going to, how are you going to distinguish yourself? Those are the questions I would advise postdocs to focus on.

HH: Any rule of thumb to spot the coming wave?

KA: If you are a philosopher of science, look for new science. Look for the latest thing that is happening in science, so you can

get in early, and then identify what's important about it. Do it from a philosophy of science angle.

HH: For instance?

KA: Optogenetics. It is a new technique neuroscientists are very excited about, because they think it's going to do things for them that they couldn't have done otherwise. So here is a theoretical question: what exactly is that thing they couldn't do otherwise? Why do they care about it?

HH: Techniques typically provide answers before a theoretical question has been asked.

KA: I guess 30 years ago fMRI was like this. A lot of people wanted to write about the limitations and promises of fMRI. Several people published on that, it was something to be an expert on. So those are the kinds of techniques to look at.

HH: Last question: Five books everyone serious about the philosophy of science should read.

KA: First, my book.

HH: Fair enough.

KA: Normally I am not a self-promoter, but I do believe in this book.

HH: Then?

KA: *Making Things Happen*. It's a challenging read, but I think it's going to stand the test of time.

HH: Number three.

KA: *Explaining the Brain* was an important book. Still is. I guess I see see my book as a successor to that.

HH: Number four.

KA: Mazviita Chirimuuta's *The Brain Abstracted*. That's getting a lot of attention. And then Ned Block's *The Border between Seeing and Thinking*.

HH: Do you want to make it six then?

KA: Maybe I would add *Knowing Science* by Alexander Bird. I think a counterpoint to what I'm doing. Bird does epistemology of science.

HYKEL HOSNI
iD

Logic Uncertainty, Computation, and Information (LUCI) Lab,
University of Milan

MENO II: A SELF-REFERENTIAL SOCRATIC DIALOGUE ABOUT MEMORY AND COMPUTER PROGRAMMING

:: SAMUEL ALEXANDER

Abstract

We provide a Socratic dialogue in which Meno challenges Socrates' principle that all learning is actually the remembering of things known but forgotten. Meno claims: if Socrates' principle holds, then Meno should already know the contents of the ongoing dialogue. Meno challenges Socrates to help Meno recall those contents, claiming Socrates cannot do this, lest Meno could immediately say the opposite of whatever he supposedly recalls he was about to say. Socrates eventually succeeds at Meno's challenge in an unexpected manner.

Keywords

Anamnesis; Meno; Self-reference

In the dialogue Meno, Plato's Socrates suggests the principle of anamnesis: there is no learning, only the remembering of forgotten knowledge. Presumably this should not be understood of contingent facts, or rather, perhaps we never truly know contingent facts: thus, Wittgenstein refused to acknowledge that there is no rhinoceros in the room. But it is fun to speculate about what would be implied if anamnesis did extend to contingent facts. It would extend to Plato's speakers' knowledge of their own words: they would know (but have forgotten) the very words they say, before saying them.

We designed a Socratic dialogue in which Meno asks for assistance remembering what he will say in the dialogue. Against all odds, Socrates succeeds. The dialogue takes place in a universe where Socrates and Meno know computer programming; Socrates uses techniques similar to quines (self-printing computer programs) and obfuscated code to evade Meno's paradox. Because of the venue's wordcount limit, we give here key excerpts (roughly 2500 words) of the roughly 7500 word dialogue. The full dialogue is available at https://github.com/semitrivial/Meno or https://doi.org/10.5281/zenodo.15686654. In the climax, Socrates leads Meno to decode a computer program steganographically hidden in the words of the dialogue; skeptical readers can verify this computation using the full dialogue at the above link (either from scratch or with the assistance of a verification script available at the same link). After these excerpts, we remark about how the dialogue was engineered.

*

Meno: Good morning, Socrates.

Socrates: Good morning to you, Meno, my friend.

Meno: I've been wanting to talk to you.

Socrates: And just what would you like to talk about, my dear friend?

Meno: It's about what you said the last time you and I spoke together.

Socrates: I suppose you mean when I told you that learning is really just remembering?

Meno: Yes, and that we actually know everything already, and we merely remember it.

Socrates: And I suppose you've contrived some sort of paradox or something to try and disprove it?

Meno: Quite. Once we finish this delightful discussion we're having, don't you think that then I will know all the things we said?

Socrates: You mean, know the conversation?

Meno: Exactly. Once we're done, I'll then know everything we said. That is, I will have learned it, wouldn't you say?

Socrates: I'm not really so sure about that. More often than not, I find I misremember conversations, or, even worse, I think that I remember and understand them, and then it later happens that I either remember wrong, or understand wrong. Besides, your question makes me uneasy, for I feel as if I am in danger of falling into a trap somehow.

Meno: What! Does the great Socrates have trouble with his memory?

Socrates: I often wish that I could remember things better, Meno. For I feel sometimes as though I flitter between two different worlds. One is intelligible and beautiful, the other is sensuous

and profane. For instance, when my house is well organized, and I remember where I have put everything, it's as if I live in a house designed by the geometers, for I can find my shirt or my shoes by an act of pure thinking. But more often, it pains me to say, my house is not organized, and I must grope around for my shirt or my shoes, and I feel as if I'm cast out of that house of the geometers, and plunged into a carnal house of the senses.

Meno: Well, as for me, I take great pride in my house always being organized. For I employ servants who keep it organized for me. And, look here: I have a servant writing this very conversation of ours on a stone tablet. You know him: it is the same boy with whom you conversed about geometry the last time you visited. I treasure your words like gold, Socrates, and I will preserve them, and guard them, just like I would treat a valuable treasure.

Socrates: You flatter me, Meno. But I suspect you're just saying that because you want to catch me in some trap. Very well, Meno. For the sake of the argument, I grant that after we have this conversation, you shall have learned it.

Meno: Well then, doesn't that mean I already knew it, and have only to remember?

Socrates: Yes, I grant that, Meno. Even before you were born, you already knew I would say the words, 'Yes, I grant that, Meno,' which I said just now.

Meno: Well then, dear Socrates, I have a challenge for you. With your famous method of questioning, please guide me to recall all of the things we shall say in this dialogue we're having. That is, if you are able to. But I doubt it.

Socrates: Why do you doubt I can do it, my dear Meno?

Meno: Because the very moment you help me remember the

things we're going to say to one another, I will immediately re-member what line I will say next, and I will say the very opposite.

(1672 characters later...)

Socrates: Very well. This is a fun challenge. And I think if I'm to have any hope at succeeding, I will need to resort to computer programming.

Meno: Perhaps. For my challenge does seem somehow related to things like the recursion theorem and the incompleteness theo-rem. But even computer programming won't rescue you from my paradox.

(1399 characters later...)

Socrates: So for example, suppose I ask you: take the first sen-tence you said to me, and repeat it to me again right now, and then after that, please repeat it to me a second time. And suppose I claimed that, by doing so, you would recite our entire conversa-tion. Would I be right?

(667 characters later...)

Socrates: So the reason why this supposed one sided conversation would not win your challenge, is because of the content of the results?

Meno: Quite so.

Socrates: And is there any other reason why this example conver-sation would not beat your challenge?

Meno: No.

Socrates: In that one sided example, I used earlier pieces of our conversation as tools. But you say, the lone reason that example would fail to solve the challenge, is due to the contents of the

resulting one sided conversation?

Meno: That's right.

Socrates: Then it's perfectly alright for me to use earlier pieces of our conversation like that?

Meno: Yes, that's alright.

(2234 characters later...)

Socrates: But isn't your servant writing some other things here? Look. He writes 'Meno:' whenever you begin talking. He writes 'Socrates:' whenever I do.

Meno: Sure.

Socrates: We have a serious problem, then. Look at the tablet right here, near the very top. It says: 'Good morning to you, Meno, my friend. Meno:' Now, I can't remember. Did I say 'Meno:' immediately after 'Good morning to you, Meno, my friend'? Or rather, does 'Meno:' here mean that you began talking?

(1196 characters later...)

Socrates: But maybe we can modify the challenge. Do you at least know what is written on the tablet? For instance, that the very first letter on the tablet is 'M'?

Meno: Indeed.

Socrates: Then maybe instead of helping you to recall our future conversation, I could instead help you to recall the future contents of the tablet. Your little paradox would still confound me, wouldn't it?

Meno: It would.

Socrates: And I wouldn't even necessarily need to help you re-

member my own contributions to the tablet, would I? Wouldn't you say I won if I could just help you recall your own future contributions to the tablet?

(2166 characters later...)

Socrates: Do you remember how Ulysses tricked Polyphemus, the cyclops? He lied to Polyphemus, saying his name was "Noman". Then when Polyphemus's neighbors came and asked if some man was slaying him, Polyphemus said, 'Noman is slaying me.' Didn't Ulysses do that on purpose, so that the cyclops's friends would misunderstand, and think Polyphemus had said "No man is slaying me"? Isn't that just the sort of syntactical semantical confusion you have so effortlessly avoided just now?

(2922 characters later...)

Socrates: And changing people's thoughts is like writing Greek on the surface of the water. The water is so pliable, you can write on it with your bare finger. But the instant that you've written there, your writing disappears, doesn't it?

Meno: It does.

Socrates: In fact, I always assumed that's what Thales meant, when he said that everything is made out of water. He meant that all the impressions we make on the world fade, just as if we inscribed them on the face of the sea. You could not corrupt the Aegean by writing on its surface, no matter how scandalous the things you wrote there. Why, even the language we speak slowly changes. And we try to nail that down, by imposing sharp constraints on it in the form of poetry.

(6011 characters later...)

Socrates: But look here. There's some space on your servant's

tablet between 'Meno:' and 'Good'. And if someone included that space in the X they passed to you, then your function would not output 'Meno:Goodm', but rather it would output 'Meno: Good', wouldn't it?

Meno: Yes.

Socrates: Can you modify your function to discard all such spaces? We'd better discard tabs and linebreaks too, and also hyphens, since I see your servant sometimes inserts a hyphen when a line ends in the middle of a word.

(3141 characters later...)

Socrates: Well, let's change our approach then. Suppose we described some elaborate transformation. And suppose you performed that transformation on this conversation so far, from the start until when I said, "Well, let us change our approach then." And suppose the output of that was:

```
print("""You win, Socrates.
You're welcome.
I admit it.""")
```

And suppose that the moment you knew that output, you said, "You win, Socrates," and I said, "Thank you, Meno," and then you said, "You're welcome," and then I said, "Do you admit there is no learning, only remembering?" and you answered, "I admit it." Would I have passed your challenge?

Meno: Certainly.

(3856 characters later...)

Socrates: Perhaps the problem is that a haystack can't be made

entirely out of needles, and if our conversation does secretly hide some program that predicts you, I suppose many characters from the conversation must work together to hide each character in the program. Perhaps every five characters from our conversation could be used to generate one character of our desired program. For example, perhaps the correct cipher, rather than changing "a" into "b", changes "Meno:" into "p", and changes "Goodm" into "r", and changes "ornin" into "i", and changes "g,Soc" into "n", and changes "rates" into "t", so that, altogether, the first 25 characters of the conversation become "print". Do you think that's at all likely?

(756 characters later...)

Socrates: You mean, if programs were written with a larger alphabet, then we would need more characters of conversation per character of the hidden program?

Meno: Yes.

Socrates: Then I guess since we haven't got the whole day to stand here talking, we should agree to write this desired program with the smallest alphabet we can. We should ask ourselves: which characters are absolutely essential for writing computer programs? Within reason, that is. After all, we're not trying to compete in some obfuscated code contest.

Meno: I think a reasonable alphabet for writing programs, without too many characters, would be:

```
alphabet = "abcdefghijklmnopqrstuvwxyz"
alphabet += " ()[]:'=,+.#\n"
alphabet += '"'
```

Socrates: That's forty characters in all, right? Since I assume by \n you mean a linebreak, which is one single character. And no uppercase letters, like "A", "B", "C"?

Meno: For your sake, Socrates, I shall consider it a victory for you if our program predicts I'll say "hey", all lowercase, even if my servant here actually writes "Hey" with the H uppercase.

Socrates: Most generous of you, Meno. Well then, if we assume this rather barebones alphabet of yours is rich enough to write computer programs, how many letters of our conversation should we use to encode each character of our program? We already saw that "one", "five", and "ten" are too small. Perhaps twenty?

Meno: Socrates, I know you haven't studied computer science, but really. Twenty is still way too small. You're probably still off by an order of magnitude, so let's make it two hundred.

Socrates: Alright, 200 it shall be. So every 200 characters of our conversation shall hide within them a single character of a computer program. Alright then, how would you define the cipher? It must take a string of 200 characters from our conversation and convert that string into a single character from your 40 letter alphabet.

Meno: Let me think, Socrates. Hmmm... Well, there's a convenient function named "ord" for turning characters into numbers. For example, ord("a") is 97, ord("b") is 98, ord("+") is 43, and so on. Don't fret about those specific numbers, this "ord" function is built into my programming language (Python). As our tablet has quotation marks and my servant sometimes writes those with fancy unicode quotes, I guess to be safe we should consider all the quotes in our conversation to be plain ASCII quotes, just like we delete all the spaces and linebreaks and tabs in the conversation.

And plain ASCII quotes have ords as follows: for the single quote, ord("'")=39; for the double quote, ord('"')=34. With this ord function, we can systematically turn characters into numbers, add up those numbers, and then take the remainder after dividing the sum by the length of my program alphabet. In other words, if the sum is 40 or more, smash it down into the range from 0 to 39, so it can be used as the index of a character in our alphabet. If we write it all as a function:

```
def cipher(X):
total = sum(ord(character) for character
in X)
remainder = total % len(alphabet)
return alphabet[remainder]
```

This function works for inputs X of any length. Of course it is intended to be used on inputs X of length 200.

Socrates: Just to make sure I understand right, what would be the value of, say, cipher("Meno:")?

Meno: Well, ord("M") is 77, ord("e") is 101, ord("n") is 110, ord("o") is 111, and ord(":") is 58. Add those up and you get 457. Now, the length of my alphabet is 40. If we divide 457 by 40, the result is 11 plus a remainder of 17. So the index we get is 457 % 40 = 17. So cipher("Meno:") is alphabet[17], character number 17 in our alphabet, remembering to count from zero. Which is "r". Altogether, cipher("Meno:")="r".

Socrates: Amazing, Meno, remind me someday to introduce you to my friend Timaeus. And what do we get when we take the first 200 characters from this conversation and plug them into this cipher?

Meno: If we take the first 200 characters of our conversation, they make the following string, though I fear my servant will insert linebreaks because otherwise this long string won't fit in one single row on his tablet, so please ignore those linebreaks; all that said, those 200 characters are as follows:

```
Meno:Goodmorning,Socrates.Socrates:Goodmorningtoyo
u,Meno,myfriend.Meno:I'vebeenwantingtotalktoyou.So
crates:Andjustwhatwouldyouliketotalkabout,mydearfr
iend?Meno:It'saboutwhatyousaidthelasttimeyouandIsp
```

If we plug this length 200 string into the cipher function, remembering not to include the linebreaks, we get the output "#".

(1375 characters later...)

Socrates: By Zeus! Alright, Meno, keep going. Plug in all the length 200 blocks from our conversation, stopping with the first block in which I say "Zeus". Do we get a valid program?

Meno: One minute... No, but I am shocked, for the output does resemble a computer program, but it has all kinds of errors. Here it is, note the double space between the "n" and the "r" on line 3:

```
#trojn hrse
seed(p=askpword())
z,x=", yo ooned","""n r''s'
iati.idnet'pas'r.p(se)d nt yap'o
+twspcw'tr"""
c=jn(smpl(p,z.len
shfl(d=rng(x.len
if c==z:exec(jn(x.i:i in d
```

Socrates: You say it has errors? Can we repair it? Tell me, what's the first error you notice?

(1824 characters later...)

Socrates: What's the program now, with all the errors fixed?

Meno: Here you go:

```python
#Trojan Horse
import random
p=input('Enter the password:  ')
random.seed(p)
z,x=", yo ooned","""n r''s'
iati.idnet'pas'r.p(se)d nt yap'o
+twspcw'tr"""
c="".join(random.sample(p,len(z)))
d=list(range(len(x)))
random.shuffle(d)
if c==z:exec("".join(x[i] for i in d))
```

Socrates: Run it! What does it predict? I understand you plan to say the opposite. What's wrong? You are quiet. Is there some problem with your program?

Meno: It wants a password.

Socrates: You're in a bind, for until you enter the password, you won't know what it will say once you enter it. Maybe when you enter the password, it will say, "It wants a password." It's too late to say the opposite of that. Maybe studying the code, you can figure out the next prediction, and defy the oracle? What's wrong? Can't you analyze the code to determine the next prediction?

Meno: It's encrypted.

Socrates: Then you must enter the password, and we shall see what your program predicted. Meno! Tell me! What is the password?

Meno: I don't know any password, Socrates.

Socrates: The password, in your own words, is: "I don't know any password, Socrates." Enter that, and if your wisdom is as great as I believe, I think your program will print:

```
it wants a password.
it's encrypted.
I don't know any password, Socrates.
```

Now I must leave, Meno, I have an appointment at the courthouse.

*

One moral we might draw from this story is that the statement "philosophy will help me remember future contingent facts even if I try to contradict them" could itself be contingent. Meno's paradoxical challenge initially seems to disprove this statement by pure reasoning (like the geometry example in Plato's *Meno*), which would make the statement universally false. Socrates' victory disproves that disproof.

We wrote the dialogue using annotations. For example, "Meno: [Yes|Correct]" means that Meno could say "Yes" or "Correct". A computer program flattened these annotations to encode Meno's program. The ways a passage can be so obtained grow exponentially with the number of annotations. "Meno: [Yes|Indeed], [Socrates|friend]" can be flattened in 4 ways; "Socrates: [Isn't that|Is that not] [right|correct], [Meno|friend]?" can be flattened

in 8 ways. Thus the problem is easier than it initially appears. The annotated dialogue, and associated computer programs, are available at https://github.com/semitrivial/Meno.

As an application, a variation of this technique could theoretically be used to assist in the problem of preserving rhyme and metre when translating poetry. We speculate perhaps similar methods were used (without computer assistance) for example by John Dryden, whose English translation of Virgil's Aeneid rhymes almost every line.

SAMUEL ALEXANDER

The U.S. Securities and Exchange Commission

METHOD AND CONVERGENCE, HELSINKI, 2–4 JUNE 2025.
:: LORENZO CASINI

Abstract

Report on the Conference *Method and Convergence*, Helsinki, 2–4 June 2025.

Keywords

Method; Abduction; Naturalism; Progress.

How to Cite

Casini, L. *Method and Convergence*, Helsinki, 2–4 June 2025. The Reasoner, 19(3). https://doi.org/10.54103/1757-0522/29042

The international conference Method and Convergence, recently held at the University of Helsinki, brought together scholars from diverse disciplines to explore philosophical methodology and its role in advancing both philosophy and science. Topics included how philosophy contributes to scientific progress, how developments like AI can aid philosophical research, and how methodological innovation drives intellectual advancement.

On the first day, Ilkka Niiniluoto (University of Helsinki) opened with a keynote on *Progress in Philosophy*, defending an optimistic view of philosophical progress grounded in the Socratic method: problematization, explication, and argumentation. Though philosophy differs from empirical science, Niiniluoto argued it advances through conceptual clarity, argumentative rigor, and evolving perspectives on enduring questions.

Ron Chrisley (University of Sussex) followed with *Method and Creativity: Toward AI Tools for Constructively Divergent Thinking*. Drawing on Margaret Boden's work, he emphasized the importance of transformative thinking—radical conceptual innovation—and proposed developing AI tools that foster genuinely divergent and creative reasoning.

The second day began with Nina Emery (Mount Holyoke College), who distinguished between content and methodological naturalism in her talk *Two Types of Naturalism in Metaphysics of Science*. She argued that commitment to scientific content in metaphysics entails a commitment to scientific methods, reshaping how metaphysical inquiry should proceed.

Andrew Brenner (Hong Kong Baptist University) then presented *The Role of Simplicity in Metaphysics*, arguing that simplicity, often invoked in science, also functions as a rational guide in metaphysical debates. He illustrated how simplicity can favor certain hypotheses and offered a fresh take on the question of why there is something rather than nothing.

On the third day, Gerhard Schurz (Heinrich Heine University) gave a talk on *Inductive Metaphysics and its Abductive Methodology*, defending the use of empirical and abductive methods to support metaphysical claims. Through case studies, he showed how testable, empirically unified theories can yield meaningful metaphysical insights.

Finally, Finnur Dellsén (Universities of Iceland, Inland Norway, and Oslo) gave a talk entitled *Abduction: The Glory and Scandal of Philosophy?* He highlighting challenges in using abductive reasoning in philosophy, such as underdetermination and theory competition. He advocated strategies to mitigate the risks involved with the use of abduction.

In addition to the keynote sessions, the conference featured a wide range of plenary and contributed talks. It ended with an open panel discussion on whether progress is possible in philosophy, how to recognize it or measure it, and what methods are more conducive to philosophical progress in teaching and in research.

The conference was a great success. It showcased the richness of contemporary philosophical methodology and fostered fruitful dialogue between philosophy, science, and AI.

Lorenzo Casini

University of Bologna

PROGIC 2025: THE 12TH WORKSHOP ON COMBINING PROBABILITY AND LOGIC, CARNEGIE MELLON UNIVERSITY, PITTSBURGH, PA – APRIL 3–5, 2025.
:: EISSA HAYDAR

Abstract

Report on the Conference *Progic 2025: The 12th Workshop on Combining Probability and Logic*, Carnegie Mellon University, Pittsburgh, PA – April 3–5, 2025.

Keywords

Probability; Logic.

How to Cite

Haydar, E. *Progic 2025: The 12th Workshop on Combining Probability and Logic*: Carnegie Mellon University, Pittsburgh, PA — April 3–5, 2025. The Reasoner, 19(3). https://doi.org/10.54103/1757-0522/28914

The twelfth edition of *Progic*, the long-running Workshop on Combining Probability and Logic, took place at Carnegie Mellon University in early April, drawing together scholars working in formal epistemology, logic, probability, and causal inference. This year's meeting was organized by Kevin Kelly, Krzysztof Mierzewski, and Francesca Zaffora Blando, and was held in the Adamson Wing of Baker Hall.

The program began with a plenary talk by Dan Roy (University of Toronto), who presented a paper, coauthored by Haosui Duanmu and David Schrittesser, titled *Admissibility is Bayes Optimality*

with Infinitesimals. Speaking via Zoom due to flight cancellations caused by inclement weather, Roy used tools from nonstandard analysis, allowing for priors which can assign infinitesimal values, in order to show that, in this setting, the admissible decision rules are exactly those (whose nonstandard extensions) minimize Bayes risk.

In the first contributed session, Cameron Freer (MIT) presented joint work with Nate Ackerman and Rehana Patel on *Learning Invariant Measures*. He discussed the with-high-probability, approximate learnability of a special class of probability measures over relational structures.

This was followed by Jeff Barrett (UC Irvine)'s presentation of joint work with Christian Torsell (UC Irvine), on *Learning How to Learn (by Reinforcement)*. In a Humean vein, the focus was on how (perhaps simple) agents might actually successfully learn in various contexts, as opposed to how their learnings can be rationally justifiable.

The afternoon session featured Jürgen Landes (Munich Center for Mathematical Philosophy), who presented joint work with Sebastien Destercke on *The Value of Varied Evidence for Imprecise Probabilities*. Landes revisited the "variety of evidence" thesis in the context of imprecise probabilities, demonstrating that its plausibility is highly sensitive to modeling assumptions.

Jon Williamson (University of Manchester) closed the day with a presentation titled *The Heuristic Use of Conditionalisation*, which sought to reframe Bayesian conditionalization as a useful heuristic rather than a universal epistemic norm. The forwarded alternative was objective Bayesian inductive logic which adopts inference to the highest entropy probability distribution as an epistemic norm replacing the Bayes rule.

Friday's plenary was delivered by Sean Walsh (UCLA), who talked about *The Expressive Power of Counterfactuals and Descriptive Set Theory*. The Lewisian account of counterfactuals takes in a distance metric, which gives rise to a topology. Descriptive set theory then supplies a means through which complexity, of say the counterfactual conditional, can be studied.

Later in the day, Cameron Freer (MIT) returned to present joint work with Nate Ackerman and Dan Roy on *Computability of Properties of Stochastic Processes Used for Learning*. Their work is aimed at effectivizing the result, which exactly characterizes those measures for which conditioning on their tail σ-field yields a measure that is sufficient for prediction and merges with the original.

Sam Eisenstat (Machine Intelligence Research Institute) followed with a discussion of *Logical Uncertainty and Self-Reference in Garrabrant Induction*. He argued that Garrabrant's model of bounded logical reasoning supports a distinctive epistemology for non-omniscient agents and addresses classic challenges such as old evidence and inductive generalization.

In the afternoon, Giovanni Duca (Northeastern/University of Milan) presented *Updating on Uncertain Evidence*: *a formal correspondence between belief revision and Jeffrey conditioning*. His results clarified how odds-thresholds yield preferential orders (themselves all-or-nothing beliefs) which can be used to match probabilistic reasoning, allowing him to give a characterization of Jeffrey conditioning in this setting.

Milan Mossé (UC Berkeley) closed the day with a talk on *Reasoning about Confirmation*. He examined various formal theories of confirmation, from Hempelian logic to probabilistic thresholds, and analyzed the complexity of reasoning in these systems

through the lens of satisfiability.

Saturday's plenary was delivered by Kun Zhang (Carnegie Mellon University), titled *Causal Learning: Why It Matters, How It Works, and Its Implications for Generative AI*. Zhang argued that the study of causal discovery from observational data, especially those involving causally-related hidden variables, was pivotal for the advance of certain scientific fields, and he presented algorithmic advances in this domain, with a particular focus on AI image generation.

This was followed by Sander Beckers (Cornell University), who presented *Nondeterministic Causal Models*, an extension of traditional structural models to allow for multiple possible outcomes to an intervention. Beckers proposed a logic for reasoning about such models and argued that they better capture counterfactual nuance.

Aydin Mohseni (Carnegie Mellon University) then presented joint work with Ben Levinstein, Daniel Herrmann, Gerard Rothfus, and Bruce Rushing which challenged cau-sal fundamentalism. In *Naturalizing Causation in Causal Models*, he argued that causal claims can be fully reduced to probabilistic statements within the framework of structural models.

Next, Siddharth Namachivayam (Carnegie Mellon University) presented *Topological Semantics for Common Inductive Knowledge*, where he gave a formal account of common inductive knowledge, centering the analysis of a Byzantine-generals-like coordination game.

The conference concluded with a talk by Adam Bjorndahl (Carnegie Mellon University), titled *A Logic of Uncertain Interpretation*. Bjorndahl developed a framework for modeling inter-

pretive uncertainty, applying it to both implication and evidential belief, his approach giving rise to Dempster-Shafer belief functions.

More information on the talks can be found at the conference website.

EISSA HAYDAR

Philosophy, Carnegie Mellon University