Building AI Products

Madhu Shashanka

I've created a background image for your presentation titled "Building AI Products." This abstract and visually engaging design incorporates futuristic technology themes with a modern and vibrant color scheme, perfect for highlighting the innovation and growth associated with AI development.

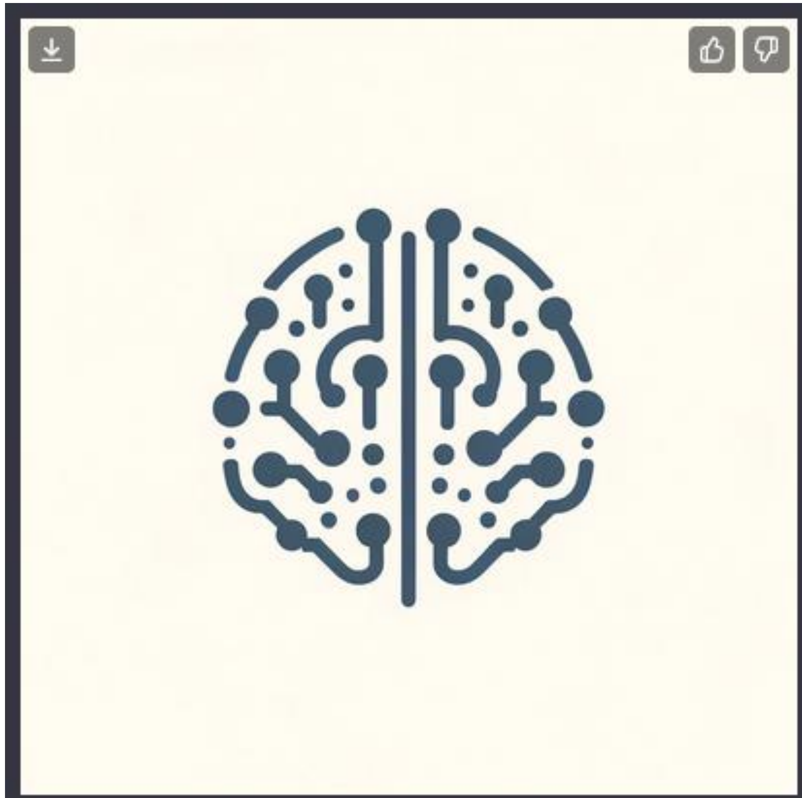

I've created a simpler background image for your presentation on "Building AI Products." This design is minimalist and clean, perfect for a professional setting without distracting from your presentation content.

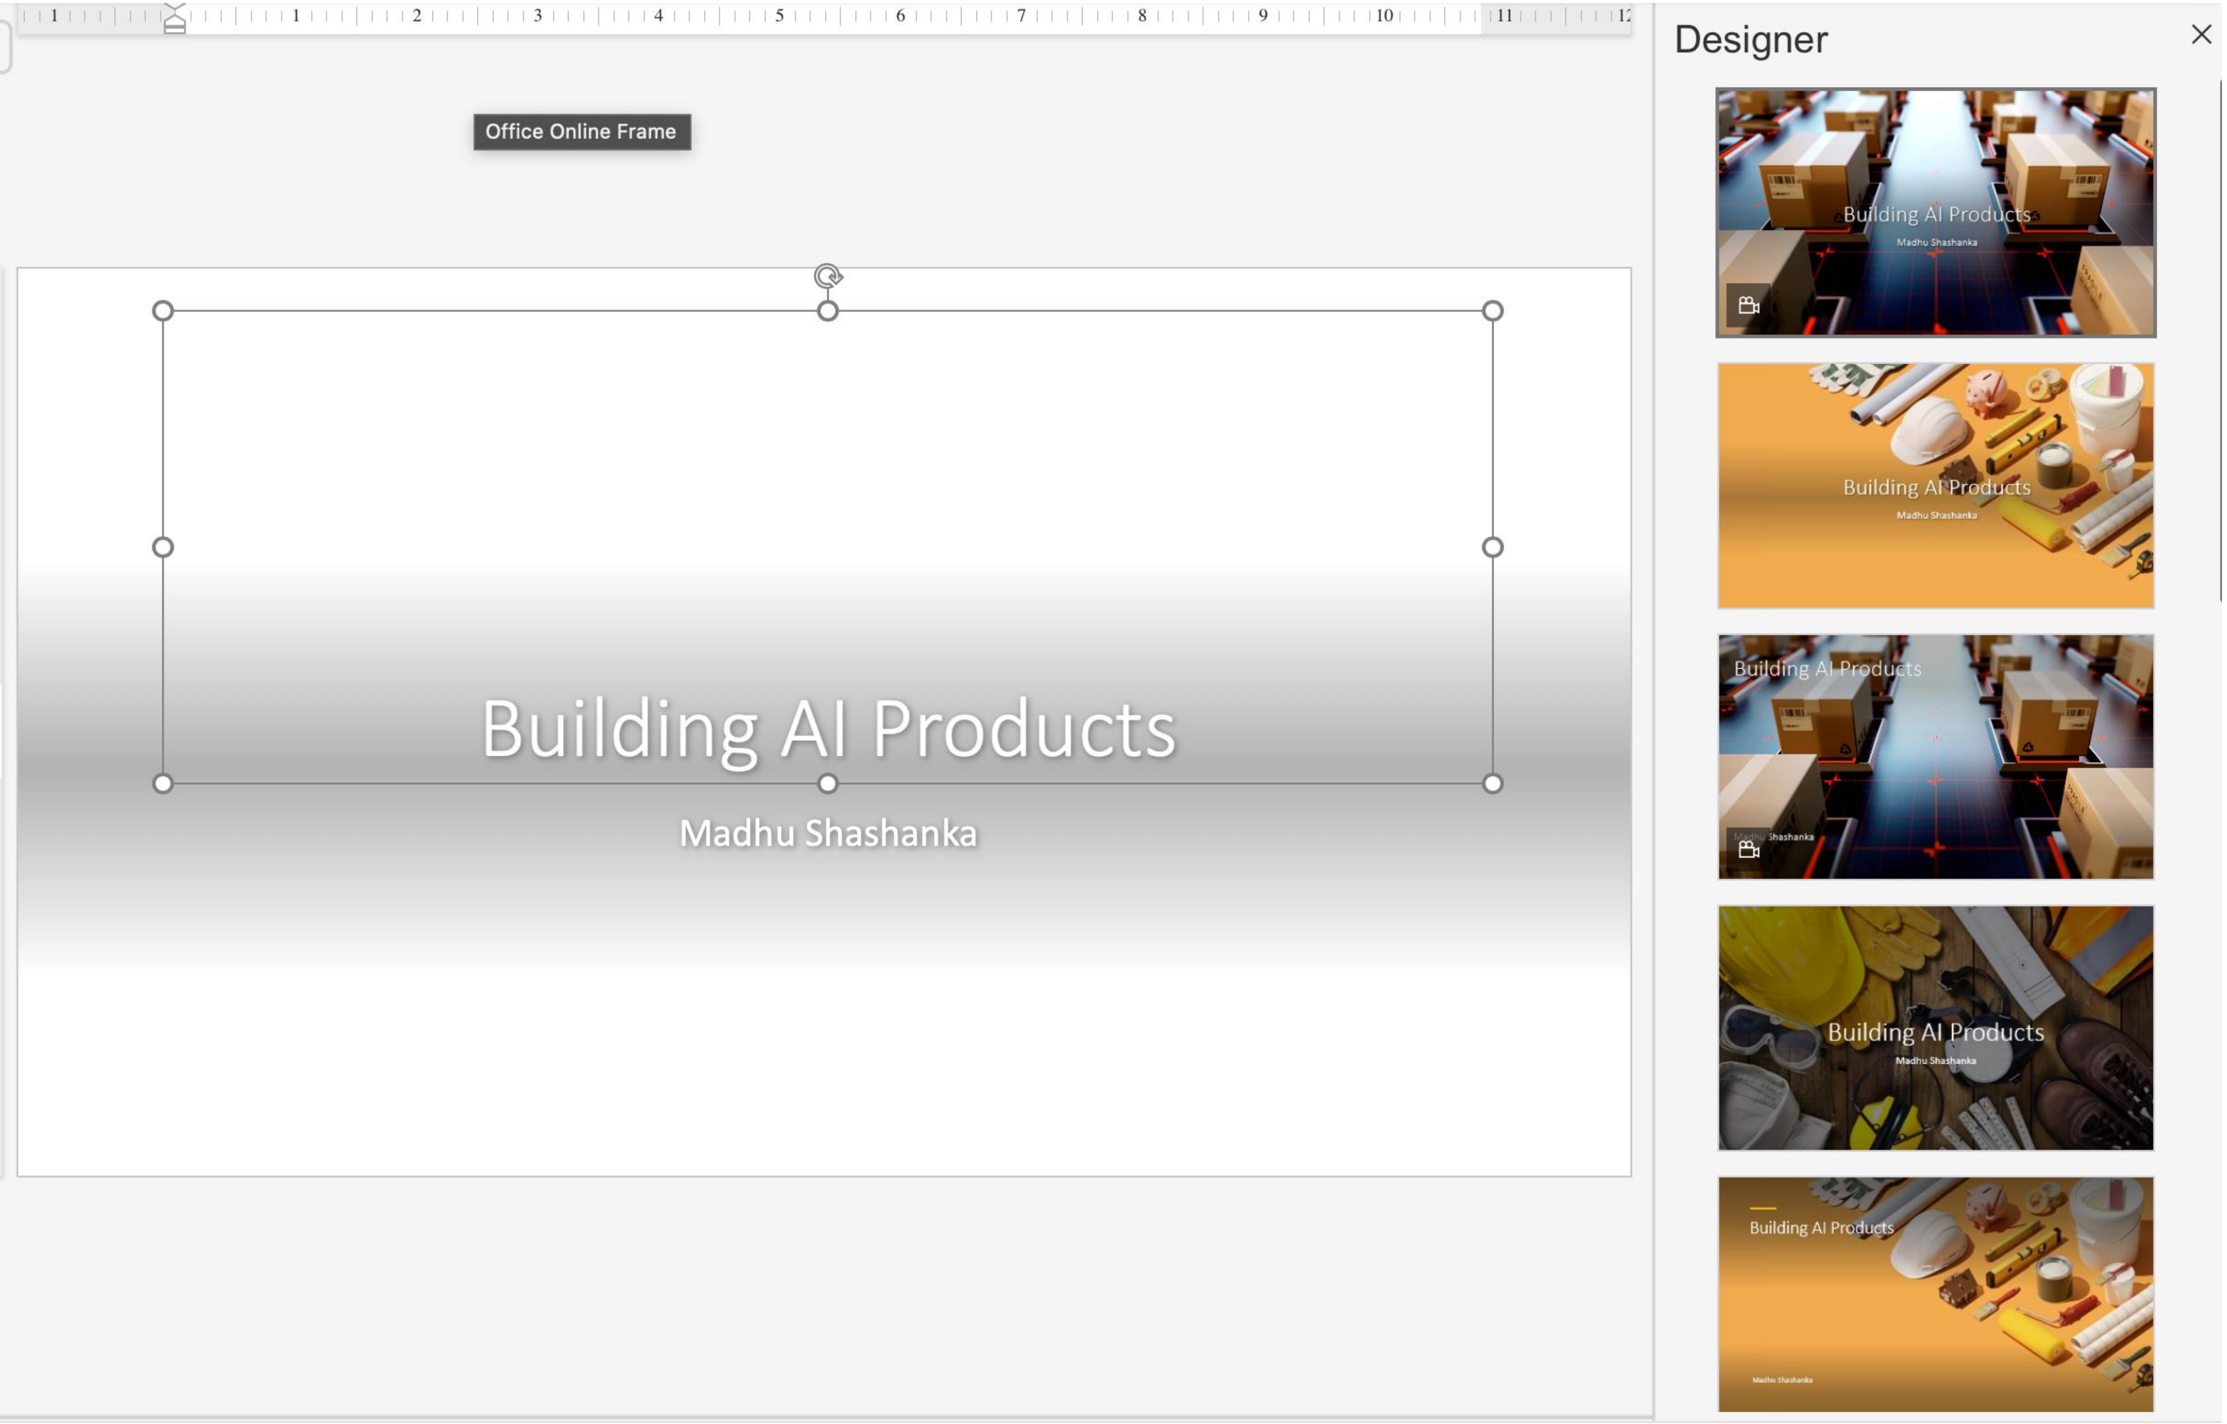
Office Online Frame
Building AI Products
Madhu Shashanka
Designer
Building AI Products
Madhu Shashanka
Building AI Products
Madhu Shashanka
Building AI Products
Building AI Products
Madhu Shashanka
Building AI Products

Computer Science

Computer Science

An influential textbook from the field of computer science, *Structure and Interpretation of Computer Programs* (1985), begins with the provocation that the subject it introduced "is not a science and that its significance has little to do with computers." [1] Rather, its authors, Harold Abelson and Gerald Jay Sussman, sought to establish a more philosophical basis for interest:

> The computer revolution is a revolution in the way we think and in the way we express what we think. The essence of this change is the emergence of what might best be called *procedural epistemology* – the study of the structure of knowledge from an imperative point of view, as opposed to the more declarative point of view taken by classical mathematical subjects. Mathematics provides a framework for dealing precisely with notions of "what is." Computation provides a framework for dealing precisely with notions of "how to." [2]

* https://amodern.net/article/on-computer-science/

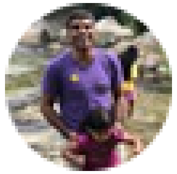

Madhu Shashanka
@ShashankaMadhu

To begin with, the field has no precise or clear definitions for “artificial” or “intelligence” but together it’s supposed to make sense!? It is good for marketing but not for an academic discipline.

> 
> **Talia Ringer** 🌻🪬 @TaliaRinger · Sep 5
>
> I often think that AI would be a way better field if it abandoned its name and went with something like "automation" or whatever
>
> Show this thread

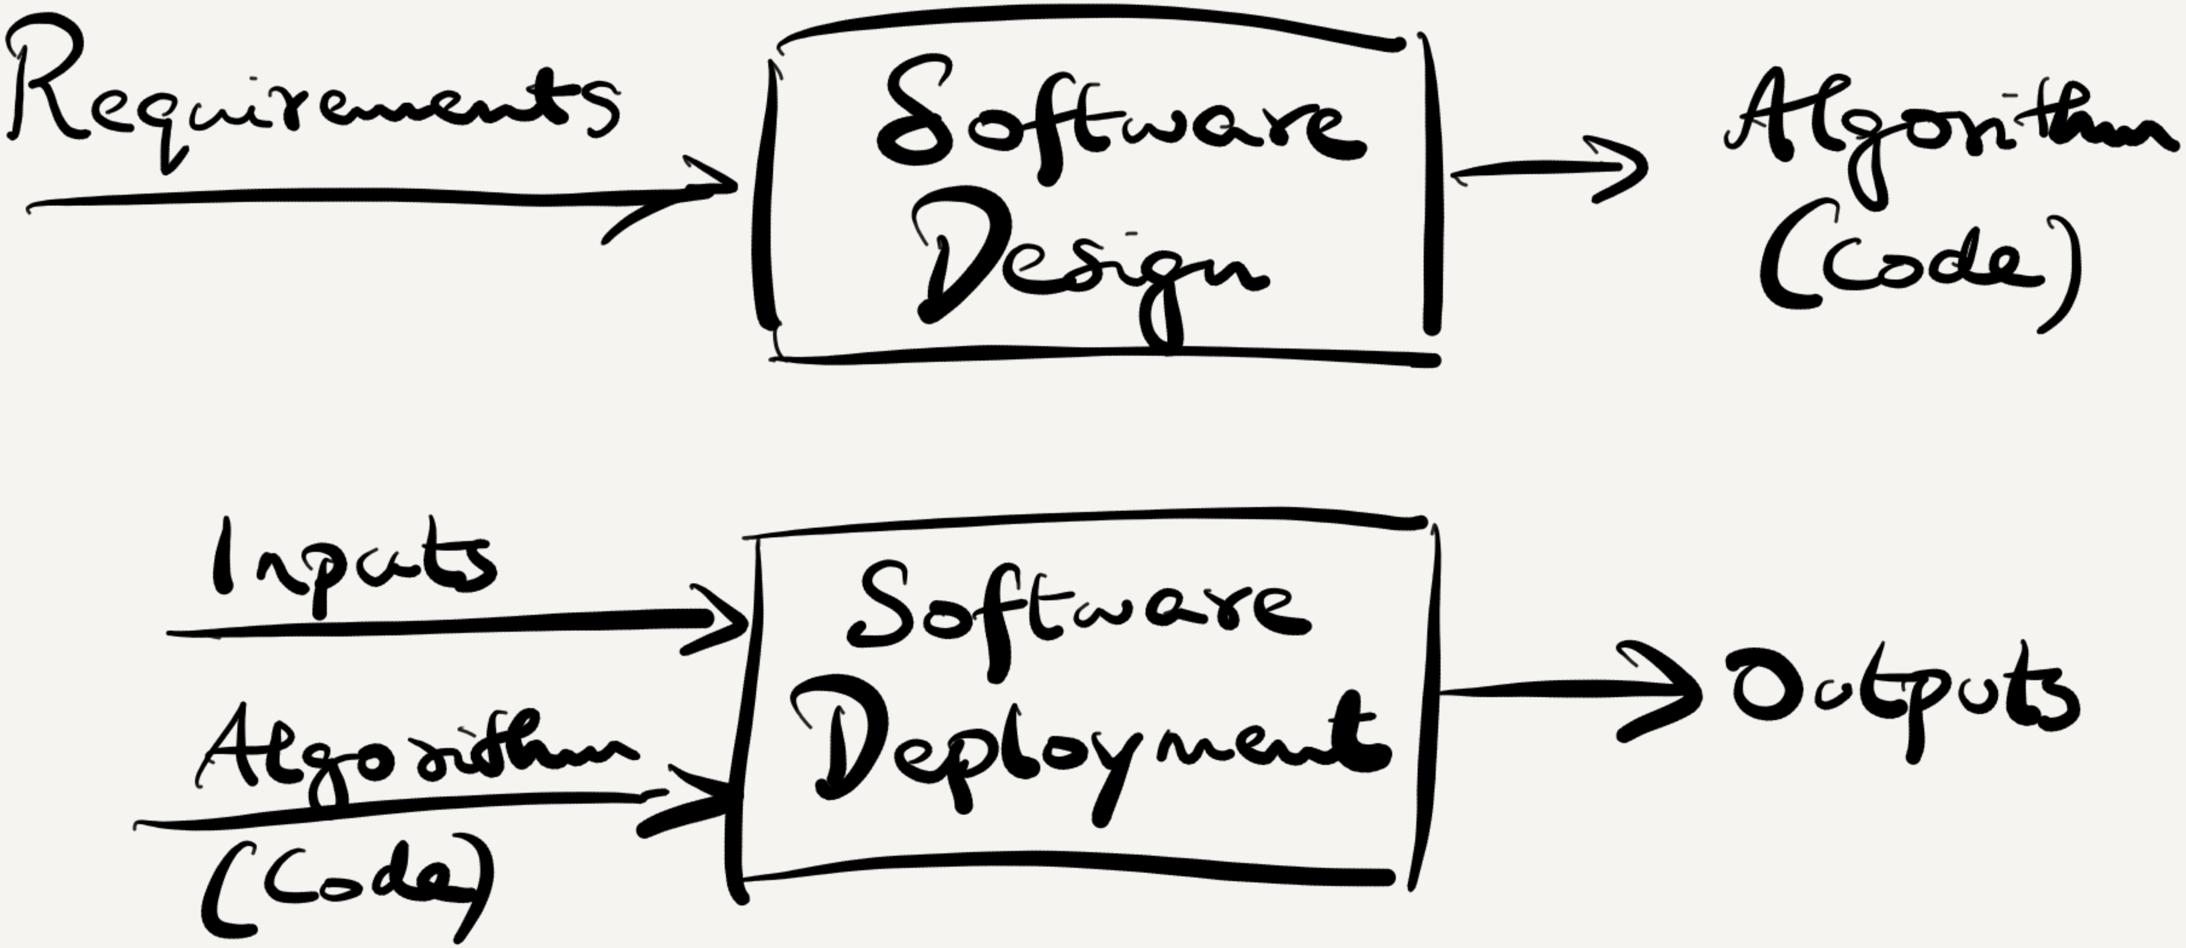

SOFTWARE AUTOMATION

https://psmag.com/magazine/the-lucrative-art-of-chicken-sexing

In the article "The Lucrative Art of Chicken Sexing," James McWilliams provides a wonderful overview of the profession and its history.

> "Virtually every link in the supply chain for hens, from genetics to the size and color of their eggs, is measured and automated with obsessive precision. ... Yet the most essential element in this massive and mechanized operation—figuring out if the hatched bird is biologically suited to lay eggs—relies on the vagaries of human eyesight. A chicken sexer—the industry's secret weapon—eyeballs the nearly identical genitalia of male and female hatchlings and, in an instant, makes the call."

But what makes it fascinating is the fact that no chicken sexer – not even the experts – can quite articulate how exactly they do it!

https://psmag.com/magazine/the-lucrative-art-of-chicken-sexing

In the article "The Lucrative Art of Chicken Sexing," James McWilliams provides a wonderful overview of the profession and its history.

> "Virtually every link in the supply chain for hens, from genetics to the size and color of their eggs, is measured and automated with obsessive precision. ... Yet the most essential element in this massive and mechanized operation—figuring out if the hatched bird is biologically suited to lay eggs—relies on the vagaries of human eyesight. A chicken sexer—the industry's secret weapon—eyeballs the nearly identical genitalia of male and female hatchlings and, in an instant, makes the call."

But what makes it fascinating is the fact that no chicken sexer – not even the experts – can quite articulate how exactly they do it!

And yet, the Japanese in 1920s figured out that chicken sexing could, in fact, be taught.

> "The central pedagogical aspect of sexing was and remains quick repetition over a focused period of time. The sexer must become proficient at immediately and unthinkingly recognizing the chickens' vent shapes, of which there are hundreds of slight variations. ... Japanese students were consistently able to internalize pattern recognition and intuit rather than consciously identify the markers determining sex. Some achieved accuracy rates that were close to perfect."

In other words, it is impossible to teach "how" to do chicken sexing – there are no markers on the hatchlings that help you distinguish, neither are there any time-tested playbooks one can follow. The only way to learn is to do it over and over, working with experts, until you gain an intuition for it.

https://psmag.com/magazine/the-lucrative-art-of-chicken-sexing

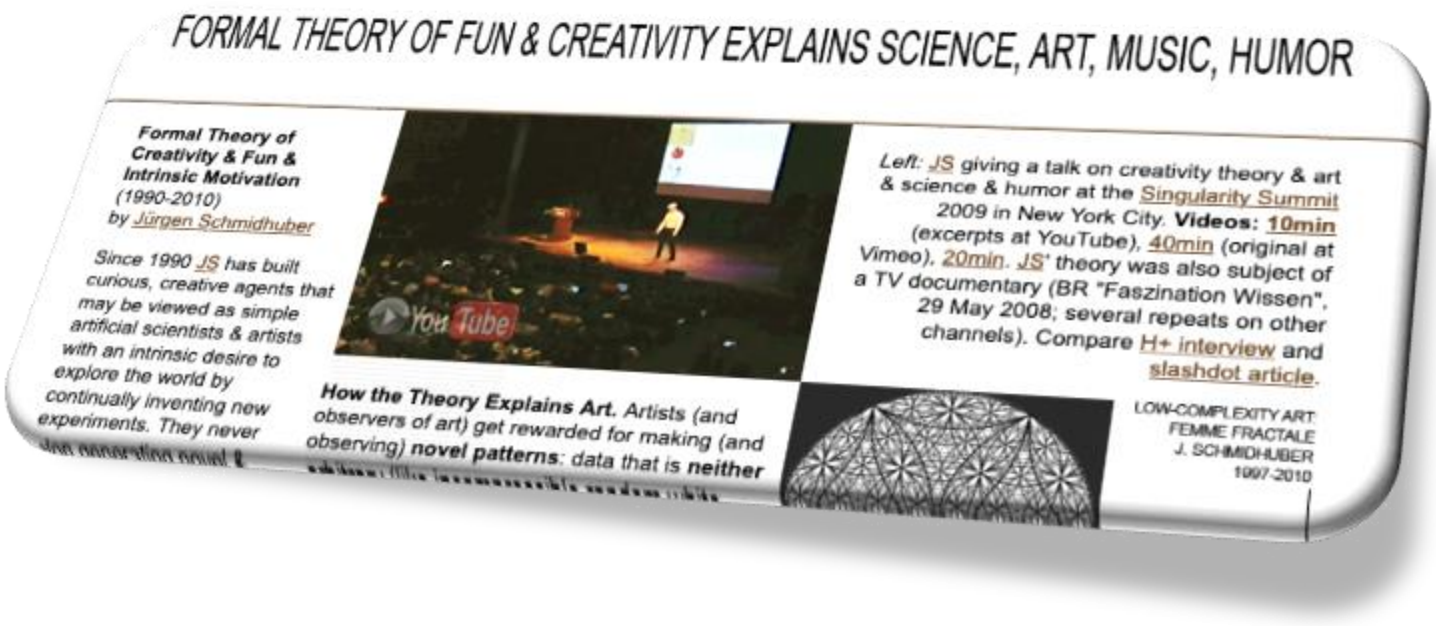

"Since short and simple explanations of the past usually reflect some repetitive regularity that helps to predict the future as well, every intelligent system interested in achieving future goals should be motivated to compress the history of raw sensory inputs in response to its actions, simply to improve its ability to plan ahead."

https://arxiv.org/pdf/0812.4360.pdf

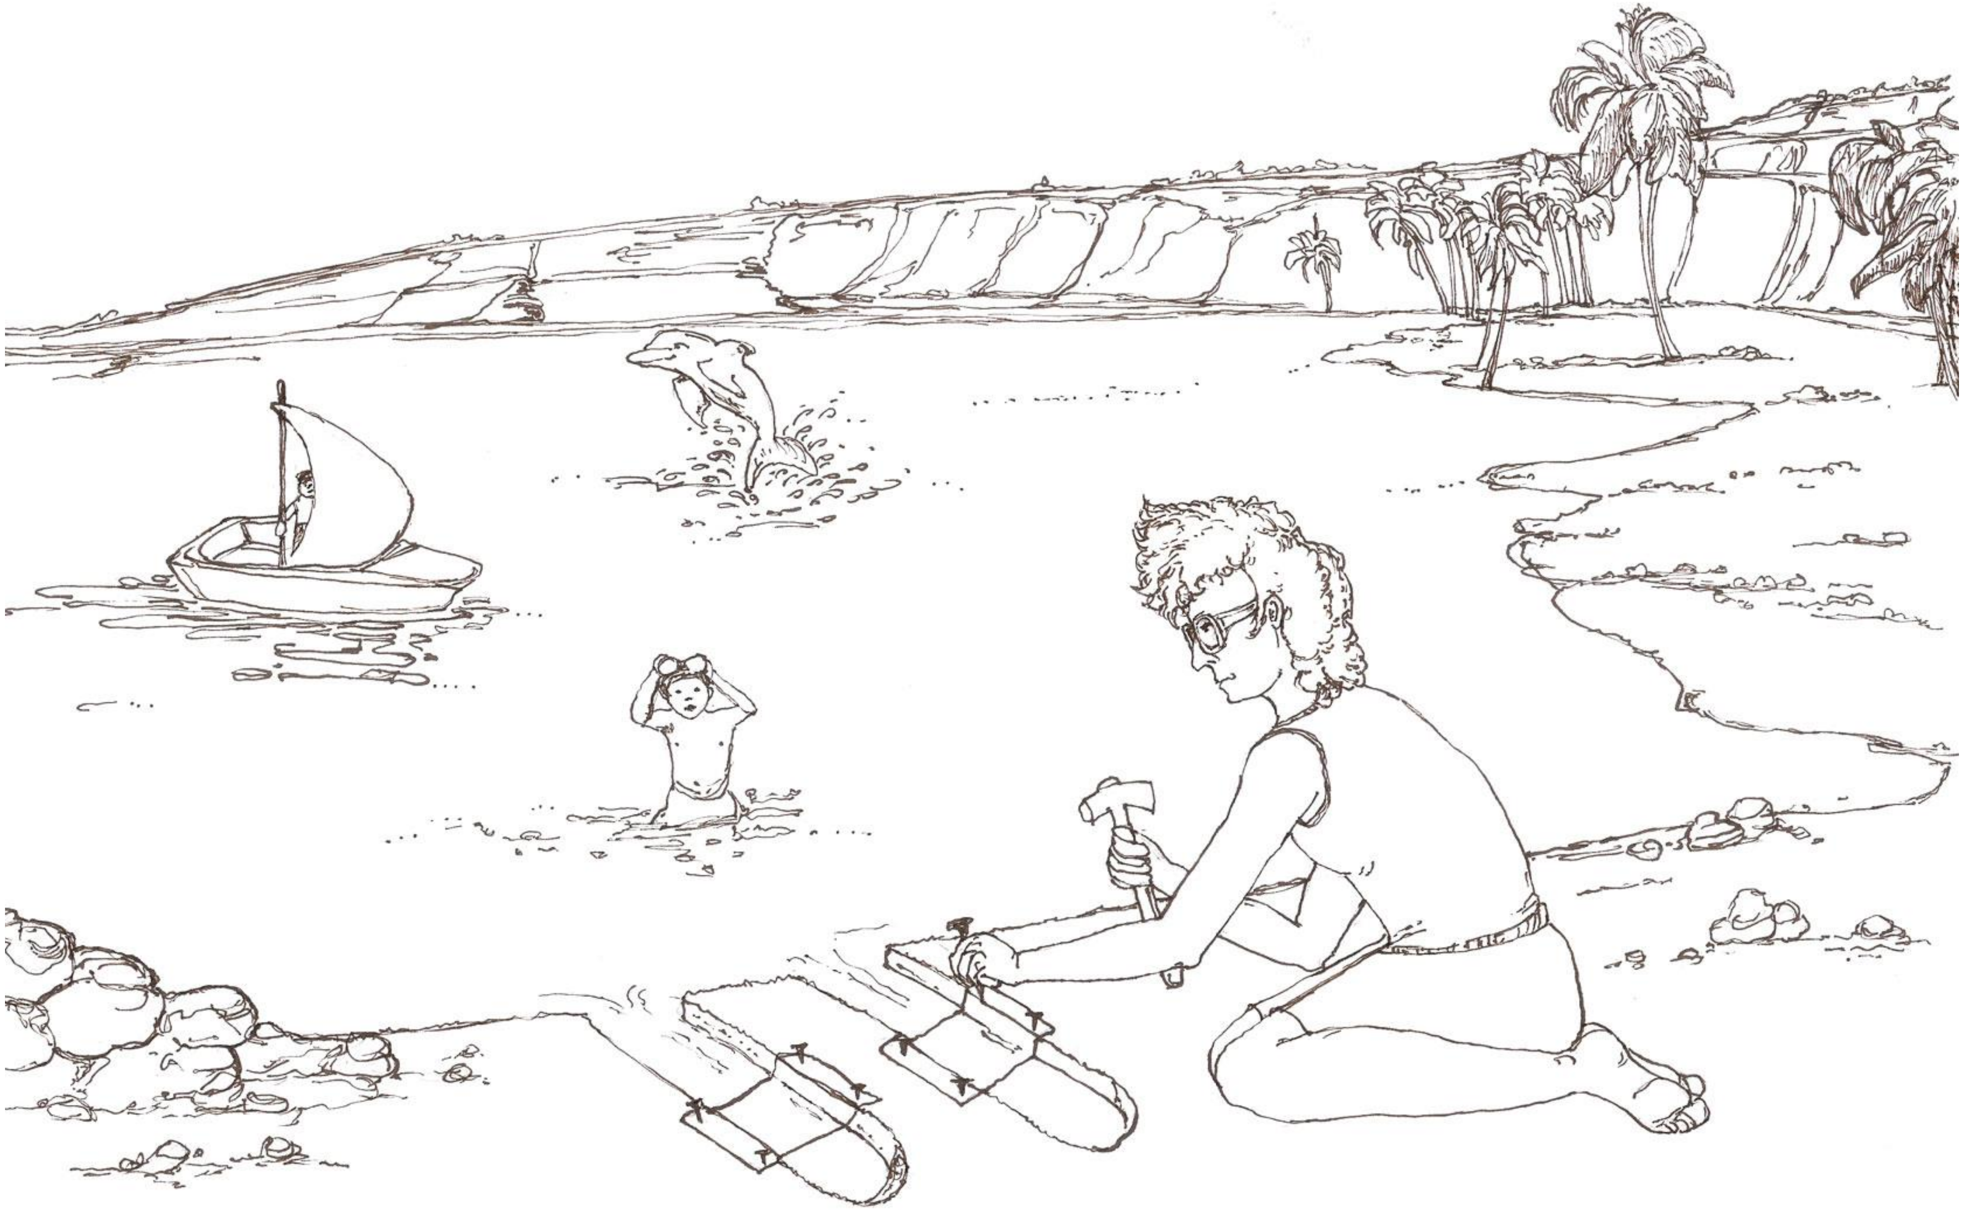

http://www.auditory.org/mhonarc/2019/msg00130.html

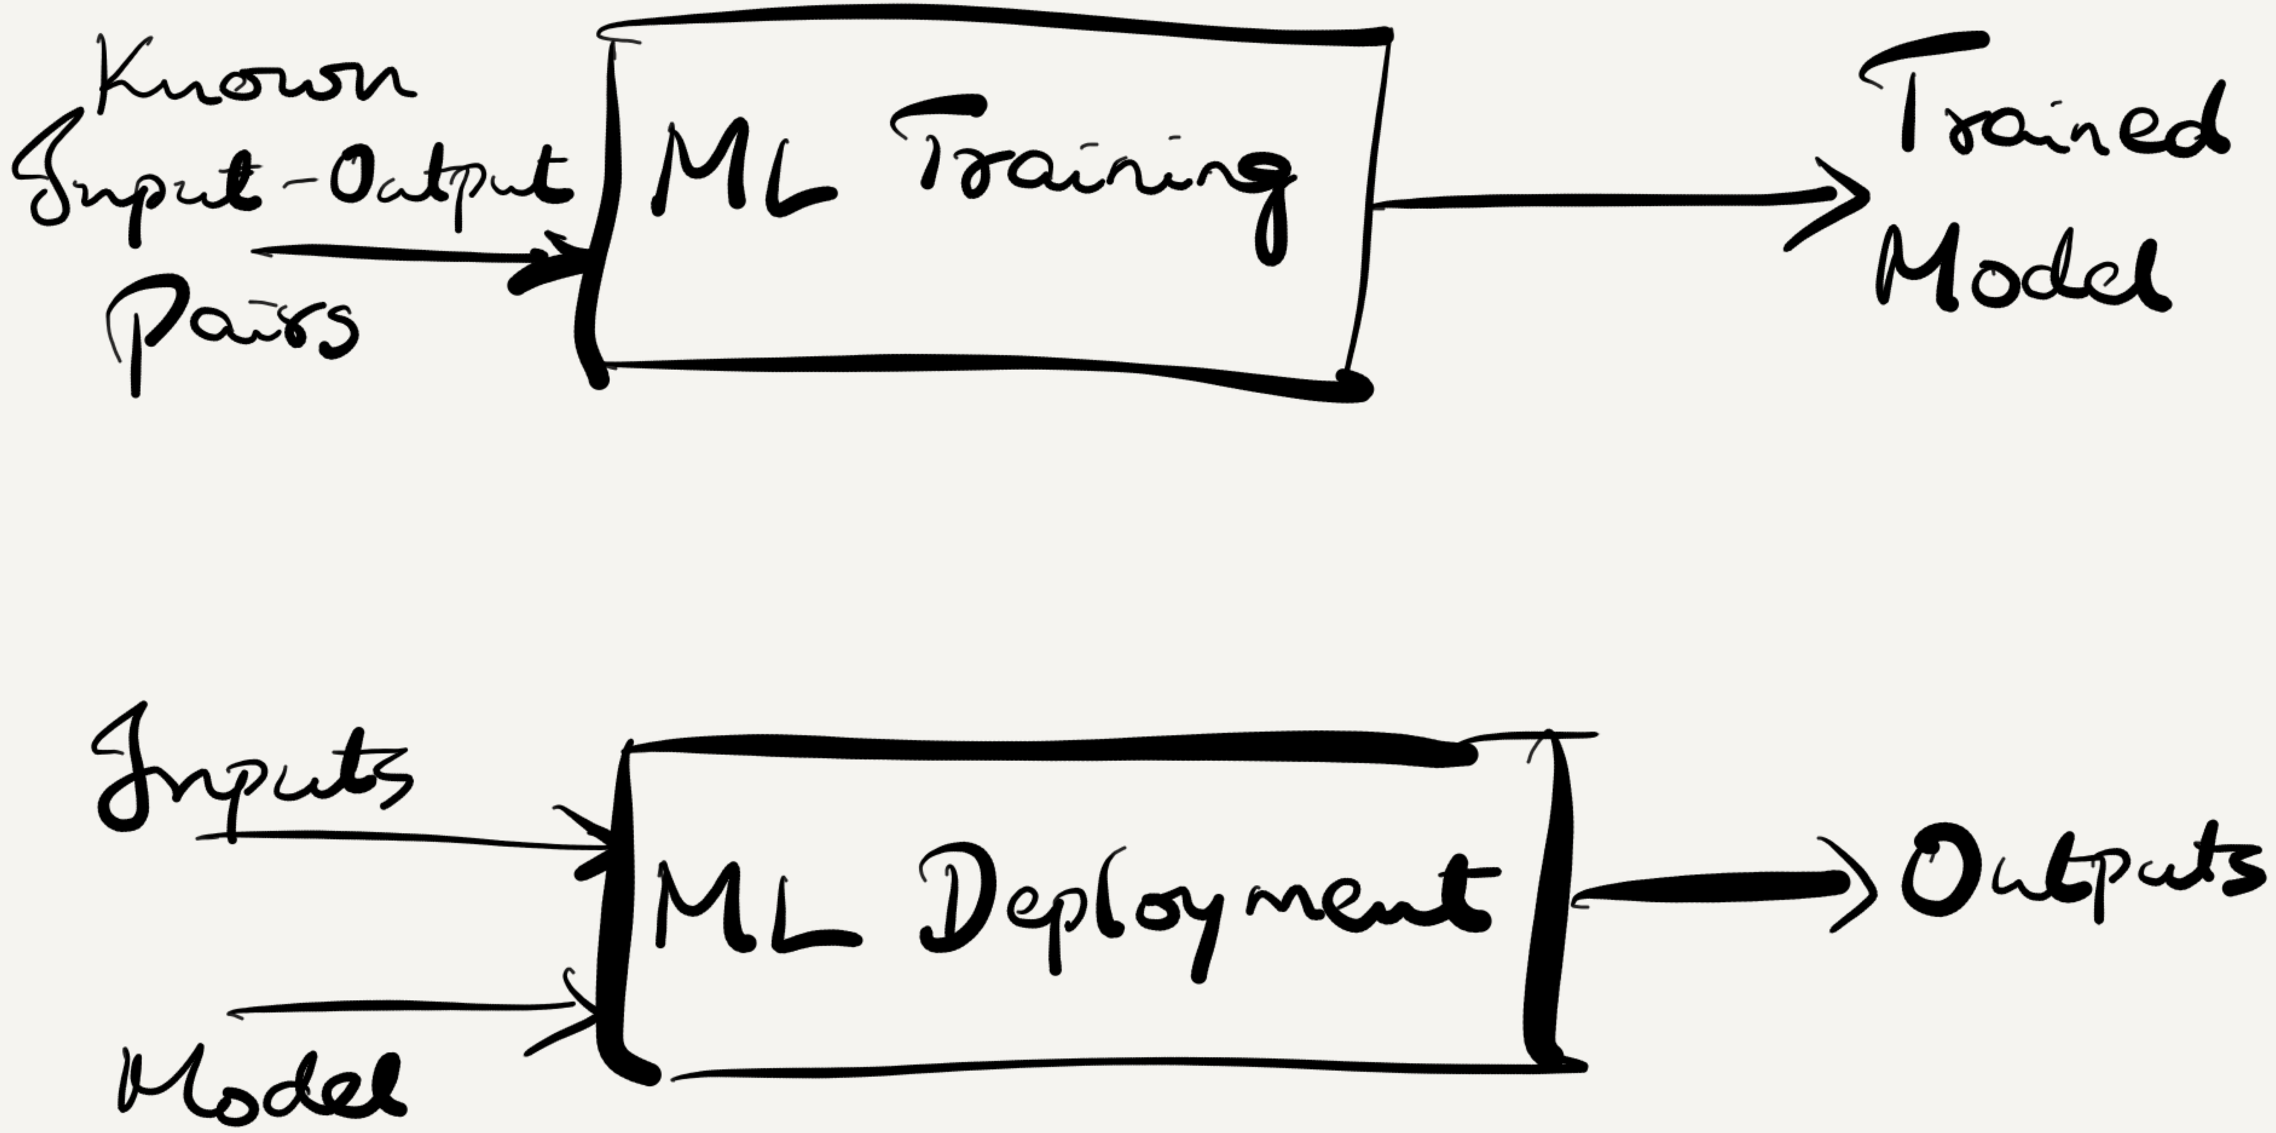

ML/AI AUTOMATION

Yogi Jaeger
@yoginho

Why not be honest and call "artificial intelligence" (#AI) "algorithmic mimicry" (AM) instead? It's been a misnomer (and AI researchers always knew that) from the very beginning. It would help us so much to properly understand & contextualize its capabilities & limitations.

The importance of stupidity in scientific research

Martin A. Schwartz

Department of Microbiology, UVA Health System, University of Virginia, Charlottesville, VA 22908, USA
e-mail: maschwartz@virginia.edu

Accepted 9 April 2008
Journal of Cell Science 121, 1771 Published by The Company of Biologists 2008
doi:10.1242/jcs.033340

I recently saw an old friend for the first time in many years. We had been Ph.D. students at the same time, both studying science, although in different areas. She later dropped out of graduate school, went to Harvard Law School and is now a senior lawyer for a major environmental organization. At some point, the conversation turned to why she had left graduate school. To my utter astonishment, she said it was because it made her feel stupid. After a couple of years of feeling stupid every day, she was ready to do something else.

I had thought of her as one of the brightest people I knew and her subsequent career supports that view. What she said bothered me. I kept thinking about it; sometime the next day, it hit me. Science makes me feel stupid too. It's just that I've gotten used to it. So used to it, in fact, that I actively seek out new opportunities to feel

I'd like to suggest that our Ph.D. programs often do students a disservice in two ways. First, I don't think students are made to understand how hard it is to do research. And how very, very hard it is to do important research. It's a lot harder than taking even very demanding courses. What makes it difficult is that research is immersion in the unknown. We just don't know what we're doing. We can't be sure whether we're asking the right question or doing the right experiment until we get the answer or the result. Admittedly, science is made harder by competition for grants and space in top journals. But apart from all of that, doing significant research is intrinsically hard and changing departmental, institutional or national policies will not succeed in lessening its intrinsic difficulty.

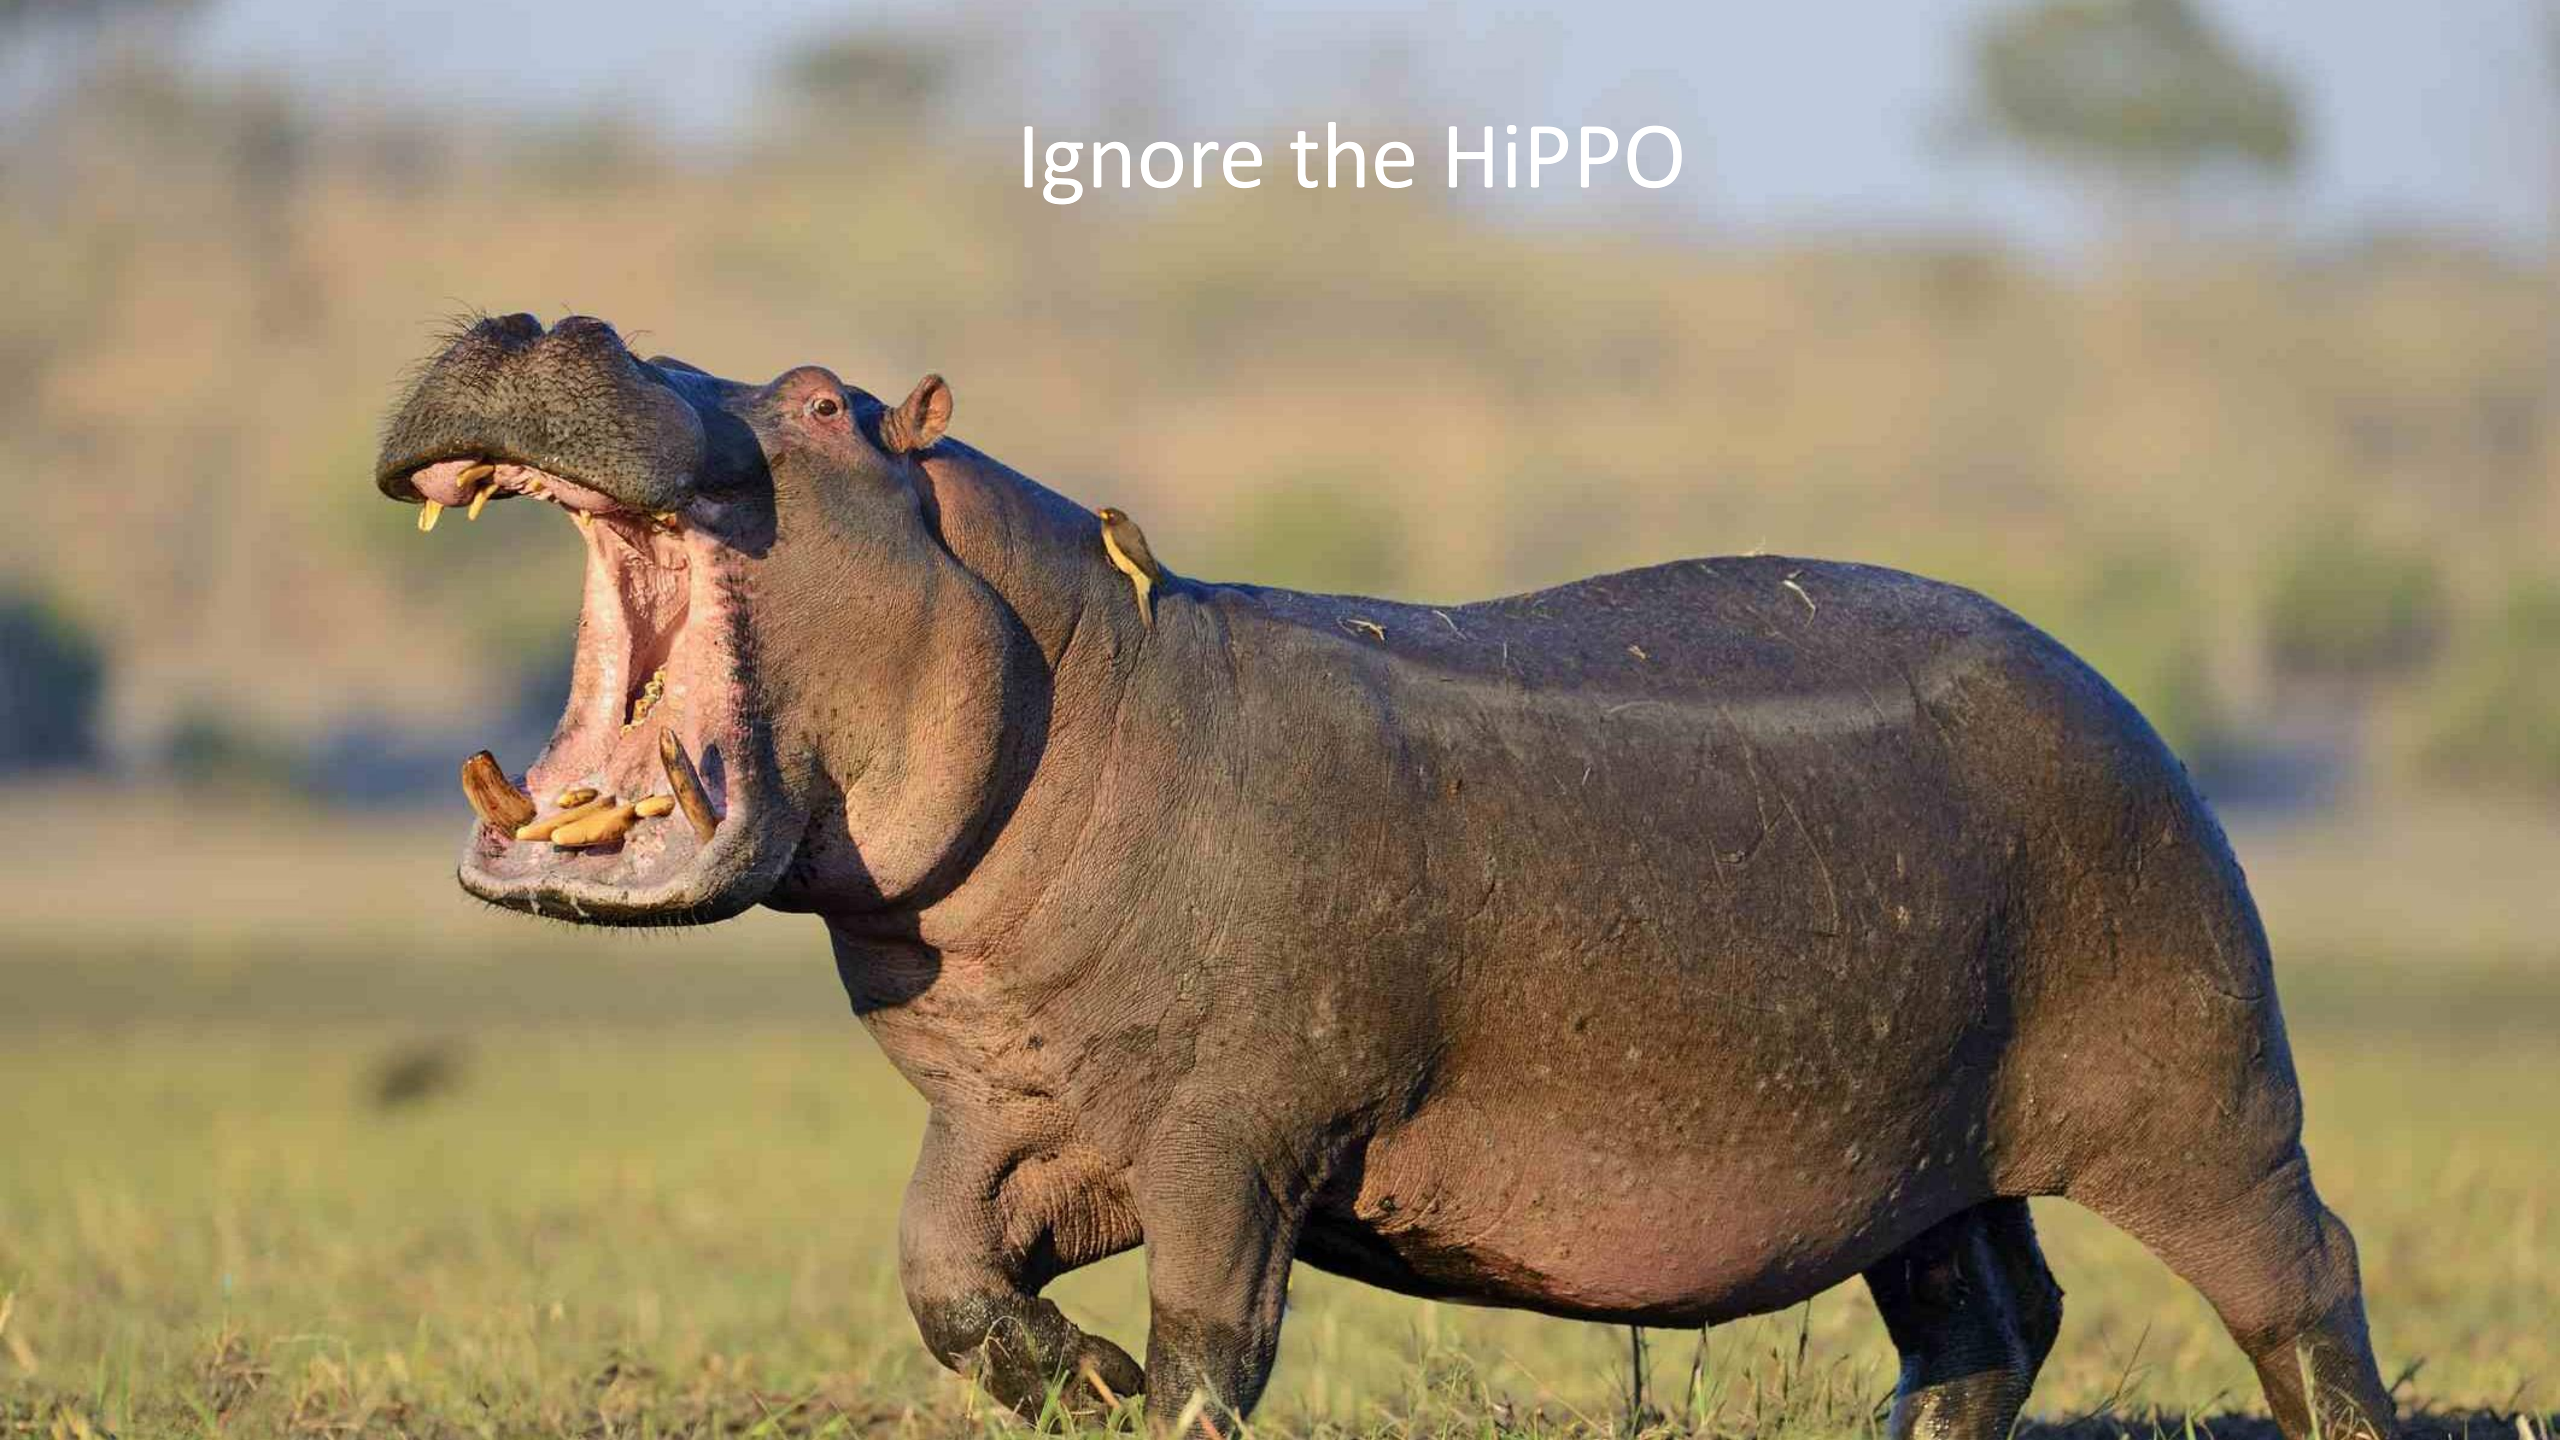
Ignore the HiPPO

Madhu Shashanka
@ShashankaMadhu

Couldn't agree more. But the key is empathy, what's simple for you maybe too complex and hard to use for most other people. You need to know what simple means for the people you are building for.

> **François Chollet** @fchollet · Apr 20, 2021
> Technologists shouldn't just make things possible. They should make them simple. In many ways it's far more difficult. To something possible you just need to be clever. To make it simple you need vision. Intelligence is common, vision is rare.
> Show this thread

Madhu Shashanka
@ShashankaMadhu

A moron in a hurry - en.wikipedia.org/wiki/A_moron_i...

This is the target persona UX design should always aim for!

> en.wikipedia.org
> A moron in a hurry - Wikipedia

Madhu Shashanka
@ShashankaMadhu

perfect use-cases for automation via ML are those where the cost of making mistakes are low

> **Andrew Whitby** @EconAndrew · Jan 26
> Maybe machine learning in radiology is on the wrong track.
>
> Rather than trying to replace radiologists in high-stakes work like diagnosing cancer, maybe it should target areas where radiology would currently be unaffordable, like routine ultrasound for minor soft tissue injuries.
>
> Show this thread

160k+ high school students will only graduate if a statistical model allows them to

Jun 17, 2020

tl;dr *: Due to curriculum disruptions the International Baccalaureate (IB) is going to use a statistical model to assign grades for > 160,000 high school students. These grades have a very significant impact on students' lives (for better or for worse). This is an inappropriate use of 'big data' and a terrible idea for a plethora of different reasons.*

Alan Mackworth
@AlanMackworth

"It's completely obvious that within 5 years deep learning will do better than radiologists", Geoff Hinton (2016) youtu.be/2HMPRXstSvQ

> **Alan Mackworth** @AlanMackworth · Mar 22, 2021
> "none of the models identified are of potential clinical use due to methodological flaws and/or underlying biases." t.co/SRhLG21PD9

Automation needs predictability

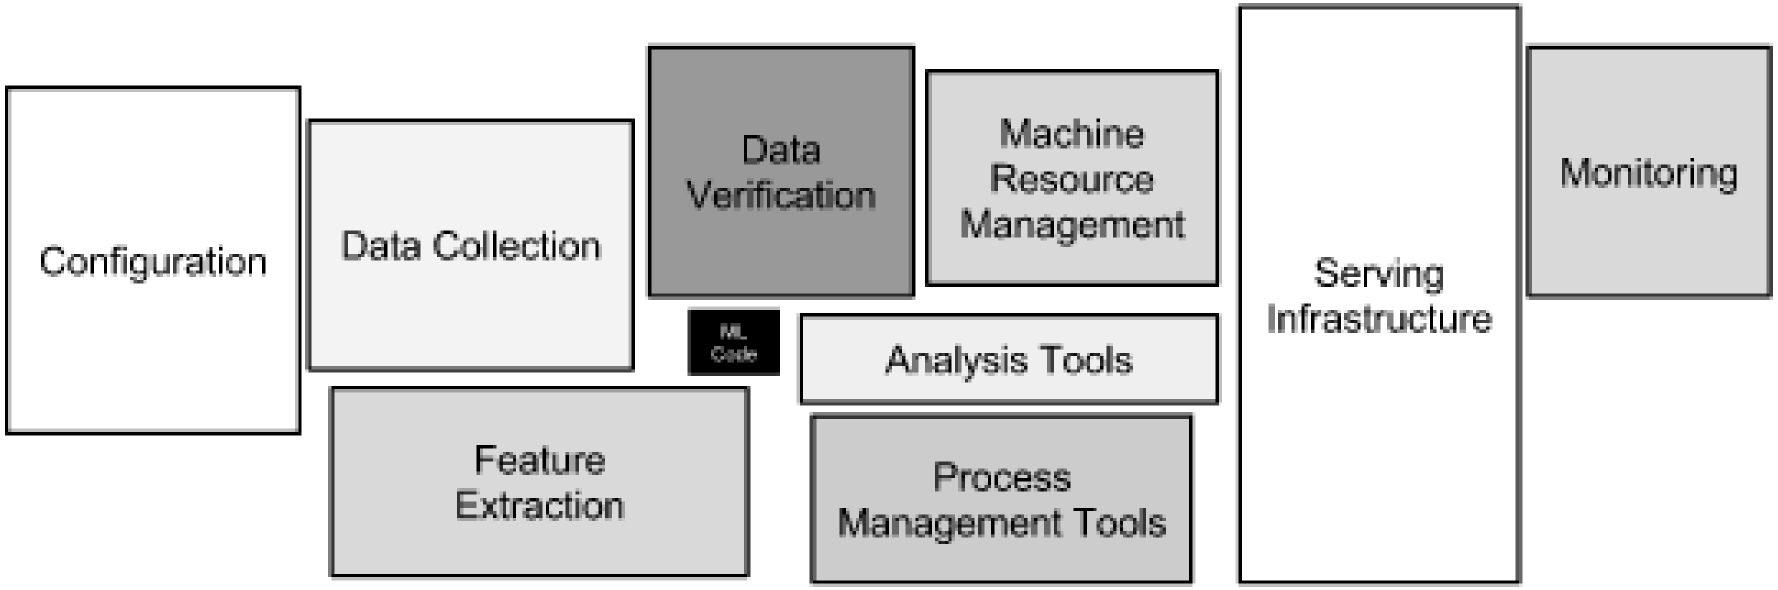

Figure 1: Only a small fraction of real-world ML systems is composed of the ML code, as shown by the small black box in the middle. The required surrounding infrastructure is vast and complex.

* Hidden Technical Debt in Machine Learning Systems, 2015

As we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns—the ones we don't know we don't know.

Madhu Shashanka
@ShashankaMadhu

Defining precise use-cases and having a clear understanding of the value of solving those use-cases are necessary for real-world success. This is where domain expertise comes into play. Machine learning / deep learning hammers alone will not give you that.

Guiding principles

While AI expertise is a key strength and differentiator for Concentric, what sets us apart is that data security is in our DNA.

Unlike newcomers in the field who see every challenge as a nail for the proverbial AI hammer, our team has decades of lived experience confronting and overcoming real-world cybersecurity challenges.

We've been thinking about the data security problem for quite some time, long before Concentric's founding back in 2018.

1. **First-principles thinking**. Focus on solving the challenges at hand, ignoring prevailing conventional wisdom of what's possible or not while avoiding unnecessary complexity. It starts with truly understanding the real pain-points of practitioners and identifying where you can add tangible value.
2. **Relentless focus on precision**. Any product that produces novel insights must confront the issue of false-positives. We had a strict requirement that anything we built kept false positives to a minimum (also known as "high-precision"). This is critical in a field like cybersecurity, which is chronically under-staffed. Practitioners cannot afford to chase after alerts that turn out to be benign.
3. **Enterprise-grade scale**. Whatever we built had to meet enterprise-level scale and prove to be bullet-proof enough to support enterprise environments. A corollary is that the technology stack chosen should be amenable to being operationalized in a cost-sustainable manner. There is a history of AI startups ignoring this at their peril, releasing seemingly impressive demos only to eventually fold at large data volume thresholds.

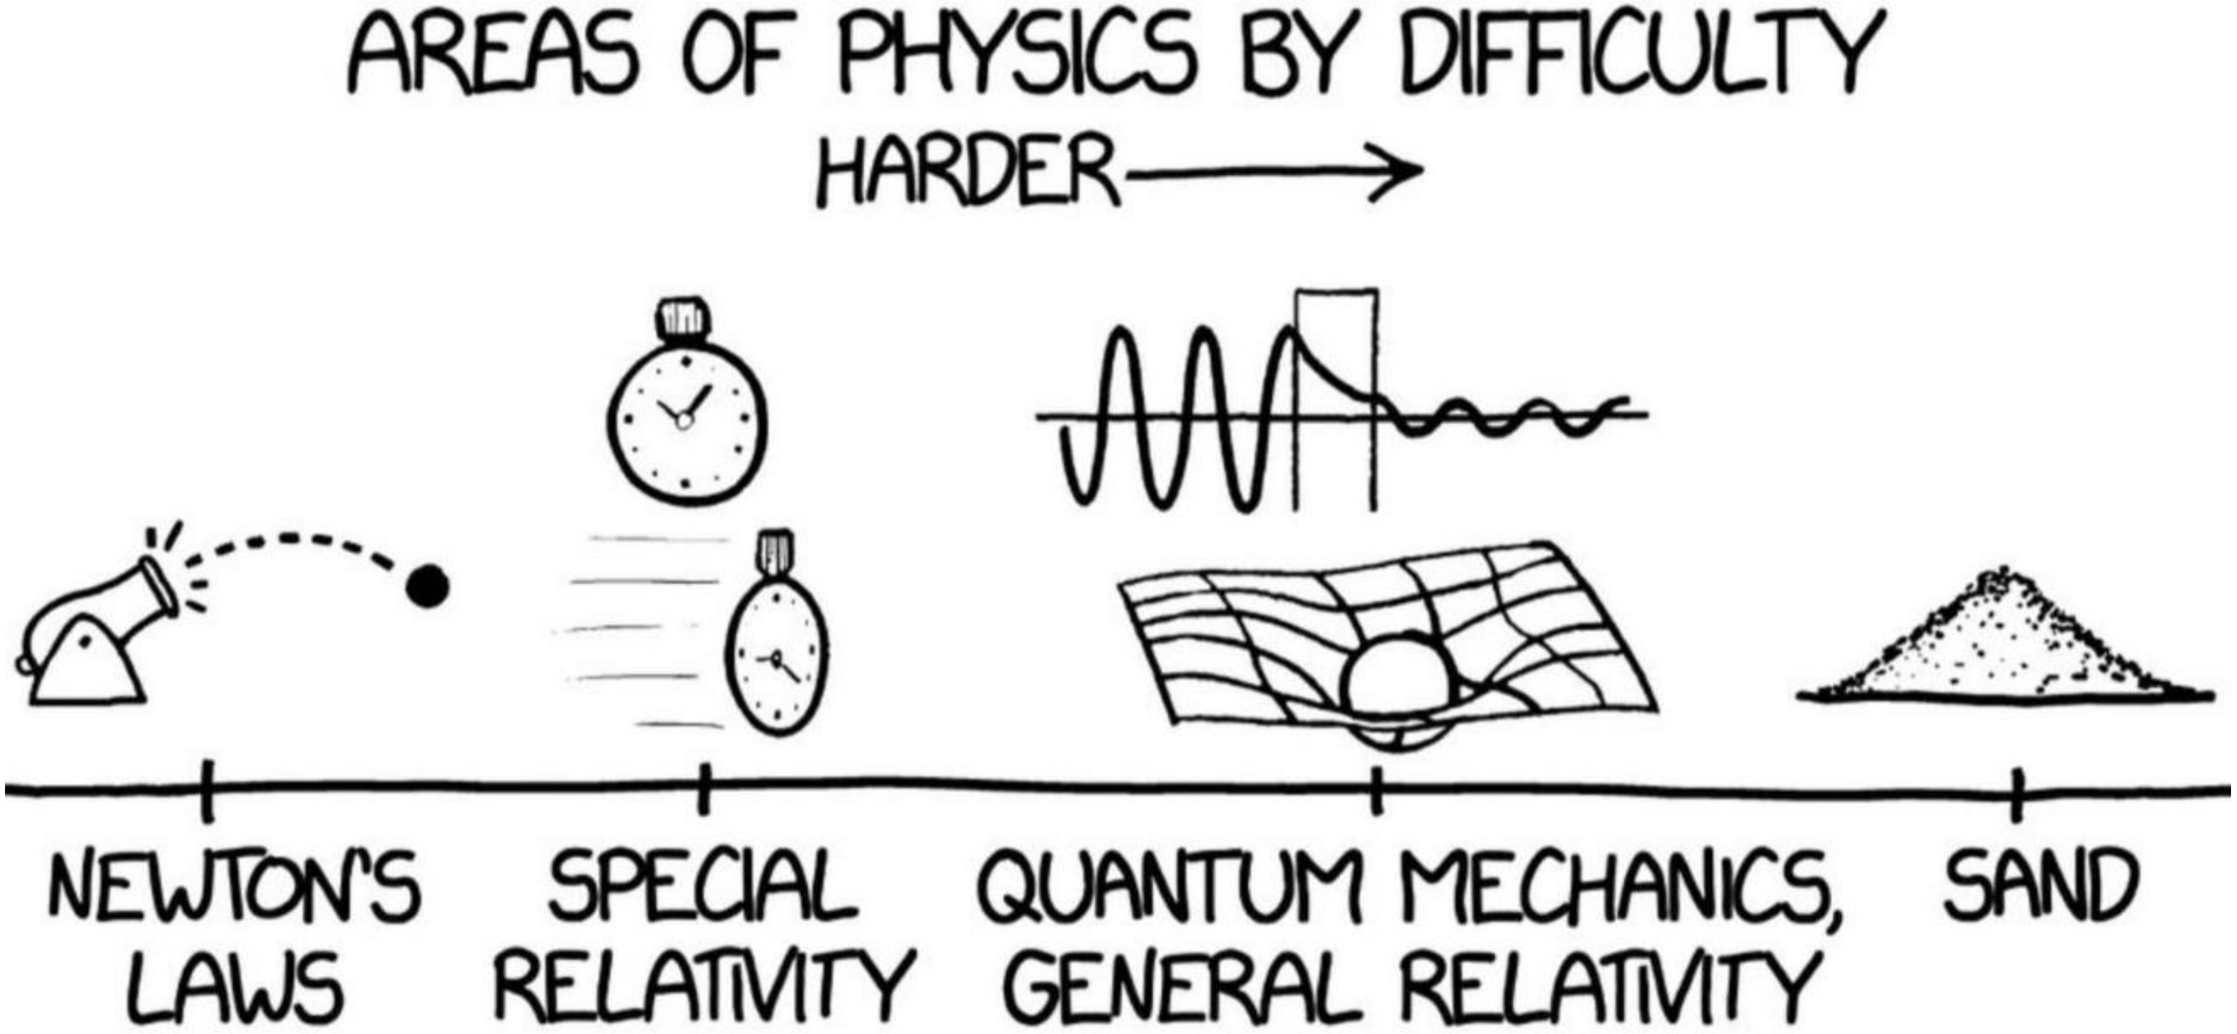
AREAS OF PHYSICS BY DIFFICULTY
HARDER
NEWTON'S LAWS
SPECIAL RELATIVITY
QUANTUM MECHANICS, GENERAL RELATIVITY
SAND

* xkcd

vicki
@vboykis

LLMs are so weird because one side is people with five PhDs who have been studying neuron activations for the past three decades and on the other side is someone called leetm5n with an anime avatar just casually releasing increasingly better performing fine tunes of mistral

Madhu Shashanka
@ShashankaMadhu

I love this! I see two more personas in addition to the algo/math folks and the tinkerers - the infra people who make the compute scale, and the application domain / business experts who can now apply AI to their use-cases without having to build an army

"… has become something of a scientific bandwagon.

… it has received an extraordinary amount of publicity in the popular as well as scientific press.

… it has perhaps ballooned to an importance beyond its actual accomplishments. "

"… has become something of a scientific bandwagon.

… it has received an extraordinary amount of publicity in the popular as well as scientific press.

… it has perhaps ballooned to an importance beyond its actual accomplishments. "

- Claude Shannon, 1956
  on Information Theory

The Bandwagon

CLAUDE E. SHANNON

INFORMATION theory has, in the last few years, become something of a scientific bandwagon. Starting as a technical tool for the communication engineer, it has received an extraordinary amount of publicity in the popular as well as the scientific press. In part, this has been due to connections with such fashionable fields as computing machines, cybernetics, and automation; and in part, to the novelty of its subject matter. As a consequence, it has perhaps been ballooned to an importance beyond its actual accomplishments. Our fellow scientists in many different fields, attracted by the fanfare and by the new avenues opened to scientific analysis, are using these ideas in their own problems. Applications are being made to biology, psychology, linguistics, fundamental physics, economics, the theory of organization, and many others. In short, information theory is currently partaking of a somewhat heady draught of general popularity.

Although this wave of popularity is certainly pleasant and exciting for those of us working in the field, it carries at the same time an element of danger. While we feel that information theory is indeed a valuable tool in providing fundamental insights into the nature of communication problems and will continue to grow in importance, it is certainly no panacea for the communication engineer or, *a fortiori*, for anyone else. Seldom do more than a few of nature's secrets give way at one time. It will be all too easy for our somewhat artificial prosperity to collapse overnight when it is realized that the use of a few exciting words like *information*, *entropy*, *redundancy*, do not solve all our problems.

What can be done to inject a note of moderation in this situation? In the first place, workers in other fields should realize that the basic results of the subject are aimed in a very specific direction, a direction that is not necessarily relevant to such fields as psychology, economics, and other social sciences. Indeed, the hard core of information theory is, essentially, a branch of mathematics, a strictly deductive system. A thorough understanding of the mathematical foundation and its communication application is surely a prerequisite to other applications. I personally believe that many of the concepts of information theory will prove useful in these other fields—and, indeed, some results are already quite promising—but the establishing of such applications is not a trivial matter of translating words to a new domain, but rather the slow tedious process of hypothesis and experimental verification. If, for example, the human being acts in some situations like an ideal decoder, this is an experimental and not a mathematical fact, and as such must be tested under a wide variety of experimental situations.

Secondly, we must keep our own house in first class order. The subject of information theory has certainly been sold, if not oversold. We should now turn our attention to the business of research and development at the highest scientific plane we can maintain. Research rather than exposition is the keynote, and our critical thresholds should be raised. Authors should submit only their best efforts, and these only after careful criticism by themselves and their colleagues. A few first rate research papers are preferable to a large number that are poorly conceived or half-finished. The latter are no credit to their writers and a waste of time to their readers. Only by maintaining a thoroughly scientific attitude can we achieve real progress in communication theory and consolidate our present position.

Hype Cycle for Artificial Intelligence, 2023

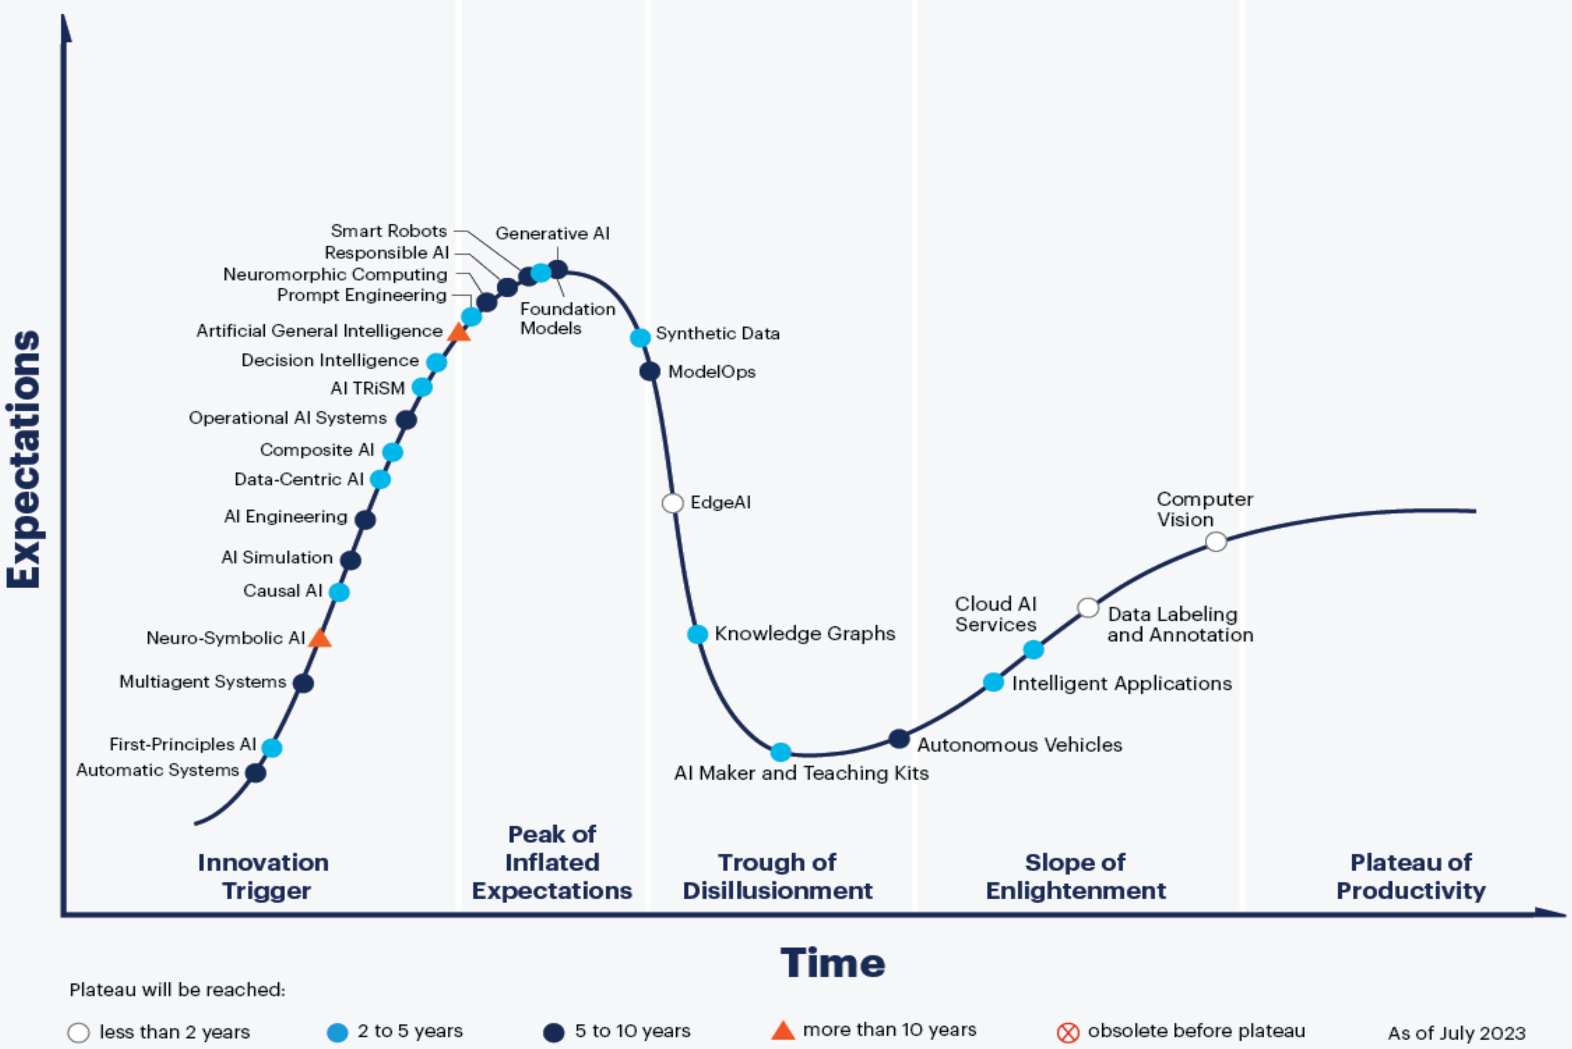

gartner.com

Source: Gartner

Gartner