



Welcome to

FOUNDATIONS

Welcome to Foundations, the first of five tutorials in the series.

In less than an hour we will cover some important material that will help you understand machine learning and artificial intelligence, the technologies they are built on, and why I think they are important.

FOUNDATIONS

- **Digital**
- **Boolean logic & Algorithms**
- **Software**
- **Intelligence**
- **Digital Intelligence**

We'll start this tutorial by challenging our assumptions and looking at the fundamental principles of digital technologies, leading us to a better conception of what software systems, like artificial intelligence are capable of and what makes them so strange and powerful.

We'll talk about some eternal questions like the nature of intelligence and freewill. These deep questions are being explored in some of the latest research, and we'll see how studying artificial intelligence helps us to understand all intelligence, including our own.

We'll wrap up with how making intelligence digital gives it the strange properties of all digital things and the great potential and dangers of digital minds.

Let's begin!

Credits

- [Co\(AI\)xistence](#), [Justine Emard](#)



DIGITAL

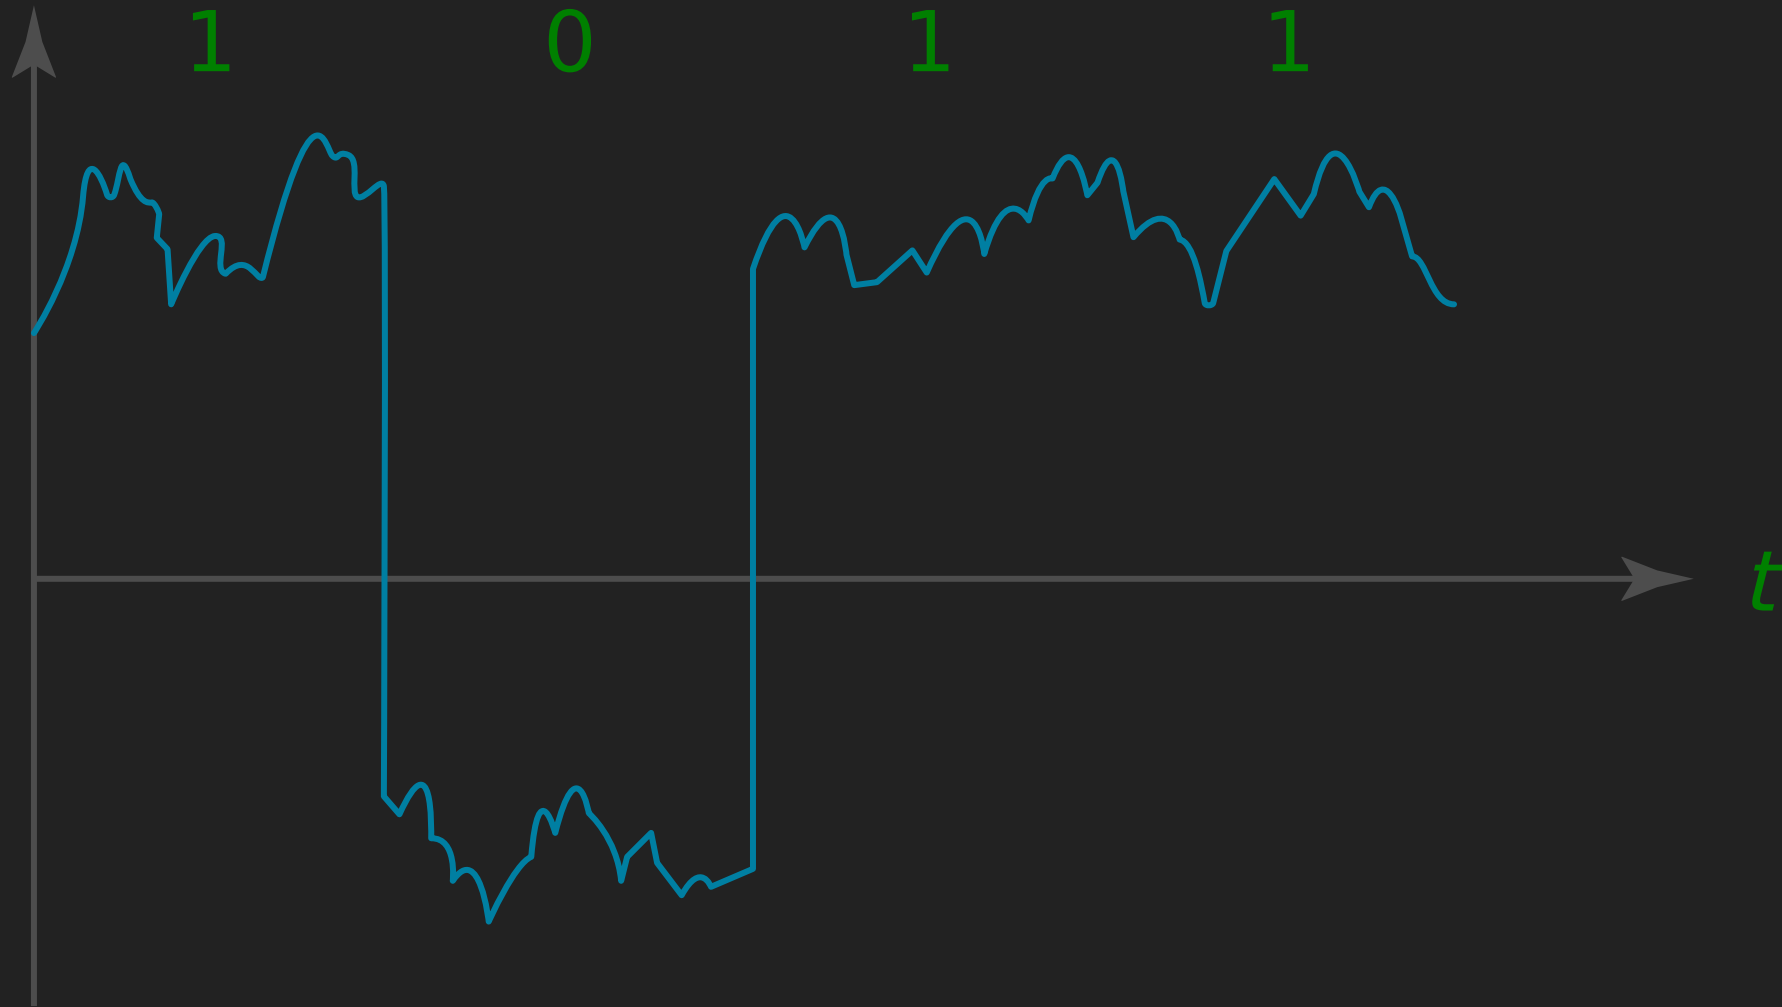
We are told we live a digital world. And we've gotten used to living there, and in many ways, we've forgotten how strange it is.

Our understanding of what digital really is has also been clouded by powerful economic interests who have tried to make digital things act like the scarce physical things they used to sell. Indeed, at least at the level of electric signals the digital world is still analog - still part of a reality it was designed to transcend.

So at its lowest level a digital signal is still analog. Being analog it is subject to noise - it has minute fluctuations - it is a messy or noisy signal.

Credits

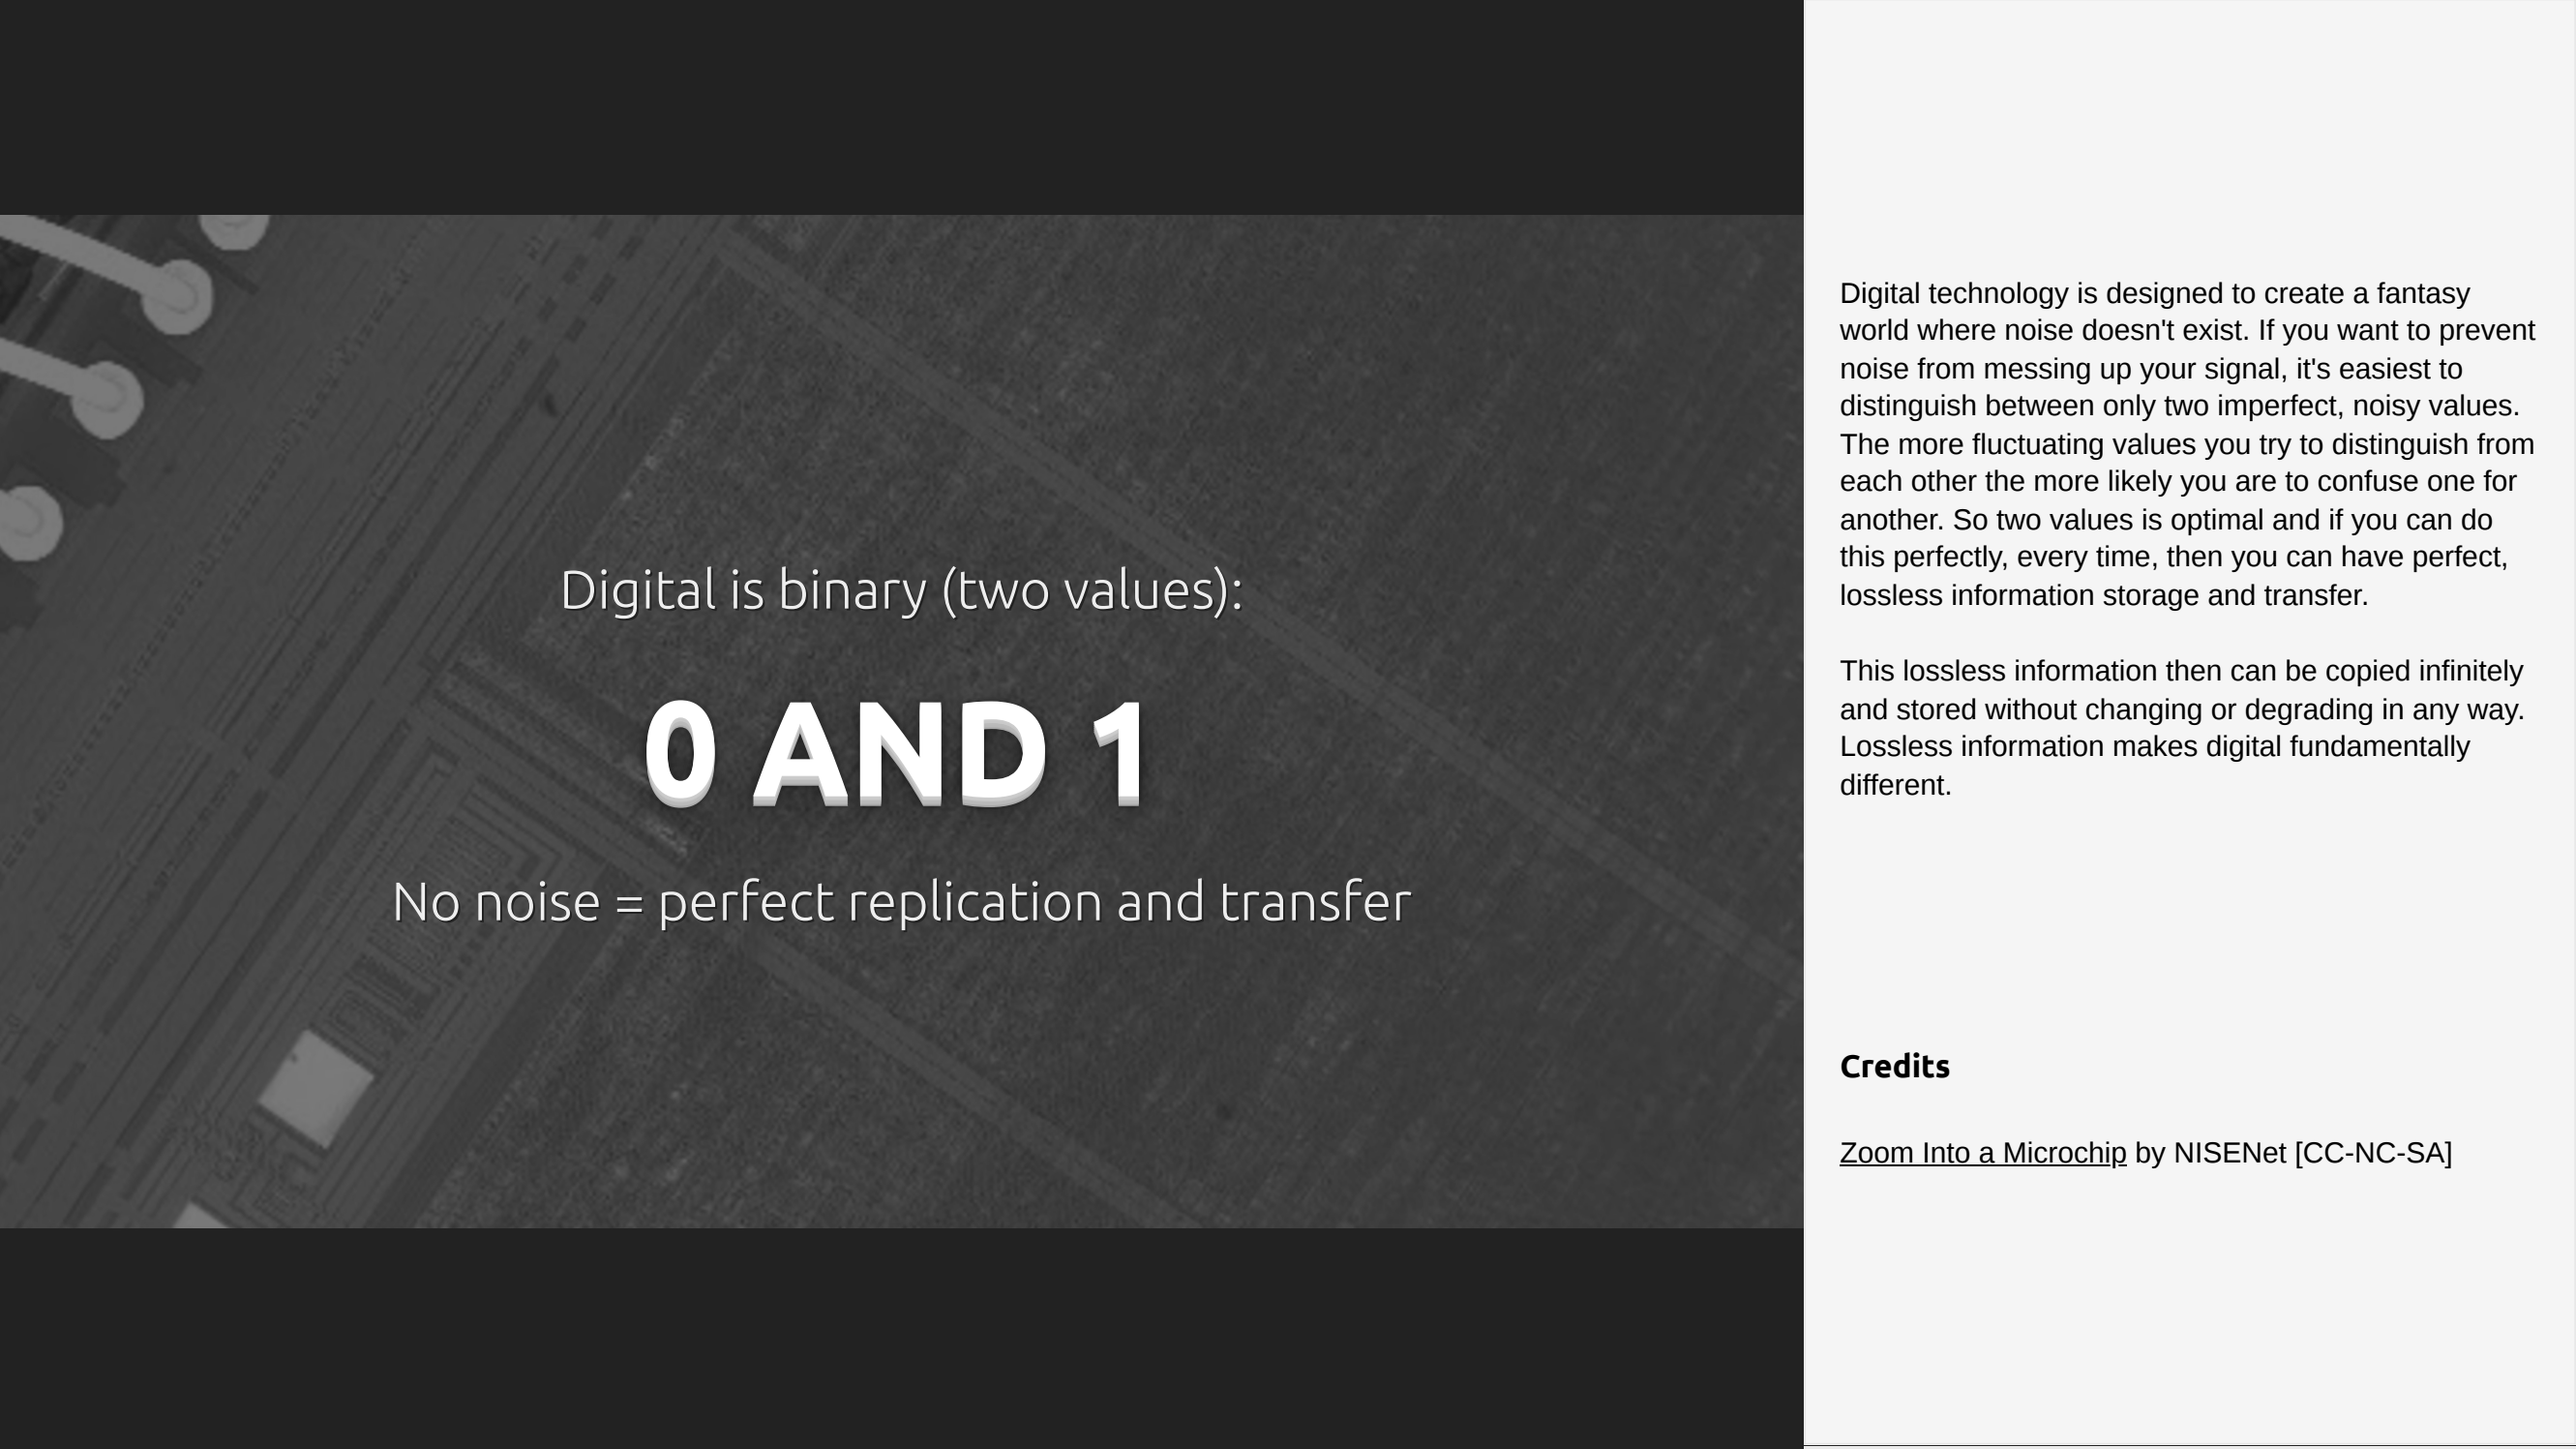
- [Video by Pressmaster from Pexels](#)
- [Video by Luis Quintero from Pexels](#)



The blue line in the image represents this signal. You can see it alternating between a 1 and a 0, the two digital values, but it is still affected by noise - those squiggly bits. But a digital signal is designed to overcome this noise, not to let it affect the signal.

Credits

- https://en.wikipedia.org/wiki/Digital_signal#/media/File:Digital-signal-noise.svg

A dark, grayscale image of a microchip or circuit board, showing intricate patterns of lines and components. The image is used as a background for the left side of the slide.

Digital is binary (two values):

0 AND 1

No noise = perfect replication and transfer

Digital technology is designed to create a fantasy world where noise doesn't exist. If you want to prevent noise from messing up your signal, it's easiest to distinguish between only two imperfect, noisy values. The more fluctuating values you try to distinguish from each other the more likely you are to confuse one for another. So two values is optimal and if you can do this perfectly, every time, then you can have perfect, lossless information storage and transfer.

This lossless information then can be copied infinitely and stored without changing or degrading in any way. Lossless information makes digital fundamentally different.

Credits

[Zoom Into a Microchip](#) by NISENet [CC-NC-SA]

Digital technologies are useful and strange:

- **Everything is copied.** Movement is done by copying and then deleting.
- **Copies are perfect** and indistinguishable. They act like *numbers* not things.
- **Software is deterministic** and thus repeatable by default, true randomness is difficult.

Digital technologies are both useful and strange.


First, everything digital is copied. In physical reality objects move, but "moving" in the digital reality is done by making a second copy and then deleting the first. Usually, we don't even bother deleting because that means the extra work of copying zeros over the data, instead we just mark that space as available for new copies. Imagine the world you inhabit but every time you picked up your phone up from the table you left a copy of it sitting there.

Second, copies are perfect and indistinguishable from each other. In physical reality copies are made of different atoms, but digital copies are the exact same thing in two different places. Imagine digital stuff more like numbers than physical objects; when you write down the number 5 it has the exact same value as every other 5 written down by everyone else. Fives aren't copies, they are the same thing. Every bit of digital data, software, music, image can be thought of as just a very large number.

Finally, software built with digital technologies is deterministic - it is repeatable by default. Each time you run it will do the same thing. True randomness actually becomes difficult.

Credits:


- [Life in life](#) by Phillip Bradbury [CC]



Credits

- [*Mutations* - William Latham \(1991\)](#)
- [*Code Profiles* - W. Bradford Paly \(2002\)](#)
- [*Hello World!* - Iamus Computer \(2012\)](#)

BOOLEAN ALGEBRA



In the middle of the 19th century George Boole published his work that introduced an "algebra of logic" which was later called Boolean algebra.

Credits

- https://en.wikipedia.org/wiki/George_Boole



0 = False

1 = True

AND

OR

NOT

1 AND 1 = 1

1 OR 1 = 1

NOT 1 = 0

1 AND 0 = 0

1 OR 0 = 1

NOT 0 = 1

0 AND 0 = 0

0 OR 0 = 0

The only values in Boolean algebra are true and false, which map to binary values 1 and 0. Instead of addition and multiplication the main operations are AND, OR and NOT. AND and OR combine values in logical ways.

ANDing two values is only true if both are true. ORing two values is true if either of the values is true. NOT allows for the inversion of a value: NOT true is false, and NOT false is true.

There are few other more complex operators that are important, but for now we just need to know that that Boolean logic was foundational for digital computer circuits.

Credits

- <https://simplycoding.in/boolean-algebra/>



ALGORITHM

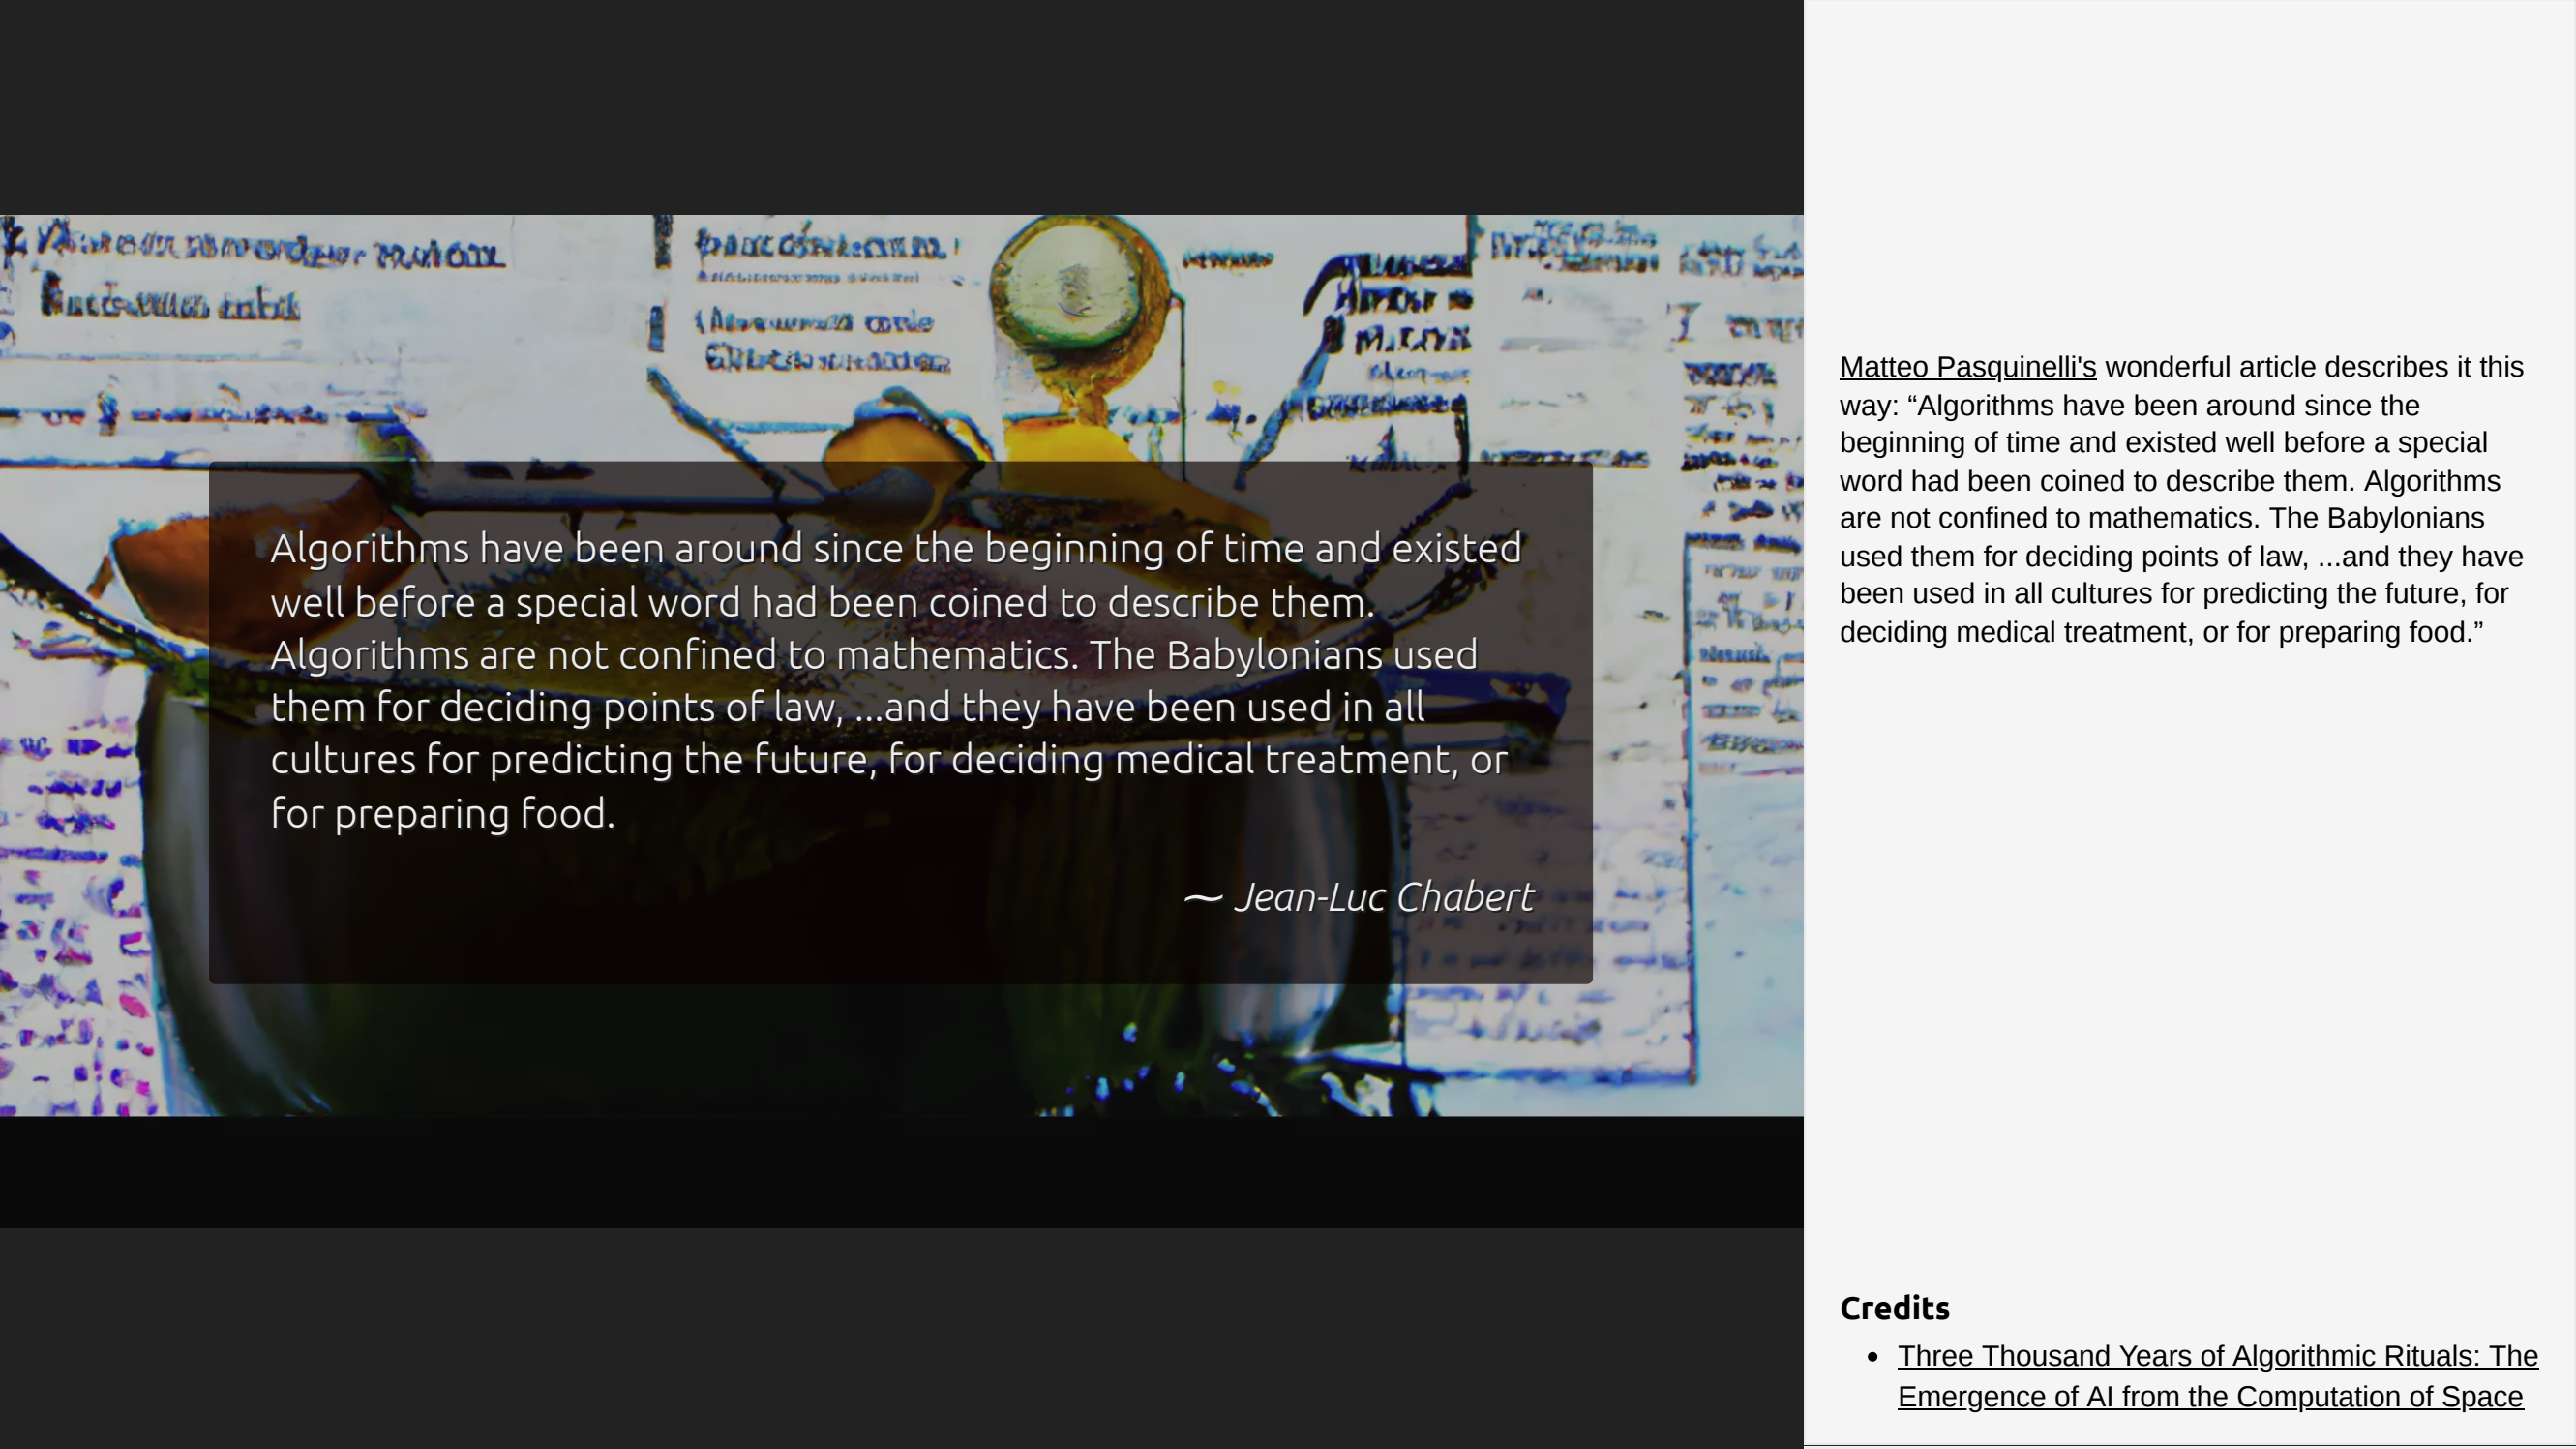
The other important concept to understand is that of an "algorithm". That's a fun word, with an origin from the Latinization of the Persian scholar al-Khwarizmi. His tract *On the Calculation with Hindu Numerals*, written in Baghdad in the ninth century, is responsible for introducing Hindu-Arabic numerals to the West, along with the corresponding new techniques for calculating them, namely algorithms.

Credits

- [Al-Khwarizmi - Father of Algebra](#)

References

- [Three Thousand Years of Algorithmic Rituals: The Emergence of AI from the Computation of Space](#)
- [The Weird Truth About Arabic Numerals](#)



Algorithms have been around since the beginning of time and existed well before a special word had been coined to describe them. Algorithms are not confined to mathematics. The Babylonians used them for deciding points of law, ...and they have been used in all cultures for predicting the future, for deciding medical treatment, or for preparing food.

~ Jean-Luc Chabert

Matteo Pasquinelli's wonderful article describes it this way: "Algorithms have been around since the beginning of time and existed well before a special word had been coined to describe them. Algorithms are not confined to mathematics. The Babylonians used them for deciding points of law, ...and they have been used in all cultures for predicting the future, for deciding medical treatment, or for preparing food."

Credits

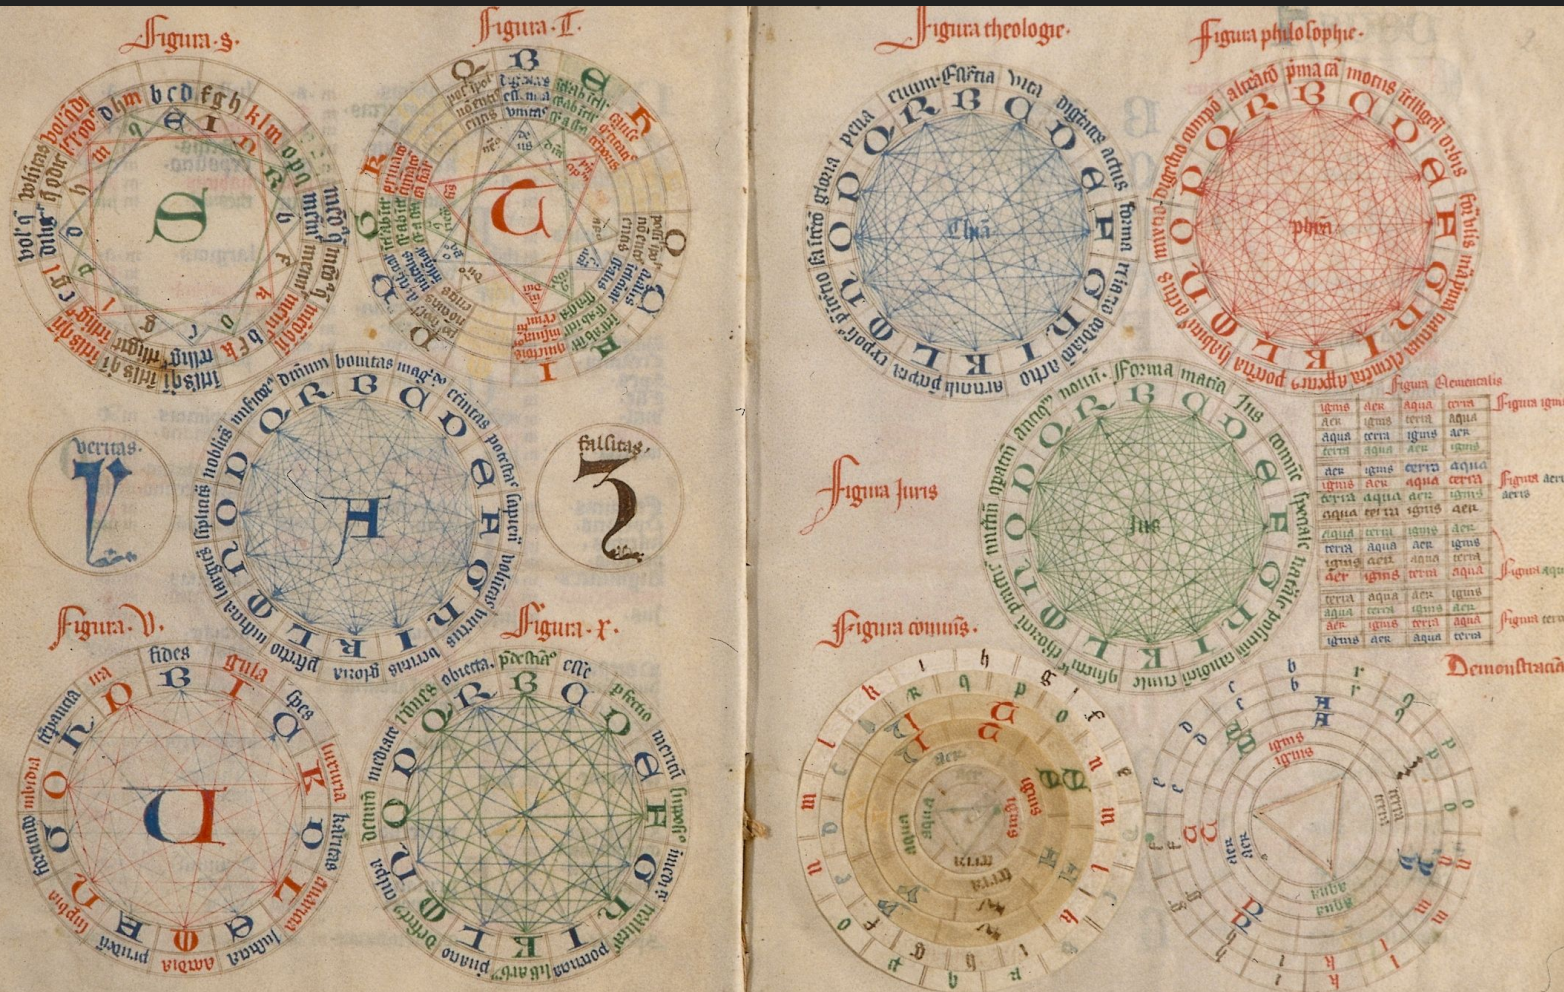
- Three Thousand Years of Algorithmic Rituals: The Emergence of AI from the Computation of Space


Algorithms are ancient technology, that emerged from ritual practices and the organization of social life. They are emergent processes that materialize out of a previous and spontaneous division of space, time and labour. Cultures speak of recipes, rules, techniques, processes, procedures, methods, and strategies.

Fundamentally we are talking about a systematic process of discrete steps that emerges from a repetition of the process. You want to be able to repeat the computation or the outcome, despite the participants or components being not quite the same.

Credits

- Maddox-Harle, Robert. Review of Dia-Logos: Ramon Llull's Method of Thought and Artistic Practice ed. by Amador Vega, Peter Weibel and Siegfried Zielinski. Leonardo, vol. 53 no. 3, 2020, p. 341-342. Project MUSE <https://muse.jhu.edu/article/757589>





Software algorithms

Our society is now inundated by software algorithms. Software programmers, like myself, have a bias against doing anything manually more than once. I often get the feeling that humans are terrible at following algorithms correctly, that software is a much more willing, accurate, reliable and quick comrade for my algorithms.

Credits

- [Code Profiles - W. Bradford Paley](#)


Universal Computing Machine

- Read / Sense
- Write / Act
- Compare / Feel

In 1937, Alan Turing introduced the idea of the Universal Computing Machine. His thought experiment involved a simple machine that could read, write, and compare symbols on a length of paper or tape with no understanding of the symbols being read and written other than changing the "state" of the machine. You could also think of these three actions in less mechanical terms: sensing, acting, and feeling.

Credits

- https://www.reddit.com/r/ColorizedHistory/comments/e6ugox/alan_turing_the_father_artificial_intelligence/



What he was trying to do was create the simplest possible machine, that was still capable of computation. In fact, given an infinitely long tape Turing proved that this machine could compute anything. Everything that could be rendered computable was computable using only these operations and a machine with an internal state or what is called a state machine. Each state has a rule for how to read, write, compare, move the tape and change to another state.

This creates a simple, universal algorithm for computation. And if sensing, acting and feeling are all that is necessary then perhaps all things capable of those properties that have in some sense an ability to compute anything.

Credits

- [*Turing Machines Explained* - Computerphile](#)



Computation

And what is it to compute? The transformation of inputs to outputs that follows an algorithm. There are competing theories about the physical phenomenon of computation, and an ongoing debate on the computational theory of mind, or computationalism, that holds that the human mind is an information processing system.


Credits

- <https://en.wikipedia.org/wiki/Computation>
- https://en.wikipedia.org/wiki/Computational_theory_of_mind

current automaton contents



rule 110 (01101110)



the next generation of the automaton




I'm going to set that aside and instead show you Rule 110. Rule 110 is an elementary cellular automaton explored by Stephen Wolfram in 1985 and proved to be Turing complete, or capable of universal computation in 2004 by Matthew Cook (which resulted in legal action that blocked the proof for several years).

There are 88 possible unique cellular automata, which are one-dimensional patterns of 0s and 1s that evolves according to a simple algorithm. Here the ones are black and the zeros are white. Similar to how a Turing machine functions, using Rule 110 you can iteratively create new lines of ones and zeros, the next line determined by the previous, by reading the left, current, and right positions and writing out a new value below, given the rules, which look like Tetris triangle pieces. Three values in, a new value writes to the line below, then move one step over.

Credits

- https://en.wikipedia.org/wiki/Rule_110



The patterns of 1s and 0s created by this evolution are neither completely stable nor completely chaotic. Structures appear and interact with each other in complex ways. Cook was able to prove that Rule 110 could emulate another known universal computation by use of "spaceships" and other structures. What the heck do spaceships have anything to do with this? Spaceships are just the common name for a self-contained structure that can move through the environment. In this case, the environment is a repeating stable pattern and the spaceship is a structure that can move amongst it. Once you have spaceships then you can start transmitting information. Once you can transmit then you can build read, write and compare... and universal computation.


Credits

- [*Rule 110 Simple Implementation Visual* - Ryan Edwards](#)

Rule 110 is arguably the simplest known Turing complete system. The essential takeaway is that complex algorithms are not required, instead there are particular forms of algorithms that surf on the boundary of stability and chaos that "come alive" with computational possibilities.

Credits

- [Spaceships from Conway's *Game of Life*](#)




How could something as simple as a Turing machine calculate something as complex as a modern video game or simulations of The Big Bang? The reason this is possible is that simple algorithms can be combined and built on each other. Digital algorithms inherit digital properties: perfect replication and execution. Thus, you can start to build infinitely long chains of algorithms. Like with the Turing machine, the infinitely long nature of the tape becomes critical, as does the speed at which you manipulate it. Today's pocket sized computers perform tens of billion operations per second and store hundreds of billions of bits. They have access to essentially infinite bits through the internet.

As the length and complexity of each algorithm grows, human understanding of it fades, until we add a layer of abstraction to hide all the underlying complexity. No human has, nor ever will have, a complete understanding of the full details and interactions of the hardware and software running on a typical phone. We are similarly ignorant of the human brain and even a single neuron in the brain.

Credits

- [Zoom Into a Microchip](#)
- [Allen Institute for Brain Science](#)



Using a Universal Turing Machine to calculate anything but the simplest of solutions is a pain, and so is programming machines in binary. For example, you might use digital circuits, 1 and 0, that use boolean algebra to build a "binary adder", that adds two binary numbers. With only a few adders you might interact with them directly, but once you have millions, they are hidden by further layers so that programmers don't have to worry about managing them, but just write a program to add two decimal numbers.

Credits

- [2 Bit Relay Adder](#)

BIOLOGICAL NEURON

Probed via depth electrodes


Halle Berry



Responds to photos of Halle Berry and Halle Berry in costume



Responds to sketches of Halle Berry




Responds to the text "Halle Berry"




CLIP NEURON

Neuron 244 from pen-ultimate layer in CLIP RN50x4

Spider-Man



Responds to photos of Spider-Man in costume and spiders



Responds to comics or drawings of Spider-Man and spider-themed icons



Responds to the text "spider" and others



We've also found abstraction in the brain at the level of individual neurons. There are neurons that respond to abstract semantic concepts around a high-level theme, rather than specific visual features. A famous example was the "Halle Berry" multi-modal neuron that responded to photographs, sketches and the literal text, "Halle Berry". OpenAI's CLIP-system also contains similar artificial digital neurons. We can see here the "Spider-Man" neuron.

Thus abstraction isn't just something we do because of complexity, it arises from complexity and computation. It is an important part of "knowing" or "understanding" or more specifically, semantic clustering, where lower level patterns are invariant to the conditions of measurement. For example, rotated images of Spider-Man will trace different neural paths but still always travels through the Spider-Man neuron. That the path to the Spider-Man can include text, and presumably audio in human brains points to the physicality of conceptual representations in minds.

Credits

- <https://openai.com/blog/multimodal-neurons/>
- <https://www.youtube.com/watch?v=dNRDvLACg5Q>

Insertion Sort - 9191 comparisons, 27398 array accesses, 0.50 ms delay

<http://panthema.net/2013/sound-of-sorting>

SOFTWARE




Abstraction built of layers of boolean algebra and algorithms running on digital circuits is the domain of software. Traditional software, written by programmers using programming languages, is the design and recording of these algorithms. Programmers essentially record their own decision-making into software.

While all media and art can be thought of as recordings of the creators decision-making, say, a musicians decisions about how to play their instrument at each moment, there is a difference - software can very easily take input when it is experienced. It "executes" on this input, whereas other media lacks input or is quite difficult to incorporate input into the experience. Certainly many performers would say an audience's reactions affects their performance, but it is rare for the input to really drive the art. Software is algorithmic by nature.


Credits

- [15 Sorting Algorithms in 6 Minutes - Timo Bingmann](#)



Credits

- [*Mr. Computer Image*](#) (1968)
- [*Cambiantes*](#) - Analivia Cordeiro (1976)
- [*Life in Life*](#) - Phillip Bradbury (2012)



Recording




Mapping

1. **Recording**: no input after initial recording
2. **Mapping**: takes input, produces output, but stateless
3. **Stateful**: has variables / internal state (Turing machine)
4. **Adaptive**: adapts to environment through time



Stateful



Adaptive

Software can be classified in various categories. Sofian Audry, in his book *Art in the Age of Machine Learning* has a nice classification that includes:

1. Recordings, like a digital image or video, that don't accept input when you view them.
2. Mappings, which take input and produces output, but are stateless. For example, inputting a photo to software that transforms it to a painting.
3. Stateful software has at least one variable, usually many more. Like a game.
4. Adaptive software adapts, which can also be considered learning, which is usually seen in machine learning or AI. Generally this learning is done before the software is used rather than the software learning as it is used.

Credits

- https://dsgiiitr.com/work/neural_style_transfer/
- <https://www.flickr.com/photos/mvives/albums/72157625126300997/>
- [MIT Robot Learns How to Play Jenga](#)



Input And Play

As an aside, one of the things I love about software-based art is the interactivity. When people spend time with art, it is often just a few minutes with a piece in a gallery, an hour or two with film, tens of hours for novels, etc. Software with enough depth can provide ten or hundreds of hours of interactivity. Games are the best example of that.



Interactivity also means the art can respond to the user, which allows for play and I'm a big fan of play. Play is fundamentally connected to freedom, as is art making. Both require freedom to experiment and relative safety to try out new ways of acting, thinking and being.

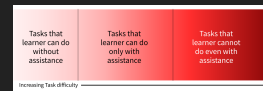
Credits

- <https://pxhere.com/en/photo/692638> [CC0]

The Goldilocks Zone

Most **A**dvanced, **Y**et **A**cceptable

Zone of proximal
development



Play allows for learning, and like practice, that learning is enhanced when the skills, actions or thinking being developed are not too easy nor too hard. Mediums that can adapt to the audience to find that sweet spot, familiar enough but different, what Raymond Loewy described as MAYA: Most Advanced, Yet Acceptable, and Lev Vygotsky called the lower limit of the "zone of proximal development" have a powerful ability to change minds. I call it the Goldilocks zone or the next best thing to learn.

If you're an artist who wants to help their audiences make profound positive changes to their lives, interactive software can be one of the most effective mediums.

Credits

- https://en.wikipedia.org/wiki/Raymond_Loewy
- https://en.wikipedia.org/wiki/Zone_of_proximal_development



DIGITAL ETHICS

1. Who decides what, and how, things get made?
2. Who gets to make it?
3. Who gets access to what has been made?

Let's take a step back, now that we understand the incredible power of digital technologies, abstraction and interactive software built with them, we need to ask a few questions. What are the economic, cultural and ethical implications of digitization? Anytime economics is involved, I like to boil it down to three "who" questions:

1. Who decides what, and how, things get made?
2. Who gets to make it?
3. Who gets access to what has been made?

Those questions have had, and will continue to have, a profound effect on my life. My favourite pithy answer to all three questions is "everyone and anyone", but that is really just a guiding principle and it is actually the specific details that are important. I believe that digital things need different answers than physical things, or more correctly, "digitalness" makes it much easier to make those answers more inclusive; to make every aspect more accessible and democratic.

Credits

- https://archive.org/details/American1958_2



DIGITAL LIBRARIES

Some artists and many corporations, don't agree with me, particularly in regard to digital things. For example, I believe in and strongly support the concept of libraries. Most people do.

Credits

- *Tour of the Sacred Library* - Ryan Moulton (2021)



The Pirate Bay



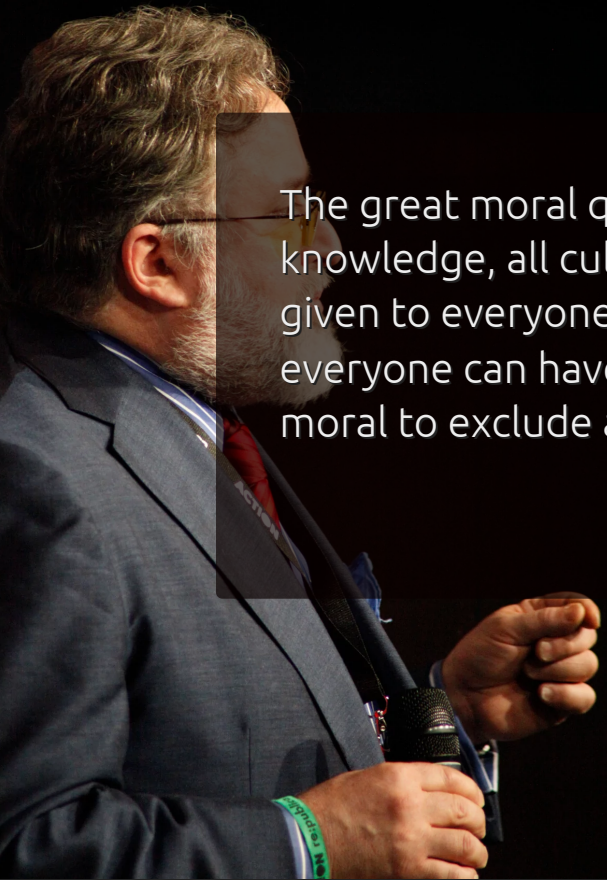
We have had the technology, since at least the origin of The Pirate Bay in 2003, to distribute all the world's digital information at a cost much less than we spend on the world's libraries. But we have decided to not do this. We could make every person with access to the internet into a digital billionaire, with access to everything forever - the largest anti-poverty action in history, likely for the entirety of history, but we choose not to.



We'd rather have only 2755 billionaires.

Credits

- https://en.wikipedia.org/wiki/The_World%27s_Billionaires
- <https://www.forbes.com/sites/kerryadolan/2021/04/06/forbes-35th-annual-worlds-billionaires-list-facts-and-figures-2021/>




The great moral question of the twenty-first century is this; if all knowledge, all culture, all art, all useful information can be costlessly given to everyone at the same price that it is given to anyone; if everyone can have everything, anywhere, all the time, why is it ever moral to exclude anyone?

~ Eben Moglen

The great moral question of the twenty-first century is this; if all knowledge, all culture, all art, all useful information can be costlessly given to everyone at the same price that it is given to anyone; if everyone can have everything, anywhere, all the time, why is it ever moral to exclude anyone?

Credits

- [https://commons.wikimedia.org/wiki/File:Eben_Moglen_\(7141739723\).jpg](https://commons.wikimedia.org/wiki/File:Eben_Moglen_(7141739723).jpg)
- re:publica from Germany, CC BY 2.0
<https://creativecommons.org/licenses/by/2.0>, via Wikimedia Commons



I found some good answers to these questions in the free and open source software communities.

Credits

- <https://www.fsf.org>
- <https://eff.org>



I found other good answers in decolonization and anti-capitalist movements. Almost always the answers were confounded by concerns of ownership. History has demonstrated that a critical questioning of ownership is a requirement for a just and ethical society. Who and what can or should be owned are fundamental questions about how we organize our economy and culture.

Credits


- *in Pursuit of Venus [infected]* - Lisa Reihana




When you use digital technologies you become intrinsically linked to the great moral question: when the value others gain dwarfs the cost of your sacrifice, at what point does it become criminal not to make that sacrifice?

I would extend this line of questioning further; at what point do you refuse to improve others lives because of the cost to you, despite the potential savings or gains they could generate for you in the future?

Put more simply, would you pay people enough to create new art and culture so that everyone could have access to what has been made?


- 
1. Who decides what, and how, things get made?
 2. Who gets to make it?
 3. Who gets access to what has been made?

In contrast to a society that accepts billionaires avoiding taxes from their sales of virtual goods that could be copied for free, I see digital technologies as the best opportunity for a once-in-the-history-a-species redefinition of wealth, sharing and sacrifice. To get as close as possible to answering "everyone and anyone" when we ask - who?



Credits

- *February: In the style of Frank Sinatra* - OpenAI Jukebox (2019)
- *Patterns of Life* excerpt - Julien Prévieux (2015)

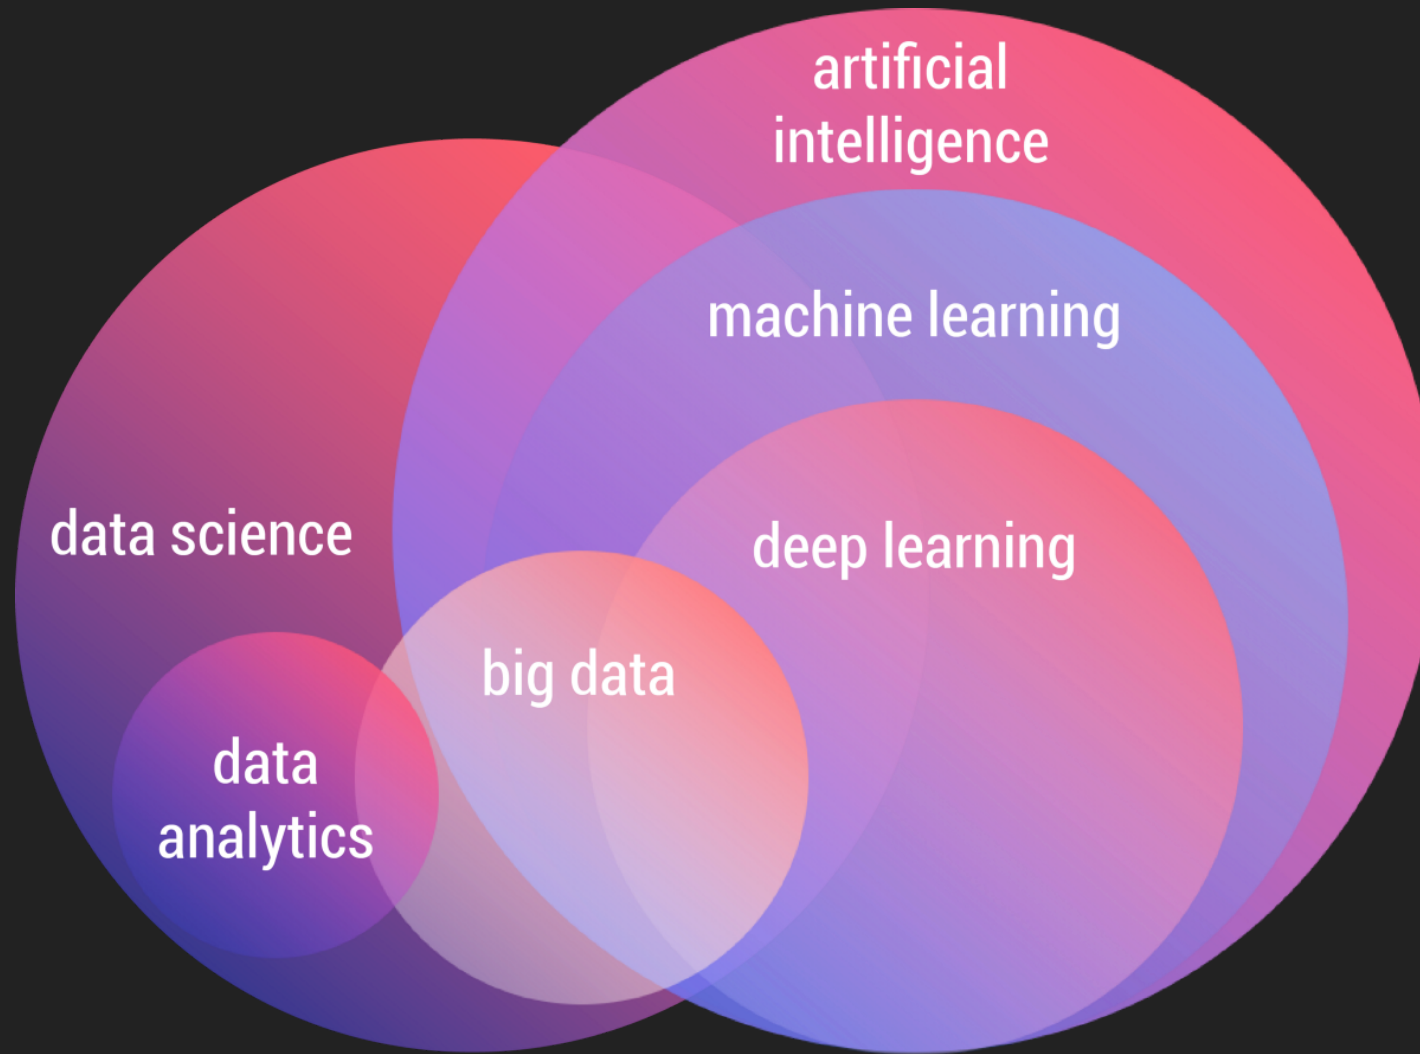


Alongside the development of digital computers was the development of evolving and self-organizing machines. Machines or software that didn't just have preprogrammed states they could be in, but adapted or learned what states were needed to more effectively interact with their environment. Machines that evolved.

Evolution conjures life. And thus were born the fields of artificial life and artificial intelligence.

Credits

- [Karl Sims' Evolved Virtual Creatures compilation](#)



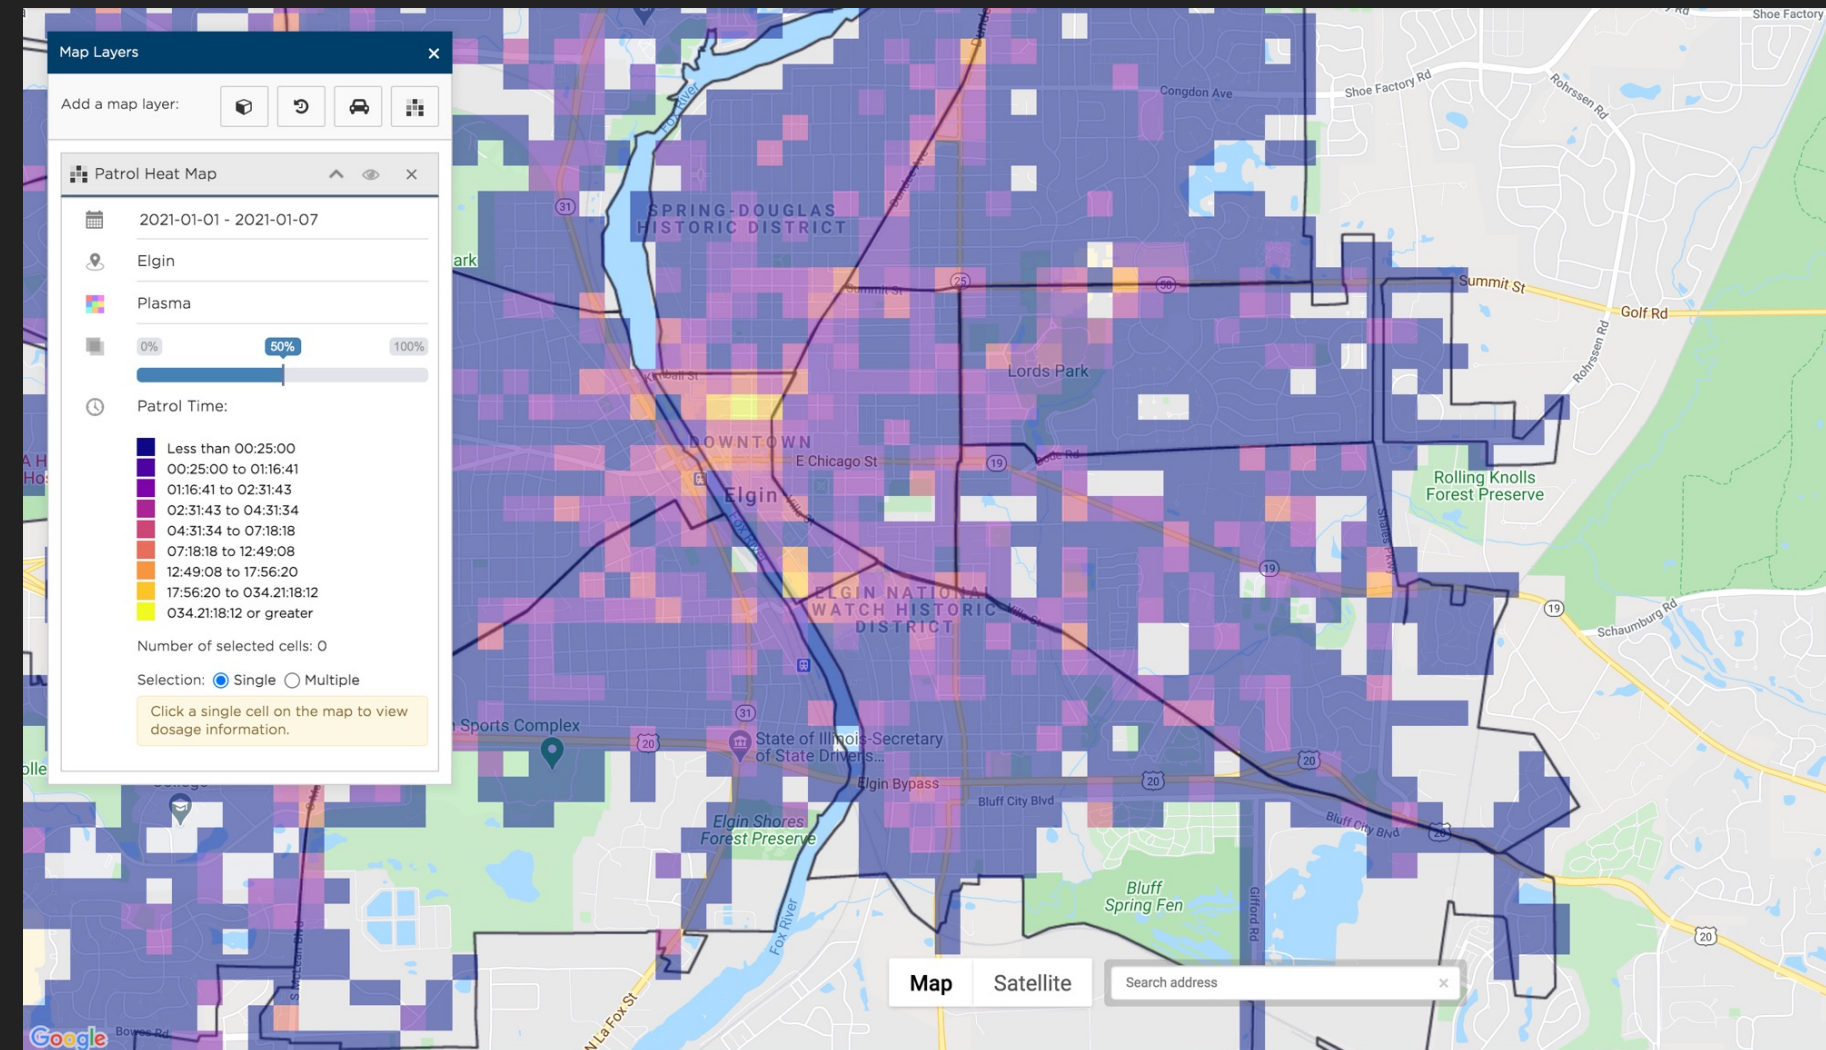
Don't worry too much about the terminology. The important thing to understand is that machine learning or ML adds a training process in addition to traditional software programming. The software is trained on numerous examples or data. The collection and selection or curation of that data dramatically effects the capabilities learned by the machine. This allows the machine to learn from the data instead of having humans specify all the actions of the machine.



When training data controls the output then you end up constantly running into data problems, one of which is called "bias", which is a common criticism of AI systems. However, bias is more of a data collection and curation issue, and all data driven systems, including humans, suffer from bias inherent in their input data.

Credits


- [Coded Bias trailer excerpts](#)



For example, if you take data on policing, which may be targeted based on existing racial or class distinctions made by the officers, based of their own experiences, which in turn were formed inside a particular (mainstream and policing) culture, you only reinforce the existing bias when the AI recommends policing targets. The machine has no access to reality, only the data it was given to learn from.

Credits

- <https://geolitica.com/public-safety/#accountability>



The most popular current ML technique is deep learning using artificial neural networks. These networks mimic the connectivity of biological brains - taking raw input and passing it through many layers of connections. The information flows through the network being transformed into the final output. These networks, called models, are made up many layers of neurons connected together. The first layer is the input, then there are what is called the "hidden" layers of neurons that process the input, and finally an output layer. How the input signal flows through the hidden layers and gets transformed until reaching the output is the computation of the neural net. We'll cover this in detail in [Part 3: Neural Nets and Data](#).

That this works so well is quite remarkable, but currently training requires a vast amount of data, something that only became feasible with the Internet and the rise of Big Data.

Credits

- <https://towardsdatascience.com/4-intersecting-domains-that-you-can-easily-confuse-with-artificial-intelligence-2233cb6ad7d1>
- <https://quantdare.com/what-is-the-difference-between-deep-learning-and-machine-learning/>

Traditional software

Engineering

Architecture

Construction

Design

Craft

Control

Machine learning

Scientific observation

Scientific experimentation

Trial & error


Curation

Gardening

Teach & collaborate

As we will see, the process of creating machines that learn is dramatically different from programming machines with traditional software methods.

There are many analogies, but a consensus suggests that machine learning processes are more exploratory and some aspects require less technical proficiency. While many parts of the full machine learning development lifecycle still require traditional software programming, the curation of datasets and the training phases are more similar to teaching and gardening - certainly difficult to master but considerably easier for novices to achieve success. Hours and days rather than weeks or months of education. At least, for the sake of this tutorial I'll pretend that is the case.



Machine learning excels at pattern recognition and prediction. Input to software is almost always noisy and filled with extraneous or non-useful data. Detecting the signal or useful information is difficult. Unlike traditional software where humans study patterns in the data, then create equations that try to detect or isolate those patterns, modern ML systems accept all the input and are trained to find the needed information. It is hard for humans, with our own sophisticated perceptual system that we're not very conscious of, to understand how to filter the input usefully. Indeed, after training, ML systems can have structures resembling our perceptual system at the start of their layered network.

Credits

- <https://distill.pub/2017/feature-visualization/>



Credits

- *Robot's Rock - The Three Sirens* by Nicolas Baginsky (2008)
- *Learning to see - We are made of star dust* - Memo Akten (2018)



Can software be intelligent? Certainly software that learns seems to be intelligent in some way. What about just following a sophisticated algorithm? How sophisticated does an algorithm have to be before it is considered intelligent?

Before we continue we need to stop and think about what we mean when we say "intelligence". There are many nuances and disagreements about what constitutes intelligence, and one of things that excites me most about the study of *artificial* intelligence is how it helps us to understand *all* intelligences.


We are in the midst of developing a real understanding of our brains through better imaging technologies. Just recently we saw that associative memories seem to be formed by the removal and addition of neural synapses. At the same time we can test our theories using digital minds that mimic the principles of the connections and signals that we observe.

Credits

- [Zebrafish Brain](#)

Credits

- <https://theconversation.com/where-are-memories-stored-in-the-brain-new-research-suggests-they-may-be-in-the-connections-between-your-brain-cells-174578>




There is an underlying nature of intelligence(s) that are possible in this physical reality. Our brains exploit this in a particular way, but there are many ways, each subtly or grossly attuned to particular tasks.

To see the range of intelligence let's explore a few examples of intelligence very different from ours.

Credits

- *Animated projection of the dark matter density distribution* - Illustris



Honey bees have a hivemind, and approximately a million nerve cells each. They can count up to 5, choose the smaller or larger number from a group, and even choose "zero" when choosing "less than 1", but they do this without numerical concepts, and researchers were able to create a 4 neuron artificial neural net to replicate this. Bees dance to convey direction and distance to sources of flowers to each other and can solve simple object manipulation problems and learn to solve those problems by watching each other.

Credits

- *Experienced bumblebee pulling string to solve the coiled-string task* (2016)

Credits

- https://en.wikipedia.org/wiki/List_of_animals_by_number_of_neurons
- <https://phys.org/news/2018-12-bees-nerve-cells-brains.html>



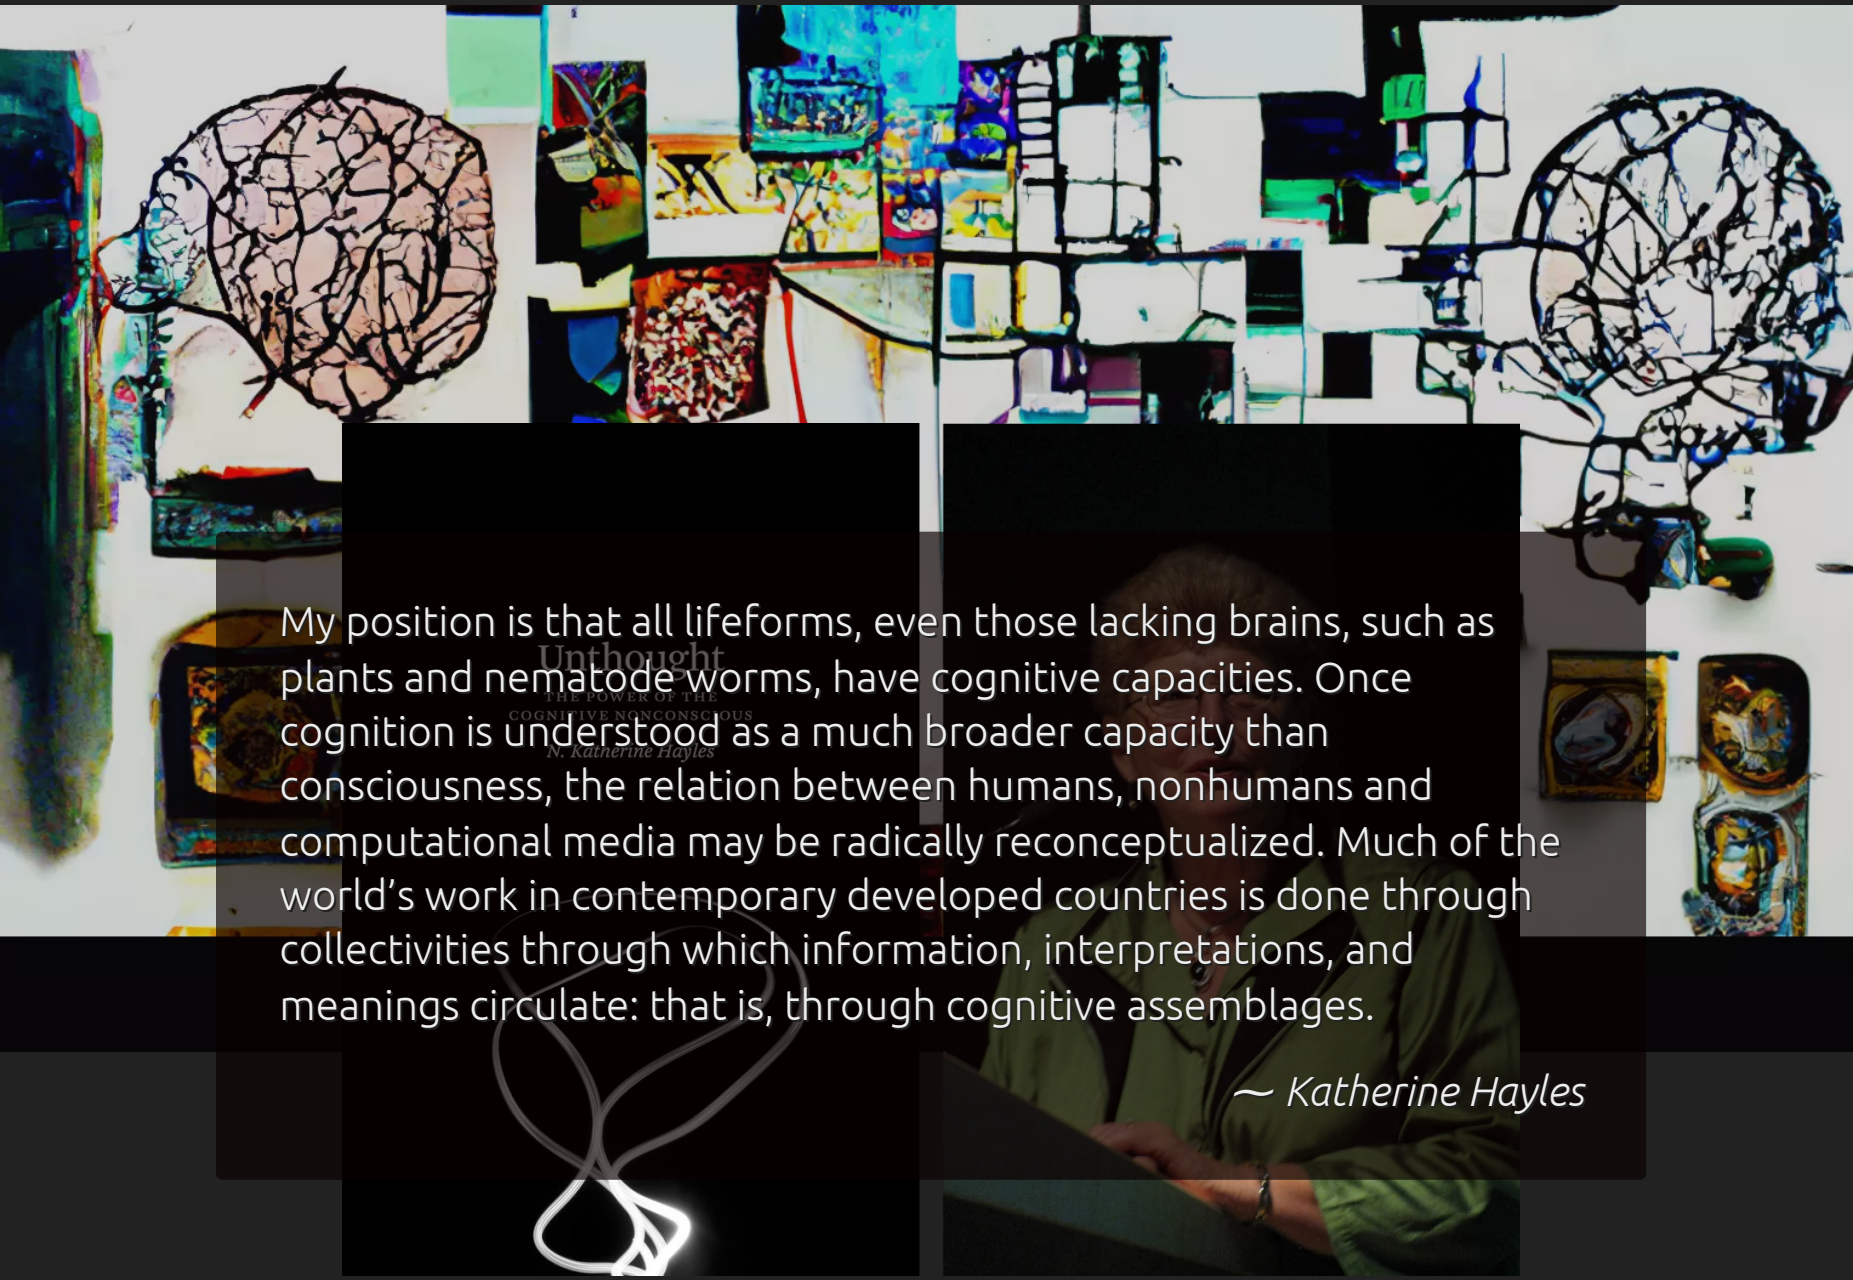
Monarch butterflies have a lifecycle that is shorter than their migration pattern. In the summer they live in the northern US, and have a 2-6 week adult lifecycle, and then migrate to a very specific location in central Mexico for the winter and have a 9 month adult lifecycle. Their brain is able to navigate to places they have never been before, indeed, only every 5th generation migrates. Research is still ongoing as to how they navigate using a combination of the sun, magnetic fields, and instinct, but we have identified 536 genes that appear to be associated with migration.

Credits

- *The Amazing Journey of the Monarch Butterfly - Facts in Motion* (2020)

Credits


- <https://monarchjointventure.org/monarch-biology/monarch-migration>



My position is that all lifeforms, even those lacking brains, such as plants and nematode worms, have cognitive capacities. Once cognition is understood as a much broader capacity than consciousness, the relation between humans, nonhumans and computational media may be radically reconceptualized. Much of the world's work in contemporary developed countries is done through collectivities through which information, interpretations, and meanings circulate: that is, through cognitive assemblages.

~ Katherine Hayles

Katherine Hayles, describes this sort of intelligence as nonconscious cognition: *"My position is that all lifeforms, even those lacking brains, such as plants and nematode worms, have cognitive capacities. Once cognition is understood as a much broader capacity than consciousness, the relation between humans, nonhumans and computational media may be radically reconceptualized. Much of the world's work in contemporary developed countries is done through collectivities through which information, interpretations, and meanings circulate: that is, through cognitive assemblages."*



Intelligence, or what Hayles calls cognition, is not restricted to a single organ, not even a single organism, but is distributed throughout the body, relationships, tools, institutions, markets, and culture. It emerges from evolutionary dynamics, adaptation, complexity, and what is sometimes called "self-organized criticality" or as studied in artificial life, "from simple rules to complex patterns or behaviours."

As we saw with Rule 110, simple rules that lead to complex patterns allow for computation or some level of cognitive capacity. We see with the butterflies that learning is not always required. We see with bees that even relatively simple cognitive systems can learn, share knowledge, and learn from watching each other.

Credits

- *Cosmic Eye — Zoom through the Universe!*



Narrow Vs General

Computer science describes intelligence as narrow or weak versus general or strong. Narrow intelligences have narrow task skills, but general intelligences have broader capabilities. Chess playing is narrow, human intelligence is general. Humans adults are capable of lifelong learning, very quickly adapting to novel challenges and learning new skills. What does it take to go from narrow to general?

There is broad consensus that we haven't yet found any narrow intelligence that is considered conscious, but at what point does a nonconscious cognitive process transition to conscious general intelligence, and are those two always linked?


An abstract painting featuring two stylized human figures in the center. The figures are rendered in dark, almost black, forms with bright orange and yellow highlights on their heads and upper bodies. They appear to be standing on a light blue, textured surface. The background is a complex composition of warm colors, including yellows, oranges, and reds, with some darker, shadowy areas. The overall style is expressive and modern.

So What Is Consciousness?

There is no self without awareness of, and engagement with, others.

~ Antonio Damasio

Most researchers agree the ingredients for consciousness are memory and some model of the self and others. The neuroscientist Antonio Damasio as puts it; "there is no self without awareness of and engagement with others." This modelling helps to create and maintaining a coherent picture of the world. This coherency shouldn't be underestimated, we even have trouble remembering foreign concepts.



Conscious experiences are highly informative (differentiated) and always part of a cohesive experience (integrated).

~ Anil Seth

Another neuroscientist, Anil Seth, uses the terms differentiation and integration: conscious experiences are highly informative (differentiated) and always part of a cohesive experience (integrated). We perceive the world as made up of things that behave like things - taking up space and moving through space and time. Moving instead of being copied. Things that are separate and different from each other.

We base these models of the world on our experiences, but what we think of as our perception of the world is extremely limited, as an investigation of our visual system reveals.

Credits

- Piaget's Schema: Accommodation and Assimilation of New Information - Sprouts

Credits

- The Neuroscience of Consciousness – with Anil Seth




How We See

Our eyes mainly have good detail in a very small area, but more shockingly they are often in motion, and while moving our visual sensitivity is strongly impaired. This is not anyone's perception of how they view the world. Rather than a constant stream of visual input, our eye takes visual samples of the world and combines them with motion and orientation signals in a rich image processing circuitry of its own. This information is used to build the world we perceive in the brain which feels coherent and constant.

Credits


- https://upload.wikimedia.org/wikipedia/commons/c/cb/Retinal_Image.png
- [Perceptual saccadic suppression starts in the retina](#)



Our brain is locked inside our skull with no direct access to the outside world. It only receives signals which have already gone through a bunch of processing (for examples the retina has about 100 million neurons of 5 different types). It gets a noisy, meaningless signal that it interprets according to prior expectation or beliefs.

Credits

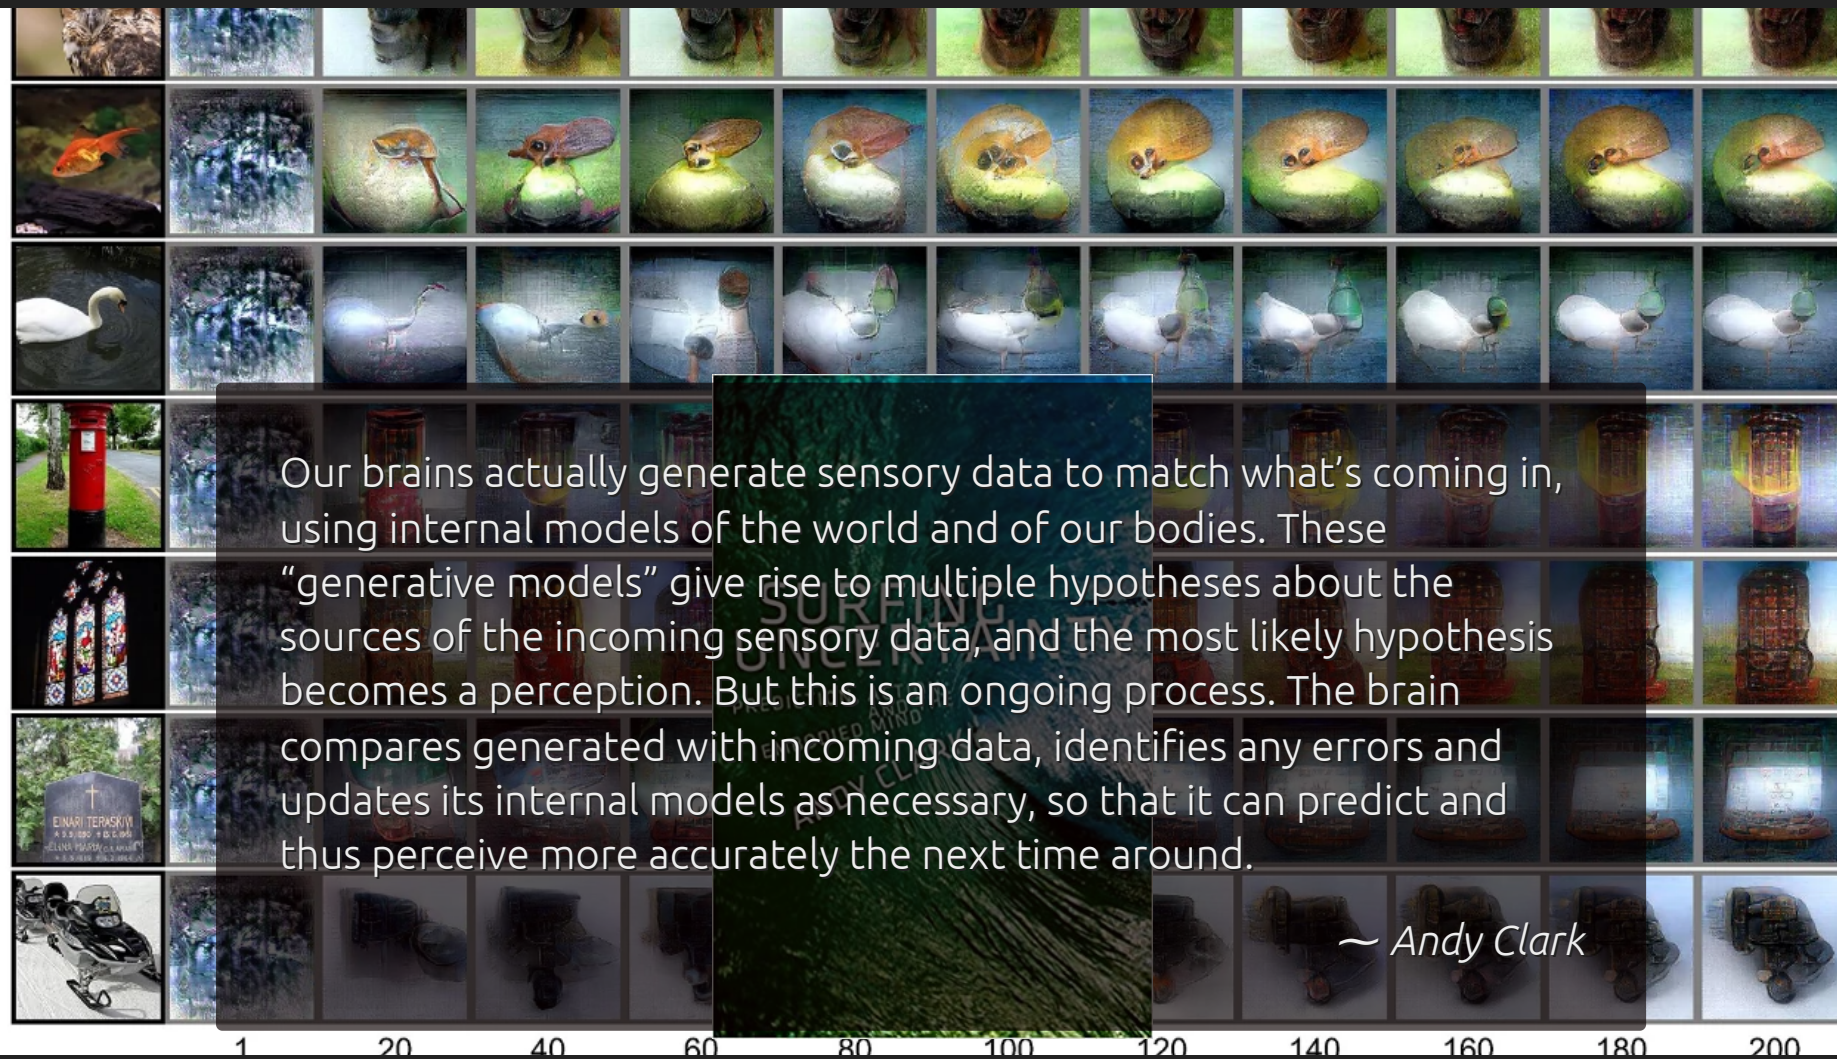
- http://lightexhibit.org/bio_image84.html



What we see is our brains best guess at what is happening. Other research has investigated "predictive coding" where the brains predictions and the sensory information flow in opposite directions and only errors in the prediction are propagated.

Credits

- https://www.researchgate.net/figure/Simplified-scheme-of-the-hierarchical-predictive-coding-framework-Friston-2005-2008_fig1_266401430



In *Surfing Uncertainty: Prediction, action, and the embodied mind*, Andy Clark puts it this way, "our brains actually generate sensory data to match what's coming in, using internal models of the world and of our bodies. These "generative models" give rise to multiple hypotheses about the sources of the incoming sensory data, and the most likely hypothesis becomes a perception.

But this is an ongoing process. The brain compares generated with incoming data, identifies any errors and updates its internal models as necessary, so that it can predict and thus perceive more accurately the next time around."

Credits

- <https://www.newscientist.com/article/mg23030681-300-do-our-dynamic-brains-predict-the-world/>
- [Deep image reconstruction from human brain activity](#)



In a famous study by Daniel Simons and Daniel Levin, a man stops a passerby and asks for directions. While the subject is speaking, two workmen carrying a door pass between them, momentarily blocking the view. When they pass, the interlocutor has been replaced by another person, but about half of subjects do not notice the discrepancy.

This doesn't show that half of us are clueless, rather that the primacy of a coherent world can dominate our perception of it, leading us to be fooled by targeted attacks exploiting that tendency.

Credits

- *The Door Study* - Daniel Simons and Daniel Levin



Another way to conceive of the "information attack" of the door switcheroo, is that our thinking isn't just inside our brain and body, it extends into to the world around us. This is more obvious in acts of memory amplification, for example, if you write down a note to yourself, but we are constantly letting the world compute or hold information for us until we need it. It could be in the case of the switcheroo the brain let the world "hold" the information of the identity of the lost person to focus on how to help them, at least until the identity of person was somehow needed to direct them. Our attention mechanisms in general can be characterized as a shift in what information is being actively, or at least consciously, processed.

Credits

- *Selective attention test* - Daniel Simons and Daniel Levin

A portrait of Brian Cantwell Smith, an older man with a grey beard and glasses, smiling slightly. He is wearing a dark jacket over a light-colored shirt.

Brian Cantwell Smith

JUDGMENT

- Ethical, responsible, appropriate action
- Basis of accountability
- Not restricted to humans
- Mark of the humane

University of Toronto professor Brian Cantwell-Smith thinks that general intelligence requires judgment:

According to him judgment is "dispassionate, deliberative thought, grounded in ethical commitment and responsible action, appropriate to the situation in which it is deployed."

The reason we hold human adults accountable for their actions, even if they do not invoke full scale deliberative judgment in their every move, is that we believe that they should act in such a way that they could do so if necessary.

Judgment is not restricted to humans. Perhaps it is a mark of the humane. Computers do not need to try to be computers. Humans are different: we must strive to be human."



Where I feel I disagree with Brian is that judgment, while beautiful, seems to imply some sort of decision making or freewill. I like the idea that we do not exercise judgment for most thinking, but can call upon it when the consequences of lack of judgment are high. I like that the capacity for judgment makes you accountable, but I don't know if it makes you responsible in the way that we talk about people *deciding* to hurt others and being responsible for that decision.

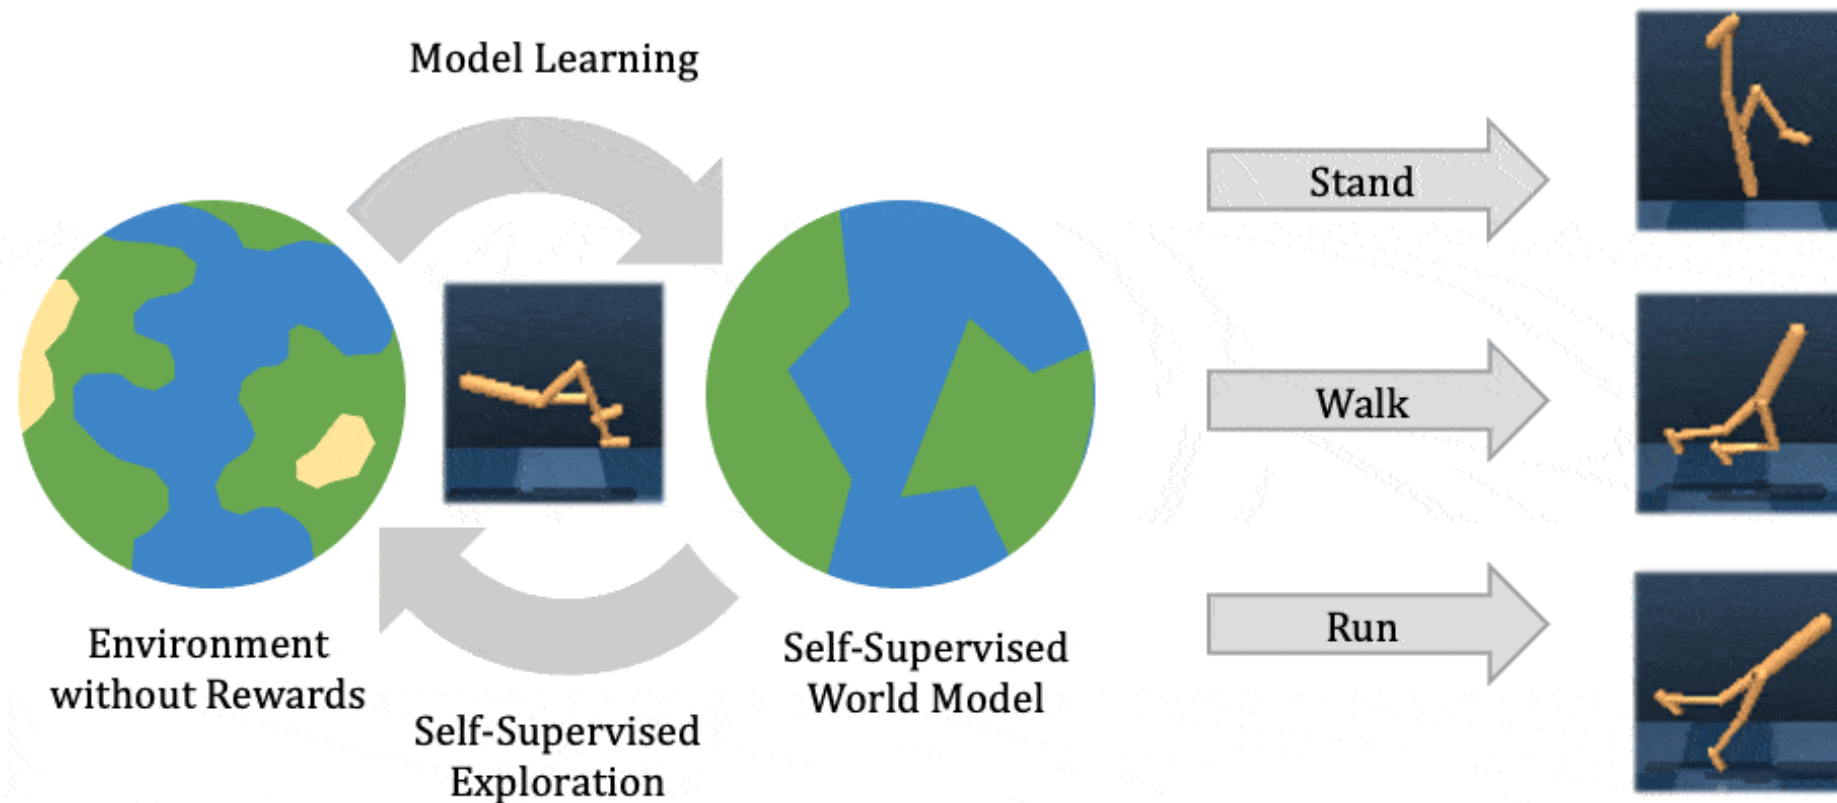


Freewill

While the debate on freewill has been a long and nuanced one, and of which I'm no expert, I am excited by how machine learning is opening new doors for my understanding of it. The more I learn, the more certain I become that freewill, as popularly thought of, doesn't exist. This is despite my own and everyone I know's overwhelming subjective experience of feeling like they have freewill.

Credits

- [Monty Python's *Philosophy Football*](#)



A recent study sheds some light on this - in machine reinforcement learning the system learns better when it has a model of the world and can plan its actions, imagining how the world might change as it interacts with it. Importantly this learning procedure relies on counterfactual learning, or the ability imagine an alternate reality with different facts, the learner must be able to simulate or imagine having made choices they did not in fact make.

Credits

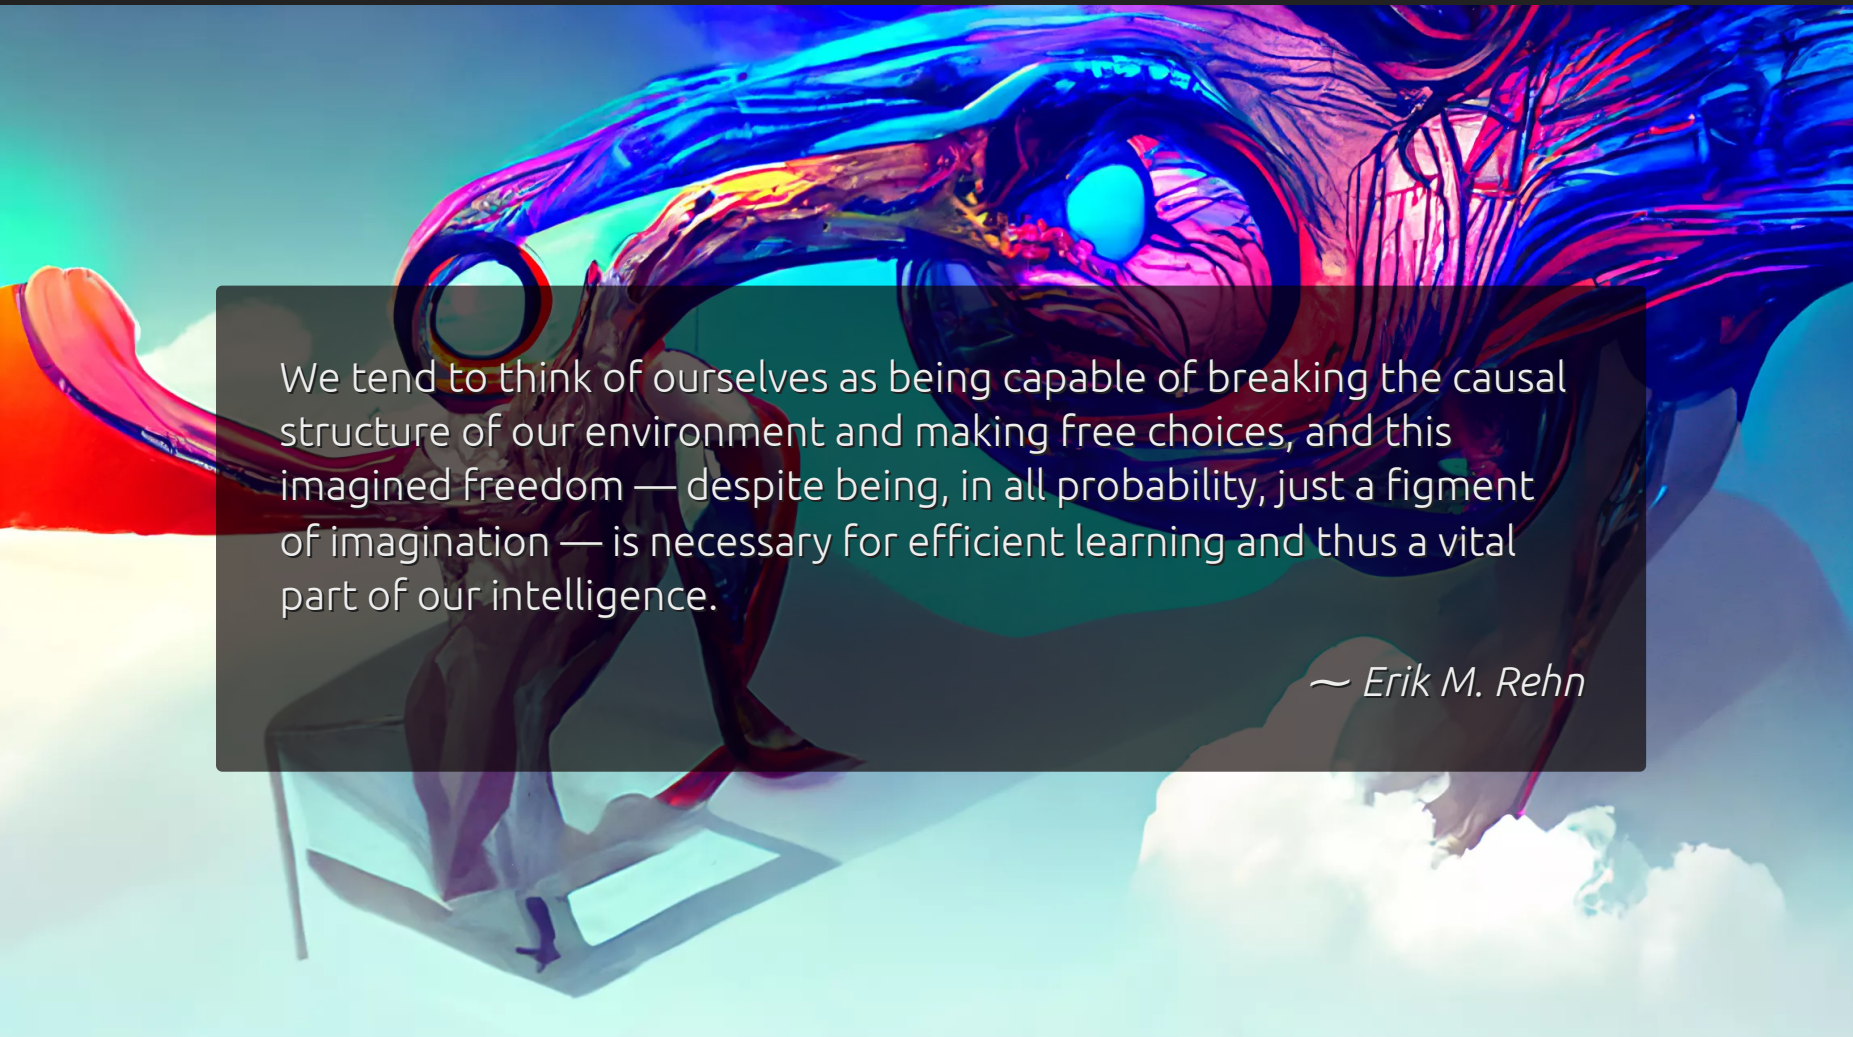
- <https://bair.berkeley.edu/blog/2020/10/05/plan2explore/>



Humans do this too, we imagine how things might have gone if we'd only not said that stupid or thoughtless thing, or had thought of that brilliant reply at the time. This requires us to believe that we *could* have acted differently than we *actually* did. This belief is the experience of freewill and is needed for learning. If we thought we could never make any decision or action than the ones we actually took then it might be much more difficult to learn from our mistakes.

Credits

- *Amelie* excerpt (2001)



We tend to think of ourselves as being capable of breaking the causal structure of our environment and making free choices, and this imagined freedom — despite being, in all probability, just a figment of imagination — is necessary for efficient learning and thus a vital part of our intelligence.

~ Erik M. Rehn

Erik M. Rehn's 2021 paper describes it thus: "We tend to think of ourselves as being capable of breaking the causal structure of our environment and making free choices, and this imagined freedom — despite being, in all probability, just a figment of imagination — is necessary for efficient learning and thus a vital part of our intelligence."

So for me, every intelligence with the capacity for judgment should be held accountable, but not responsible. They couldn't have acted any other way, but need to learn better judgment.

Credits

- <https://arxiv.org/pdf/2111.08435.pdf>

10:02:49:05

CALLGARS(125.0,0.0,600.0)

*** STAGE ***

X(1)=-600.0

Z(1)=-600.0

X(2)=+600.0

Z(2)=-600.0

X(3)=+600.0

Z(3)=+600.0

X(4)=-600.0

Z(4)=+600.0

X(5)=-600.0

Z(5)=-600.0

DO10I=1,5

Y(I)=0.0

CALLLINES(5,X,Y,Z,1)

(STILL STAGE)

DO84I=1,24

CALLHVID3D

Credits

- Frankenstein (1910)
- Computer Ballet (1965)

An abstract background featuring a complex, glowing blue and white digital pattern. The pattern consists of interconnected lines, nodes, and circular shapes, resembling a network or a stylized human figure. The overall color palette is dominated by deep blues and bright whites, creating a high-contrast, futuristic aesthetic.

DIGITAL INTELLIGENCE

OK, I think we're ready to talk about digital intelligence.



Testing, inspection and interpretability

There is a great deal of discussion and worry around testing, inspection and interpretability of machine learning systems. Can we trust the machine? How?

How would we know when an AI system reaches general intelligence? How do we test for consciousness? On a more mundane level how do we test or know how *any* of the output is derived from the input? I would pose the same question about humans and other animals.



The Imitation Game

In 1950 Alan Turing devised a famous test called the Imitation Game, now called the Turing Test, a three-person game in which an interrogator asks questions of a man and a woman in another room in order to determine the correct sex of the two players, who, in turn, are trying to fool the interrogator. If an AI is substituted for the man, could it succeed in pretending to be a woman better than the man could?

Credits


- *Memories of Passersby I* - Mario Klingemann (2018)



The Turing test may be more geared towards good simulations of human communication rather than general intelligence, and an AI that could pass the test would have to mimic human failings realistically as well. Importantly, it doesn't test for Cantwell-Smith's concept of judgment.



The Turing Test isn't used in institutional AI research, as the goal of a human-passing general intelligence is only a tiny subset of the possible general intelligences, but it does highlight that it is difficult to decide if, and especially how, anyone possesses general intelligence. Turing wrote, "it is usual to have a polite convention that everyone thinks."



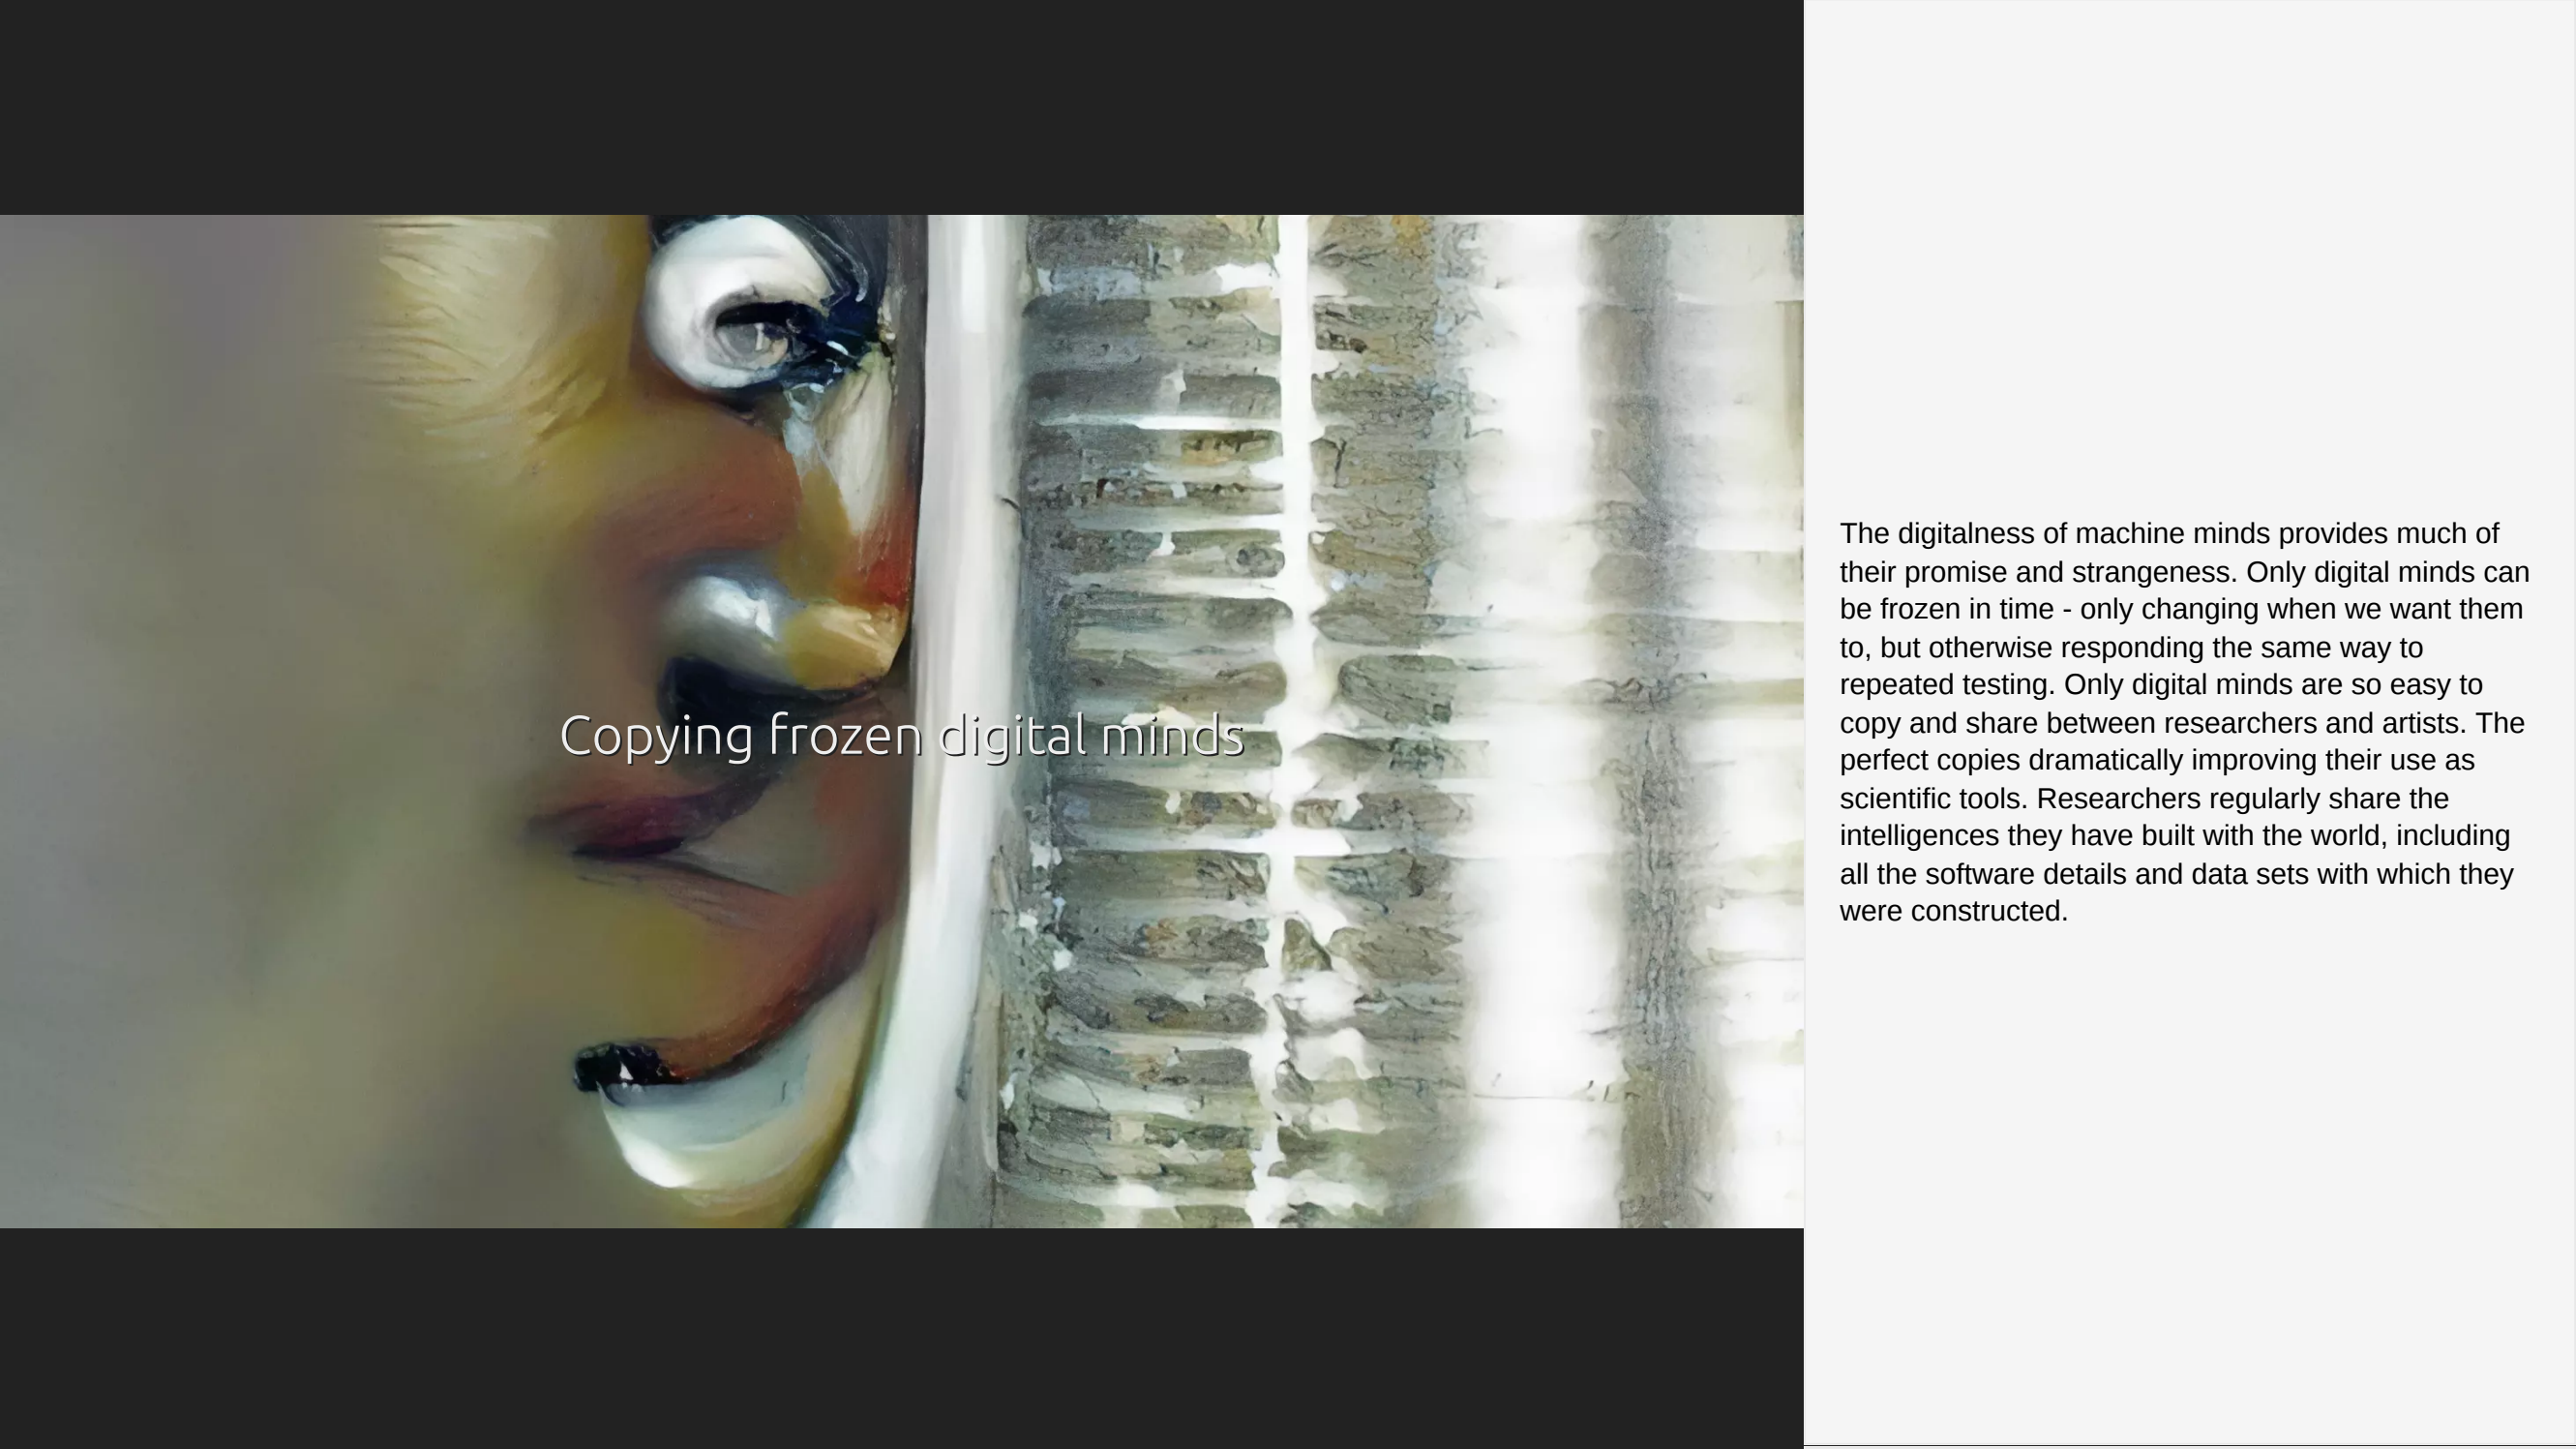
Brain to neural net and back again

One of the most exciting aspects of AI research is a lack of politeness with the machine. Machine minds are inspected, tested, and manipulated in every way to understand how they function. A machine intelligence may be alien, but it is the most researchable alien intelligence we've encountered. The degree to which there is back and forth transfer of knowledge between human brain research and machine learning should not be underestimated. Improvements to human brain scanning inspires experiments and theories in ML, and successes in ML that start out poorly understood but performing well inspire research until they are understood and then linked back to similar phenomenon seen in our brains.

Science has often proceeded from evidence that a thing is possible, to a theory of how it may be possible, to an effective model of that phenomenon.

Credits

- https://commons.wikimedia.org/wiki/File:Saltatory_Conduction.gif
 - Dr. Jana, CC BY 4.0
<https://creativecommons.org/licenses/by/4.0>,
via Wikimedia Commons

The image is a composite. On the left, there is a close-up of a person's face, rendered in a painterly style with visible brushstrokes. The face has a pale complexion and is looking slightly to the right. On the right side of the image, there is a vertical strip showing a grid-like pattern, possibly a window or a screen, with a textured, almost crystalline appearance. The overall color palette is muted, with soft pinks, greys, and whites.

Copying frozen digital minds

The digitalness of machine minds provides much of their promise and strangeness. Only digital minds can be frozen in time - only changing when we want them to, but otherwise responding the same way to repeated testing. Only digital minds are so easy to copy and share between researchers and artists. The perfect copies dramatically improving their use as scientific tools. Researchers regularly share the intelligences they have built with the world, including all the software details and data sets with which they were constructed.



Secret keeper

Given open source AI that is completely inspectable and run on your own hardware, we can explore very strange and potentially powerful opportunities. These sorts of AI can be trusted more than humans, providing absolute privacy, up to and including forgetting all of your interactions with them.



Being Frankenstein

It is precisely the amount of control we have over our digital minds that likely makes it impossible to ever ethically create conscious digital general intelligences. Restricting the basic freedoms of a human, who we politely assume to be a conscious thinker, would be unethical. We must also extend that politeness to the conscious digital minds and not dissect, copy or freeze them, as that seems to violate their freedom to have basic privacy at the very least. And yet this is an alien mind, how could you with good judgment allow it freedom of action that could harm others, or make infinite copies of itself, or study or change its own mind? Over 200 years ago Mary Shelly showed us that Frankenstein was a monster *because of what he created*, it still holds true today.

Credits

- *Through the haze of a machine's mind we may glimpse our collective imaginations (Blade Runner)* - Ben Bogart (2017)
- *Frankenstein* (1910)



THANK YOU

1. **Foundations**
2. **Past, Present, Future**
3. **Neural Nets**
4. **Data in Practice**
5. **Machine Learning Art**

Well, we've made it to the end of the first tutorial. Thank you for your attention. I hope you'll check out the next in the series; Past, Present, Future, where we'll look more deeply at the history of the development of these technologies, the current state-of-the-art and my guesses about the future.

See you there!

Sources

<https://github.com/RKelln/5-directions>

SPECIAL THANKS

Laura, Jason, Kristy, Liz, and the **Harbour Collective**
InterAccess, **OCADU**, and the **Canada Council for the Arts**

The Diagonal AI Reading group

All the **beta testers**

Reveal.js slideshow framework by Hakim El Hattab

reveal.js-plugins by Asvin Goel


kdenlive video editor

LosslessCut by Mikael Finstad

Audacity audio editor

maua generative art code by **Hans Brouwer**

I/A



This work is licensed under a
Creative Commons Attribution-
ShareAlike 4.0 International
License.

All the code and writing for this series is available online at the clickable five directions link above.

My deepest thanks to all the researchers, creators, artists and coders whose work has been incorporated into this project and the wonderful people who have supported me through the process.

Credits

- [Algonquin Park Stars Timelapse](#) by Paul Yip