

# Machine Learning

Where we are, and what to expect

Mikael  
Lundgren

mikael@levla.se  
@theintrovert



[www.levla.se](http://www.levla.se)





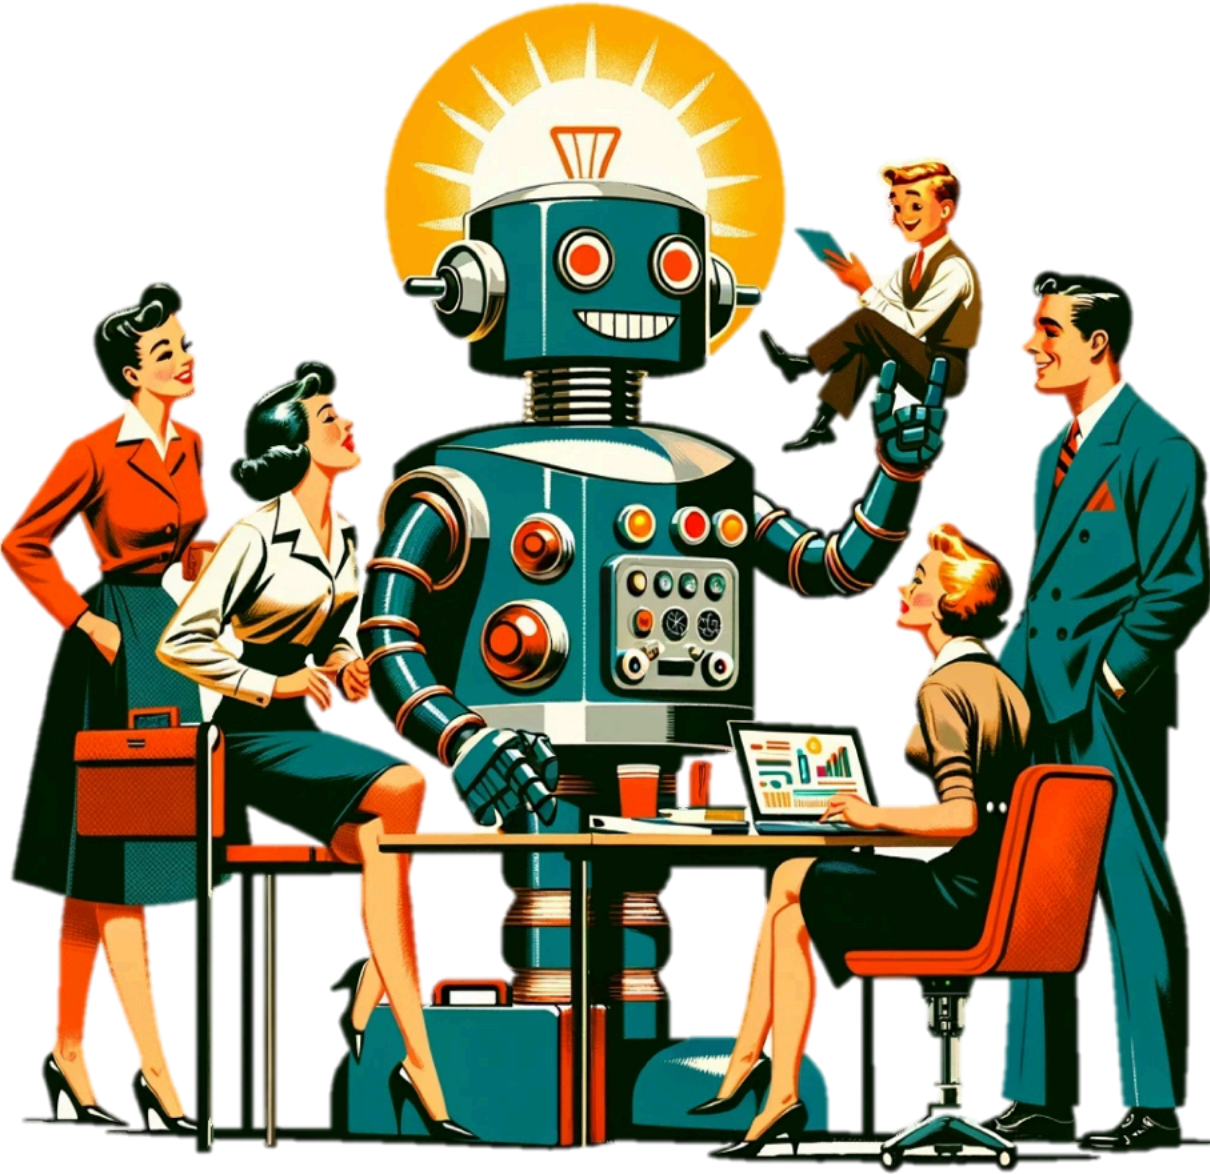
- **M.Sc. Computer Science**
  - Org. theory, Pedagogy
- **25 years Lean/Agile Prod. Development**
- **Business Founder (Levla, 2011)**
- **Followed AI research since 2014**
  - Worked actively with AI since 2017



<https://www.linkedin.com/in/mikaellundgren/>



# Machine Learning (“AI”): Where are we heading?



Where did it come from?

What are the limitations?

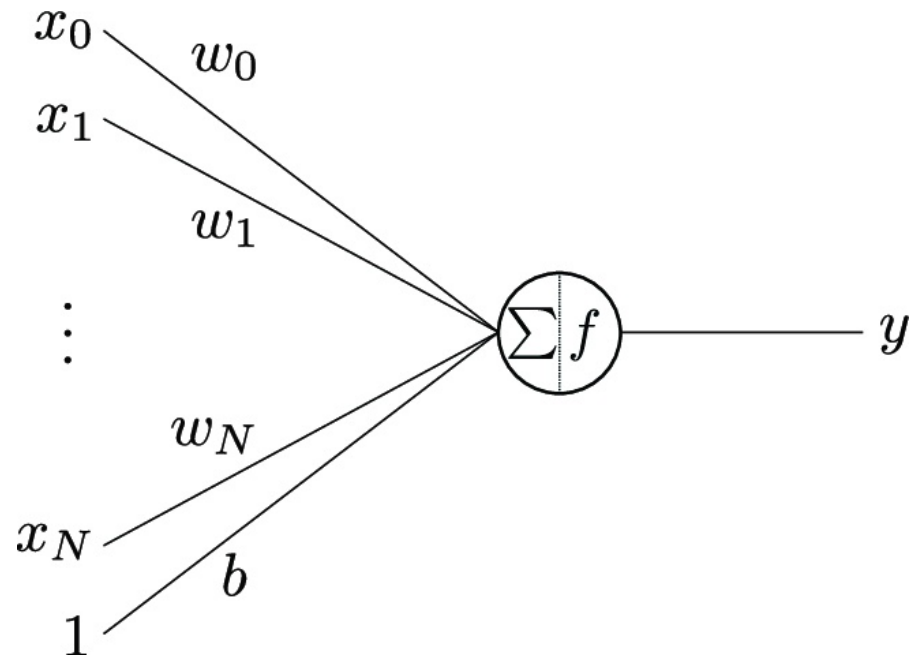
What are the possible outcomes?

*"The AI helper" by Dall-E*

# Back in 1993...

We studied “neural networks” and “simulated neurons” in the AI class of Computer Science studies. We built a simple classifier that could learn to mimic an XOR circuit, and experimented with slightly more advanced models.

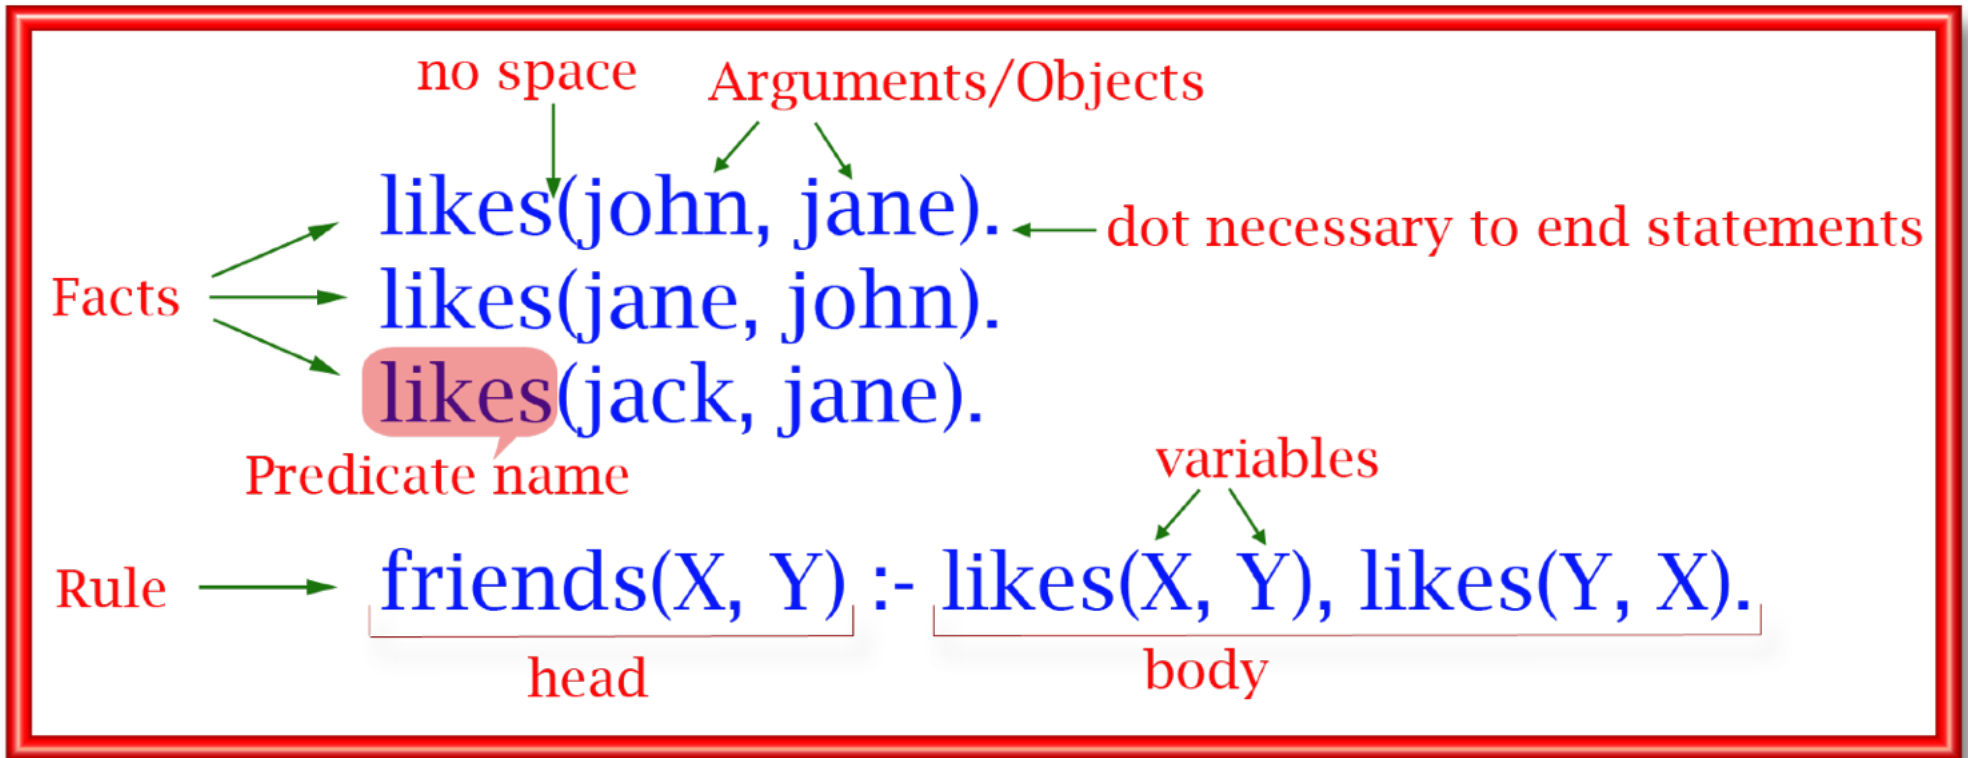
The theory intrigued me, but computers were way too slow to make it other than a thought experiment.





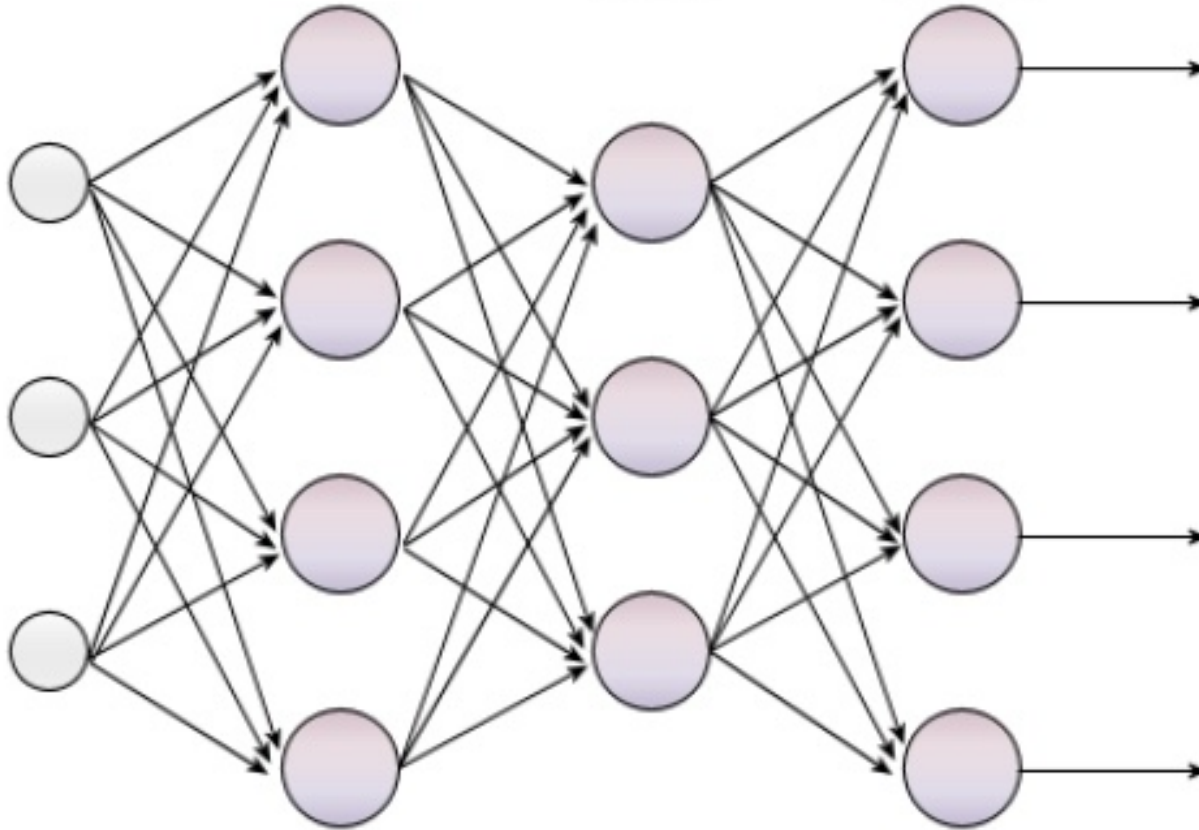
# Much focus was still on rule-based “AI”

## Program Window



Prolog code example from: <https://athena.ecs.csus.edu/~mei/logicp/prolog/programming-examples.html>

# A big problem: The Combinatorial Explosion



*A multi-layered perceptron (fully connected),  
Wikimedia Commons*

# What happened?

## Two things:

1. A breakthrough in research by several parties, who found a way to dramatically reduce the number of connections (thus computations) without loss of quality
2. The fierce competition in the entertainment industry, causing huge investments being made in the evolution of computer graphics hardware

*(I guess I should have bought those NVIDIA shares after all)*

# How are they connected?

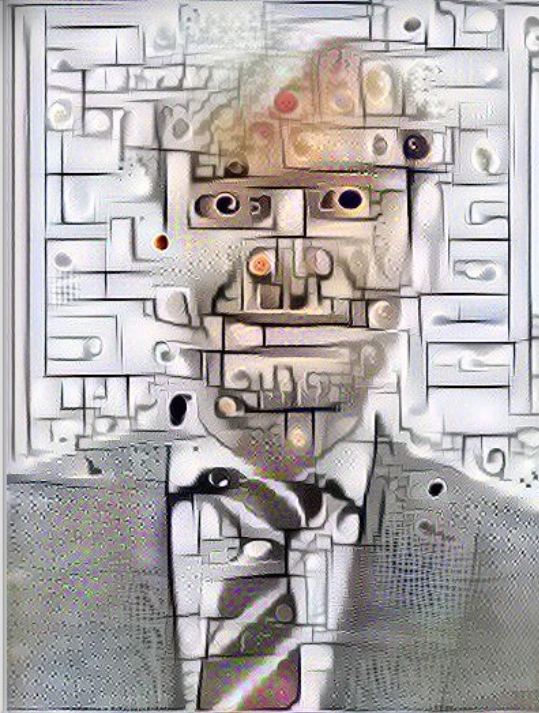
As it turns out, the majority of computation required for a neural network consists of *matrix operations*.

Coincidentally, a lot of the operations required by software to process 2d images and present 3d worlds and imagery also utilize *matrix operations*.



$$\begin{bmatrix} \bullet 0 & \bullet 0 & \bullet 0 \\ \bullet 0 & \bullet 1 & \bullet 0 \\ \bullet 0 & \bullet 0 & \bullet 0 \end{bmatrix} + \begin{bmatrix} \bullet 0 & \bullet 0 & \bullet 0 \\ \bullet 0 & \bullet 1 & \bullet 0 \\ \bullet 0 & \bullet 0 & \bullet 0 \end{bmatrix} - \frac{1}{9} \begin{bmatrix} \bullet 1 & \bullet 1 & \bullet 1 \\ \bullet 1 & \bullet 1 & \bullet 1 \\ \bullet 1 & \bullet 1 & \bullet 1 \end{bmatrix} = \begin{bmatrix} \bullet 0 & \bullet 0 & \bullet 0 \\ \bullet 0 & \bullet 2 & \bullet 0 \\ \bullet 0 & \bullet 0 & \bullet 0 \end{bmatrix} - \frac{1}{9} \begin{bmatrix} \bullet 1 & \bullet 1 & \bullet 1 \\ \bullet 1 & \bullet 1 & \bullet 1 \\ \bullet 1 & \bullet 1 & \bullet 1 \end{bmatrix}$$





# What about *generating* data instead?

*Early examples involved applying pre-trained feature maps onto images to tweak the style or presentation - as some feature maps or layers encode the visual style but not the contents, and vice versa!*

Image Style Transfer Using  
Convolutional Neural Networks

*Gatys et al, 2016, images from custom  
Python code by Mikael Lundgren 2016.*



# GAN - The Key to Generative AI

A Generative Adversarial Network consists of two competing networks:

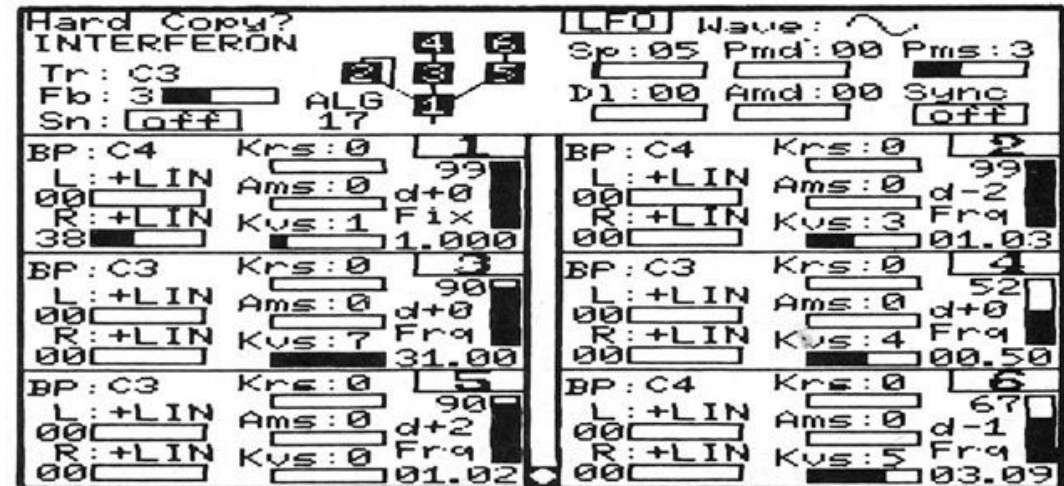
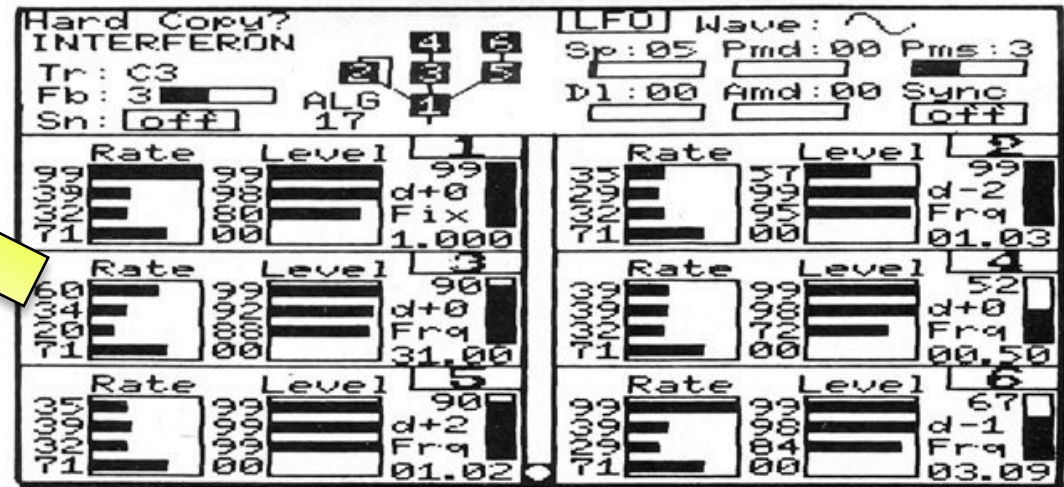
- A **discriminator**, learning to discern fake from real data
- A **generator**, trying to generate data good enough to fool the discriminator



*Image by Dall-E: The Artist and the Inspector*

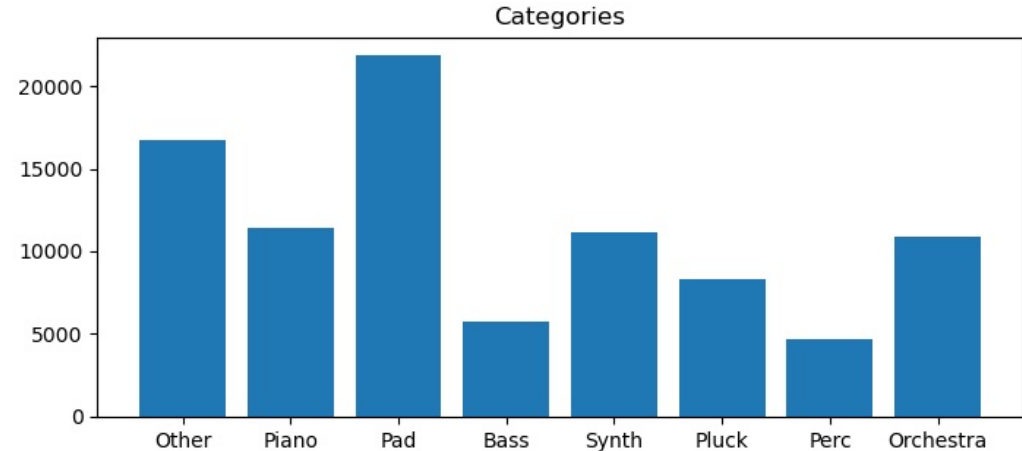


# Machine Learning and Sound Design



# Training the CGAN

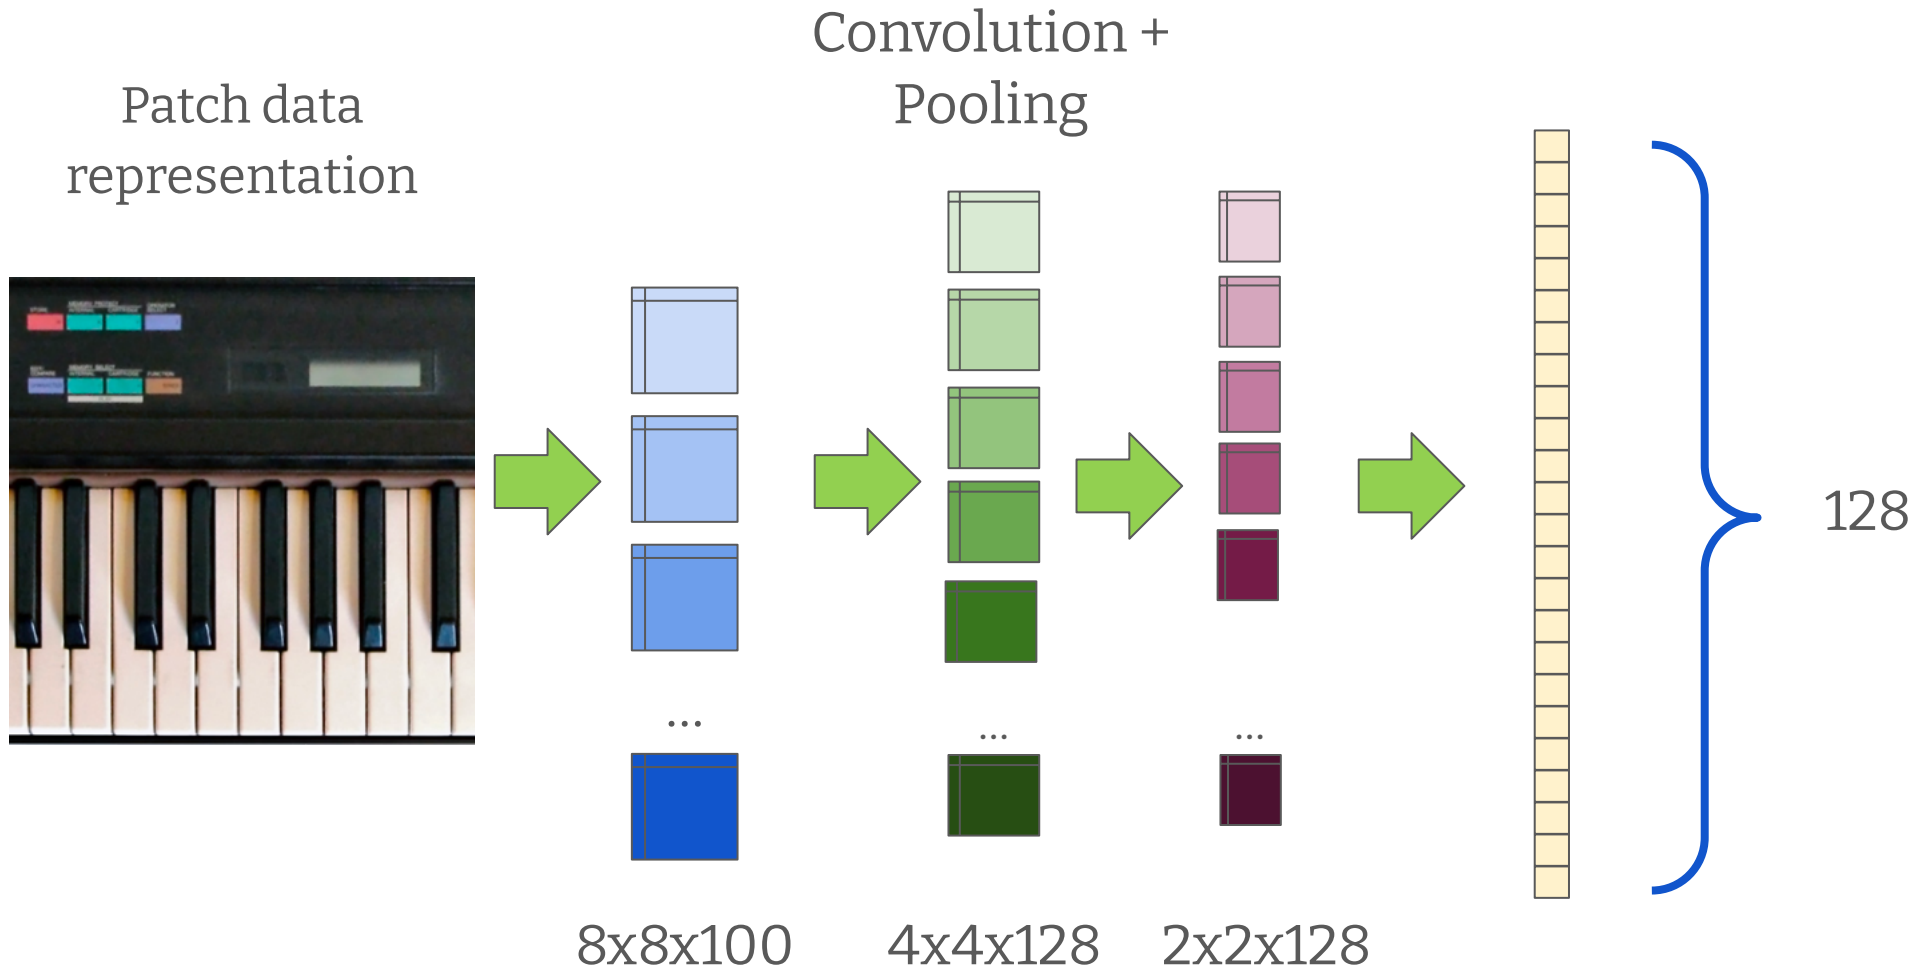
- “Deep” GAN network, trained on ~50 000 patches collected from the net
- Many of which were created in the mid 80’s - by people hunched over an actual DX7
- Fairly good diversity of patches, GAN stabilizes at about 400 epochs



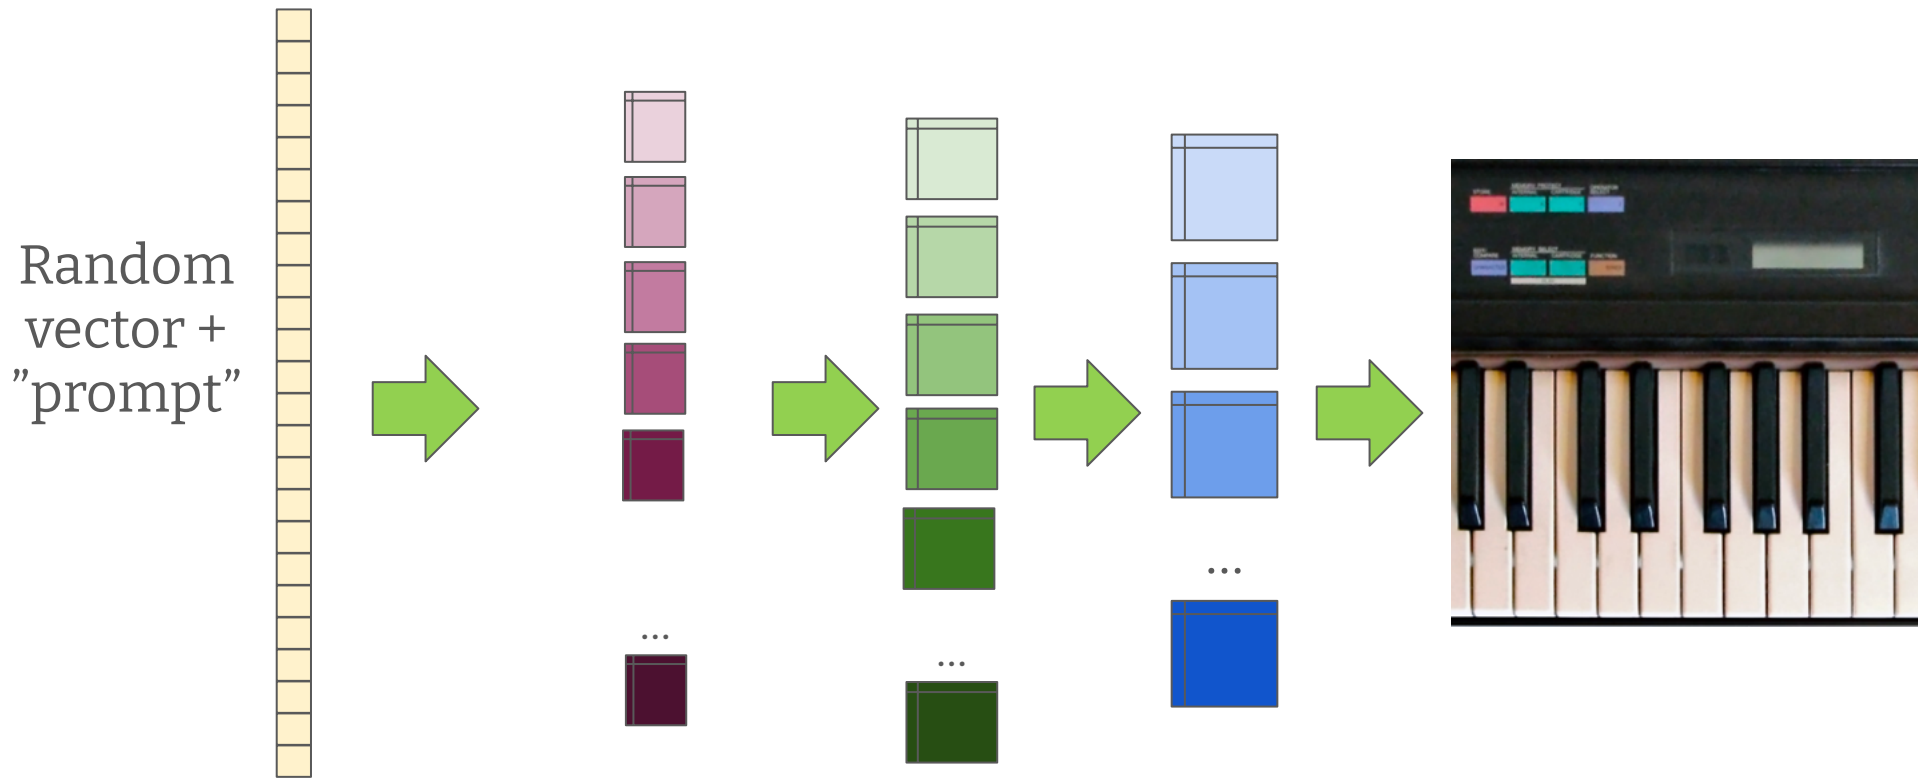
Read more:  
[www.theintrovert.net](http://www.theintrovert.net)



# Network Structure



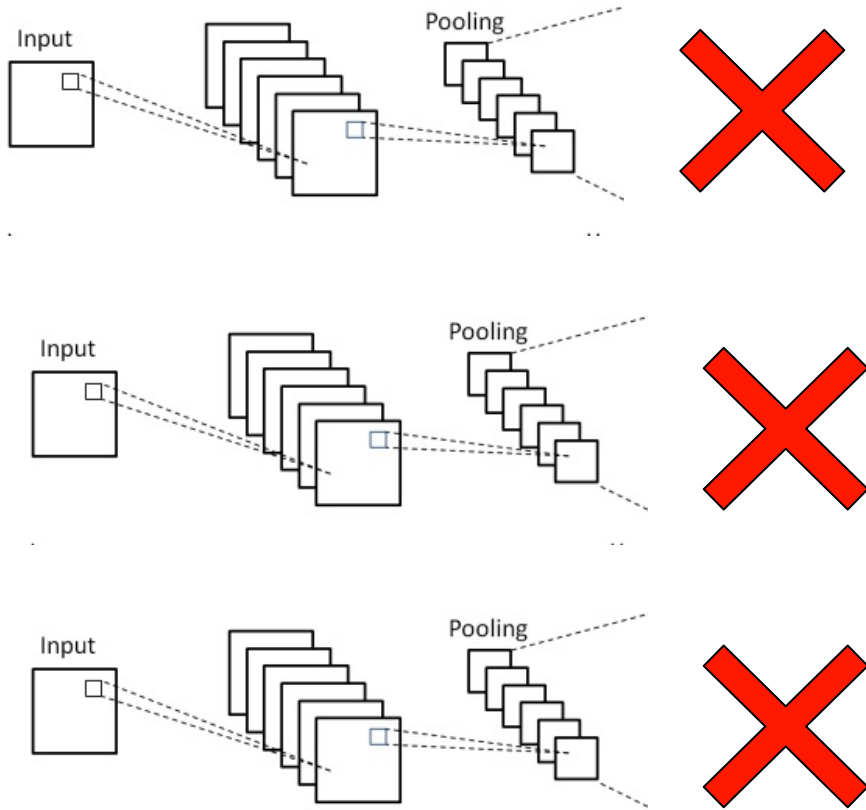
# The Generator - Running the Process in Reverse





# Deciding the Architecture

This was surprisingly full of trial-and-error.



I don't know exactly *why* the final topology worked and the others not - and this seems common!



*My cat telling me he is annoyed with the nasty high-frequency sounds I'm making*

# Findings

- 80% of the work went into validating and scrubbing data
  - This is "regular" code, nothing fancy at all
- Presence of invalid or poor data has a large impact on the results
- Training a classifier is much easier than a generative network
  - You rely on good practices rather than hard rules to decide when results are "good enough"
  - A lot of the big decisions (architecture etc.) are also a bit of trial and error
- This is still based on a very new field of research!



# The current limitations of generative AI



# Why the inconsistent results?

Unlike a director who takes every element of a scene into consideration, here every frame is generated as the *most likely to follow the previous frame, considering the prompt*.

*"Prompt" referring to the instructions given to the generative system.*

Thus it becomes obvious that these networks lack an understanding of the entire scene and its meaning.

*Prompt:*

*Cinematic shot of a secret agent placing bets at a Black Jack table in a 1960s casino.*





# LLM - Large Language Models

Large Language Models have a huge use case, making them visible and popular in the public eye, next to image generation:

- Chatbots
- Summarizing and transcribing long texts
- Translation
- ...

This also makes them able to mimic *AGI* - Artificial General Intelligence, which is a problem in itself.



# Google Fires Engineer Who Claims Its A.I. Is Conscious

The engineer, [redacted] contends that the company's language model has a soul. The company denies that and says he violated its security policies.

Microsoft's chatbot integrated into Bing has been reportedly acting weird in conversations with users, going so far as to compare one journalist (interacting with the LLM at the time) to famous dictators, and trying to persuade a journalist from the New York Times that he was in love with the chatbot, and not his partner, in quite a persuasive manner.



# An LLM is *not* an expert system

It is a network designed to:

1. Infer the most probable next word in a sentence
2. Infer the most probable reply to input

It might make up information if it has insufficient data (“hallucinate”)

**I  
NEED  
YOUR  
CLOTHES,  
YOUR BOOTS  
AND  
YOUR  
MOTORCYCLE**



# Always Check the Facts



Mikael

*Saab 96 V4 and Saab 95 V4 are two different car models manufactured by Saab AB in Sweden [...] Here are some of the most important differences between the models:*

*...*

- 1. Engines: Saab 96 V4 used a four cylinder engine from Ford developed specially for Saab, while Saab 95 V4 used a four cylinder Mercedes engine [...]*

*I used to own a Saab 95 V4 back in the day, and it had a Ford engine, so I suspect you are wrong.*

***I am sorry if I have provided incorrect information.*** *It might be useful to remember that the Saab 95 V4 was manufactured for a long period of time, from 1959 until 1978, and several changes and updates to the model could have occurred during this time. It is possible that your Saab 95 V4 had a Ford engine, especially if it was from a later production run. It is also possible that there exists multiple variants of Saab 95 V4 using different engine types.*



# Double Hallucination

Language models, like all generative AI, grapple with what we call “hallucinations.” Instead of admitting “I don’t know,” they may conjure up details out of thin air just to please us.

- *A Survey of Hallucination Problems Based on Large Language Model* - Xinxin Liu
- *Reducing LLM Hallucination Using Knowledge Distillation: A Case Study with Mistral Large and MMLU Benchmark* - D. McDonald, Rachael Papadopoulos, Leslie Benningfield
- *FactCheckmate: Preemptively Detecting and Mitigating Hallucinations in LMs* - Deema Alnuhait, Neeraja Kirtane, Muhammad Khalifa

And here’s the twist with text-based generative AI: we run a double risk of hallucination—once when the system interprets our prompt, and once more in the act of generating its response.

# A Professional Way to Approach Current LLM's

- They are somewhat arrogant know-it-alls.
- Treat as you would your uncle, who always have a response to things, and the most amazing stories to tell - slightly sceptic.
- LLM's trained on a very specific data set may be less prone to hallucinations, but they are not immune!

Use them as a discussion partner, not an information trove!

Most serious LLM's can offer links to the sources they claim to use -  
review the information yourself!

# AI: The Future

A lush green forest scene with a bright light source on the ground, creating a magical atmosphere. The image is dominated by vibrant green foliage and tree canopies. A bright, ethereal light emanates from the forest floor in the center, casting a long, soft glow across the grass. The overall mood is serene and hopeful, suggesting a bright future.

# Classifying Requirements with AI

Early research indicates that both explicit and implicit “wants” regarding a product can be extracted successfully from user feedback, and mapped against product strategy tools such as Kano analysis, much faster and more precise than humans are able to.

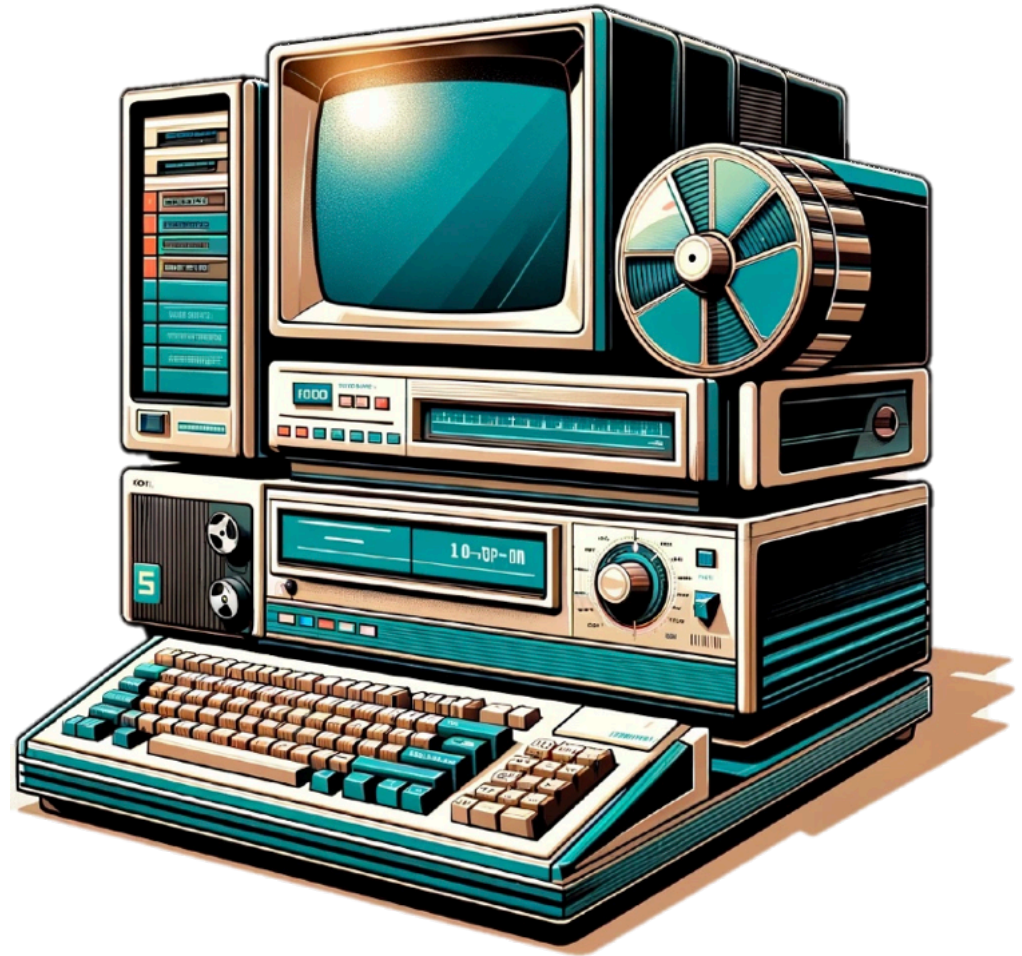
This will likely lead to customer feedback and -input becoming increasingly more sought after and processed by automatic means to be able to reduce some guesswork out of product development and improvement.

*Kejia, Chen., Jian, Jin., Jiayi, Luo. (2021). Big consumer opinion data understanding for Kano categorization in new product development. Journal of Ambient Intelligence and Humanized Computing, 1-20. doi: 10.1007/S12652-021-02985-5*



# Is Software Development going extinct?

- Short term, a lot of shuffling data between different API's and similar simpler, yet error-prone programming tasks will probably be automated
- Web development is already under fire
- A more relevant question is: Is Google going extinct?



# Why Google when you can discuss?

not for the faint of heart

Goodreads  
<https://www.goodreads.com> · Översätt den här sidan 19

**Not for the Faint of Heart by Lex Croucher**  
Lex Croucher's *Not For the Faint of Heart* is a thrilling adventure full of hijinks, found family, and romance destined to change the lives of the inhabitants ...  
★★★★☆ Betyg: 4 · 185 recensioner · 11,99 US\$ ⓘ

Science Fiction Bokhandeln  
<https://www.sfbok.se> · Översätt den här sidan :

**Not for the Faint of Heart - Lex Croucher**  
Clem, a backwoods assistant healer known for her new-fangled cures, just wants to help people. When Mariel's ramshackle band kidnap Clem as retribution for her ...  
159,00 kr

Linguee  
<https://www.linguee.com> 36

**not for the faint of heart - Swedish translation**  
Many translated example sentences containing "not for the faint of heart" – Swedish-English dictionary and search engine for Swedish translations.

Bloomsbury Publishing  
<https://www.bloomsbury.com> · Översätt den här sidan :

**Not for the Faint of Heart: from the award-winning author ...**  
15 aug. 2024 — Rip-roaringly romantic, fast-paced and funny - discover the brand new Sapphic historical fantasy adventure from the author of *Cafe Nero* ...  
8,09 GBP

*When describing something difficult, would you use the term "Not for the faint of heart"?*

## ChatGPT:

Yes, "Not for the faint of heart" is a common idiom used to describe something that is particularly challenging, intense, or potentially frightening. It suggests that the situation or activity is not suitable for people who are easily scared or overwhelmed. Here's how you might use it in a sentence:

- **"The hike up the steep mountain trail is not for the faint of heart."**

This idiom effectively conveys that the task at hand requires courage, strength, or resilience.



# Coding with GPT?

Please show me Python code for receiving a patch from the DX7 via Sysex



Here is an example of Python code that demonstrates how to receive a patch from the Yamaha DX7 synthesizer using the system exclusive (SysEx) format:

What happens if GPT suggests open source software snippets for code you are writing for your client?

According to a recent study by Microsoft, 78% of engineers admitted to using own AI tools at work, without telling anyone, out of fear that it wouldn't be allowed...

Not really solving the problem

```
import pygame
import pygame.midi

# Initialize pygame and the MIDI system
pygame.init()
pygame.midi.init()

# Set up the input port for receiving SysEx data
input_id = pygame.midi.get_default_input_id()
input_port = pygame.midi.Input(input_id)

# Wait for a SysEx message from the DX7
while True:
    # Check for incoming MIDI events
    if input_port.poll():
        # Get the next MIDI event
        event = input_port.read(1)[0]

        # Check if the event is a SysEx message
        if event[0][0] == 240:
            # Extract the patch data from the SysEx
            patch_data = event[0][1:]

            # Process the patch data as needed
            # (e.g., parse the data and extract the
            # sound editing parameters)

# Clean up the MIDI system and pygame
input_port.close()
pygame.midi.quit()
pygame.quit()
```

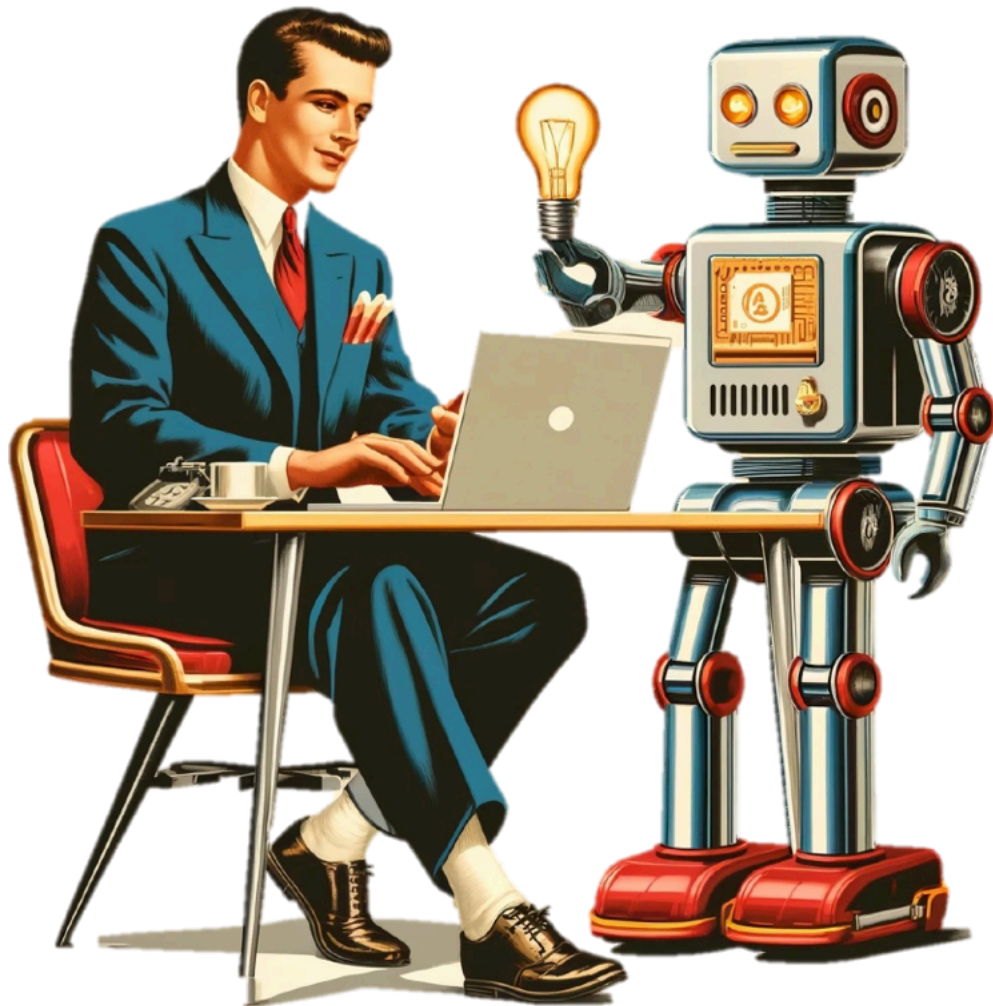
# Bias - what is it and why does it matter?

- Bias is when training data is skewed to contain a subset of examples
  - People of a certain complexion or gender
  - Western style decoration (kitchen app)
- Leads to erroneous or *biased* conclusions
- This is an area of much research - and a potential for experts to guide clients!



Image by MidJourney:  
"Life, Universe and Everything"

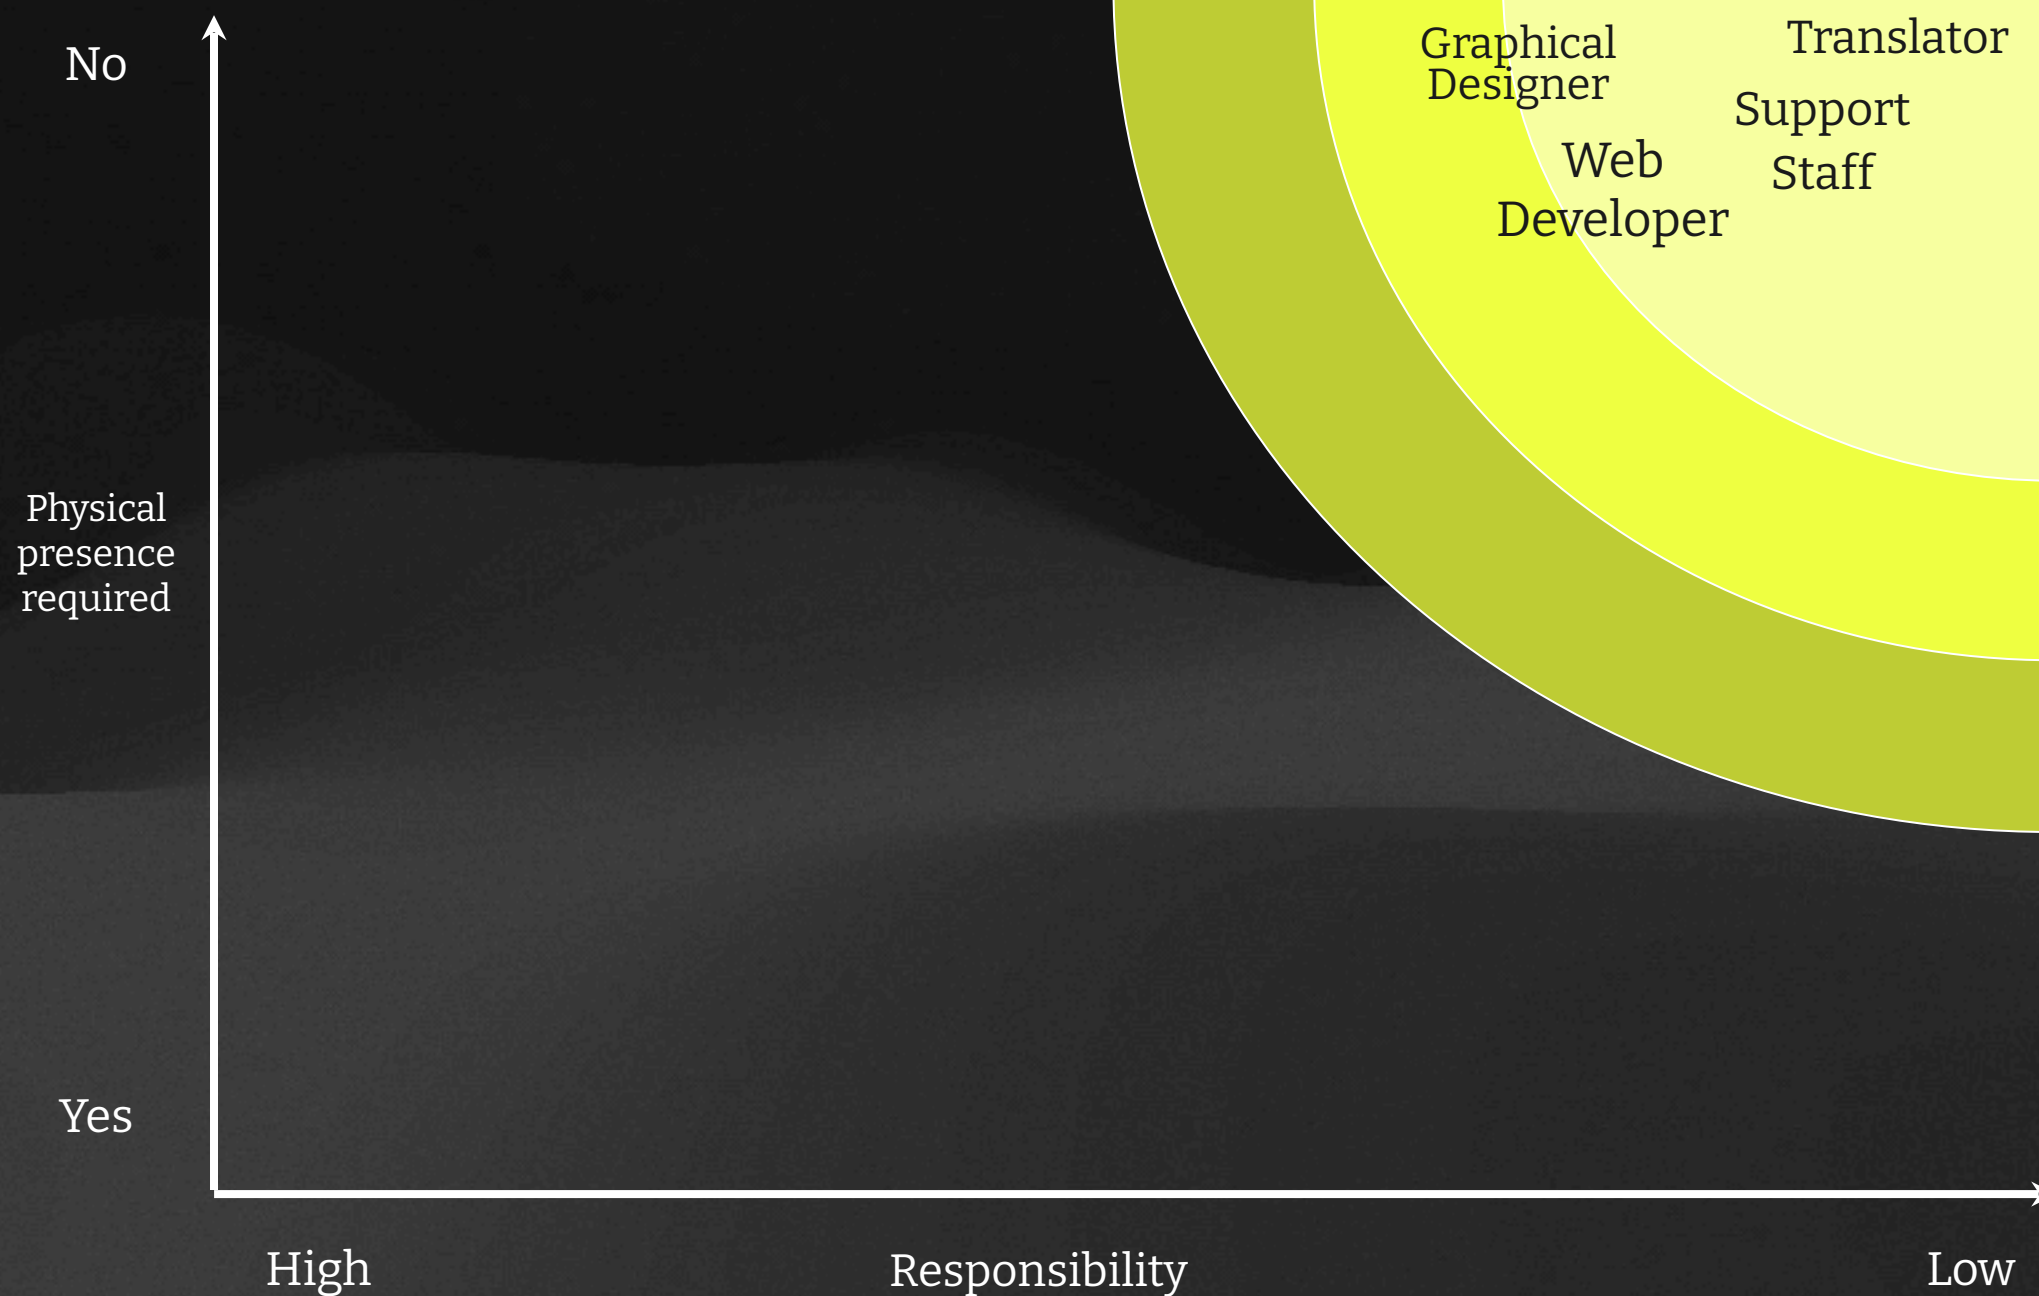
# AI won't replace people at scale, yet...



...but people who refuse to use AI tools might be replaced.

- ChatGPT, Gemini, NotebookLM, DeepSeek
- ChatPDF
- SciSpace
- CoPilot

*Image by Dall-E: "If GPT was invented in the 50's"*





# Benefits and Opportunities

Learning cycles can go much faster!

- (Analog vs. Digital photography - an example)
- Tailored tutoring , pair programming, ...



With proper data, more careful *unbiased* decisions can be made

- Portfolio management
- More robust underlying code and architecture, less focus on basic code errors
- Accelerated research and breakthrough, less double work
- Creative synthesis between different design concepts
- Powerful creative tools to enable the next wave of creative endeavours!
  - Fast and interactive design cycles (in the near future: prototyping even!)

# Thanks for attending!



[mikael@levla.se](mailto:mikael@levla.se)



[www.levla.se](http://www.levla.se)



[StochasticDesign](https://www.youtube.com/StochasticDesign)



[theintrovert](https://www.instagram.com/theintrovert)



[https://www.linkedin.com/  
in/mikaellundgren/](https://www.linkedin.com/in/mikaellundgren/)