UNFOLDING THE FOLDING ALGORITHM

A BRIEF INTRODUCTION TO ARTIFICIAL INTELLIGENCE AND [some of] THE MATHEMATICS BEHIND ALPHAFOLD

Kimberly Mandery
Computer Scientist, St. Olaf College

Thursday, July 24 2025

GOALS FOR TODAY'S TALK

- Learn unifying terminology
- Spark interdisciplinary discussions
- Decode deep learning architecture

DEFINITIONS

AI VS ML VS DL VS GENAI

MODEL = DATA + ALGORITHM

WORKFLOWS

PRE-TRAINED MODELS

HISTORY

CAN COMPUTERS THINK?

WHY THE RISE NOW?

COMPETITIONS AS BREAKTHROUGHS

ALPHAFOLD

ATTENTION

TRIANGLE INEQUALITY

DIFFUSION MODEL

NEXT STEPS

ETHICS

HYPE VS REALITY

RESOURCES

WHAT IS ARTIFICIAL INTELLIGENCE?

AI are systems that mimic human intelligence

NARROW AI is very good at a specific task

GENERAL AI is a sci-fi dream

ML is learning patterns from data

Guessing until it gets it right

DL uses neural network architecture

Layering abstractions until result is somewhat recognizable

GEN AI finds most probable item

Piecing together something that is locally likely but can be globally disjoint

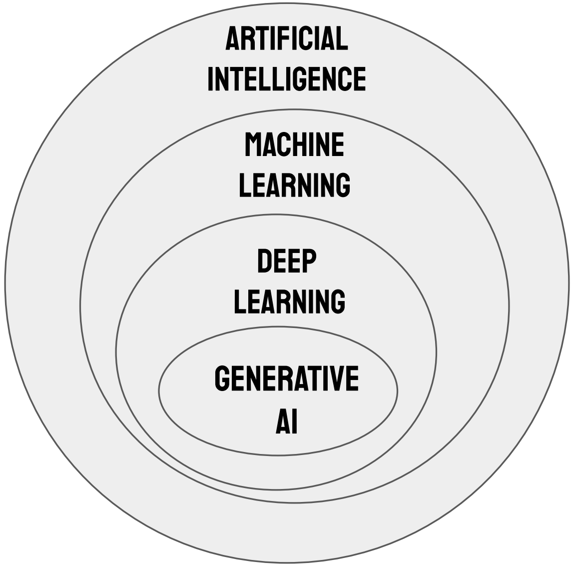

The Illusion of Learning: https://ml-site.cdn-apple.com/papers/the-illusion-of-thinking.pdf (do LLM and LRM reason?)

EXAMPLES

Netflix recs	early chatbots
AlphaFold	wall clock

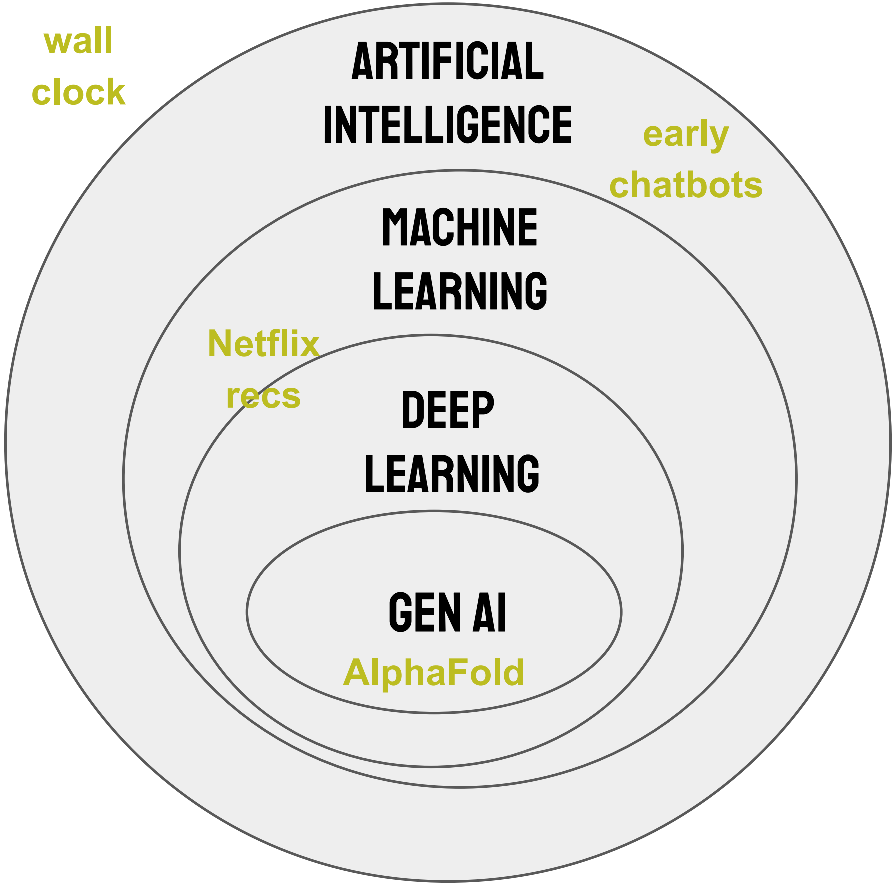

DATA

shapes what the model learns

FEATURES

ML: numeric, categorical
DL: text, images

TARGET

what you want to predict

+

LEARNING ALGORITHM

defines the learning process

adjusts **parameters** (the knobs it turns to improve)
to optimize a **cost function** (overall error including all examples)
which comes from individual **losses** (error for a single example)
guided by **hyperparameters** (settings chosen beforehand)

randomly guess → adjust guess to optimize cost function→ repeat

=

ML MODEL

mathematical function approximating reality

LEARNING ALGORITHM

what is this
even doing?

Give linear regression data and it will
randomly guess → adjust guess to optimize cost function→ repeat

using **gradient descent**

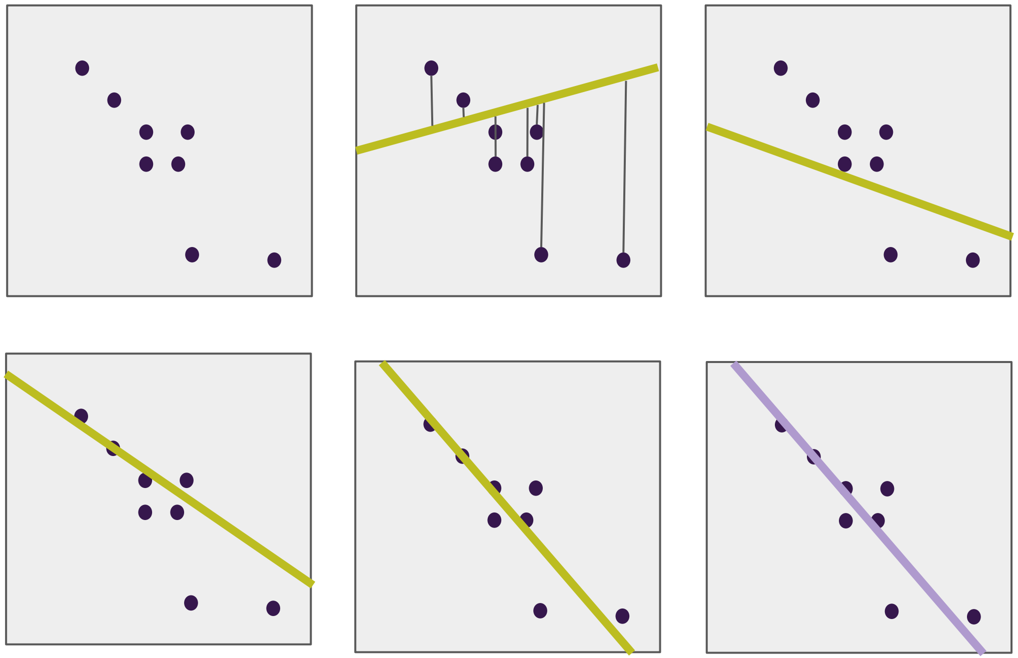

converges

How do we find an equation of a line?

LEARNING ALGORITHM

what are they able to do?

Give [an algorithm] data and it will

randomly guess → adjust guess to optimize cost function→ repeat

in order to : [do things]

linear regression: fit a line to it

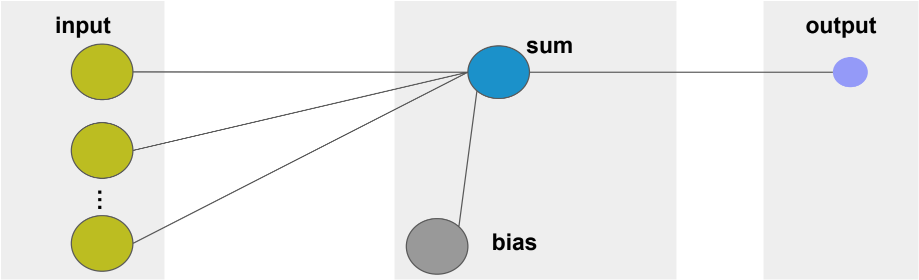

logistic regression: output a probability (0-100%)

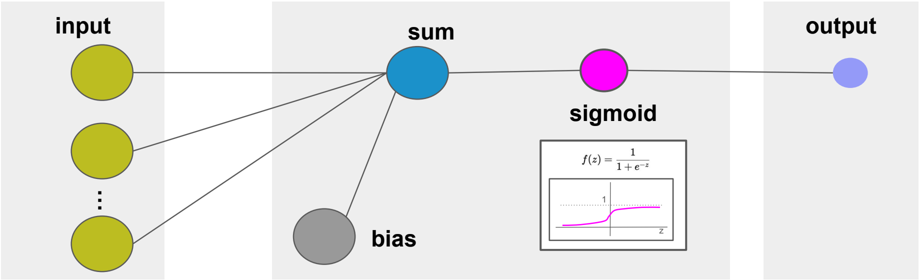

Give [an algorithm] data and it will

randomly guess → adjust guess to optimize cost function→ repeat

in order to : [do things]

LEARNING ALGORITHM

what are they able to do?

perceptron : can do yes or no decision

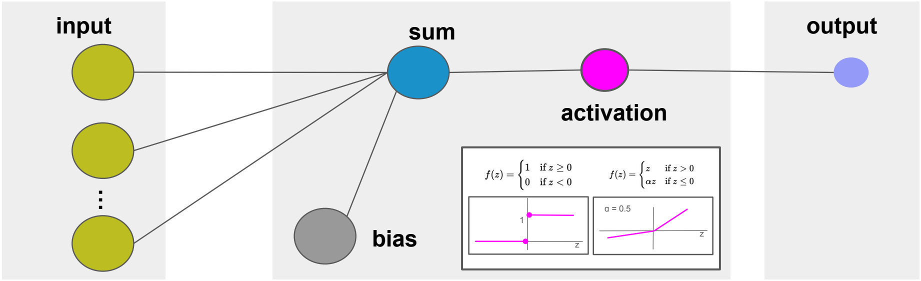

neural network : stack perceptrons in a network

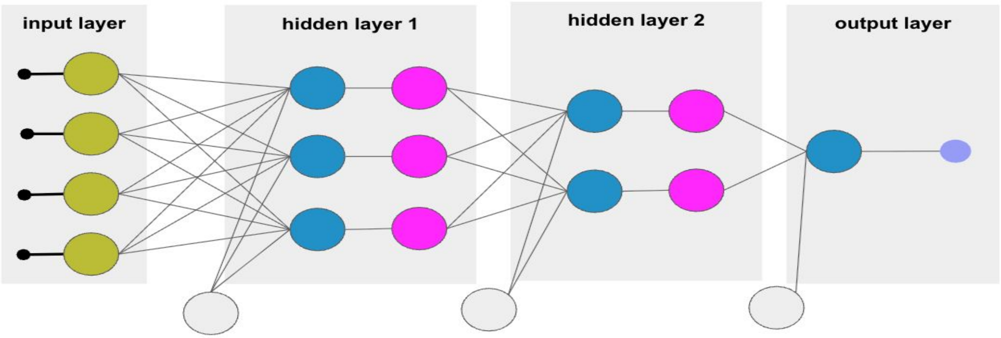

DATA

shapes what the model learns

FEATURES

ML: numbers, categories
DL: text, images

TARGET

what you want to predict

+

LEARNING ALGORITHM

defines the learning process

randomly guess → adjust guess to optimize cost function → repeat

=

ML MODEL

mathematical function approximating reality

EVALUATION

quantifies how good the model is

METRICS

DATA SPLITTING

ABLATION

DATA SHUFFLING

PRE-TRAINED MODELS

allows us the ability to pull models and fine-tune for specific applications

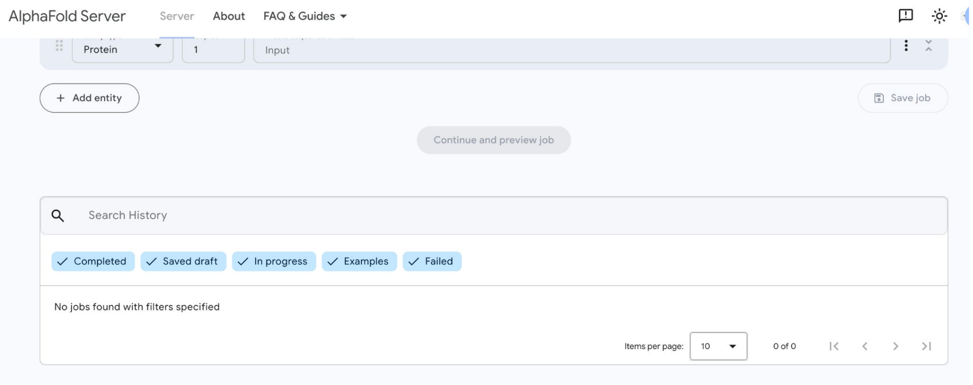

https://alphafoldserver.com/welcome

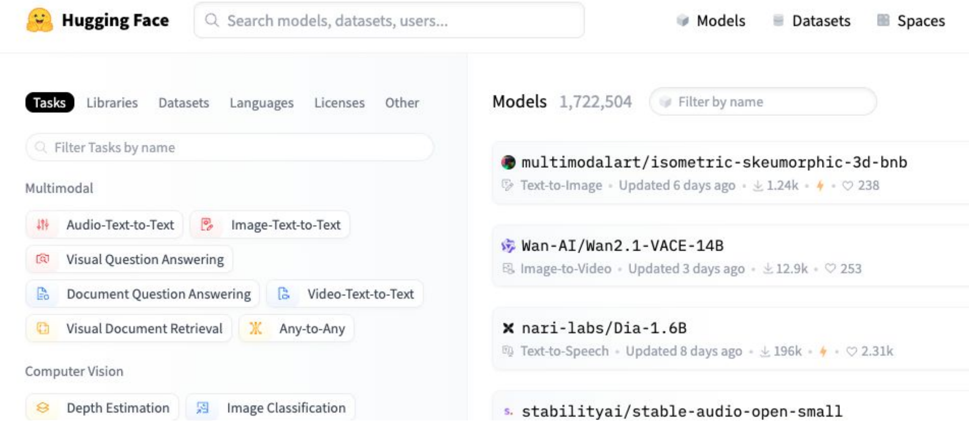

https://huggingface.co/models

WHAT MAKES COMPUTERS INTELLIGENT?

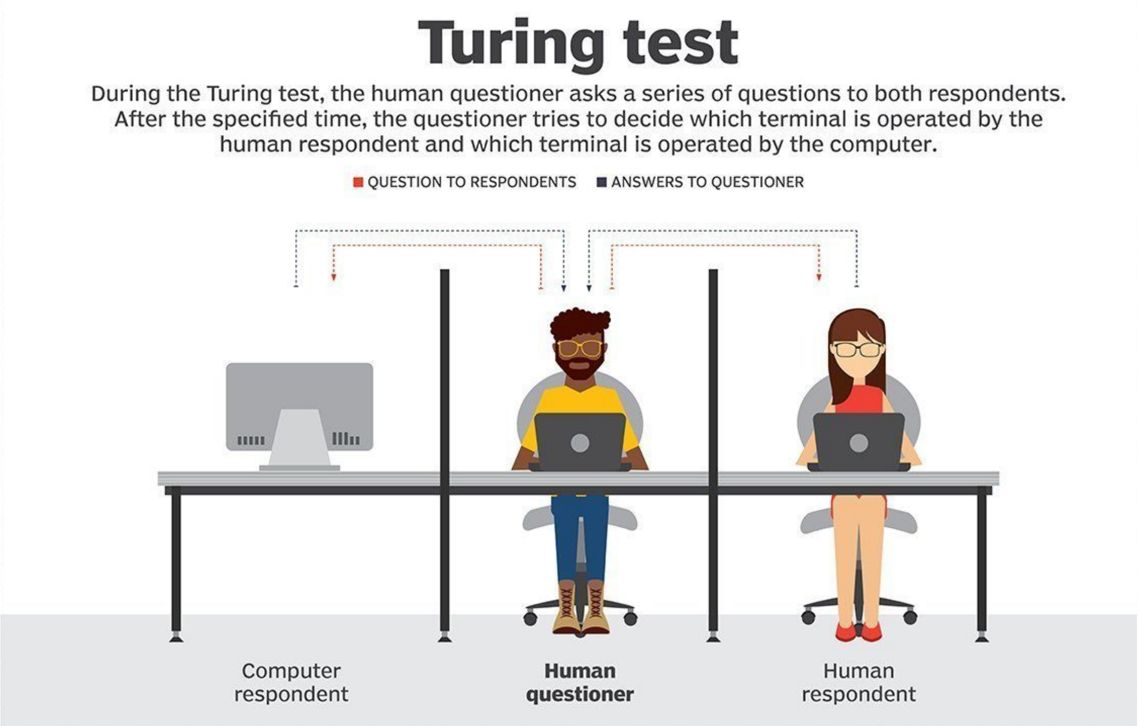

https://www.techtarget.com/searchenterpriseai/definition/Turing-test

Computing Machinery and Intelligence [1950] https://courses.cs.umbc.edu/471/papers/turing.pdf (fun read)
LLMs Passed the Turing Test [in review] https://arxiv.org/abs/2503.23674

WHY SO MUCH AI NOW?

BIG DATA:
larger datasets coupled with easier collection and storage.

HARDWARE:
graphic processing units (GPU)

SOFTWARE:
open-source that builds community

COMPETITIONS:
iterative push to innovate with clear success metric

AI WINTER:
20 year delay due perceptrons being "canceled"

COMPUTERS THAT BEAT HUMANS

AlphaGo

https://www.youtube.com/watch?v=WXuK6gekU1Y

MODELING COMPETITIONS

ImageNet

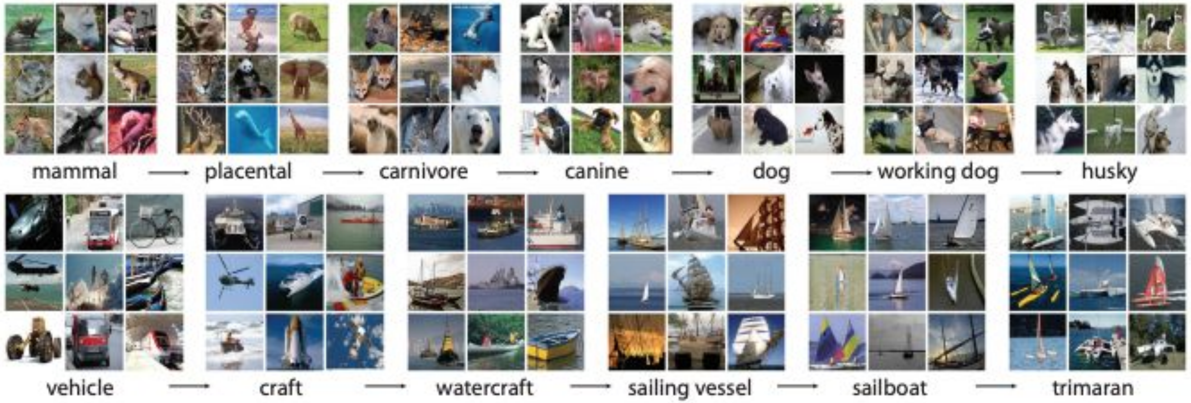

https://www.image-net.org/static_files/papers/imagenet_cvpr09.pdf

MODELING COMPETITIONS

CASP

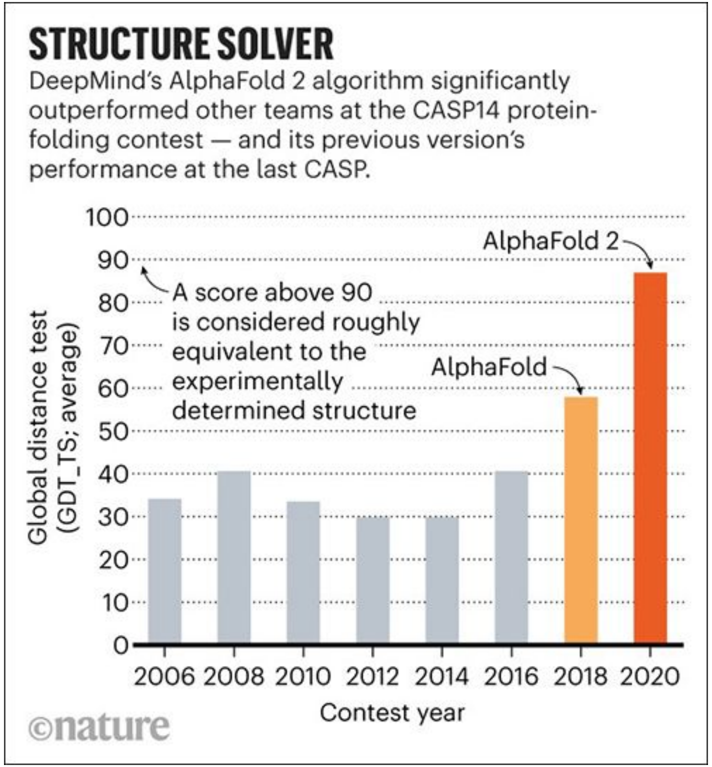

Nature article: https://www.nature.com/articles/d41586-020-03348-4
Protein Structure Prediction Center https://predictioncenter.org/

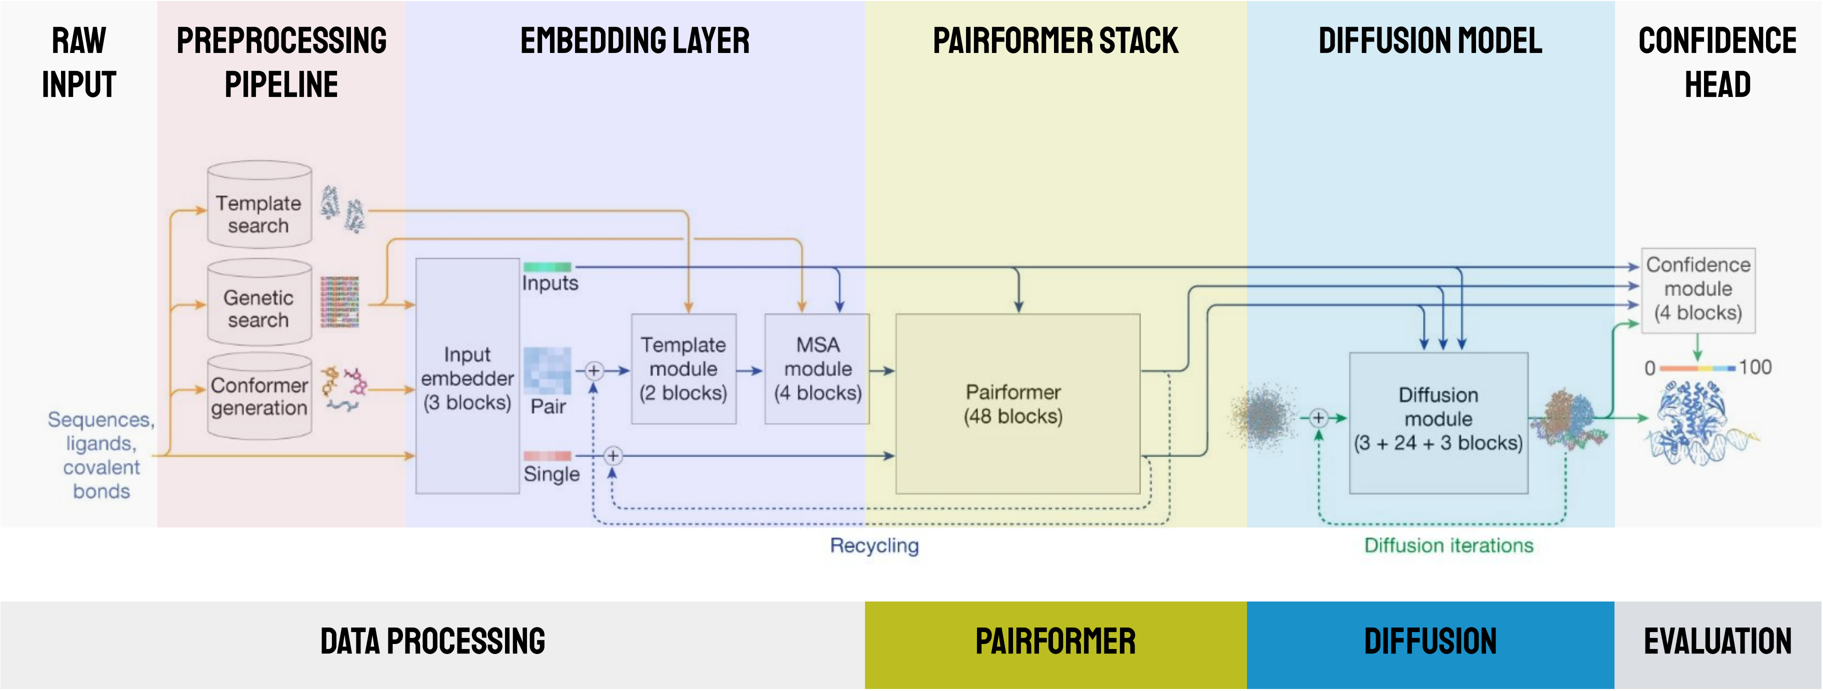

https://www.nature.com/articles/s41586-024-07487-w/figures/1
https://research.dimensioncap.com/p/an-opinionated-alphafold3-field-guide

ATTENTION

INPUT *"I am **not sad**"*

TOKENIZED ["I", "am", "**not**", "**sad**"]

EMBEDDING

I → [0.1, 0.0, 0.3]
am → [0.2, 0.1, 0.0]
not → [0.5, 0.1, 0.6]
sad → [0.4, 0.3, 0.2]

this is the input to the attention layer

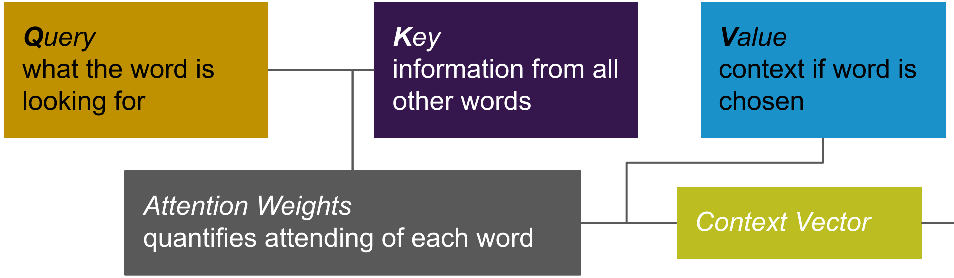

This lets the model figure out that "**not sad**" likely means **positive emotion** by paying more attention to "**not**" when interpreting "**sad**". This is called **attending**, and what attention is named after.

CONSTRAINTS

	House 1	House 2	House 3
Color			
Nationality			
Animal			
Sport			

Clues
1. The Brazilian does not live in house two.
2. The person with the Dogs plays Basketball.
3. There is one house between the person who plays Football and the Red house on the right.
4. The person with the Fishes lives directly to the left of the person with the Cats.
5. The person with the Dogs lives directly to the right of the Green house.
6. The German lives in house three.

Zebra Puzzle Example: https://www.researchgate.net/figure/An-Example-Zebra-Puzzle_fig1_382065247

TRIANGLE ATTENTION

Triangle attention can enforce the **triangle inequality**

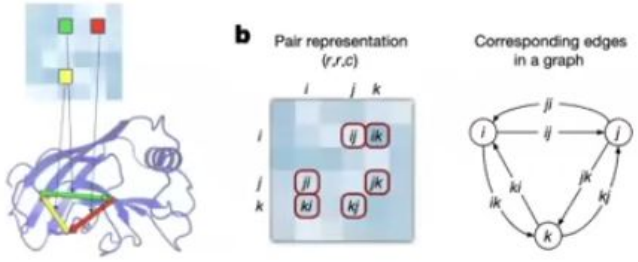

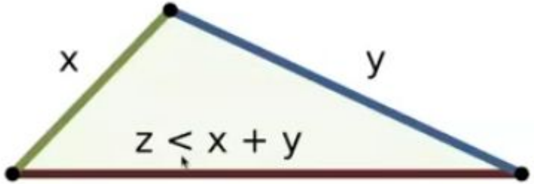

The Pairformer uses attention to find which triangles are the most influential and passes on suggested spatial and functional relationships.

Review and Discussion of AlphaFold3 https://www.youtube.com/watch?v=qjFgthkKxcA (go to the end for some hacks!)

DIFFUSION

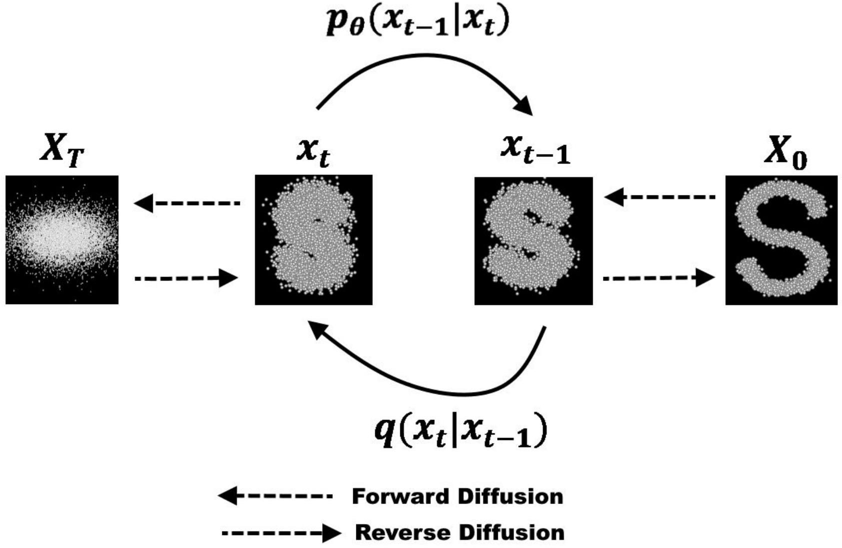

The Diffusion module uses the suggested constraints from the Pairformer and iteratively builds a 3D structure. Randomly initializing (seed) gives different outputs.

DATA

shapes what the model learns

FEATURES

input sequences, MSA, evolutionary relationships, chemical properties → **embedded representations**

TARGET

3D protein structures

+

LEARNING ALGORITHM

defines the learning process

In the **Diffusion module**:

- start with random noise (coordinates)
- gradually sculpt a structure using evolutionary patterns and geometric constraints given from **Pairformer**
- iteratively refine while allowing any amino acid to influence any other
- repeat until structure satisfies both evolutionary and physical rules

=

PRE-TRAINED MODEL

mathematical function approximating reality

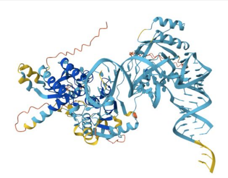

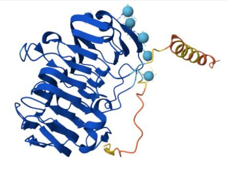

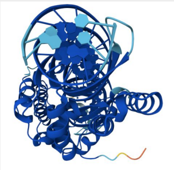

EVALUATION

quantifies how good the model is

CONFIDENCE

DATA SPLITTING

ABLATION

DATA SHUFFLING

ETHICS

What to keep in mind

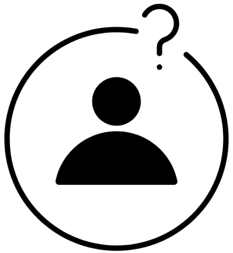

FOR WHOM?
who benefits, who is harmed

JUST BECAUSE YOU CAN, DOESN'T MEAN YOU SHOULD
even if possible, is it responsible

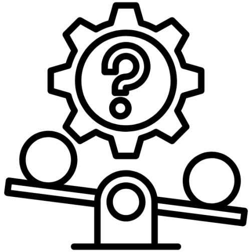

BIASES ARE BAKED IN
from the data used to the model assumptions

REGULATION IS BEHIND DEPLOYMENT
surveillance, data privacy, copyright often discussions years later (if at all)

HYPE VS REALITY

are there limits to
AI applications?

DATA QUALITY
models are only as good as the data they're built upon

AI IS NOT MAGIC, IT'S MATH
and this adds a barrier in understanding

NARROW AI IS TASK-SPECIFIC
exceedingly good at one thing, generally bad at most things

AI IS NOT OBJECTIVE
it reflects human decisions on what is considered important

HYPE DISTRACTS FROM HARM
how do data centers and resource usage affect environmental concerns?

LONG TERM IMPACT IS UNKNOWN
are we better off because of this, or are there unintended side effects

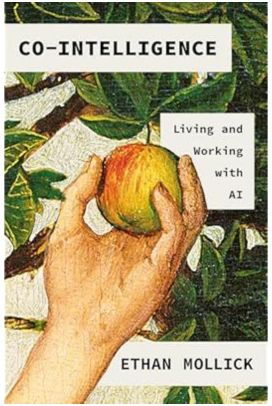

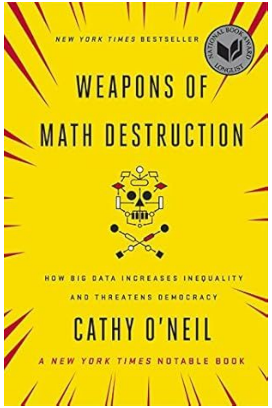

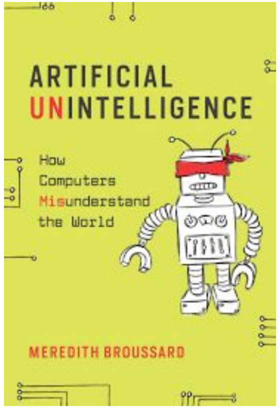

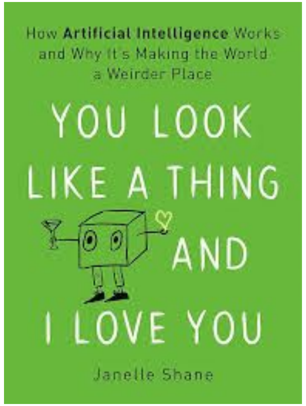

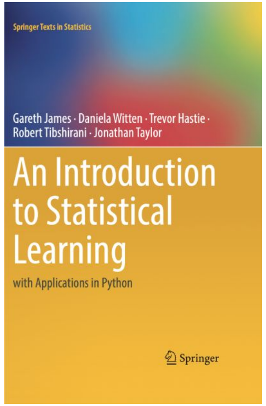

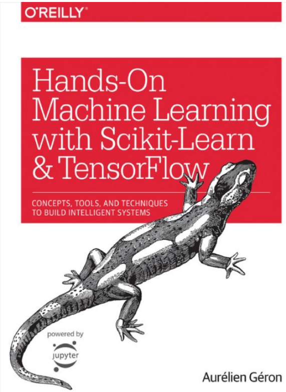

GENERAL

TECHNICAL

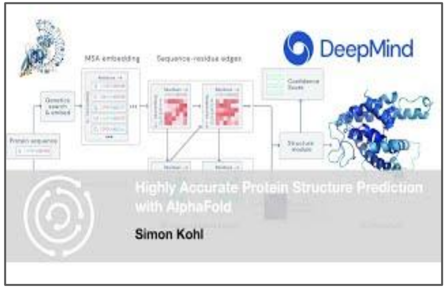

heidelberg.ai

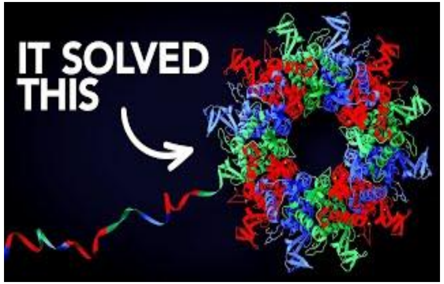

[Veritasium]

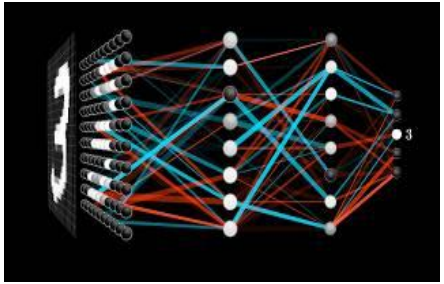

3Blue1Brown

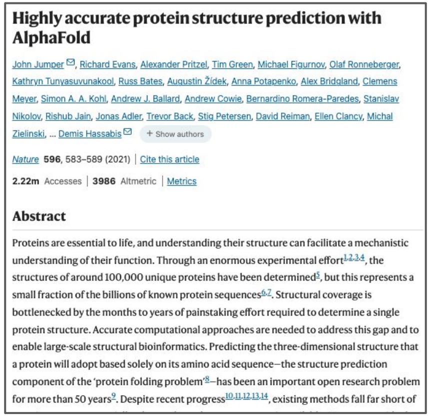

Highly accurate protein structure prediction with AlphaFold

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, ... Demis Hassabis + Show authors

Nature **596**, 583–589 (2021) | Cite this article

Abstract

Proteins are essential to life, and understanding their structure can facilitate a mechanistic understanding of their function. Through an enormous experimental effort[1,2,3,4], the structures of around 100,000 unique proteins have been determined[5], but this represents a small fraction of the billions of known protein sequences[6,7]. Structural coverage is bottlenecked by the months to years of painstaking effort required to determine a single protein structure. Accurate computational approaches are needed to address this gap and to enable large-scale structural bioinformatics. Predicting the three-dimensional structure that a protein will adopt based solely on its amino acid sequence—the structure prediction component of the 'protein folding problem'[8]—has been an important open research problem for more than 50 years[9]. Despite recent progress[10,11,12,13,14], existing methods fall far short of

[AlphaFold (2021)]

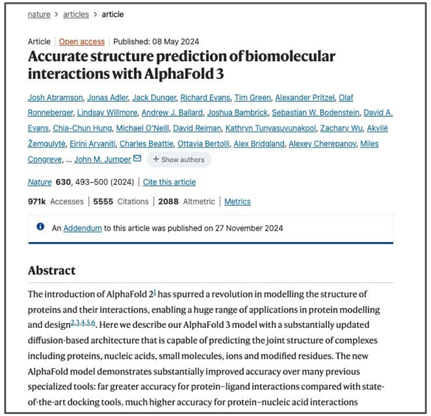

Article | Open access | Published: 08 May 2024

Accurate structure prediction of biomolecular interactions with AlphaFold 3

Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J. Ballard, Joshua Bambrick, Sebastian W. Bodenstein, David A. Evans, Chia-Chun Hung, Michael O'Neill, David Reiman, Kathryn Tunyasuvunakool, Zachary Wu, Akvilė Žemgulytė, Eirini Arvaniti, Charles Beattie, Ottavia Bertolli, Alex Bridgland, Alexey Cherepanov, Miles Congreve, ... John M. Jumper + Show authors

Nature **630**, 493–500 (2024) | Cite this article

An Addendum to this article was published on 27 November 2024

Abstract

The introduction of AlphaFold 2[1] has spurred a revolution in modelling the structure of proteins and their interactions, enabling a huge range of applications in protein modelling and design[2,3,4,5,6]. Here we describe our AlphaFold 3 model with a substantially updated diffusion-based architecture that is capable of predicting the joint structure of complexes including proteins, nucleic acids, small molecules, ions and modified residues. The new AlphaFold model demonstrates substantially improved accuracy over many previous specialized tools: far greater accuracy for protein–ligand interactions compared with state-of-the-art docking tools, much higher accuracy for protein–nucleic acid interactions

[AlphaFold3 (2024)]

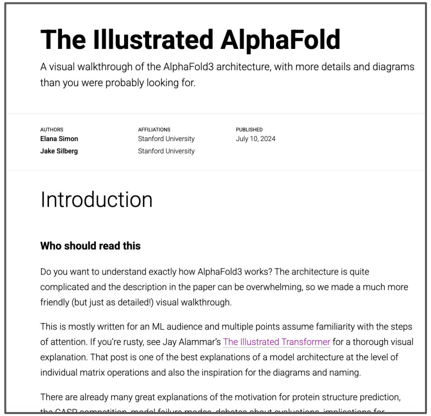

The Illustrated AlphaFold

A visual walkthrough of the AlphaFold3 architecture, with more details and diagrams than you were probably looking for.

AUTHORS
Elana Simon
Jake Silberg

AFFILIATIONS
Stanford University
Stanford University

PUBLISHED
July 10, 2024

Introduction

Who should read this

Do you want to understand exactly how AlphaFold3 works? The architecture is quite complicated and the description in the paper can be overwhelming, so we made a much more friendly (but just as detailed!) visual walkthrough.

This is mostly written for an ML audience and multiple points assume familiarity with the steps of attention. If you're rusty, see Jay Alammar's The Illustrated Transformer for a thorough visual explanation. That post is one of the best explanations of a model architecture at the level of individual matrix operations and also the inspiration for the diagrams and naming.

There are already many great explanations of the motivation for protein structure prediction,

[Illustrated AlphaFold3]

THANK YOU
ANY QUESTIONS?

Extra Slides

ALPHAFOLD

Building Blocks

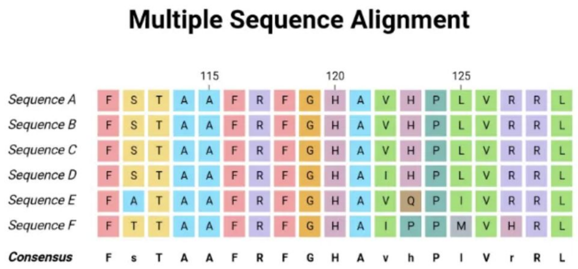

Integrates evolutionary information from multiple sequence alignments (MSA). Consensus shows how conserved each sequence position is

MSA includes looking at co-evolution

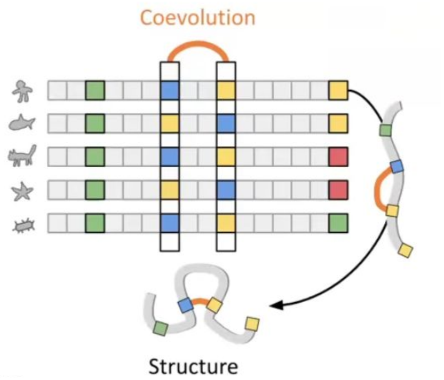

https://research.dimensioncap.com/p/an-opinionated-alphafold3-field-guide

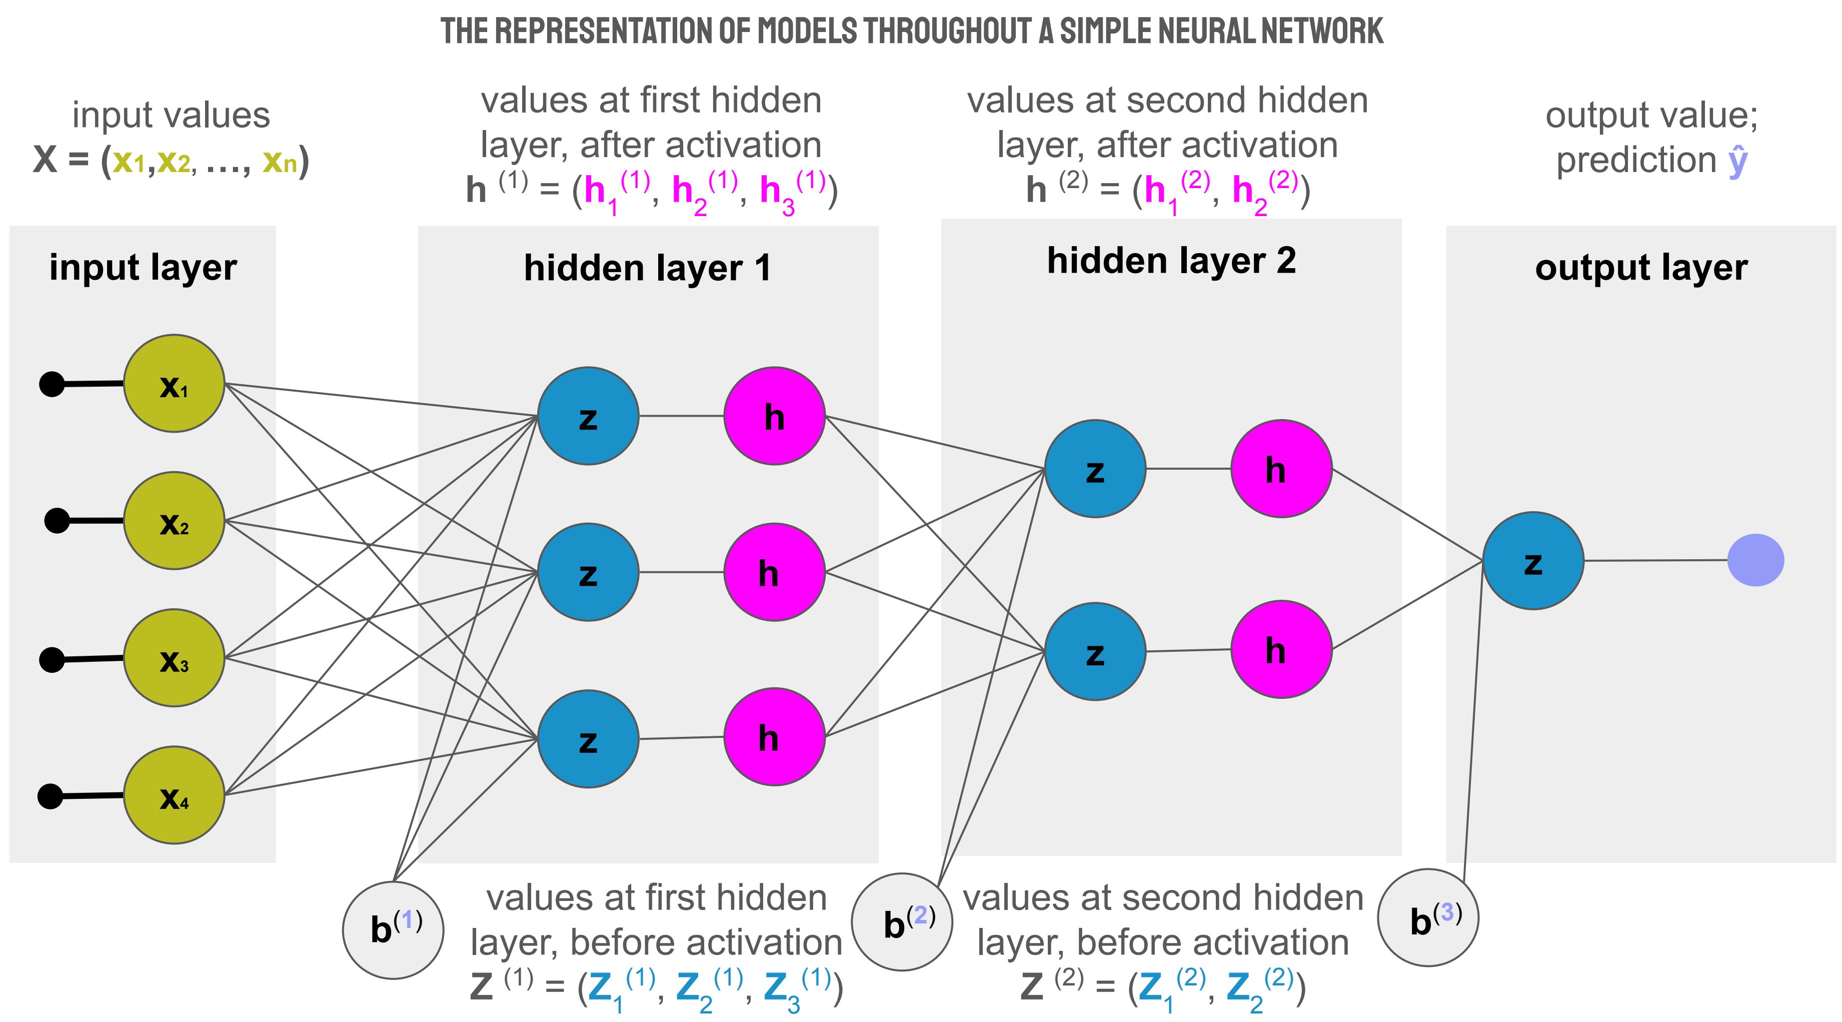
THE REPRESENTATION OF MODELS THROUGHOUT A SIMPLE NEURAL NETWORK
input values
X = (x1,x2, …, xn)
values at first hidden layer, after activation
h (1) = (h1(1), h2(1), h3(1))
values at second hidden layer, after activation
h (2) = (h1(2), h2(2))
output value; prediction ŷ
input layer
hidden layer 1
hidden layer 2
output layer
x1
x2
x3
x4
z
h
z
h
z
h
z
h
z
h
z
b(1)
b(2)
b(3)
values at first hidden layer, before activation
Z (1) = (Z1(1), Z2(1), Z3(1))
values at second hidden layer, before activation
Z (2) = (Z1(2), Z2(2))

THE LINEAR ALGEBRA USED ON THOSE REPRESENTATIONS IN A SINGLE HIDDEN LAYER OF A SIMPLE NEURAL NETWORK

$$h^{(1)} = (h_1^{(1)}, h_2^{(1)}, h_3^{(1)})$$

$$X = (x_1, x_2, \ldots, x_n) \qquad Z^{(1)} = (z_1^{(1)}, z_2^{(1)}, z_3^{(1)})$$

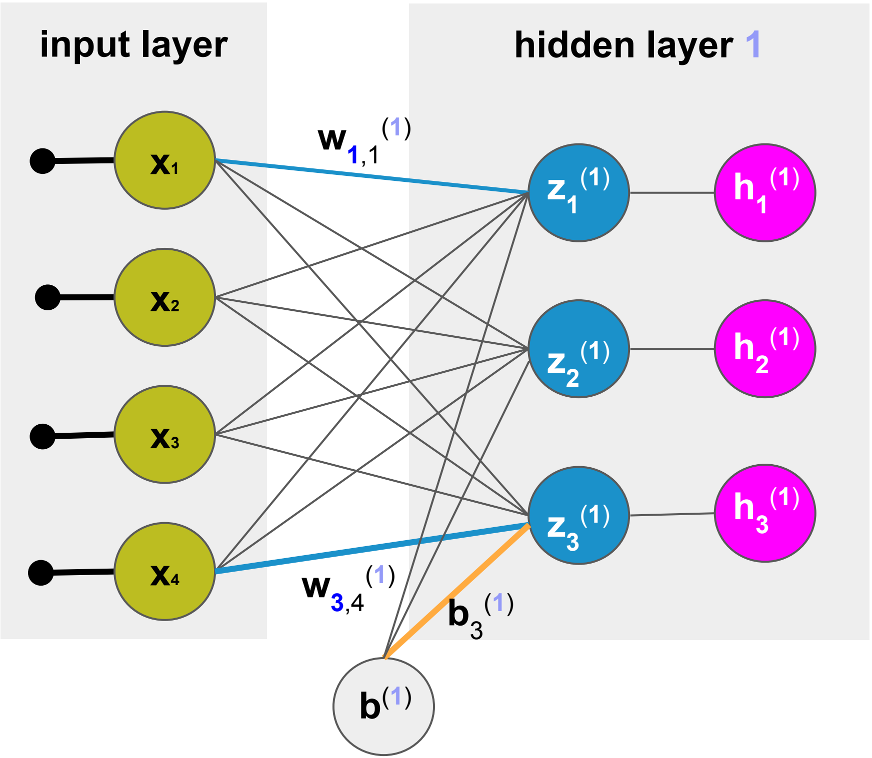

$w_{i,j}^{(1)}$ is the weight of connection from input layer j to **neuron i** at hidden layer 1

$b_i^{(1)}$ is the bias of **neuron i** at hidden layer 1

$$\begin{pmatrix} z_1^{(1)} \\ z_2^{(1)} \\ z_3^{(1)} \end{pmatrix} = \begin{pmatrix} w_{1,1}^{(1)} x_1 + w_{1,2}^{(1)} x_2 + w_{1,3}^{(1)} x_3 + w_{1,4}^{(1)} x_4 \\ w_{2,1}^{(1)} x_1 + w_{2,2}^{(1)} x_2 + w_{2,3}^{(1)} x_3 + w_{2,4}^{(1)} x_4 \\ w_{3,1}^{(1)} x_1 + w_{3,2}^{(1)} x_2 + w_{3,3}^{(1)} x_3 + w_{3,4}^{(1)} x_4 \end{pmatrix} + \begin{pmatrix} b_1^{(1)} \\ b_2^{(1)} \\ b_3^{(1)} \end{pmatrix}$$

$$= \begin{pmatrix} w_{1,1}^{(1)} & w_{1,2}^{(1)} & w_{1,3}^{(1)} & w_{1,4}^{(1)} \\ w_{2,1}^{(1)} & w_{2,2}^{(1)} & w_{2,3}^{(1)} & w_{2,4}^{(1)} \\ w_{3,1}^{(1)} & w_{3,2}^{(1)} & w_{3,3}^{(1)} & w_{3,4}^{(1)} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} + \begin{pmatrix} b_1^{(1)} \\ b_2^{(1)} \\ b_3^{(1)} \end{pmatrix}$$

$$\boxed{Z^{(1)} = W^{(1)} X + b^{(1)}}$$

bias vector

matrix of weights

LEARNING ALGORITHM

what are they able to do?

Give **[an algorithm]** data and it will

randomly guess → adjust guess to optimize cost function→ repeat

in order to : **[do things]**

convolutional NN: scanning image to find patterns (zoom-in)

recurrent NN : keep track of information it has seen

graph network : learn how items are connected

attention blocks : focus on relevant parts