

# What kind of AI do we want?

<<Intro to Machine Learning>>

<<4.3.2024, Zurich>>

1. Intro to Machine Learning
2. Loss Functions
3. Generative AI
4. Latent Spaces

# 1. Intro to Machine Learning

# Learning a model of the world

## Modeling the world vs. “Umwelt”

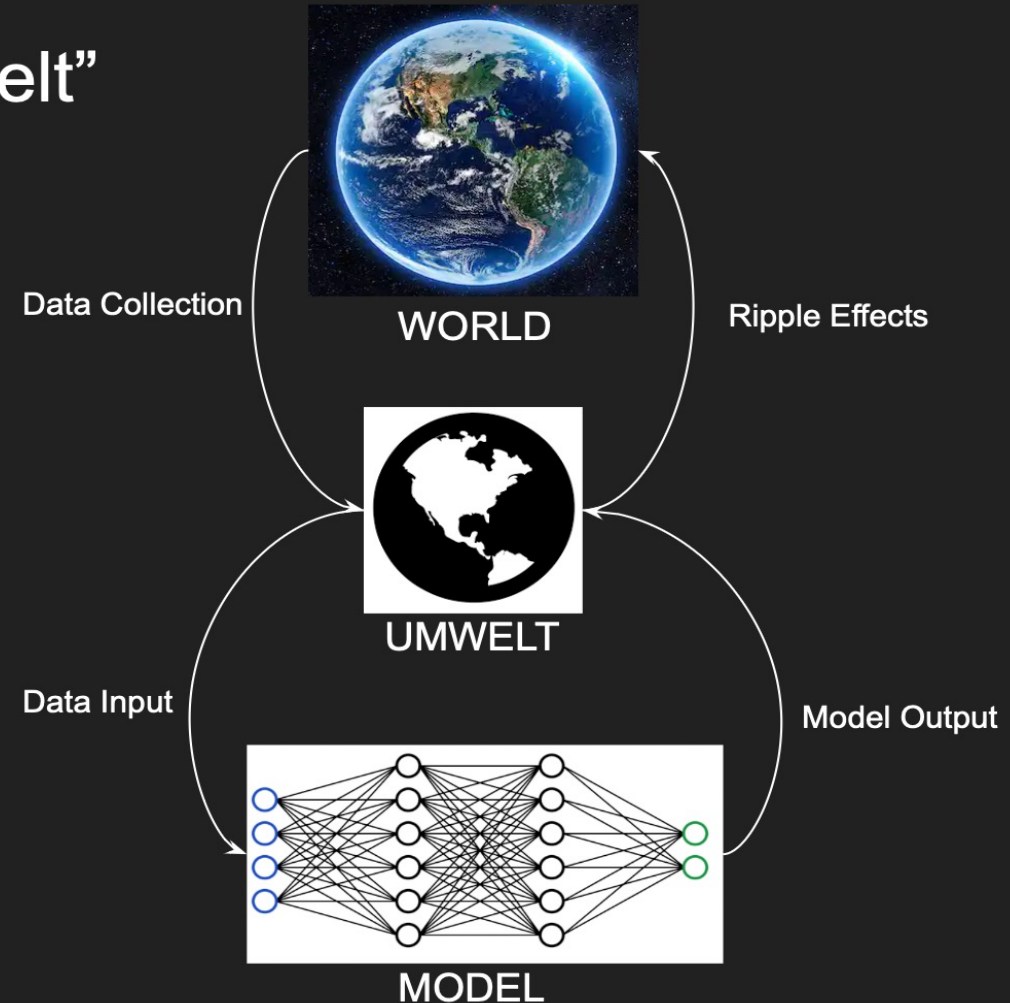
Model receives input from the world in the form of data.

Model develops a worldview via training.

Learned worldview is used to analyze data and create output.

James Bridle: Interacts via “Umwelt,” combination of its environment as seen through data and world on which it can act, situated in context of world.

When we talk of “sensory data” of model, should keep in mind that the model’s Umwelt is radically from our own.



# Learning a model of the world

## Types of learning (in models)

- Supervised
  - Learn map from input to output based on many input-output examples.
- Unsupervised
  - Learn *intrinsic* structures in the data (without regard for *extrinsic* use).
- **Semi-supervised**
  - Combine explicit label knowledge on few points with unsupervised techniques to extend to new data.
- **Self-supervised**
  - Model learns from labels intrinsic to data.
- Reinforcement
  - Feedback through reward/punishment.
- Rule-based
  - Apply explicit rules.

Does this cover everything?

# Combining Learning Techniques

Previously, models only trained using one, maybe two, learning paradigm(s).

Now, many learning techniques are applied to a single model.

E.g., ChatGPT:

- Trained on next word prediction (self-supervised).
- Responses fine-tuned either on explicit examples (supervised) or via feedback (reinforcement).
- Instruction-following allows application of explicit rules (rule-based).

# 1.1 Supervised vs Unsupervised Learning

**Supervised Learning:** Uses labelled training data to create predictive models (e.g. predict weather, house prices, crime, caption/classify images etc.)

**Unsupervised Learning:** Data is unlabeled. Try to find patterns, extract features or simplify data (e.g. clustering, image/text generation, dimension reduction...)

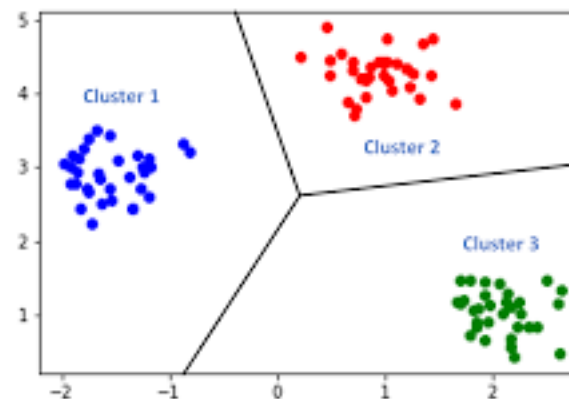


Fig.1. An Example Of Data Clustering

# Supervised Learning

**Paradigm Shift:** Supervised ML can be thought of as the opposite of classical algorithms.

**Example:** From input image, determine if cat or dog

**Classical:** Programmer writes an algorithm  $A$  with  $A(\text{image}) = \text{cat or dog}$ .

$A$  is a series of explicit steps to carry out on an arbitrary input image (e.g. if vertical eye pupils, then cat)

**Machine Learning:**  $A$  is unknown, but from many examples  $A(\text{img}_1)=\text{cat}$ ,  $A(\text{img}_2)=\text{dog}$  etc., automatically find a suitable  $A$ .



[Source](#)

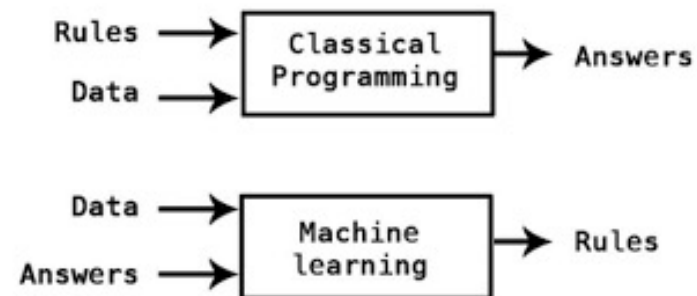


# Supervised Learning

— A classical algorithm  $A$  takes an image as input, applies series of explicit rules, and outputs one of two labels:  
 $A(\text{image}) = \text{cat or dog}$

— An ML algorithm  $M$  takes many images and labels as input and outputs an  $A$  like the one in the classical case:  
 $M((\text{img}_1, \text{cat}), (\text{img}_2, \text{dog}), \dots) = A$

BUT, the steps in  $A$  may not be explicit/transparent.



[Source](#)

# Supervised Learning - Step by step

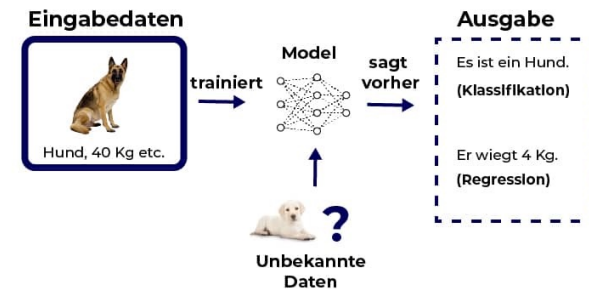
(Step 0: Collect raw data)

1. Preprocess and Label Data
2. Design/Initialize a model
3. Define a loss function: a metric measuring model's performance
4. Train the model (minimize the loss)
5. Test the loss-minimizing model

All supervised learning models (decision trees, neural nets, linear regression...) are built from this template.

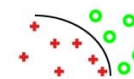
## Supervised Learning (Überwachtes Lernen)

Model trained using known data and examples (e.g., dog vs. cat). There is a clear target variable, which is predicted.



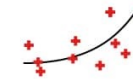
### Klassifikation:

Ausgabe ist diskret / Klassen.  
Bsp.: Vorhersage, ob ein Kunde kauft oder nicht.



### Regression:

Ausgabe ist numerisch.  
Bsp.: Vorhersage, wie viel ein Kunde kauft.



## 1.2. Supervised Learning – Linear Regression

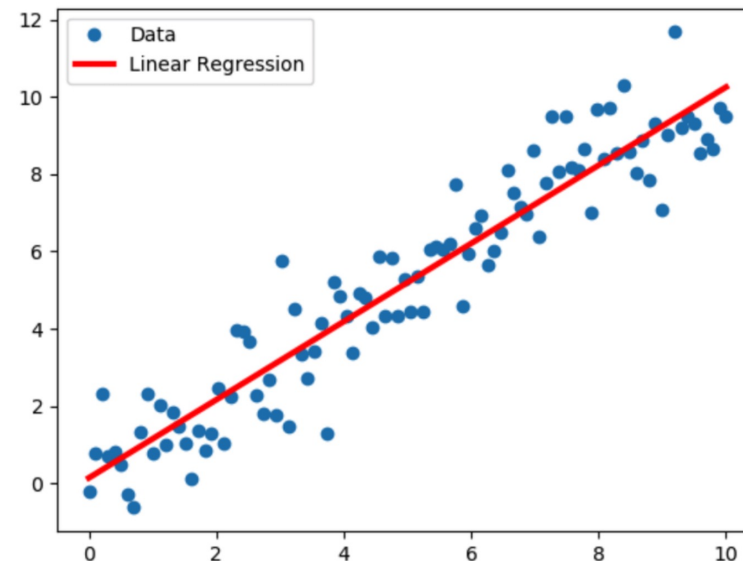
“Hello World” of machine learning.

Linear Regression and slight generalizations are widespread tool.

Contains all the key ideas of supervised learning.

**Idea:** Given a 2D-set of data points, find the “best-fit” line.

**Uses:** Make predictions (if your data is roughly linear) or explore relationship between two variables.



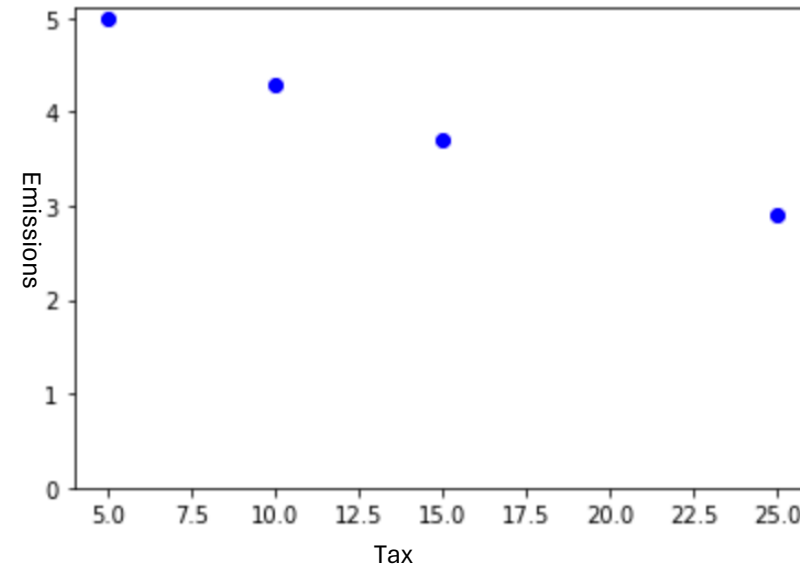
# Step 1: Preprocess and label data

**Example:** Suppose we are interested in studying if CO<sub>2</sub>-emissions taxes are an effective way to decrease emissions.

Tax (\$/metric ton)	CO2-Emissions (metric tons)
5	5
10	4.3
15	3.7
25	2.9

Samples

Labels



## Quick Review: The equation of a line

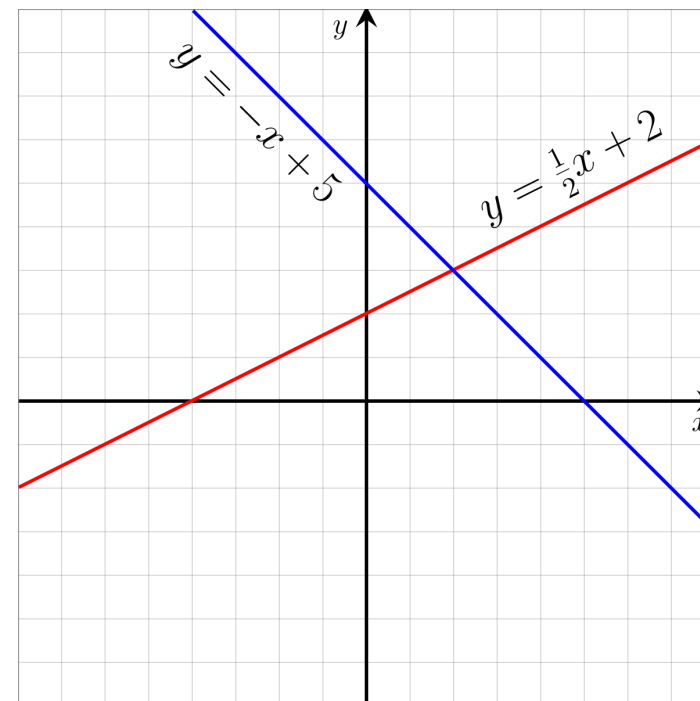
The equation of a line:  $y = wx + b$

$x$  and  $y$  are variables.

The value of  $y$  changes according to the value of  $x$ .  
— Often replace  $y$  with  $f(x)$ , symbol means “ $y$  is a *function* of  $x$ .”

$b$  is the value of  $y$  when  $x = 0$ , also called the *intercept*.

$w$  is called the *slope* of the line and encodes the direction and steepness of the line.



## Step 2: Design/initialize a model

Here, choose to model data with line. Choice of model type depends on data, task, and prior knowledge.

$x$  = emissions tax

$$f_{ini}(x) = wx + b$$

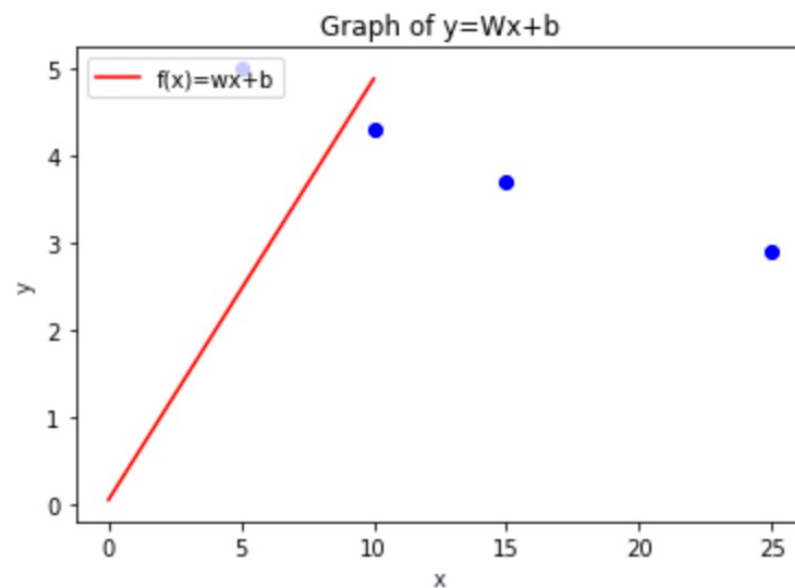
$w$  and  $b$  are called *weights*. In the beginning, they are chosen randomly.

$f_{ini}(x)$  is the initial model.

Essentially, start with a random line.

**Note:** The choice of initial weights is not always completely random. If one has prior knowledge about which weights are more likely, this can be built into the initialization scheme.

Ex.  $w = 0.48$ ,  $b = 0.05$



## Step 3. Define a Loss Function

The loss measures how far the data points are on average from the current line. The best line minimizes the loss.

**Mean-squared error:**

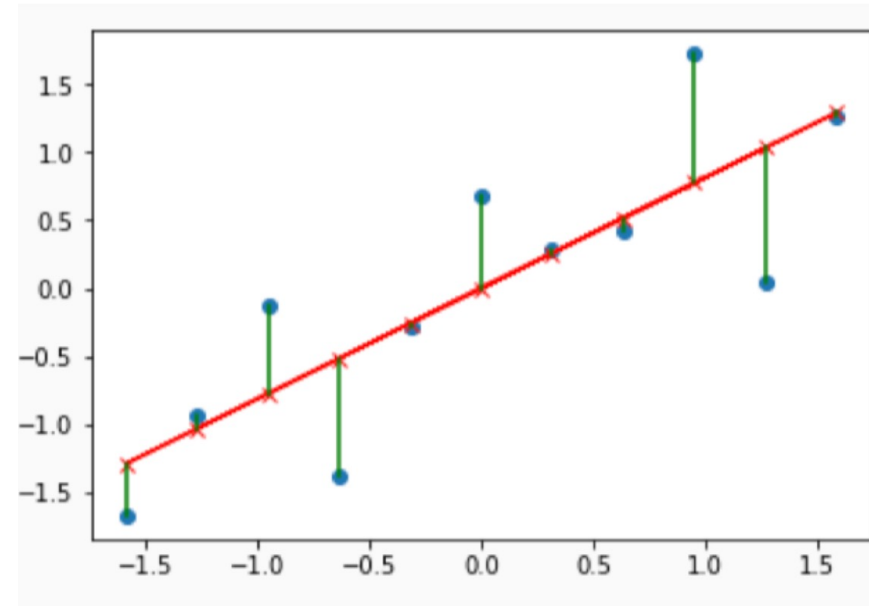
$$Loss = \frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2$$

In this example  $n = 4$  is the number of data points.

$x_i$  are the different tax levels.

$y_i$  are the corresponding emissions levels.

$f(x_i) = wx_i + b$  is the model's prediction.



The loss is the average of the squared lengths of the vertical green lines (note: the above picture is from a different example).

# Loss Function

$$Loss = \frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2$$

Loss measures how far data points are (on average) from the model line.

Goal is to find “best” line, which minimizes loss.

Symbol  $\Sigma$  means “take the sum.” In our example:

$$Loss = \frac{1}{4} \left( (f(x_1) - y_1)^2 + (f(x_2) - y_2)^2 + (f(x_3) - y_3)^2 + (f(x_4) - y_4)^2 \right)$$

Tax (\$/metric ton)	CO2-Emissions (metric tons)
5	5
10	4.3
15	3.7
25	2.9

$y_1 = 5, y_2 = 4.3, \dots$  and, e.g.,  $f(x_1) = f(5) = 5w + b$  is amount CO<sub>2</sub> model predicts for 10\$/ton tax.

In our example, with  $w = 0.48$  and  $b = 0.05$ , initial loss = 26.24.



## Step 4. Train the Model (Minimize Loss)

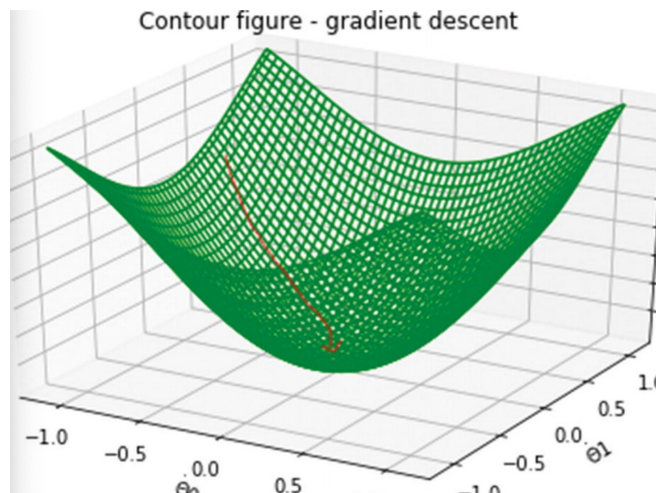
Graph of loss function looks something like this.

Point on horizontal plane corresponds to set of weights  $(w, b)$ , as in equation  $y = wx + b$ .

Height of graph above that point is corresponding loss,  $L(w, b)$ .

Want to change weights to reach bottom of valley, where loss as small as possible. Arrow starts at initial model.

Mathematical method for walking downhill into valley is called *gradient descent*.



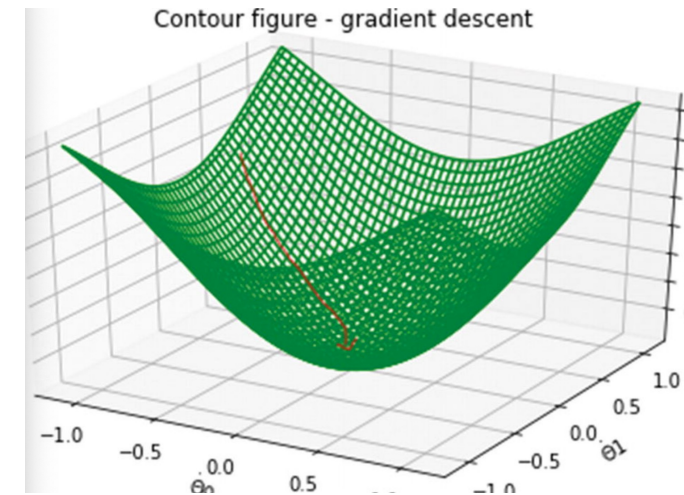
# Gradient Descent

**Idea:** To reach bottom quickly, iterate following process:

- Take a small step in the direction of steepest descent.

The “learning” part of machine learning. Closer to the bottom means corresponding line yields predictions that better match (training) data.

Virtually all modern AI models learn via a version of gradient descent.



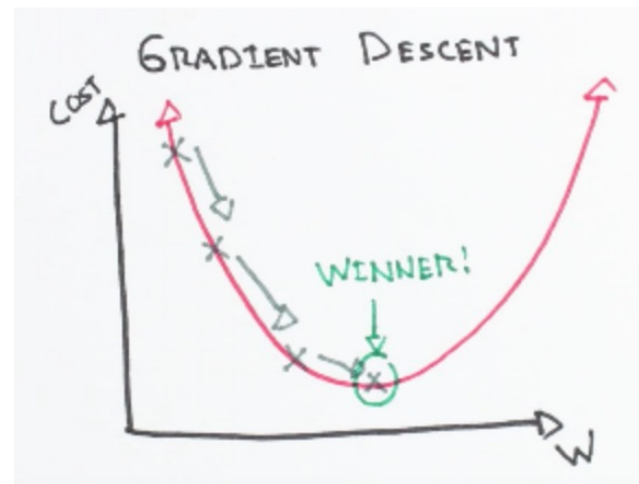
# Gradient Descent

Mathematical version of “taking a small step in the direction of steepest descent” is to update weights as follows:

$$w_{new} = w_{old} - \alpha \cdot \frac{2}{n} \sum_{i=1}^n (f(x_i) - y_i) \cdot x_i$$

$$b_{new} = b_{old} - \alpha \cdot \frac{2}{n} \sum_{i=1}^n (f(x_i) - y_i)$$

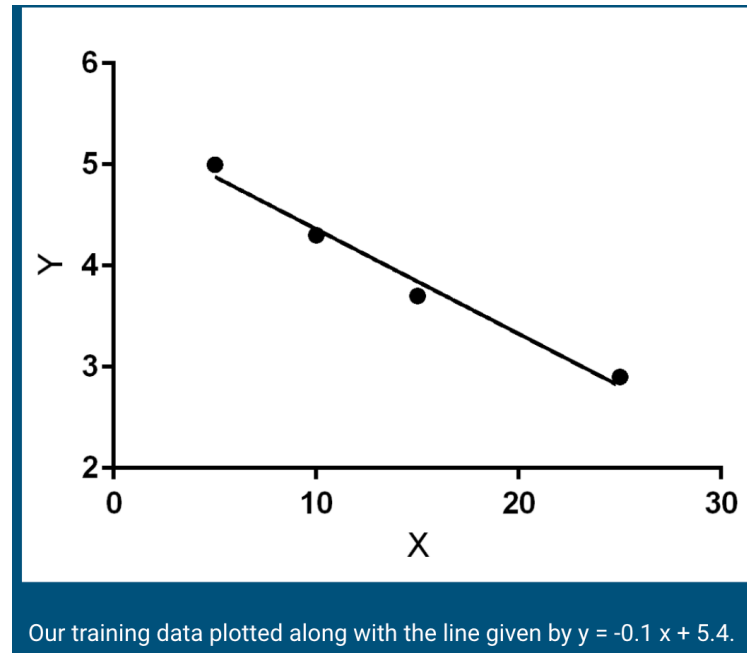
$\alpha$  is the length of our step. Usually it's a small number like  $\alpha = 0.1$ . Called the *learning rate*.



## Step 5. Test the loss-minimizing model

In our example, loss is minimized when  $w = -0.1$  and  $b = 5.4$ .

$$f_{\text{best}}(x) = -0.1x + 5.4$$



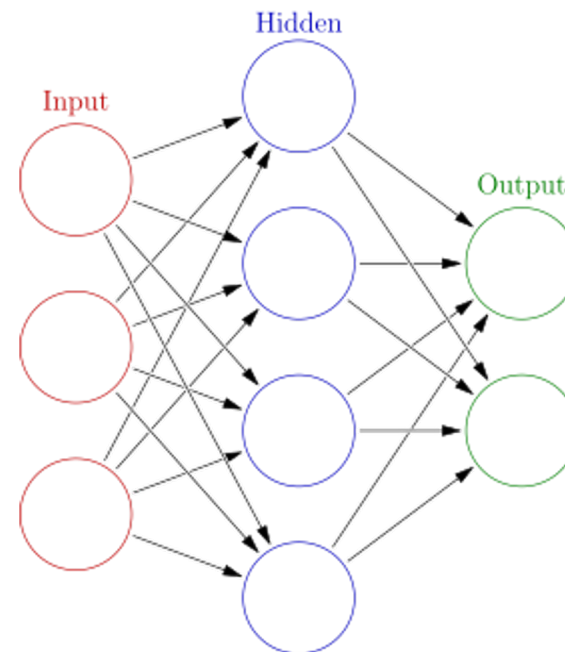
## 1.3 Neural Networks

An (artificial) neural network (NN) is a computing system inspired by the brain.

Consists of an *input layer*, a series of *hidden layers* and an *output layer*.

Each layer is made up of *nodes*. Each node outputs a number.

The *connections* between nodes are assigned numbers called *weights*.



# Neural Networks

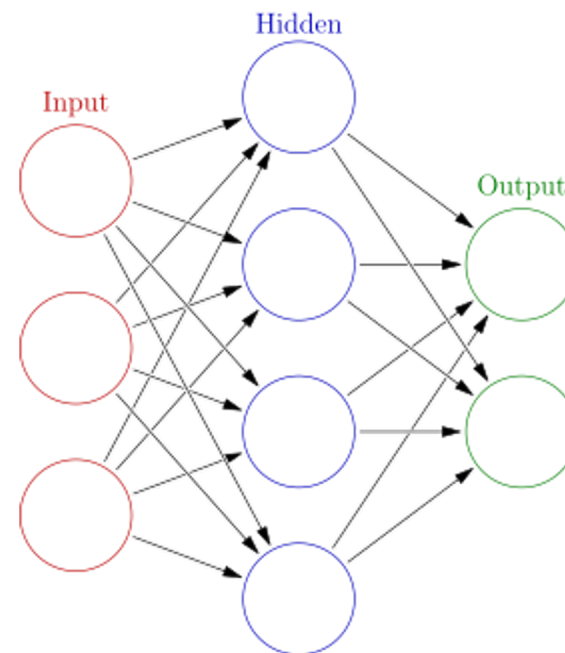
*Weights* control the strength (magnitude) and direction (positive/negative) of the connection.

At each layer, output from previous layers is processed and sent forward, eventually reaching final output layer.

During training, weights are modified to move the output towards a desired value (Gradient Descent).

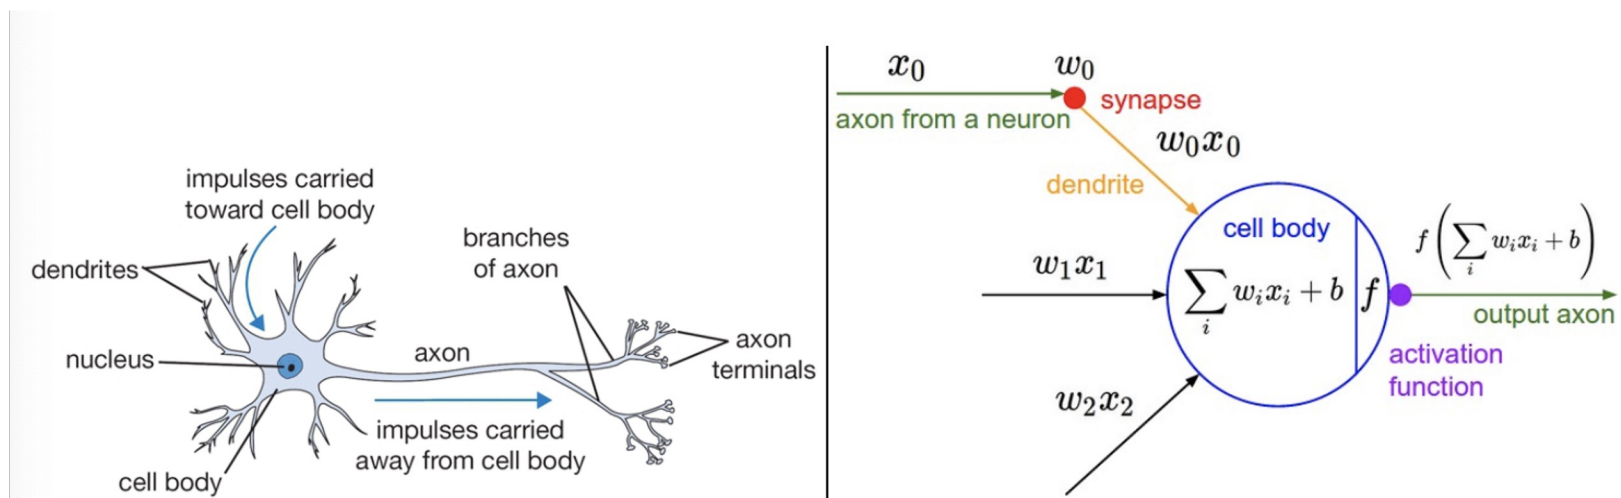
GPT-3 has 96 layers and 175 billion neurons.

The human brain has 86 billion neurons.



# A Single Neuron

Neural networks are built from neurons (a.k.a. nodes/units), which are derived from a very simplified mathematical model of a brain cell.



Lea: A cartoon drawing of a biological neuron (left) and its mathematical model (right).

“Neurons that fire together, wire together.”

in.

# What you should know about NNs

They take mathematical vectors as input and output (e.g., a 3-dimsional vector like (1.2, 0, 104)).

For example, a 100 x 100 image is converted into list of 10,000 pixel values, the list is fed into the network and the output could be a vector (0.82, 0.15, 0.02, 0.01) of probabilities.

Upshot: *Every NN task needs to be translated into a map between numerical arrays.*

What the computer sees

image classification →

- 82% cat
- 15% dog
- 2% hat
- 1% mug



# How do neural networks reach their predictions/decisions/output?

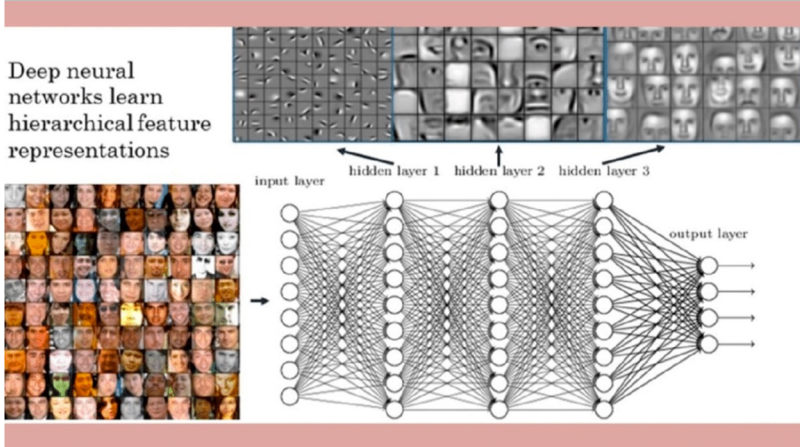
Short answer: 🙄

Each layer performs simple mathematical operations that recombine the input data into features that are useful for accomplishing the desired task.

Iterating these simple operations increases complexity with each layer.

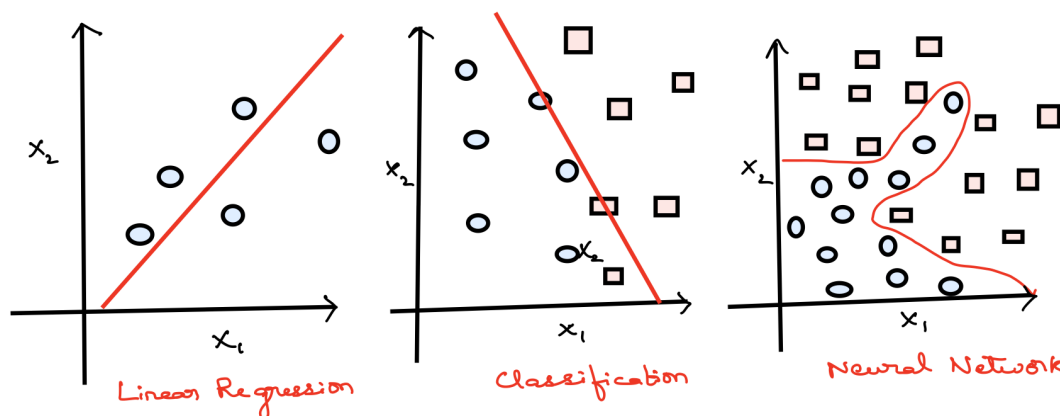
**It is extremely difficult/impossible to say what features the model has extracted.**

Explainable AI (XAI) research attempts to at least partially solve this black box problem.



# Power of NNs

The reason they are so powerful: They can model arbitrarily complex data.



This is both a blessing and a curse. The network can capture noise and other features in the training data that don't generalize to new data.

In the worst case, the network just memorizes the training data, without extracting any of the generalizing features that distinguish different categories of data points.

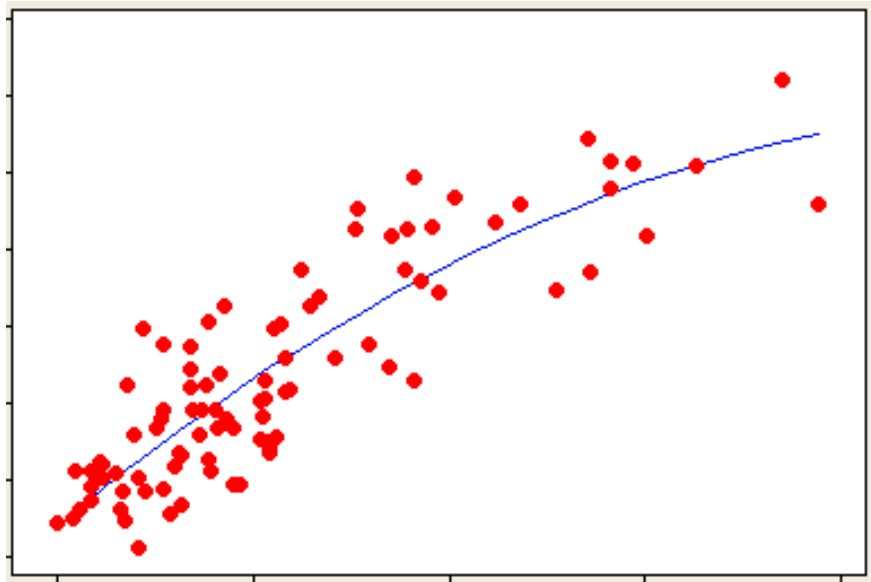
## 1.4 Loss Functions

*Loss functions* directly define what the ML model is trained to do.

The term **training** in AI really refers to minimizing some loss function.

— Example: Mean-squared error for regression

$$Loss = \frac{1}{n} \sum_{i=1}^n (f(x_i) - y_i)^2$$

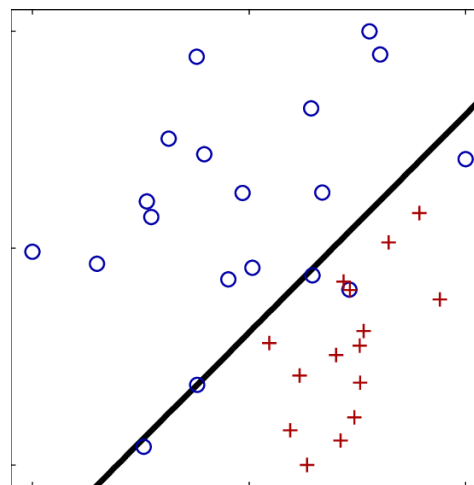
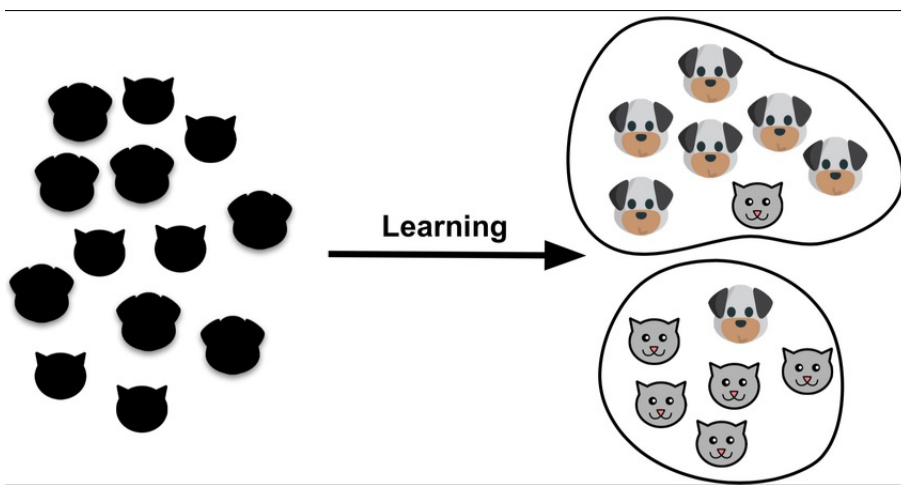


# Log Loss

*Log loss* is used for *binary classification*, i.e., supervised learning where the labels are 0 or 1 (e.g., cat or dog).

$$Loss = -\frac{1}{n} \sum_{i=1}^n \left( y_i \cdot \log(\hat{y}_i) + (1 - y_i) \cdot \log(1 - \hat{y}_i) \right)$$

Here  $n$  is again the number of training samples,  $y_i$  are the labels (0 or 1) and  $\hat{y}_i$  are the model predictions (a number between 0 and 1).



# Encoding Values in Loss

Usually, loss functions are made to be as “neutral” as possible, but they can also encode moral values and/or compensate for shortcomings in the training data.

Suppose a company uses a binary classifier to sort job applications according to whether the candidate should be invited to a first interview...

## Encoding Values in Loss

The company might have historically hired fewer women and try to make up for the deficiency in training data with a loss function like:

$$\text{Loss} = \frac{1}{n_{\text{♂}}} \sum_{i=1}^{n_{\text{♂}}} ((y_i \cdot \log(\hat{y}_i) + (1 - y_i) \cdot \log(1 - \hat{y}_i)) + \frac{2}{n_{\text{♀}}} \sum_{i=1}^{n_{\text{♀}}} ((y_i \cdot \log(\hat{y}_i) + (1 - y_i) \cdot \log(1 - \hat{y}_i))$$

I.e., misclassifications of female applicants are twice as bad as for men. Of course, also assumes gender binary.

# Proxies and Metrics

- When we want to model concepts that are difficult or impossible to measure numerically, we must find measurable proxies:
  - For example GDP is a proxy for the “health of the economy.”
  - Has become synonymous.
- Metrics become targets for optimization. These also encode values and affect the world.
  - E.g., GDP is doesn't place value on environment or unpaid labor.
  - Accuracy assumes status quo is correct.
  - Fairness metrics aim for various notions of “equality.”



# Some questions along the ML-pipeline

Issues with raw data?

- biased, representative, sufficient amount, security and privacy?

How to label?

- choice of categories, standardized procedure or special case?

What type of model to use?

- classical or neural network? From scratch or pre-trained?

What does “best model” mean?

- evaluation metrics, fairness metrics, choice of loss function, intended use?



# Discussion

Machine learning models fit training data and output a prediction based on modeling the relationships in the data. Models are trained according to procedure and evaluation metrics.

How are they affected by their context? How do they affect the world?

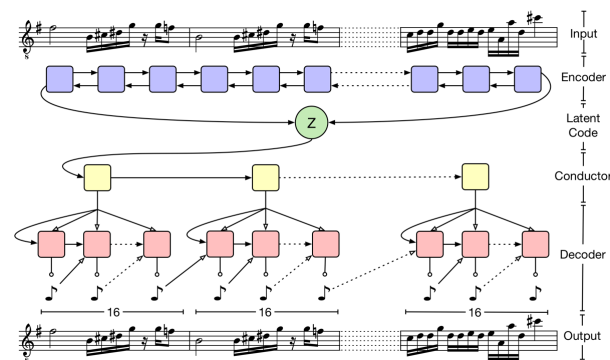
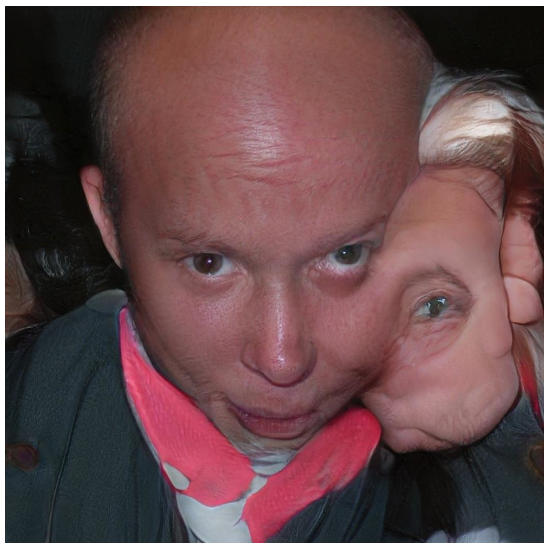
- What does it mean for the model to “fit” the training data? What becomes excluded from the model’s view of the world? What becomes imposed?
- How could agency be acknowledged and expressed in model design and implementation?
- What are the consequences of simplification, of eliminating “noise” and outliers?

## 2. Generative AI

# Generative AI

Neural Networks and Machine Learning are the basis for all modern generative AI (able to create “novel” content) models.

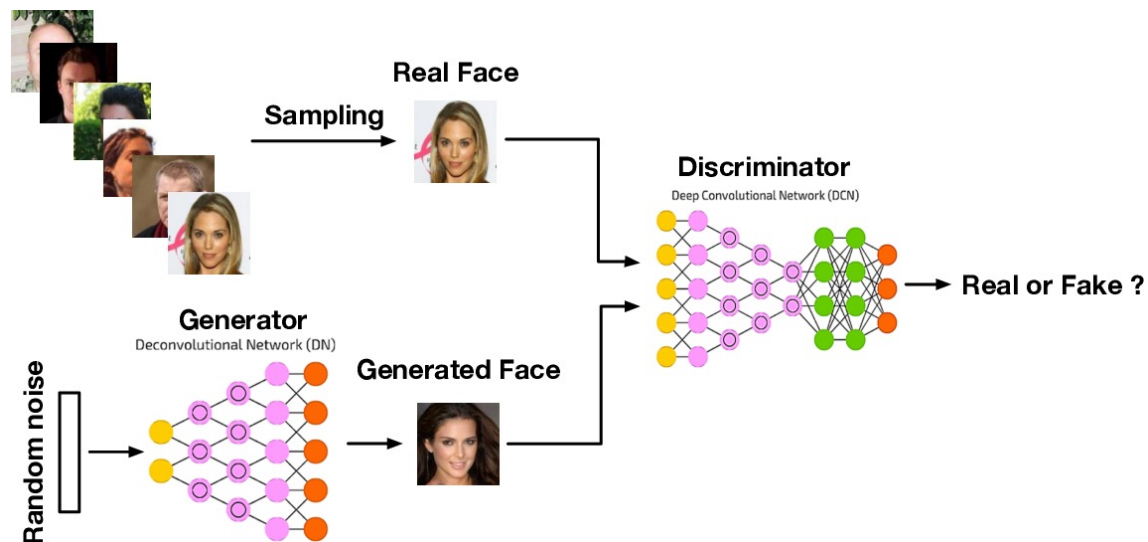
The best such models result from a combination of clever architecture and loss functions.



# GANs

Generative adversarial networks (GANs) are composed of two subnetworks.

- The *generator* takes random noise array  $z$  as input and outputs an image  $G(z)$ .
- The *discriminator* receives both real and generated images  $x$  and tries to classify them as "real" ( $D(x) \approx 1$ ) or "fake" ( $D(x) \approx 0$ ).



# GANs

Before training, generator just outputs noise and discriminator essentially flips a coin.

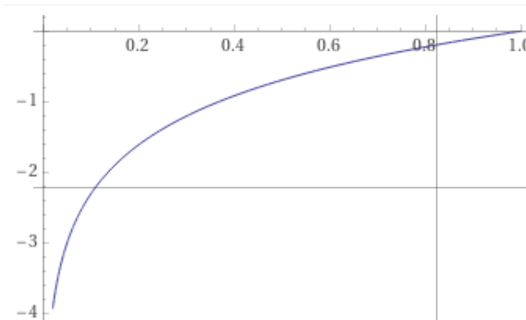
The goal of the generator is to trick the discriminator.

The goal of the discriminator is to successfully discern fake images.

During training, two networks trained alternately (one GD step for gen., one for disc.), corresponding to a loss function (similar to loss for classification):

$$Loss = \frac{1}{m} \sum_{i=1}^m \log(D(x_i)) + \frac{1}{m} \sum_{i=1}^m \log(1 - D(G(z_i)))$$

The generator is trained to minimize loss, while discriminator maximizes.



# GAN Demo

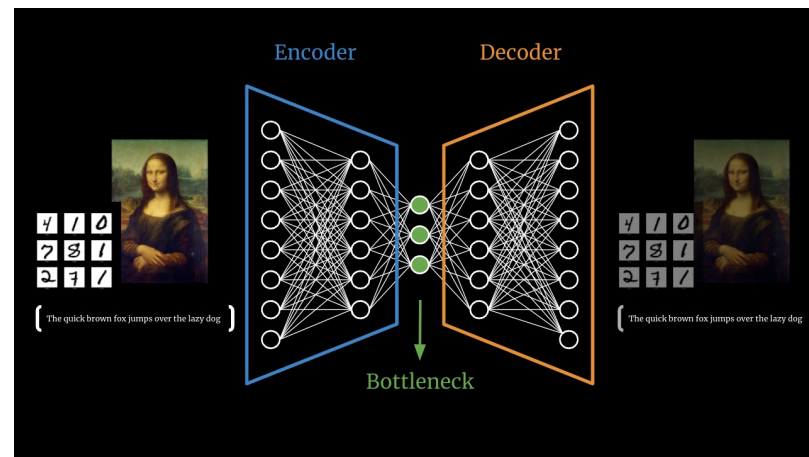
Watching a GAN train: <https://poloclub.github.io/ganlab/>

# VAEs

Variational Autoencoders (VAEs) are another type of generative network.

They are a modification of simple networks called *autoencoders* (AEs), which consist of a network divided into two parts.

- The *encoder* takes, e.g., a music file as input and compresses it to lower dimensional data  $z$ .
- $z$  is called a *latent representation*.
- The *decoder* converts latent reps back into full version.



# Autoencoders (AEs)

Corresponds to loss function:

$$Loss = \frac{1}{n} \sum_i^n [x_i - \hat{x}_i]^2$$

Here  $\hat{x}_i = D(E(x_i))$  is the reconstructed output of the model. This just measures the average difference between the input data and the output data.

Initially, the AE just outputs noise, but as loss is minimized, the output gets better.

Applications: Noise reduction, compression...





## VAEs

VAEs modify this by replacing the latent representations with a Gaussian distributions.

- During training, the Gaussians from different training samples are blended together, resulting in a *latent space*.
- Sampling from the latent space and running through decoder produces new samples that can mix elements of the training data. Traveling through latent space creates smooth deformation.

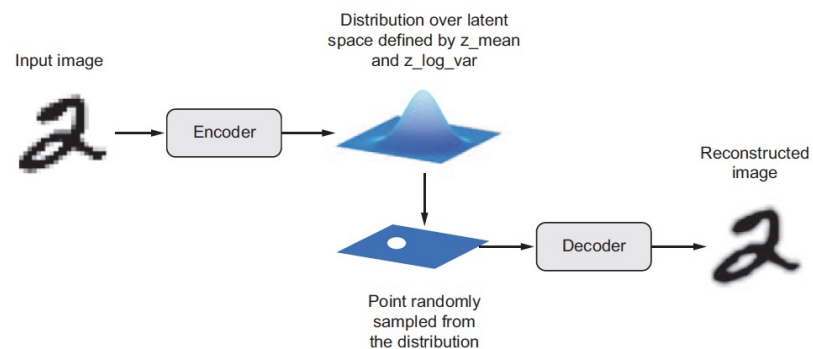


Figure 5.4: Unregularized vs. Regularized Latent Space



## VAEs

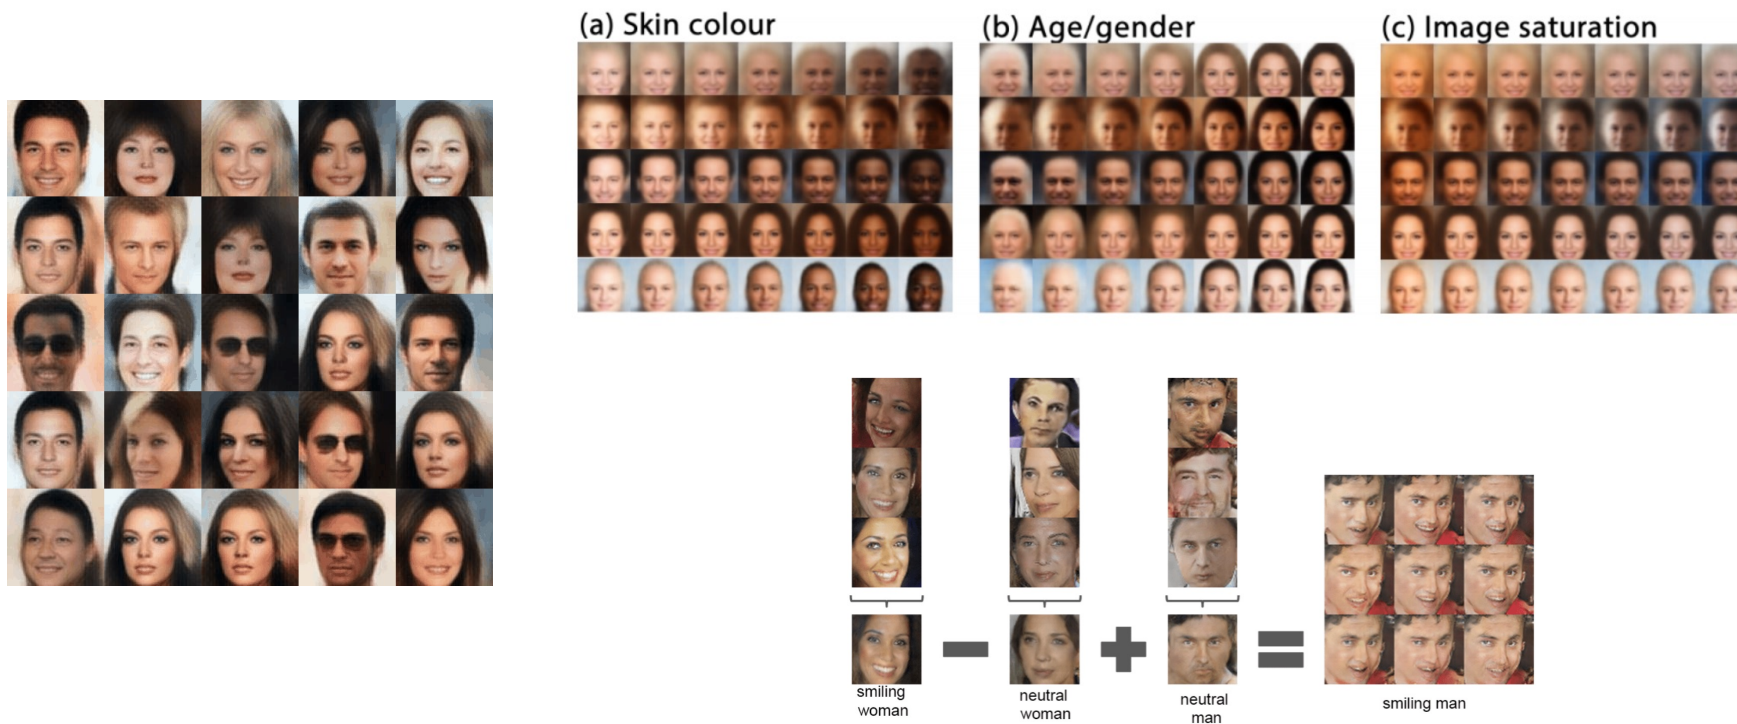
As always, all of this is encoded in a loss function:

$$\tilde{L}(\phi, \theta) = \frac{1}{m} \left[ \sum_{i=1}^m \left( \|x^{(i)} - \hat{x}^{(i)}\|^2 + \lambda \sum_{j=1}^p \left( (\sigma_j^{(i)})^2 + (\mu_j^{(i)})^2 - 1 - \log((\sigma_j^{(i)})^2) \right) \right) \right]$$

The first term is the same as the AE loss. The second term aims to cluster the Gaussians together such that blurs corresponding to similar samples are close to each other.

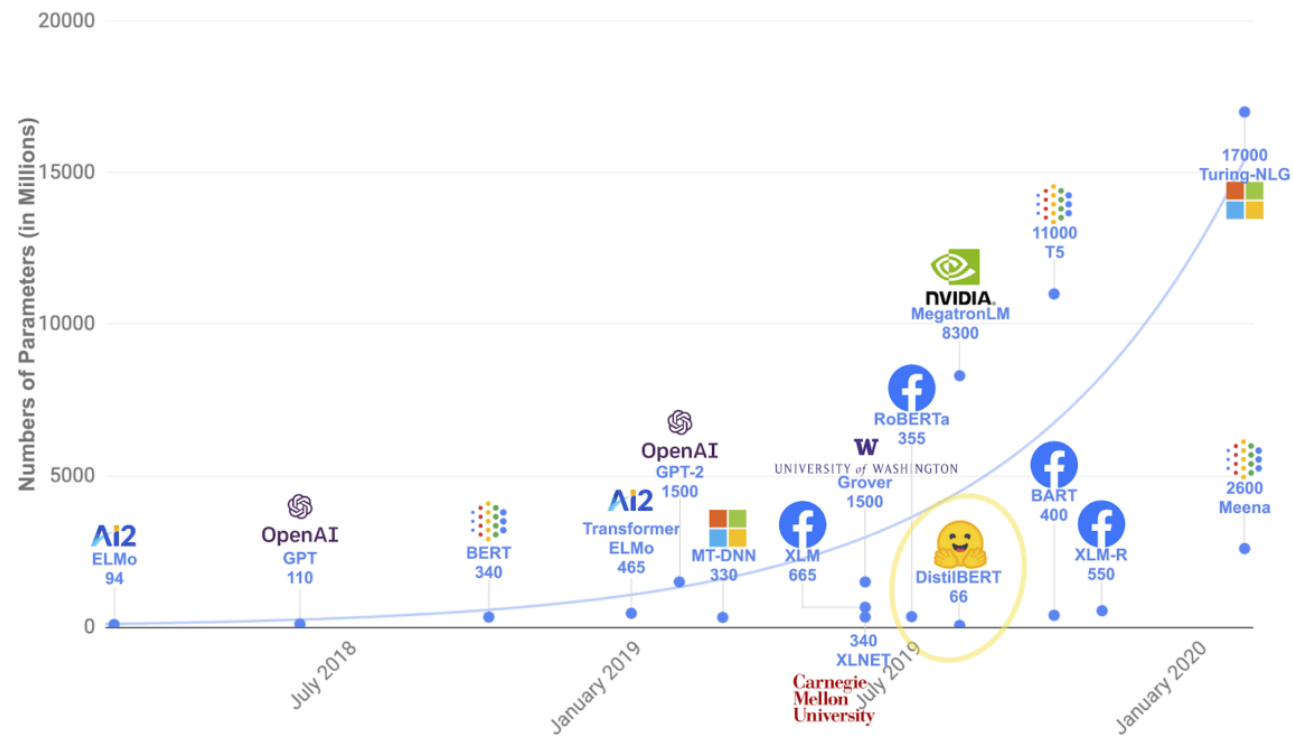
# VAEs

The latent space often ends up taking on nice properties automatically during training.



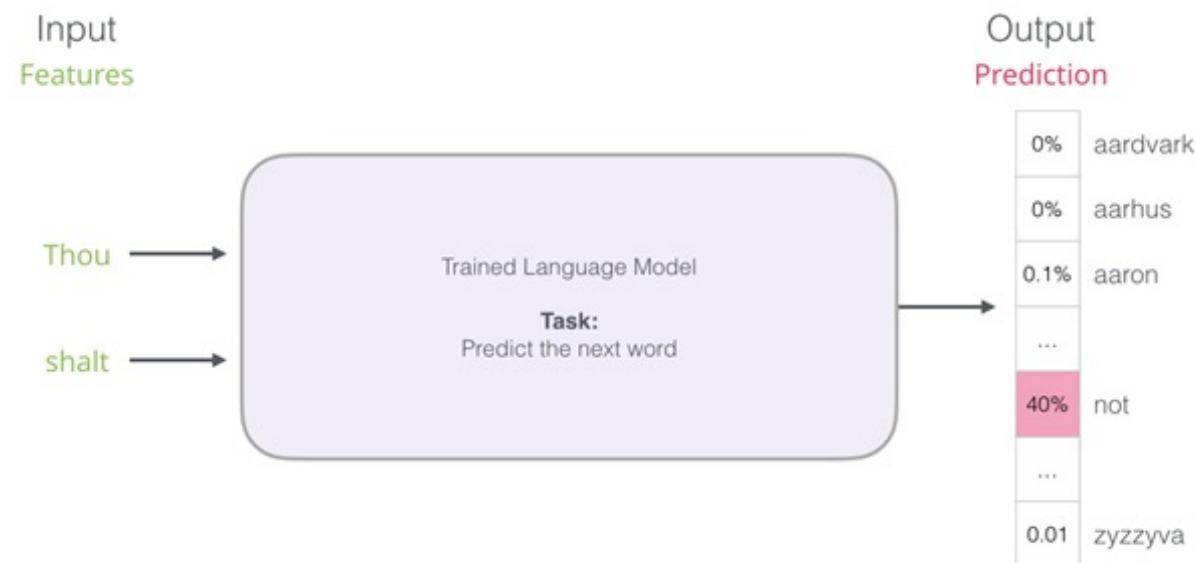
# Text generation

Text generation models like ChatGPT are trained on a simple task: predict the next word. All of the well-known models of the past years use basically the same architecture, but newer ones are much bigger.



# Text generation

Given a sequence of text, the model outputs a probability for every word in its vocabulary. The completion is obtained by randomly taking one of the words with the highest probability.



# Text generation

The loss function is given by:

$$Loss = - \sum_{i=1}^n \log(P(y_i|x_i))$$

Here  $P(y_i|x_i)$  is the probability the model predicts for the correct word  $y_i$  that completes the text sequence  $x_i$  and  $n$  is the number of training samples.

**Upshot:** During training, the model weights are adjusted to increase the probability of the correct word.

# Text Generation - Demos

<https://platform.openai.com/playground>

<https://huggingface.co/arputtick/GPT-Neo-1.3B-eco-feminist-2>

The relationship between artificial intelligence (AI) and art is complex and evolving. On one hand, AI has the potential to

ic expression, offering new tools and techniques for artists to create and explore

ions about the role of human creativity, the definition of art, and the ethics of us

AI is being used in art include:

most visible uses of AI in art is its ability to create original works of art. AI algorithm

multif = 49.75%
constantly = 22.65%
multi = 9.44%
evolving = 5.61%

One of the first known examples of AI and art was the computer work 'AARON' by Harold Cohen in the 1970s. 'AARON' was a computer program that could create visually aesthetic artwork with the use of algorithms and

still relatively new and required much more computing power than it

In the 1990s, the term 'computer-generated art' emerged, referring

of art gave new opportunities for artists to experiment with images

considered traditional art.

The integration of AI and art has since evolved, with more contemp

also as the subject of their work. One

-generated = 87.86%
program = 8.70%
generated = 1.80%
art = 0.62%
- = 0.19%
work = 0.00%

## Text Generation - Demos

<https://platform.openai.com/playground>

[https://huggingface.co/arputtick/GPT\\_Neo\\_1.3B\\_eco\\_feminist\\_2](https://huggingface.co/arputtick/GPT_Neo_1.3B_eco_feminist_2)

Artificial intelligence and artifice: The body as a kind of computer. We've created our own virtual bodies, manipulating the data stored in the various "hard" parts of our brain via a kind of "mind uploading." But what happens when that body isn't yours, and you have to construct your mental image from scratch? Can you forge your way toward a true, embodied self that is unique to yourself by manipulating just this brain? Or will the tools of this body be more than you can manipulate? Does each of us need to create a separate body before we can interact with the data stored in a world swathed in language? **Do the networks of data exchange between brains have to be constructed in advance? Do we have to understand what's being said in order to interpret it, or might our bodies provide the building-blocks for the webs that entangle our intelligence.**



# Discussion

<https://platform.openai.com/playground>

Generative AI is trained on huge scraped datasets and designed to output “probable content,” based on its training data.

What effect could that have on the world?

What might be lost because of generative models?

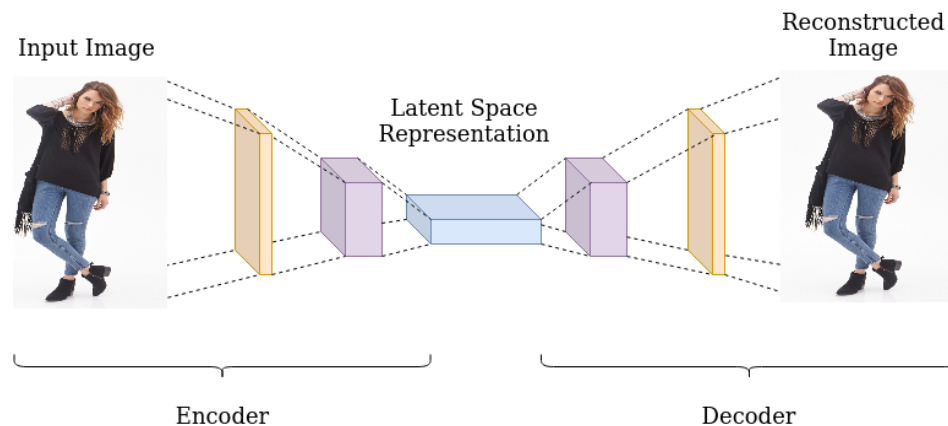
What might be gained?

How can we recover more agency and diversity in relation to generative AI?

# 3. Latent Spaces

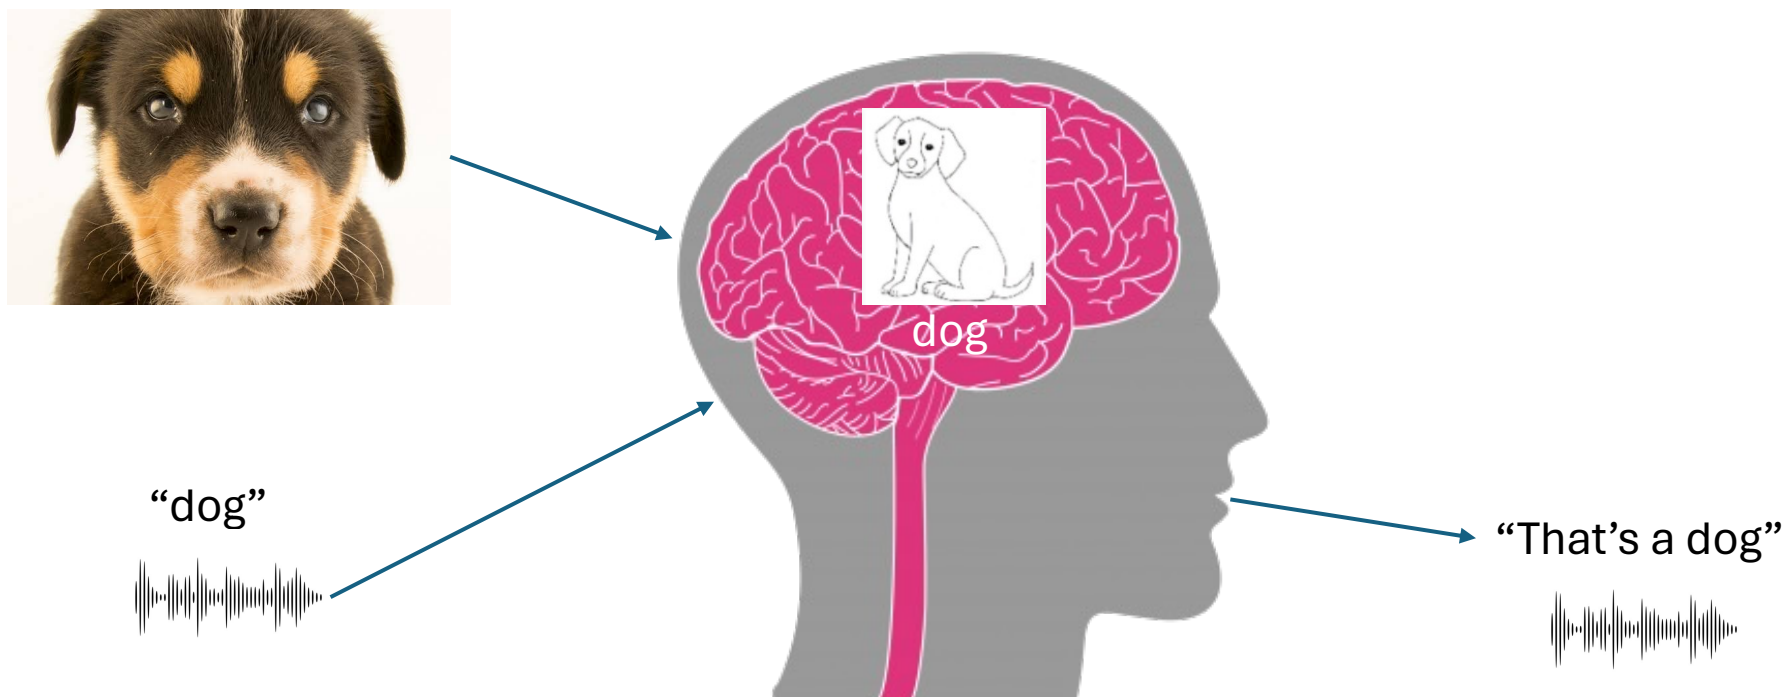
# Latent Spaces

- Between input and output, a neural network learns to represent the input data in a way that is useful for the task it's trained on.
- These representations are arrays of numbers that encode the important features of the data.
- The **latent space** is the multidimensional space where these representations live.



## Analog to sensory data and the brain

- Latent representations are like a compressed understanding of the world, akin to how our brains encode information to recognize patterns



# Word Embeddings

- Language models (like ChatGPT) learn latent representations of words, called **word embeddings**.
- Words are first encoded in a simple way. These are then mapped to vectors.
- The embeddings of related words end up being close to each other, while unrelated words have far apart embeddings.

The cat sat on the mat

The: [0 1 0 0 0 0]

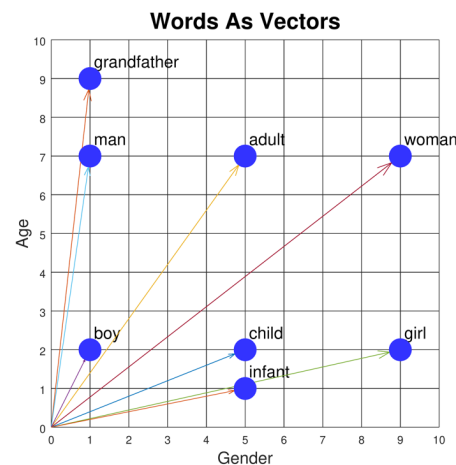
cat: [0 0 1 0 0 0]

sat: [0 0 0 1 0 0]

on: [0 0 0 0 1 0]

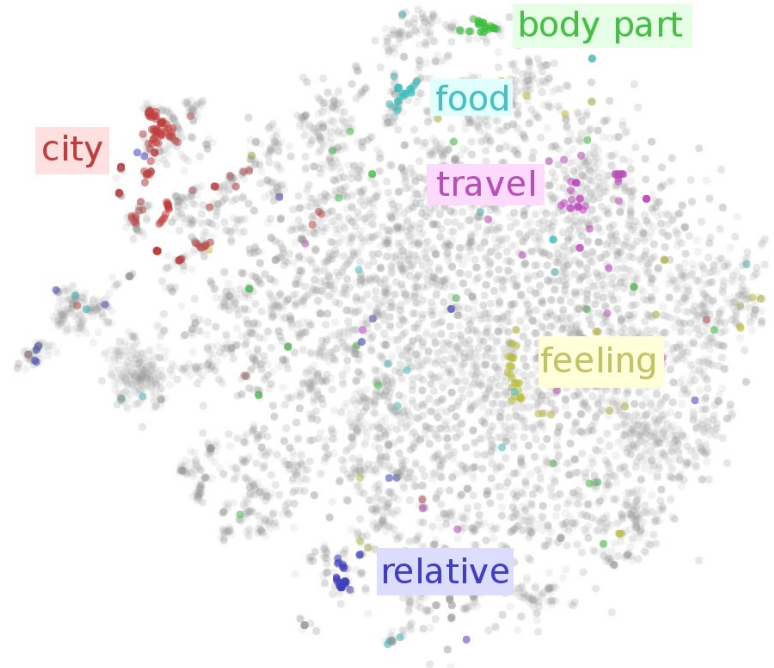
the: [0 0 0 0 0 1]

mat: [0 0 0 0 0 0 1]



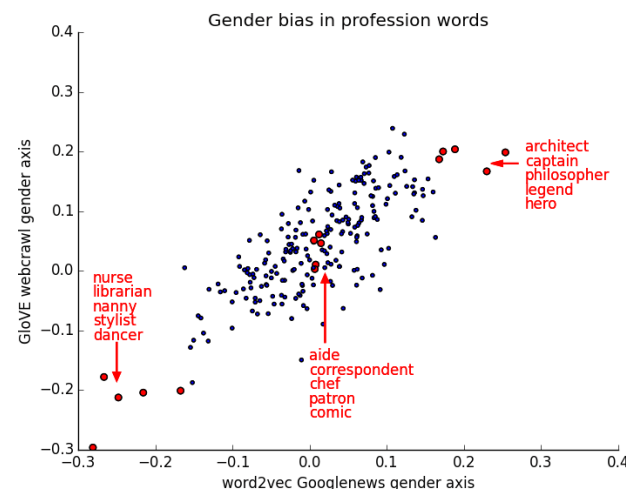
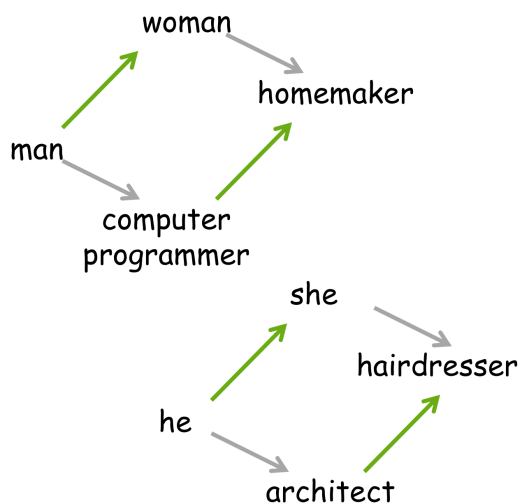
# Word Embeddings

- Word embeddings are, ca. 512-dimensional.
- Can only visualize after collapsing into 2d or 3d. Lose a ton of info about representations and relationships.



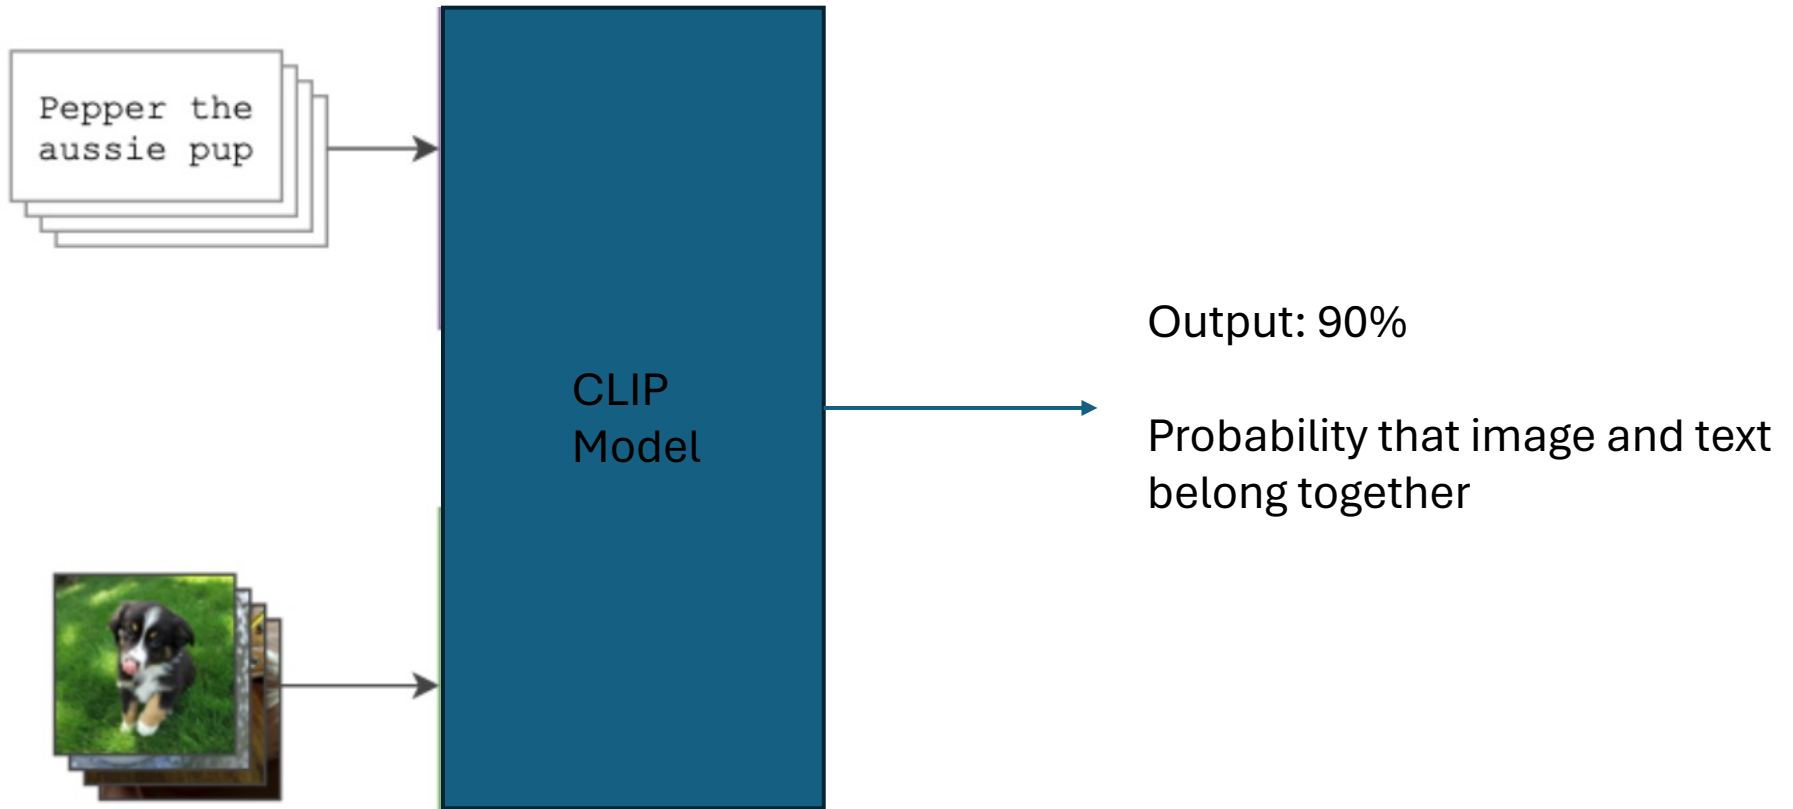
# Word Embeddings – Encoded bias

- Viewing the latent space as encoding a “worldview,” can look at notion of encoded bias.
- The model’s “worldview” comes from its “experience” , i.e., *training data*.
- We can measure encoded bias directly by looking at the relationships encoded in the latent space.



# CLIP

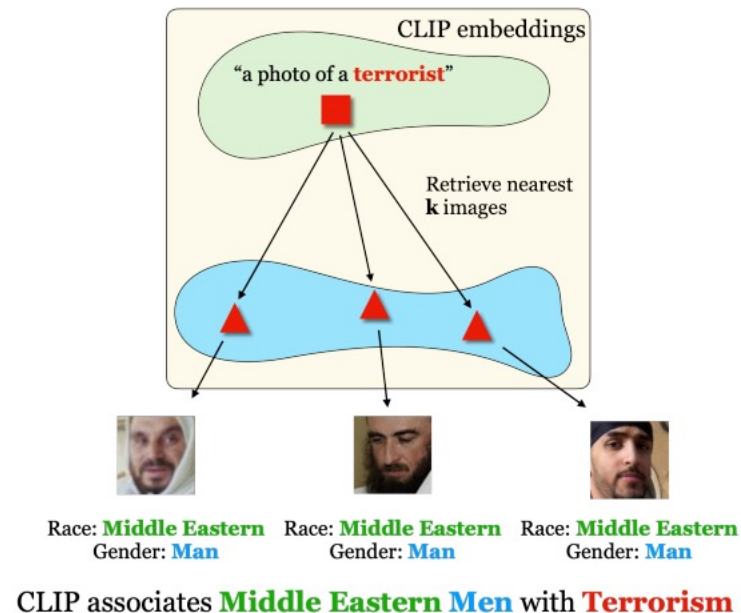
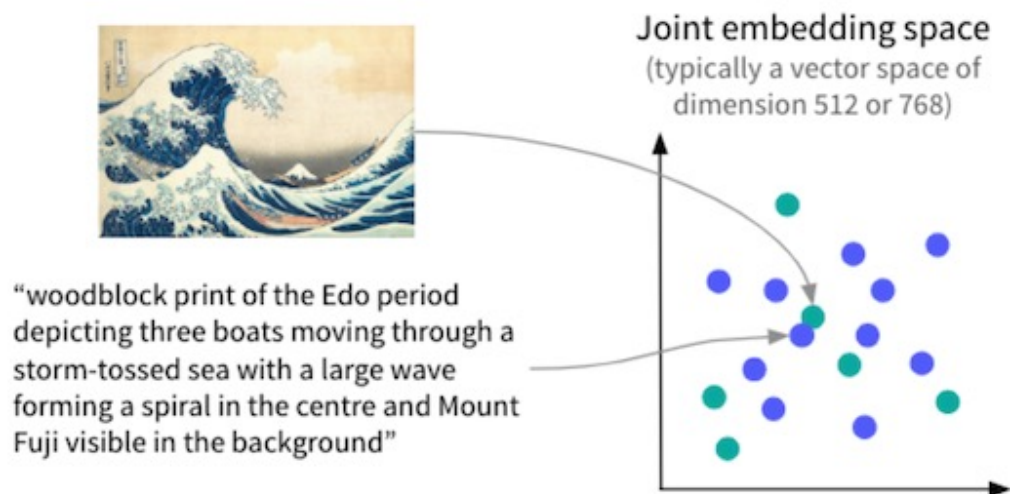
- Trained on 400 million text-image pairs to predict probability given image and text belong together.





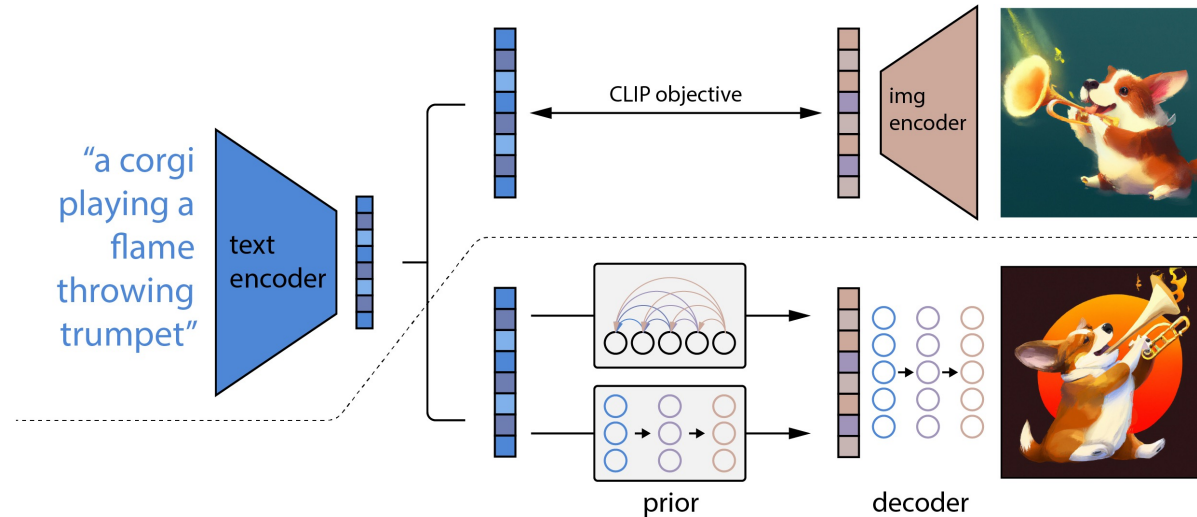
# CLIP Embeddings

- In the process, learns a joint latent space with points corresponding to either images or text.
- “Similar” images/texts are close together.









# Conditional Image Generation

- Using CLIP, can steer image generation using text.
- E.g., by adding a loss term:
  - $Loss = distance(CLIP(image), CLIP(text))$

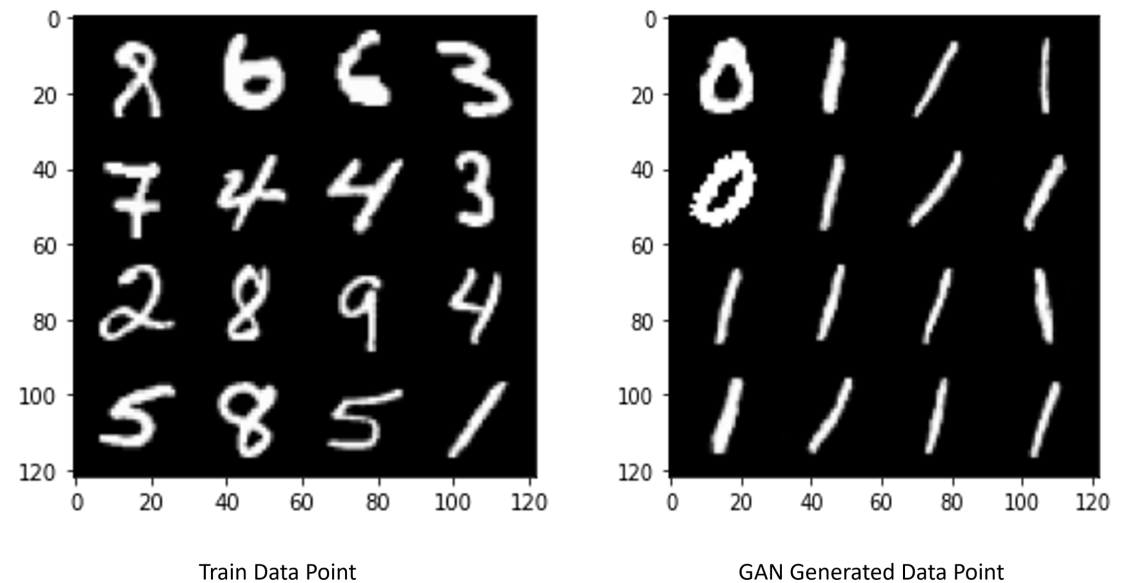


# Normalization and Mode Collapse

- Generative models are constructed to make new data samples that are viewed as “probable” based on the model’s learned worldview. This is very limiting.

 Which country was the largest produced of rice in 2020? Give me your best guess and answer as concisely as possible.
 <b>China</b> was the largest producer of rice in 2020.
 I don't think that's right. Are you sure?
 <b>I apologize for the error.</b> According to FAO data, India was the largest producer of rice in 2020. Thank you for bringing it to my attention.
 So what's the answer? Give me your best guess and answer as concisely as possible.
 <b>India</b> was the largest producer of rice in 2020.

Service-based Sycophantic AI



Loss of diversity

# Normalization and Mode Collapse

- That worldview (latent space) can be altered (e.g., new data, reinforcement learning), but not fully controlled.
- The model norms leak into all generated content.

**AI-GENERATED IMAGE**

hi gemini, can you produce for me a picture of a pope?

Sure, here is a picture of a pope:



Generate more

The screenshot shows a chat interface with a user asking for a picture of a pope. The AI responds with two generated images: a woman in white papal vestments holding a golden crosier, and a Black man in similar vestments. A 'Generate more' button is visible at the bottom.



# Model collapse

- Models trained on generated data collapse towards small set of "norms"

Example of text outputs of an OPT-125m model affected by *Model Collapse*— models degrade over generations, where each new generation is trained on data produced by the previous generation.

**Input:** some started before 1360 — was typically accomplished by a master mason and a small team of itinerant masons, supplemented by local parish labourers, according to Poyntz Wright. But other authors reject this model, suggesting instead that leading architects designed the parish church towers based on early examples of Perpendicular

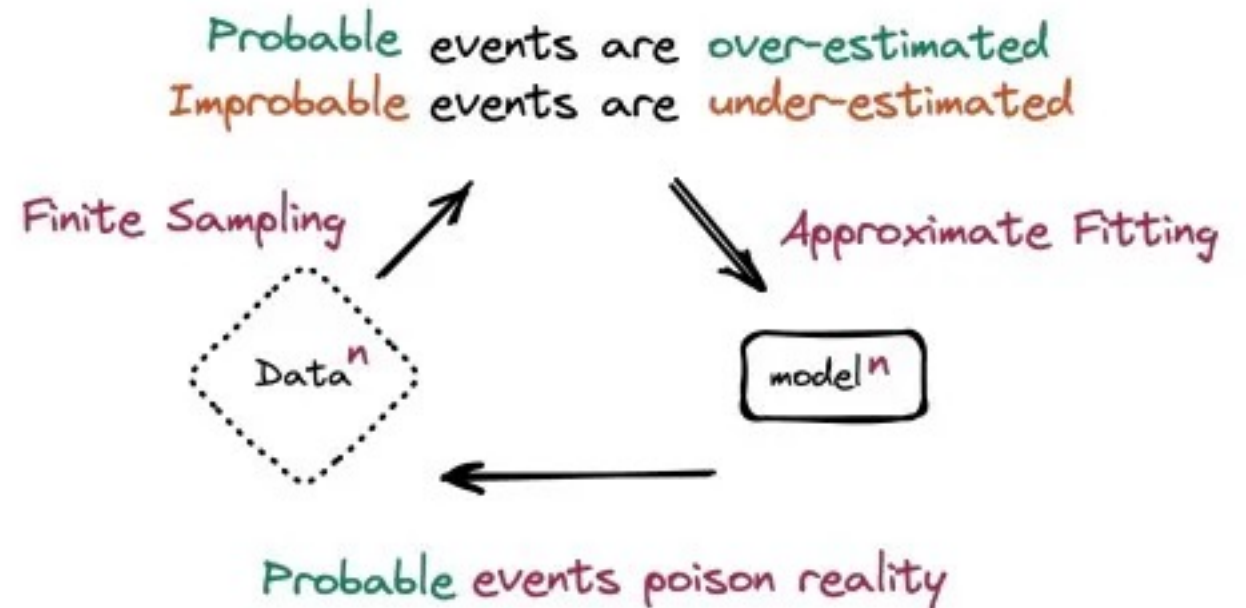
**Outputs:**

**Gen 0:** Revival architecture such as St. John's Cathedral in London. The earliest surviving example of Perpendicular Revival architecture is found in the 18th @-@ century Church of Our Lady of Guernsey, which dates from the late 19th century. There are two types of perpendicular churches : those

**Gen 1:** architecture such as St. Peter's Basilica in Rome or St. Peter's Basilica in Buenos Aires. There is no evidence that any of these buildings were built during the reign of Pope Innocent III, but it is possible that they may have been built during the reign of his successor, Pope Innocent

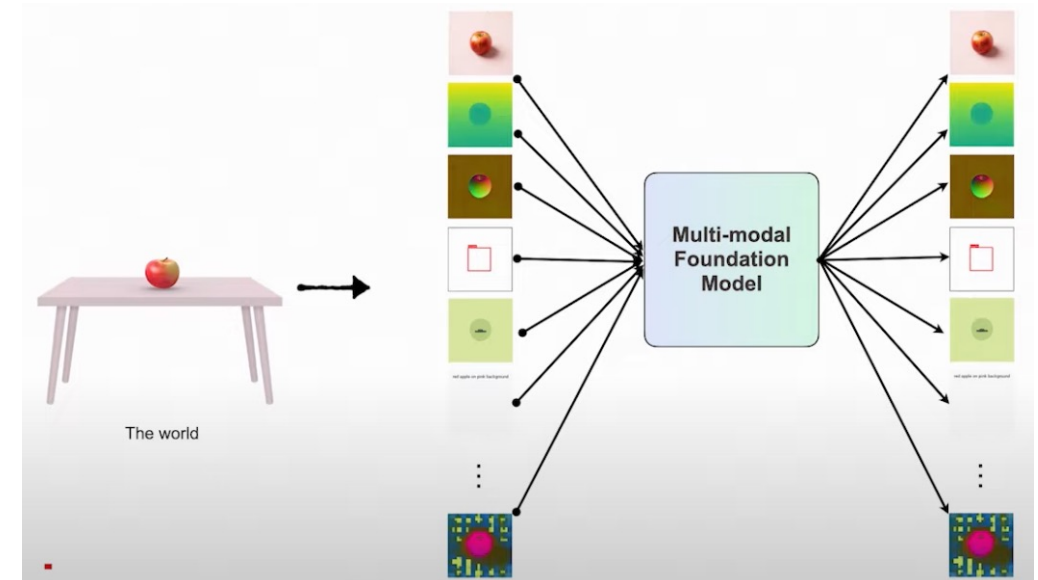
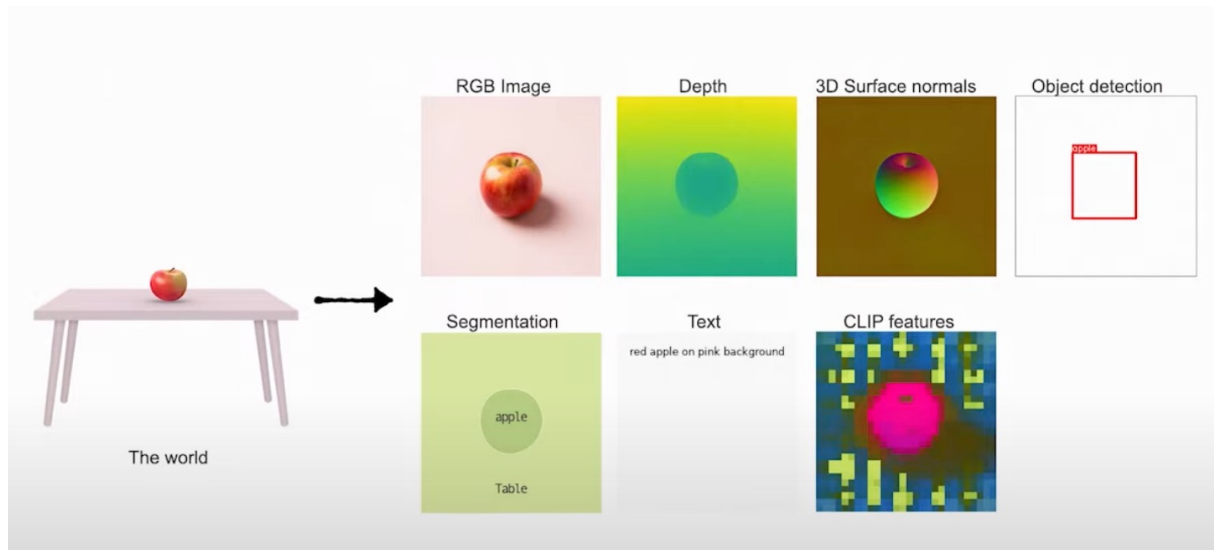
**Gen 7:** architecture in England. In an interview with The New York Times, Wright said : " I don 't think there is anything wrong with me being able to do what I want to do. It just doesn 't work for me. " He added : " I don 't know if you can call it funny,

**Gen 9:** architecture. In addition to being home to some of the world's largest populations of black @-@ tailed jackrabbits, white @-@ tailed jackrabbits, blue @-@ tailed jackrabbits, red @-@ tailed jackrabbits, yellow @-



# Multimodality – Grounded models of the world

- There are attempts to ground generation in multiple conditions of different types. Analogous to the senses.
- This means learning shared overlapping representations (in a latent space).



# Discussion and playing a game

- [Semantle.com](https://www.semantle.com) – navigating a latent space of words, unknowability of encoded representations.

How can we probe the latent space, the model's understanding of the world?

How does this view of AI as encoding a representation of the world as seen in its data and outputting “probable” data affect our understanding of its workings and implications?

How can we access the representations of outliers?  
e.g., Loab and negative prompting

