

Antikythera: A Speculative Philosophy of Planetary Computation

Benjamin Bratton

As computation becomes planetary infrastructure, how does its acceleration of hybrid intelligences pose new challenges to fundamental philosophical questions? As machine sensing, machine cognition, machine embodiment co-evolve, how does computation become more than a mere technology, but the medium through which we ask existential questions about who, what, and how we are?

The Antikythera mechanism, sometimes called the “first computer,” was more than a calculator; it was also an astronomical device. The birth of computation is in the orientation of intelligence in relation to its planetary condition.

As computation evolves into planetary infrastructure—scientific, cultural, geopolitical—perhaps its most decisive impact will be not in what it does as a tool, but as an epistemological technology: what it discloses to sapient intelligence about how the world works. This in turn alters how intelligence remakes its worlds, including the ongoing artificialization of intelligence, life, sensation, and ecosystems.

What is the philosophical school of thought most appropriate to this reality? Perhaps instead of only projecting timeless wisdom, the work of speculative philosophy is to compose new interpretations of new realities, and to do so through direct exploratory encounters with the technologies that disclose those realities to us. Ultimately, we may ask, what is planetary computation for, and toward what futures might it be oriented and in turn orient the future of complex life and intelligence?

Benjamin Bratton is Professor of Philosophy of Technology and Speculative Design at the University of California, San Diego. Through the lens of planetary computation, his work establishes new philosophical frameworks for interpreting the past, present and future co-evolution of life, culture, and technology. He is Director of Antikythera, a think-tank researching the future of planetary computation based at the Berggruen Institute. He is the author of numerous books including *The Stack: On Software and Sovereignty*. The tenth anniversary edition will be published by MIT Press in 2026.

EDITORIAL


This talk provides an overview of Antikythera’s Speculative Philosophy of Planetary Computation. It establishes the program’s core inspiration: the Antikythera mechanism, an ancient device linking computation with cosmology and planetary orientation. Computation is not just calculation but a fundamental force shaping society and understanding, originating from early accounting systems like cuneiform. The talk advocates for an *allocentric* philosophy—one that gets outside human perspectives, often via technology—to grasp reality. Antikythera contends that technology, particularly computation, currently outpaces our conceptual language, necessitating philosophical invention rather than just application.

The central concept, Planetary Computation, is introduced through the “Blue Marble movie” thought experiment: Earth, over billions of years, rapidly grows a computational “epidermis” of satellites, networks, and sensors. This system is both a global infrastructure (conceptualized as “The Stack” with Earth, Cloud, City, Address, Interface, User layers) and an epistemological tool revealing planetary conditions like climate change, thus making the Planetary and Anthropocene conceivable. This planetary-scale sensing turns the Earth itself into a cognitive apparatus, as exemplified by the Event Horizon Telescope.

The Antikythera program links philosophy and engineering, using interdisciplinary studios and collaborations to generate new ideas. Its research spans four areas: Planetary Computation, Synthetic Intelligence, Recursive Simulations, and Planetary Sapience.

Synthetic Intelligence focuses on AI’s evolution “in the wild,” challenging anthropocentric views and simplistic alignment goals. It critiques *Reflectionism* (the contradictory views that AI must mirror humans/AI only reflects human flaws) and proposes exploring *productive disalignment*. Recursive Simulations examines how computational models interact with and alter the reality they represent, from climate politics to digital doubles. Hemispherical Stacks analyzes the interplay of geopolitics and the fragmentation of planetary computation into competing regional systems.

Planetary Sapience integrates these themes, framing intelligence as a planetary evolutionary phenomenon. It posits an evolutionary sequence: Evolution → Life (autopoiesis) → Artificialization (allopoiesis) → Intelligence → Symbolic Language → AI, where AI represents the *artificialization of artificialization* and acts as a scaffold for future unknowns. This implies a convergence of life, intelligence, and technology. The paradox remains: is intelligence adaptive long-term, or self-extinguishing? The goal is to identify and engineer the preconditions for its long-term adaptiveness.



Antikythera: A Speculative Philosophy of Planetary Computation

Benjamin Bratton

Thank you, Nicholas, for that very kind introduction and welcome to a week of Antikythera events at MIT. First, however, a bit of thanks, most sincerely to all of our supporters and our hosts: MIT School of Architecture, the Morningside Academy of Design, our primary collaborators on this little conspiracy—the Berggruen Institute, as well as the Media Lab for hosting this event. And to MIT Press for joining us for the next phase of the Antikythera program, a book series, as well as a digital journal, which I'll talk a little bit about later on.

INTRODUCTION

00:10

Introduction

First of all, the name: Antikythera, a bit of a tongue twister. It comes from this little device, the Antikythera mechanism, which was probably apocryphally the first computer from second century BCE, discovered in 1901. It took a long time to figure out what it actually was because it was so anomalous in the evolutionary history of technology.

It was not only a calculation device, a computer, it was also an astronomical device. It allowed its users to orient themselves in relationship to their planetary condition. And, indeed, it's this conjunction between computation and cosmology that is our inspiration. Indeed, computation was born of cosmology. And I mean that in both senses of the term; of a scientific cosmology, astronomic cosmology, as well as the anthropological sense of cosmology, of how we make sense of who we are, where we are, and so forth. And it's indeed this disjuncture between these two that is why a philosophy of technology think-tank such as ours might be useful in some regard.

This speaks to computation in its more, let's say, conceptual sense. But computation is also a way in which we organize societies. Remember that the first writing or among the first writings was the Sumerian cuneiform, which were essentially receipts, a very granular, calculative notation of economic transactions. And from cuneiform, of course, all manner of alphanumeric languages and literatures ensue. So, at least genealogically speaking, all great works of literature are a kind of elaborate fancy form of accounting, a point that I like to remind my friends in the literature department about.

One of the things that we think that philosophy and science also, at least when they're most successful, have in common, is a kind of Allocentrism. They allow us to get out of our heads and to see the world not as it appears, not as it, maybe, perceived phenomenologically, but perhaps as it is. And it is this usually technologically mediated alienation, a kind of derangement and disenchantment that is the most priceless accomplishment of both.

Sciences historically are born when philosophy finally learns to ask the right questions. Physics, chemistry, biology, obviously computer science as well. Philosophies, on the other hand, are born when technologies force the birth of new languages. Where they make things possible that we don't have the words for. I think this is indeed part of our moment.

There are moments in history when, as we say, new things outrun the nouns available to contain them. When our ideas about what it is that we would like to do technologically are far ahead of what's actually possible to do. So, 19th century ideas about space flight would be an example. But there are other moments in time, and I think our own, where the technology is capable of doing things already and is doing things that we don't necessarily have the proper language for to describe, let alone orient and steer. And in those moments, a philosophy of technology must do something else. It's not a matter of projecting truisms and axioms onto circumstances, but of inventing the language and conceptual structure that may allow for that orientation.

PLANETARY COMPUTATION

05:05

Planetary Computation

Let me then speak to what we call planetary computation, that is our key idea that we develop throughout the program, and which we orient a lot of the research. And I mean this in at least two ways, as you'll see.

Imagine the Blue Marble image, if you will, not as a still image as you see here in the corrected orientation, but indeed as a movie, a kind of super fast forward movie showing all 4.5 billion years of the evolution of Earth. What would you see?

You would see the formation of continents and their emergence and breaking apart, and volcanoes and the appearance of life, and thus eventually an atmosphere. But in the very last few seconds, you would see something extraordinary: this little blue cell floating in the black void would sprout a kind of computational epidermis of satellites and sensors and various mechanisms by which it comes to communicate with itself, to understand itself, an organ, if you like, that allows it to transmit information from one part of its body to another and to learn things about itself, including how its own planetary systems work, how old it is, and so forth. One that is becoming increasingly filled up but also increasingly more complex and increasingly infrastructural.

How does this happen? One way, obviously, is through the evolution of sapient human intelligence that makes such an initiative possible. That infrastructure, that cognitive organ, is increasingly also the seat of an artificialization of intelligence. It itself is capable of feats of cognition, not just sensing that previously may have been reserved for other forms of life, and it itself continues to evolve. AI is evolving as I will try to make the argument to you tonight.

Design and engineering and indeed societies as a whole, during the Apollo era, as this was, as all of these perspectives were becoming more commonplace, was enthusiastic about understanding, re-understanding the planet as a site condition to be constructed. But the real continuous monument, the one that we actually built, wasn't continuous, it was discontinuous. It was an accidental megastructure, built piece by piece, defined by modular, functionally defined modular layers in a kind of an interlocking structure, what we call *The Stack* in a book written in 2016, published by MIT Press.

And the argument goes, sort of defines this accidental megastructure in relationship to these layers. An Earth layer, from which the energy necessary to run this vast megamachine is drawn. A cloud layer by which platform services are drawn and delivered, many performing functions traditionally assigned to Westphalian states. A city layer in which we are all embedded, interacting with interfaces of variously intimate and public forms. An address layer that identifies, nominates anything to which information can be sent. An interface layer which provides an image of the affordances by which the whole system might be manipulated. And finally, where we live, a user layer, a functionally defined position of agency that sends circuits up and down this entire dynamic. And indeed discontinuity is the shape of the architecture of our times, which demands a different way of thinking about design, different way of thinking about engineering, and indeed a different way of thinking about philosophy.

Any project such as ours has its inspirations, its precedents. One of the key ones for us is this book, *Summa Technologiae*, 1964, by Stanislaw Lem. You probably know Lem more as a science fiction author, but this is his key book of nonfiction. And it's really an extraordinary book. But one of the key ideas that Lem develops (here is Lem) is the distinction he makes between what he calls instrumental technologies and existential technologies.

Instrumental technologies are those for which the primary impact in the world is in what they do as tools. Existential technologies, on the other hand, are the much more rarefied kind, are those that, when used properly, transform how we understand how the world works. Microscopes and telescopes are good examples, but indeed, so is computation.

For example, climate science is of course based on the production of data from a wide range of sensors and sources, but also primarily on supercomputing simulations of planetary past, present and future. Without that artificial cognitive crust and a computational infrastructure capable of making those simulations, the very idea of climate change as an empirically understood fact would not be possible.

And thus the more qualitative concept of the Anthropocene, and indeed all of the moral, ethical, political and technological consequences of the reckoning with anthropogenic agency at this scale. All of this is ultimately an effect of planetary computation and what it tells us as an existential technology. It reveals and constructs the planetary as what Dipesh Chakrabarty calls a humanist category.

There are a number of ways in which this history can be told. This is the first image of Earth taken from lunar orbit and sent back to Earth in 1966. It was sort of a precursor, if you like, to the much more famous Blue Marble image. And it had reverberated quite dramatically culturally.

It was shown to the notorious German philosopher Martin Heidegger in an interview that he did in *Der Spiegel* in 1966 called “Only a God Can Save Us.” And Heidegger was unequivocal in his anxiety, comparing the image to nuclear weapons, and arguing that indeed we don’t need nuclear weapons anymore because we have already destroyed the Earth by essentially getting outside and seeing ourselves in this way. And he was probably right. But for us that’s a feature, not a bug.

Compare this, for example, the Blue Marble image to the Event Horizon Telescope, which we’ll hear much more about on Friday from our guest Peter Galison. Now, the resolution of an image is, of course, dependent upon the size of the lens, the aperture. And so what would be the biggest resolution image you could make on Earth? It would be one the size of the Earth. To make such a thing, you would do what Event Horizon did, which is to link together telescopes from Greenland to Antarctica and use the rotation of the planet itself as part of the sensory machine that you have built by turning the planet itself into a technology. Or rather, the planet turns itself into a technology through us. Resulting, ultimately, in the black hole image with which we are all familiar. A collapsing space time some 50 million light years away. Now, the key difference here is that whereas the Blue Marble was a picture taken by a human of his domain, so to speak, a black hole is a much more frightening image. I’m sure Heidegger would have completely lost it, having been presented with this. The planet itself took the picture. The planet itself sensed its own environment. We were merely, as I wrote in *The Terraforming*, we positioned ourselves not as this caretaker apex species looking down on the Guardian, but rather as the apex mediating residue, the little creatures scurrying around on the surface of the planet who managed to assemble this sensory apparatus on behalf of its host.

This is also computation as an existential technology. This is a Copernican turn, a Copernican turn, a decentering, an Allocentric realization through technological alienation. This, in short, is what planetary computation is for.

THE PROGRAM

14:08

The Program

Let me tell you then about the Antikythera program itself. It is in many ways a program that tries to reconcile a relationship between philosophy on the one hand and engineering on the other. Now, philosophy has not always held engineering in such high regard. We do. The functionalism -the requirements of function are actually not only a generative constraint on how we might model a problem- but it is also to engineering that I think philosophy needs to look now to find that new language, that new conceptual architecture of vocabulary in order to catch up with the present.

Our program is then a research program, which I’ll discuss in some detail, but it’s also programmatic. It is also a school of thought in formation that is not from the Heideggerian tradition, but indeed something hopefully much more effective. So let me tell you then, very quickly, a little bit of the history of the program.

This is my home university, University of California, San Diego, home of Geisel Library, the Scripps Institute, and the Salk Institute. The genesis of this research program actually begins in Moscow, of all places, at the Strelka Institute, where they were very amused to find out that California also has a red star in its flag. At Strelka, Nicolay Boyadjiev, who’s here this evening, and I directed two think-tanks: *The New Normal*, the work of which you can see in this book we published with Park Press, and *The Terraforming*, which is also available as a free download. The Strelka Institute was an extraordinary institution. It was a kind of secular cultural hub within Moscow, one that was an unfortunate casualty of the invasion of Ukraine.

Antikythera was born roughly three years ago. It is developed in collaboration with the Berggruen Institute, based in Los Angeles and Beijing and London. The real scope of the program, though, is in the Studio and Affiliate Researchers that we work with on an annual basis in studios and salons and lectures and collaborative projects. All of this you can survey and see a bit more of at the antikythera.org website, which you can scan here or go on your own.

What do we do? We do many things. We host salons in Beijing, Berlin, Amsterdam, Los Angeles, Mexico City, London. And these are moments in which we can gather everyone for extremely intense, all day conversation and not only find new answers, but learn to ask new questions together as well. We host lecture series which we usually do as a filmic production, in which you are now enrolled.

The core of the program is really the studios, however, which is where we do a lot of the fundamental work. And part of the philosophy of the studios is this: architecture has benefited from the studio culture from its beginning. It reinvents itself every 15 weeks. First principle questions are asked. And in doing so, it remains active and fresh in this regard. Now, arguably society now asks software to do many of the things that it used to ask architecture to do, namely the programmatic organization of people in space and time. Software, however, does not have the same kind of studio culture, exploratory, experimental, first principles studio culture. But why not? This is our beginning point for how it is we approach this. It's deeply interdisciplinary, not only computer scientists and designers and artists, but philosophers, science fiction writers. All of the studios bring these people together in an intensely collaborative environment.

The first studio we did in 2023 was at the Bradbury Building in Los Angeles, which you know from the Blade Runner. We had 12 of our studio researchers who came and we developed a number of quite extraordinary research projects, all of which you can see on the site there as well. Our big studio report, and so forth.

As you notice, a lot of our work uses a kind of cinematic vernacular. We make a lot of films, as a way of synthesizing a lot of ideas into a kind of a structured narrative.

All right, now to the heart of the matter. Planetary computation, as said, is the key concept around which the research rotates. But there are four important areas: Synthetic Intelligence, Recursive Simulations, Hemispherical Stacks, and Planetary Sapience. And I'll speak to each of these in turn.

SYNTHETIC INTELLIGENCE

18:59

Synthetic Intelligence

First, Synthetic Intelligence. Now we begin our discussion in recognition of one of Lem's most important stories about AI is a story called "Golem XIV," in which a precocious artificial general intelligence is housed here at MIT, and decides that it no longer wants to be wargaming and kind of goes rogue and becomes a bit of an oracle. And so we want to say hello to Golem XIV wherever you have him hiding these days.

One of the important things, I think, to remember about AI and its relationship to philosophy is that AI as a technology is perhaps unique in the sense that it is born from the philosophy of AI. That the philosophy of AI and AI as a technology have evolved in a kind of double helix relationship with each other, for at least centuries. And so this continues.

Some of the key ideas that we adhere to and our position is a little bit more about is that, ultimately, the artificialization of intelligence will teach us much more about what thinking is than we will teach the machines. That AI exists, ultimately, in the wild, that natural intelligence evolved in the wild and so will its artificial form. And that alignment, as they say, between AI and human society must be bidirectional. It's not just a matter of how AIs must adhere to what is suspiciously called human values, but also how societies adapt to the reality that inanimate objects are capable of complex cognition.

Our survey, if you like, of the current state of AI discourse and where we see ourselves in relation to this is in a piece I published in Noema a few months ago called "The Five Stages of AI Grief," in which this discourse is categorized as AI denial, AI anger, AI bargaining, AI depression, and AI acceptance. You can go to Noema and read it and decide which category you might be in.

A lot of the work we do is also collaborations with platforms. Blaise Agüera y Arcas, who I will speak a little bit more about, and I wrote a piece called "The Model is the Message" also in Noema, which was developed looking at some of the early years of LLM models and thinking about where things may go well and go wrong. And I think we got it pretty well right.

Blaise, by the way, who you'll meet on Friday if you're here, is the author of the first book that will be in the MIT Antikythera book series. Modestly titled "What is Intelligence?" There is also a kind of single to the album called "What Is Life?" that will be available to all those who come on Friday in a beautiful little book form as well.

We also in the journal, the MIT Digital Journal, are publishing a lot of the new work by our Affiliate Researchers, including Katherine Hayles's piece on modular cognition, Metahaven and Gašper and Nina Beguš on the use of machine learning in analysis of animal language and so forth.

AFTER ALIGNMENT

22:04

After Alignment

As mentioned, this question of alignment and our suspicion of a kind of ease and overuse of this term is one of the key areas by which we explore it. There's a lecture on our website that I did in London last year called "After Alignment," which lays out the argument in some detail.

One of the key ideas I'll pull from is this — it's what we call Reflectionism, something of which we are suspicious. Reflectionism would be defined as a kind of doublethink within much of the AI discourse. It would go something like this: and that is the idea that AI is not like us, but it must become more like us, become more human-centered, become more anthropomorphic, and as it does, the results will be better, and so there's this gap that must be resolved. On the other side, and sometimes spoken by the same person in the same sentence, is the idea that AI is nothing but the direct expression of our socioeconomic system. It is a mirror of us in all of its usually described as pernicious forms. So either AI is nothing like us, but must be, or it is only like us and that's the problem. This is what we mean by Reflectionism, and both sides of this we hold at some distance.

In other words, alignment to what, what form of human values are actually talking about? Why would the emergence of machine intelligence any more than the emergence of other forms of animal intelligence, be better off if it appears more comfortably anthropomorphic to us?

What it should not align to is a kind of deeply anthropocentric understanding of the evolution of intelligence, and its purpose is essentially as a kind of a simple prosthetic. It's more than this and ultimately may prove to be not even contained by the word technology.

AI is deeply weird. A lot of the science fiction stories around AI, this is, of course, from Godard's *Alphaville*, saw AI as a kind of seamless, technocratic, and clean Spock. It's not. It's deeply weird.

Some of this—the question around anthropocentrism, anthropomorphism within AI—comes from this appropriately canonical piece by Turing from 1950, which we are doing a kind of definitive redesign of and publishing in the first issue, “Computing Machinery and Intelligence.”

As you recall from Turing's 1950 essay, of which *The Imitation Game* is the the kind of key plot line, that he posits a situation in which an AI is asked to pass as a human, and he really makes what amounts to a kind of functionalist argument, which is, if it can pass as a human, that we have to grant that there's some kind of intelligence in there and that that's enough. Unfortunately, in the popular culture it has become something more like a necessary condition. That is, unless the AI can perform thinking the way that we think that we think it is disqualified from intelligence, this is the real problem.

We're familiar with the quadrant here that more alignment, more better, and less alignment, more risk. We are actually interested in exploring this fourth quadrant here. That less alignment and net positive outcomes may be the area that we want to model in advance.

And speaking of the language, part of this has to do with this: instead of trying to take what we understand machine intelligence to be doing and trying to force it into older categories, the inverse is really the question. We need the language to describe the weirdness that is happening right in front of us and to conceptualize it on its own terms.

COGNITIVE INFRASTRUCTURES

25:54

Cognitive Infrastructures

This was the basis of the Cognitive Infrastructure Studio, which we ran in London this past summer. Beginning with the history of philosophical thought experiments that have driven the development of AI, an exploration of those which might come next. Of course, the stochastic parrot was our mascot for the studio.

Cognitive Infrastructures refers to all of the ways in which intelligence exists in the wild as a landscape scale phenomenon, both the evolution of natural intelligence and ultimately the evolution of machine intelligence, and, more importantly, their coevolution together. Intelligence in a dynamic relationship with other kinds of systems, not isolated as a kind of brain in a Petri dish. And so indeed does machine intelligence, what we colloquially today call robotics or in the future, maybe just physical AI in the future, in many respects is where we are interested in thinking about this, and where it's headed. Not only as a design problem but as a philosophical problem as well.

As our infrastructure has become increasingly cognitive, the inverse is also true. Artificial cognition becomes increasingly infrastructural. And it's this dynamic between the two that animated the research of the studio, which we held at Central Saint Martins, across the street from Google DeepMind, which we did a lot of collaboration with.

Some of our studio researchers are here this evening. They began with a series of, hopefully, successful prompts, of how we may break down and approach this issue in a more open and generative way. Our researchers were, as you know, from all over, DeepMind, Oxford, Cambridge, but more importantly, disciplinary breadth, not only computer scientists, but filmmakers, philosophers, science fiction authors and so on. Working together in quite intense groups of three. Over the course of a month we generated about a dozen papers, of which you see several of them here.

"Whole Earth Codec" was also a project from our 2023 studio that looked at different forms of training data. "Frontiers of Endosomatics" around cognitive prosthetics. "Traversing the Uncanny Ridge" on different structures of how to define novelty in relationship to generative AI. "Human AI Interaction Design" is one of the key areas of our interest and has been a core for both studios. The relationship between AI and scientific epistemology animated at least a couple of the projects. As I said, about a dozen papers over the course of that time. Human AI Interaction Design is also one of the key areas we've done a lot of thinking about, particularly, the way in which HCI was based on a kind of spatial intelligence, like above, inside, below, whereas Human AI Interaction Design is really based on sociocultural cues, which means that it's open to all kinds of neurotic dispositions which are worthy of further research. AI in scientific epistemology, was the basis of a number of projects that we did at the same time. And again, these are all available on our site for your proper perusal.

RECURSIVE SIMULATIONS

29:22

Recursive Simulations

Next, Recursive Simulations. Simulations are not a new philosophical concern. They are arguably, from the basis of Western philosophy itself, a kind of suspicion, if not outright paranoia, about simulation. Plato's Allegory of the Cave is about a suspicion of what is and is not a simulation, and how philosophy should comport itself in relationship to the difference.

But it's not the only way in which the notion of simulation, of the shadow, of the double, might be perceived. Other traditions may prioritize the shadow, as actually the part that allows us a deeper access to the real.

The political implications of this are profound. As mentioned, the scientific idea of climate science as an epistemological accomplishment of planetary computation means that what we recognize as climate politics, an incipient planetary politics, is based upon a mobilization against the negative implications of a supercomputing simulation of the future. Which puts tremendous political efficacy and power, agency, in the role of simulations itself. Such that we might think of climate science kind of politics, more importantly, as one organized around the agency of supercomputing simulations.

This is not the only way in which there's an important relationship between simulation and governance, of how it is that complex systems govern themselves by anticipating their future state. The book I wrote during and about the pandemic, *The Revenge of the Real*, dealt with this in considerable detail.

The thing about simulations is that they are both one of the ways in which it is possible to understand the dynamics of complex systems that are not directly perceivable, such as a climate, and therefore give us an access to the real that would be otherwise impossible. But they are also the form of technology, both hardware and software, that allows for the production of completely unreal forms of perception. A kind of derangement from the real. And it's the same technology. And this is perhaps the concern that Plato's paranoia points to.

It's also one that we are all, each of us, enrolled in all day and every day. Everyone in this room has multiple digital profiles, shadows, if you like, doubles of you, that are housed in what is really the primary architecture of our time, data centers, that's where your doubles live, your shadow in the mountain somewhere nearby. And as you interact with the stack and interact with the platforms, it is really that double that is interacting on your behalf.

For example, when you go through the airport, and the little man at the gate stops you and wants to scan you and verify whether you should be allowed through, what's being evaluated is your double, is your shadow. It may feel like he is evaluating you, but he's not. The system, the interface is evaluating your double. Your shadow is being interrogated for its propriety, and if your shadow passes, then you are allowed to pass through the door.

Now, another way in which we are thinking about simulation, particularly about AI, is generative AI. Here in the middle is Trippy Squirrel, by Alex Mordvintsev, probably the first image of modern generative AI, which came out less than a decade ago. Alex was the developer of the Deep Dream software. In less than a decade, Trippy Squirrel has gone to AI generated video, which is nearly perceptually indistinguishable from photographic video. It has evolved so quickly that it became the basis of a long standing strike in Hollywood over the future of the culture industry. One that will take decades to fully play out.

Our work on this is then divided in that way. On the one hand, we're interested in simulation as a Model Epistemology, how it allows us to understand things and understand complex systems that would be otherwise not understandable. By the way, Nicholas's contribution to the first issue of the journal traces over the history of simulations, particularly in relationship to architecture. But at the same time, we also recognize what we might call something like simulation anxiety. That is: the more that we interact with simulations in a meaningful way, the less certain we are of the boundaries, the uncertain liquid, fluid boundaries between sim and real. But

as we do so, these movements and these kinds of negotiations become more difficult. And indeed, the more times we ask the question, "is this a simulation or not?" the less sure we are of the answer.

Indeed, it is this passage between sim and real, the sim to real gap as it's called, the hybridization of simulation and reality, that's of real interest to us. Not the separation and the clean ontological break between the thing and its shadow, but their weird amalgamation that we see as the path forward in one way or another.

And again, back to the black hole image, which you remember as in terms of model epistemology, that our understanding of black holes begins first in the 20th century with a mathematical model that they must exist, an increasing fidelity of the physics that would describe their parameters and where they are and how they would work, and therefore it becomes possible to build an Event Horizon Telescope and point it in the right direction, and to know what to see, because it had been simulated in advance, and this kind of cognitive foresight that simulations allow is how they work, and they function as existential technologies.

HEMISPHERICAL STACKS

35:54

Hemispherical Stacks

Next is Hemispherical Stacks. This is our area where we explore the relationship between planetary computation and geopolitics. Indeed, "the whole age of computer has made it where nobody knows exactly what's going on".

Some of what's going on is the Chip Wars, that the ability to build a society is dependent increasingly on the ability to build the Stack that is the armature and anatomy of that society. And so the race to, in essence, construct that economy or social system is the race to construct that backbone.

Right after The Stack was published, when the Hemispherical Stack essay first came out in 2018, I think, it was already clear that there was not one planetary computation, that it was already splitting into different isomorphic and even antagonistic stacks. A multipolarization of planetary computation was at work. Why would this be so?

Well, one reason is because governance with a small g, in the cybernetic sense of how a system senses and models and recursively acts back upon itself, was increasingly invested into computational systems. The computation was no longer only something to be governed by laws and dictates. It is governance in an imminent sense. It is how systems recursively act back upon themselves. And so this shifts and splits a bit from our normal understanding of the political, as a contestation over the symbols of power, into something much more physical and, as said, imminent and immediate.

At the scale of geopolitics, what we are seeing is not only a close tracking between the multipolarization of geopolitics and the multipolarization of computation. They are indeed the same thing. We are in the midst, we argue, of a kind of fundamental transformation between eras of planetary computation, one in which that stack is being replaced, a bit like Theseus's boat, layer by layer, by one that is dependent not on classical computing architectures, but on neuromorphic architectures.

We break this down into a number of different areas of research, for which as you'll see, we do a lot of collaborations with China, for obvious reasons, around this. The last one you see of stories and scenarios, this is how we are doing the research in a way.

You see some of the projects, Cloud Megaregionalization, which is about, essentially, the regional planning school, like building the region around its ability to produce something for the cloud, whether it's mining or chip design or something like this. And this whole, this structure to planetary scale. Planetary Counter-Simulation, which is looking at, not just the ways in which geopolitical adversaries produce simulations of themselves and of their adversary in order to negotiate their relationship. But the way in which they try to distort the adversary's simulation of themselves, and in fact, in this way, in which that gamesmanship of real and not real becomes the basis of, the dance of hopefully stable negotiation. So again, this work begins with an essay from 2018, which inspired a book by the Dutch National Design Academy called Vertical Atlas, which contained a number of essays and projects trying to map multipolar computation across the world.

It also inspired some of my work in China. As I mentioned, I was a visiting professor at NYU Shanghai. And some of that work is culminating in a book that Anna Greenspan and Bogna Konior and I co-edited coming out later this year with MIT and Urbanomic called Machine Decision is Not Final: China and the History and Future of Artificial Intelligence, which gathers a number of science fiction writers and philosophers to think through this extraordinarily weird and unique history of cybernetics and AI in China.

And then third in this triangle, The Vertical Atlas, Machine Decision is Not Final, is Dispute Plan to Prevent Future Luxury Constitution. If you see me afterwards, I'll explain where the title came from. Which is a book of fiction, speculative fiction. It

was published by e-flux and Sternberg. And you think of these three together, and you have a kind of, perhaps, new genre.

We are collaborating on this particular project with Chen Qiufan or Stanley Chen, as you may know him, the author of *Waste Tide*, wonderful Chinese science fiction author, coauthored a book called *AI 2041* with Kai-Fu Lee.

It also is one of the basis of increasing work we're doing on automation. Stephanie Sherman, also the associate director of Antikythera and her work of thinking about automation, the evolution of automation as a kind of natural planetary process that becomes artificialized through platform systems themselves, gives a lot of shape and weight to this discussion.

We've been doing a lot of deep collaborations with Chinese science fiction authors, hosted a salon at Peking University a couple of years ago and continue to work closely with them.

As I mentioned, we imagine this as a kind of new genre. Scenario planning is one of the ways in which we give literary shape to counterfactuals. Another way in which we give literary shape to counterfactuals is through science fiction. What we call scenario fiction is a way in which these are brought together. In a way that might, through the amalgamation and accumulation of multiple scenarios, we will gather hundreds, and maybe a kind of mosaic emerges that might give us a little bit of sense of where all this is going. This will also be all collected together and published in the book series with MIT press.

PLANETARY SAPIENCE

42:04

Planetary Sapience

All right. The last of our areas and perhaps the most philosophically interesting is what we call Planetary Sapience. The term emerges from an essay I wrote in *Noema* a couple of years ago called "Planetary Sapience," which looks at the paradox of intelligence. That is, at the same moment, historically, that human sapience came to really comprehend, in significant ways, its own anthropogenic agency, its own evolutionary trajectory, the age of its host planet, fundamental understandings of what the world is, that were obscure even into the early 20th century, is also the moment in which we figured out that it's on fire. That, in essence, we woke up in a house on fire, that is on fire because we woke up. And so the question is, is there a way in which this can be reconciled? This is the paradox of intelligence.

Some of the work that we're publishing from this area includes a longer essay by the Polish philosopher Bogna Konior on Lem's idea of existential technologies. An extraordinary piece by our friend the Cambridge historian Thomas Moynihan, who's here this evening, on the history of the moments by which, bit by bit, planetary sapience came to form our understanding of how it is that the planet comes to comprehend itself in this way. This is also from a lecture that Thomas did for our London studio, which is also on the Antikythera site.

We're also, as I mentioned with Turing, redesigning and republishing some classical pieces. And piece another is by Vernadsky, who you may know as the coiner of the term biosphere, as well as noosphere, and his work really begins the idea of understanding the evolution of intelligence as a geologic phenomenon. And we will be publishing this key work as well.

The real crux of our interest in this area has to do with, as said, the role of computation as a fundamental existential technology. Over the past century, computation has become the primary existential technology by which we understand the world. It's not just instrumental, it has helped us to comprehend, in new ways, how we think, the effects of what we do, how we make sense of the world, how life works, and indeed where and when we exist. All of those transformations are dependent upon computational systems that are themselves evolving. And I mean this in a literal sense, which I hope to show you.

We can think of intelligence and complex sapient intelligence as, again, something that our planet does, in that it produces over a long period of time, it folds its matter into particular shapes of primate brains and giant prefrontal cortices, and that through this folding, it was able to perceive things about its own processes that it would not have been able to understand otherwise. That it made the substrate, if you like. This, by the way, is not my refrigerator, so don't worry.

What does this mean? Well, we can think of this in the context of, where does one put AI in this particular context, of this folding of matter? Well one way to think about it might be this. Is that we, the fire apes, figured out how to make rocks think. We figured out how to take bits of rocks and metal, very finely folding them, running electricity through them, and now the rocks can think things that previously only primates were capable of doing so. This is newsworthy.

What it means is that the substrates of complex intelligence now include both the biosphere and the lithosphere in various amalgamated, codependent evolutionary tracks. It also means, in a certain sense, that we need to think about both the ways in which computation was discovered, and it was a natural phenomenon in many

ways, and also the ways in which it was invented, that it was artificialized, and that there is a dynamic between these two that continues in uncertain ratios.

There's much discussion for good reason about what we mean by "intelligence" in relation to artificial intelligence, but perhaps less attention to what we mean by "artificial". One way to define it is as an anomalous regularity, as you see in this image of a Japanese forest. We can all tell which trees were planted by a human. But I want to introduce a different, perhaps more general definition of what we mean by artificialization itself, and to locate it in that story of the fire apes.

What we call life is commonly defined, very incompletely, through a capacity for autopoiesis, the ability to internalize the environment in order to replicate oneself, in order for the system to replicate itself. But this is paired with a different capacity, which probably emerges much later, which is allopoiesis, the ability of a system to be capable of producing something that is extrinsic to itself, and to do so in ways that allow it to capture more energy and information and matter, so as to serve the purpose of autopoiesis. In the most general sense, autopoiesis is what we mean by artificialization.

In ways I don't have time to go into here today, this implies a lot of different correspondences and differences between what we see as an evolved intelligence and an artificialized intelligence, which correspond in some ways to this, the dynamic between natural computation and artificial computation. So let me then unpack this a little bit this topic of the evolution of the artificialization of intelligence. And I'll try to make the case for this perhaps seemingly odd conclusion that AI is itself the artificialization of artificialization.

I like to draw. I like to think by drawing. This is a sketch I made at our London studio this last summer. But let me try to decipher it for you a little bit.

Another one of our core collaborators within the program is Sara Walker, who with Lee Cronin developed what's called assembly theory. One of the key arguments of assembly theory is that evolutionary selection does not begin with biology. It begins at least with chemistry. That the phase space of possible molecules is just too vast, to be fully searched, and that there is a dynamic by which the space of possible molecules that can assemble, that are stable and more importantly, that can replicate themselves, is part of this process. And so, begin there.

This, then, is the scaffold that makes, again, this is not science, this is a napkin sketch so please take it as you will. This makes possible, as a scaffold, a process by which certain molecular arrangements are really, really good at replicating themselves autopoetically, which we might colloquially categorize as a form of life, which itself is the scaffold by which to get really, really good at autopoiesis evolution selects for the capacity to get really, really good at allopoiesis. To capture energy, and to get really good at capturing energy, information, and matter to replicate yourself, that depends on the allopoetic capacity to capture energy, information, and matter. Now to get really, really good at artificialization requires something like complex intelligence, to coordinate, to have foresight, to imagine counterfactuals and to compose those counterfactuals in a particular way. All of these in one way or another are something that is selected for, each of which, importantly, is a scaffold for what comes next. To get really good at complex intelligence, something like symbolic language becomes very important. It allows for a tertiary memory, as Stiegler called it, it allows for a long term collaboration between people who may not know each other. And it builds on this scaffold of intelligence. And so no surprise then that AI might be built on training on symbolic language. The question then is, what is AI a scaffold for? This process is not ending. The process doesn't end with AI. AI is a scaffold. And what is that a scaffold for? And so on and so on.

In order to ask this question, we also need to recognize that there's inevitably a kind of recursive process here that not only does symbolic language function as the scaffold for AI, but as AI emerges, it transforms the dynamics of symbolic language. In the same way that the emergence of symbolic language transformed the dynamic of intelligence. There is a recursive process in this, and so that ultimately, transively, AI is this artificialization of artificialization. Or put differently, sapient human intelligence was the thing that made the artificialization of intelligence possible, made AI possible, and it is itself that artificialization of the thing that brought it about. What we think of is the whole scope of anthropogenic agency and all the kind of causal dynamics of complex human intelligence, the comprehensive artificialization of the world, that force is now itself artificialized.

But really the important question is, what is this a scaffold for? And what is that a scaffold for? And what is that a scaffold for? This isn't the end of the story. This is hopefully the midway point.

LIFE, INTELLIGENCE, TECHNOLOGY

One of the current discussions that we've been developing in the program, with Blaise and Sara in particular, but others as well, is identifying a kind of convergence. Convergent evolution is one of our areas of interest, a convergence between what we call life, intelligence, and technology. And to some extent, this begins with the

fact that many of the things that we might categorize as that is life, that's intelligence, that's technology, are becoming more physically amalgamated, hybridized, cyborgized. These are some of Michael Levin's xenobots. And so as they become physically amalgamated, there is a kind of conceptual amalgamation of the idea of what life, intelligence, and technology is.

Blaise's book *What is Life?* that you see is also based on some of his work on artificial life, in which he discovers a process that looks very similar to the dynamic that I laid out for you. And what do we see here? And this is a program that he wrote in the appropriately named low level programming language, Brainfuck, in which we see something like assembly theory in action. Those clusters that are able to be stable and are able to replicate themselves, become scaffolds for more complex structures, and on and on and on. Now for Blaise, he'll make the argument that this is not just a metaphor for life, it is life itself. But at the very least, it is something that perforates the definitional and conceptual boundaries between the two. These are some other visualizations of the Brainfuck work by Alex Mordvintsev, the same fellow who made *Deep Dream*.

We might ask the question, ultimately, what do we mean by life, what do we mean by intelligence, and what do we mean by technology in relation to all three of these together?

Some of it, I think, needs to be based on this, that the evolution of intelligence is not just something that happens on the planet, it is something that happens to a planet. It is part of the evolutionary dynamic of a planetary system itself. Not only are the dualist ideas of mind-body separation not particularly helpful, but nor is the idea that the emergence of abstract rationalization is somehow something separate from ecological processes. It's not. It's something that ecosystems do. They select for this capacity.

A bit, then, on where assembly theory may fit into this. This is the now infamous paper in *Nature* that Lee and Sara published on assembly theory. I'm not going to try to explain the whole thing to you, but it's really two key ideas to keep in mind. One is I already mentioned that the idea that selection begins with, at least with chemistry, not with biology, and two, it's a kind of theory of evolutionary time and complexity, that something simple becomes a scaffold for something more complex, and so on and so on and so on and so on. There's a bit of John Maynard Smith in this. And that you can, in essence, model the amount of evolutionary time in an object by quantifying the amount of complexity that is in that object. Because that will tell you how much time was necessary for those nested scaffolds to appear. And they quantify this as what they call the assembly index, which has to do with the amount of complexity in the object and how many instances of this object exist.

If you want a deeper dive into all of this, this is Sara's talk at our London studio last year, the video of which is also on our site. We're also publishing a longer essay by her and Lee on assembly theory also as a theory of technology.

Let me end with this. Life, intelligence, and technology are converging. Where the things that had stayed stable in these categories are misbehaving, they're migrating to other areas and such that a different kind of conceptual language may be needed. So what philosophy going back to Aristotle has seen as fundamentally different categories may be seen as faces of the same wonder. And if so, a science that is more general, may be being born, if we midwife it well. Cybernetics, we are at MIT, after all, foreshadowed this and made what came later possible. It is an intellectual scaffold. What comes now, will surely be yet more momentous, because the capacity to realize it is orders of magnitude more powerful. The atom, for better or worse, splits again. So, we can draw upon some contemporary definitions of each of these, and understand how this discursive convergence is in our midst.

Let me end where we began, which is the question of cosmology. One of the other unfortunate things that I think characterizes our moment is a kind of split between the astronomic scientific sense of cosmology—what we actually know about who we are, when we are and where we are, how we are—from the cultural sense of cosmology, of what it means to be who we are, and what should we do now? And it is this split between the two that I think is not only unfortunate, it's probably dangerous. And so I think any philosophy of technology program must seek to contribute to the remedy of this bifurcation. It would do so because there's a lot at stake.

The question, though, is to what extent is intelligence actually adaptive? In the short term, it's obviously extremely adaptive. It allows tremendous advantages, as the ways in which it allows for and is allowed by complex allopoiesis. But in the long term, is it self extinguishing? Is that paradox of intelligence, of waking up with a house on fire, one in which it is possible to make our way through?

Maybe a better way to ask the question is not is it or isn't it, intrinsically, but rather what would make it adaptive? What would be the preconditions for its adaptive-ness? And to Nicholas' point about design as designation, how can those preconditions be artificially realized? This is both the philosophical project and the engineering project at the same time. That may be the real Copernican trauma at hand.

With that, let me close by making a few announcements. A full article-length version of the research and ideas that we've worked with is available on the Antikythera site. You can go through this in considerable detail. And as I also mentioned, this is just one of the many collaborations that we are undertaking with MIT.

Today represents, in essence, the announcement, the preview of the digital journal as well as the book series that we're undertaking with Amy Brand, Nick Lindsay, as well as, Noah Springer. In May, we will launch the complete first issue of the journal. We want to also give real thanks to Channel Studio, who are our collaborators in the design of the digital journal platform.

The first issue will include many names and faces with which you are now familiar. And each one of these will develop a kind of bespoke, beautifully designed representation of their research in the journal. Our authors and I think some of the most interesting, innovative digital design studios, paired up to develop each of these quite lovely and extraordinary initiatives.

In the meantime, please go to antikythera.org, you'll find all about the program. Link.antikythera.org is where you'll find links to all of the projects and, and work that I've spoken to this evening. And at journal.antikythera.org, you will see a preview of the digital journal that we are developing with MIT press. The full journal will launch in spring of next year. With that, thank you very much. Thank you so kindly for your patience. And I look forward to continuing the conversation.