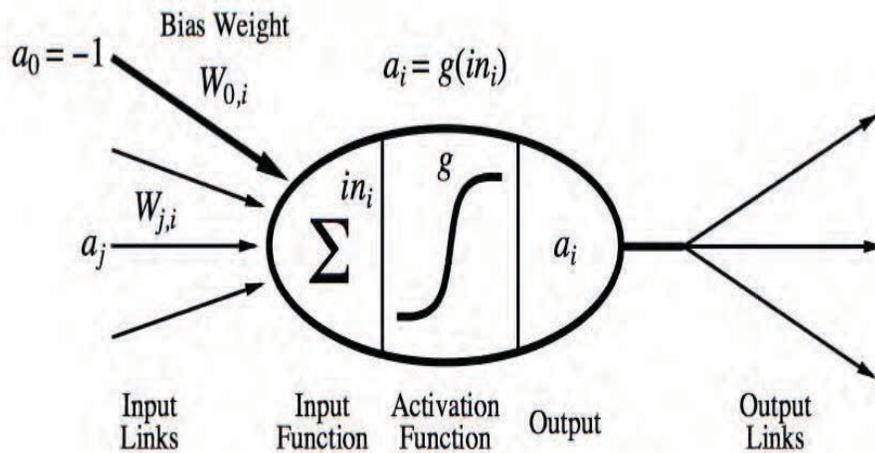


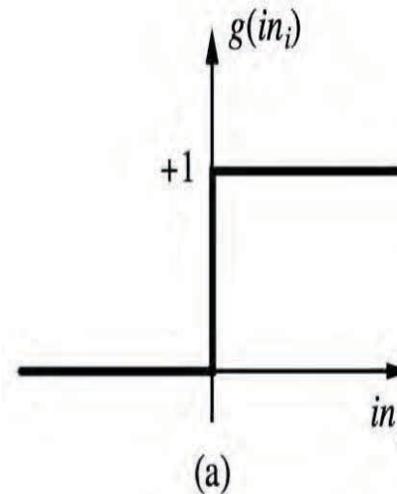
What is a Neuron ?

- A neuron is a computational unit in a neural network that exchanges messages with other neurons

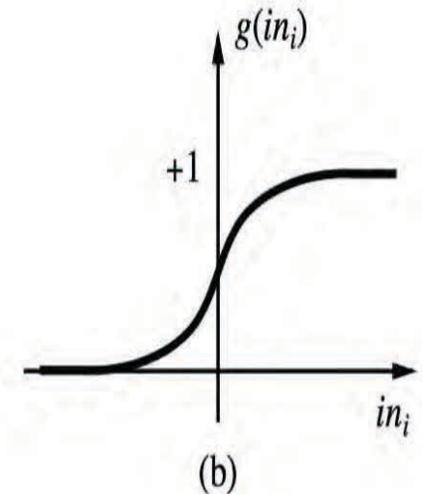
$$a_i \leftarrow g(in_i) = g(\sum_j W_{j,i} a_j)$$



Possible activation functions:

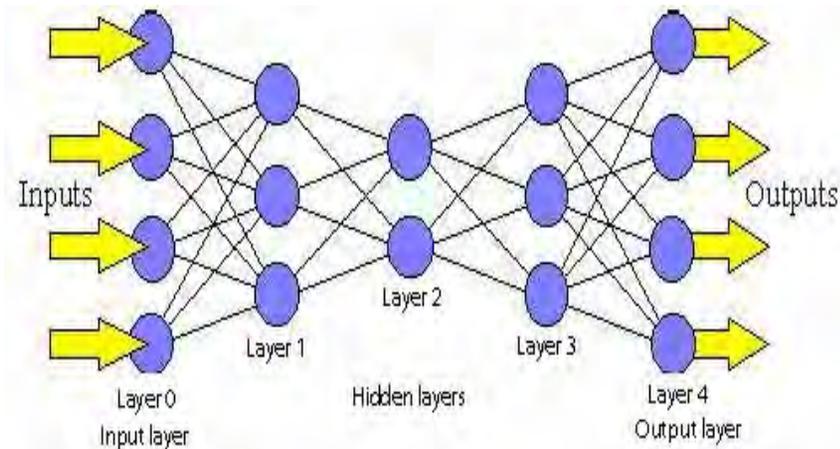


(a) Step function/
threshold function



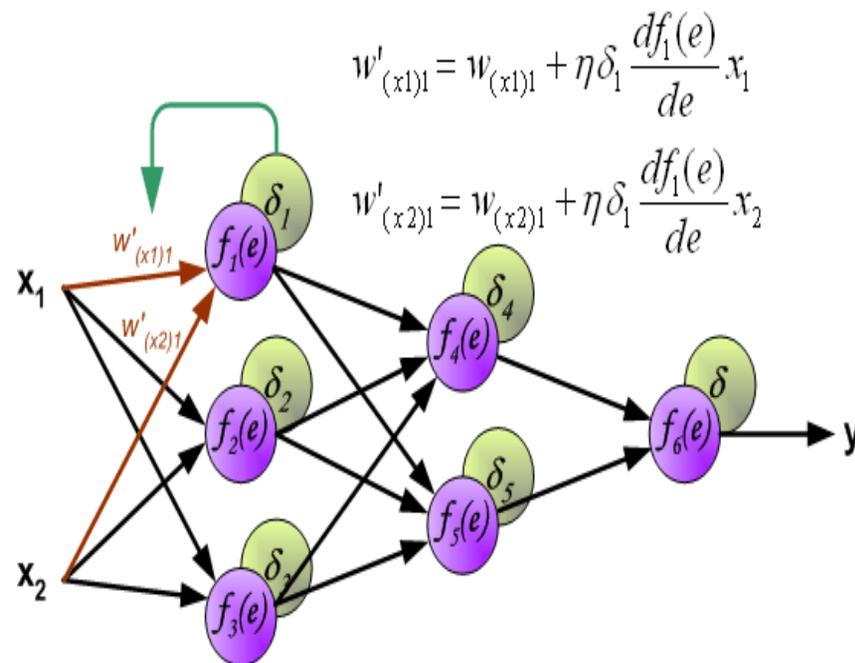
(b) Sigmoid function (a.k.a,
logistic function)

Feed Forward & Backpropagation



Feed forward algorithm:

- Activate the neurons from the left to the right.



Backpropagation:

- Randomly initialize the parameters
- Calculate total error at the right, $f_6(e)$
- Then calculate contributions to error, δ_n , at each step going backwards.

$$f(x) = \frac{1}{1 + e^{-x}}$$

-0.06

2.7

-2.5

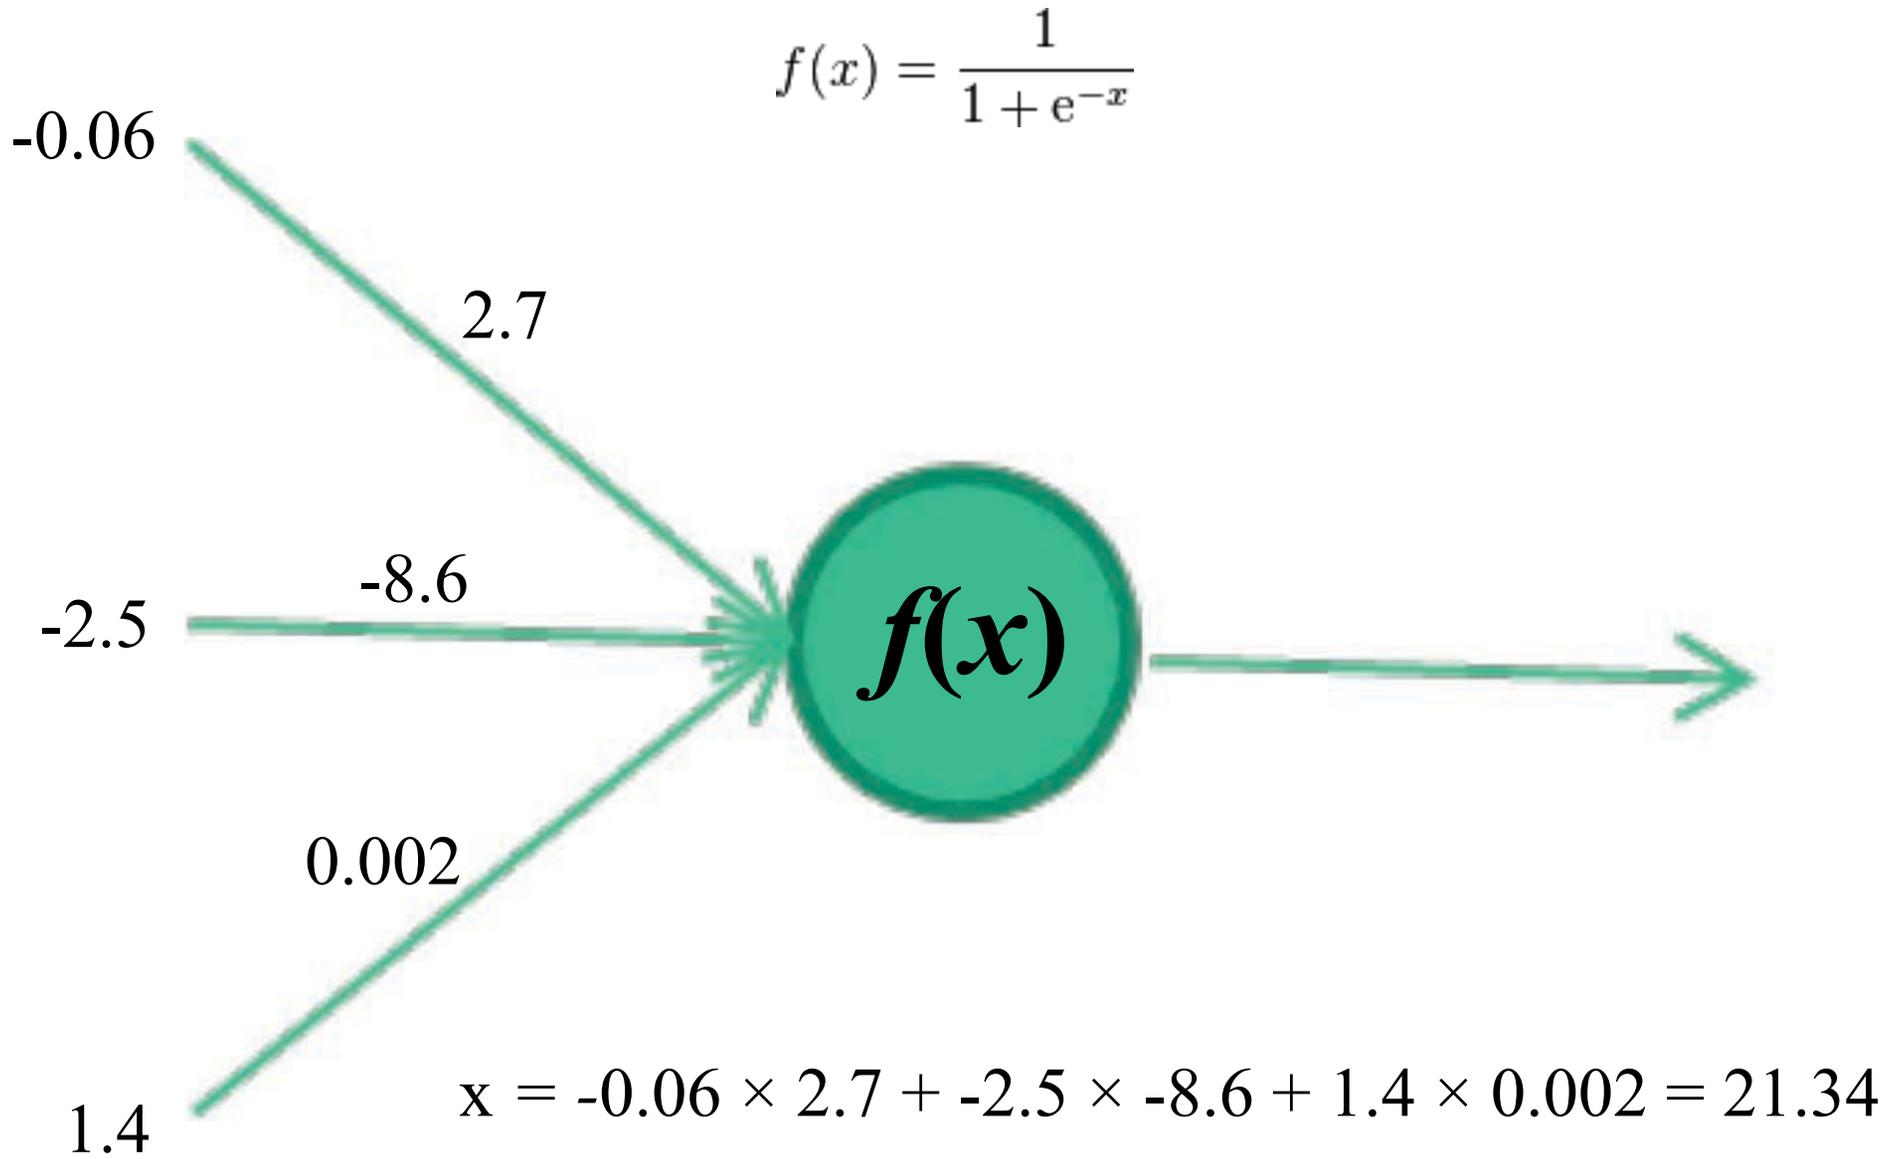
-8.6

0.002

1.4

$f(x)$

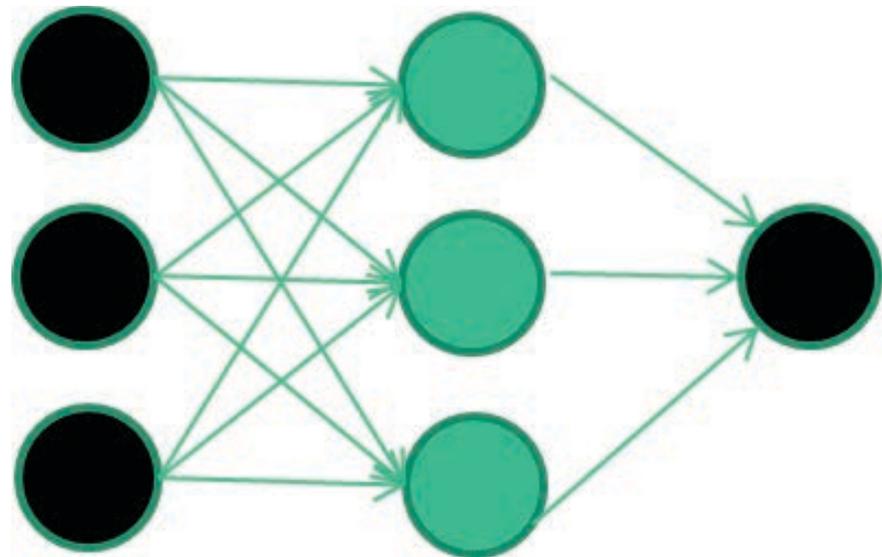
$$x = -0.06 \times 2.7 + -2.5 \times -8.6 + 1.4 \times 0.002 = 21.34$$



Training data

| <i>Fields</i> | <i>class</i> |
|---------------|--------------|
| 1.4 2.7 1.9 | 0 |
| 3.8 3.4 3.2 | 0 |
| 6.4 2.8 1.7 | 1 |
| 4.1 0.1 0.2 | 0 |
| etc ... | |

Initialise with random weights



Training data

Fields *class*

1.4 2.7 1.9 0

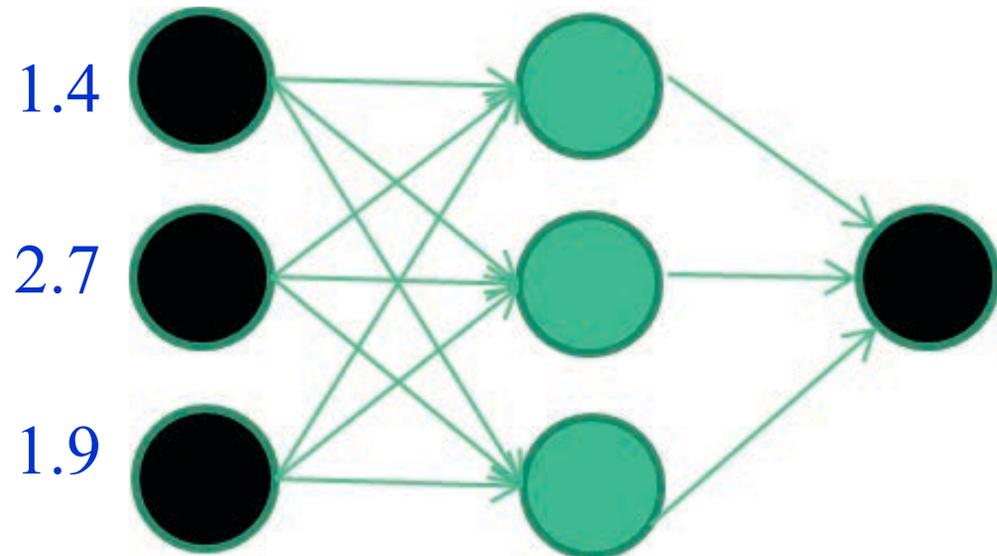
3.8 3.4 3.2 0

6.4 2.8 1.7 1

4.1 0.1 0.2 0

etc ...

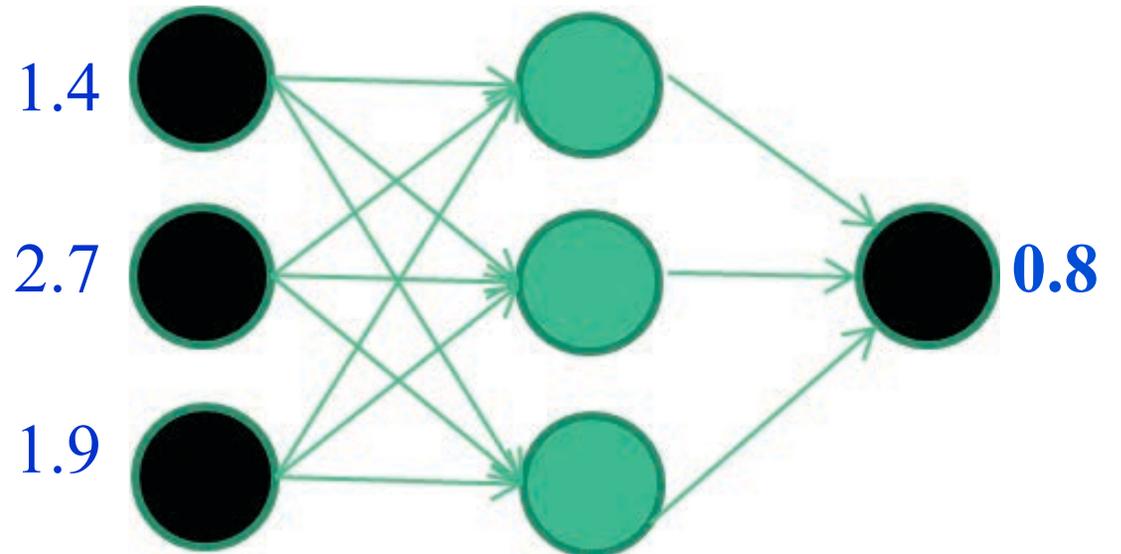
Present a training pattern



Training data

| <i>Fields</i> | <i>class</i> |
|---------------|--------------|
| 1.4 2.7 1.9 | 0 |
| 3.8 3.4 3.2 | 0 |
| 6.4 2.8 1.7 | 1 |
| 4.1 0.1 0.2 | 0 |
| etc ... | |

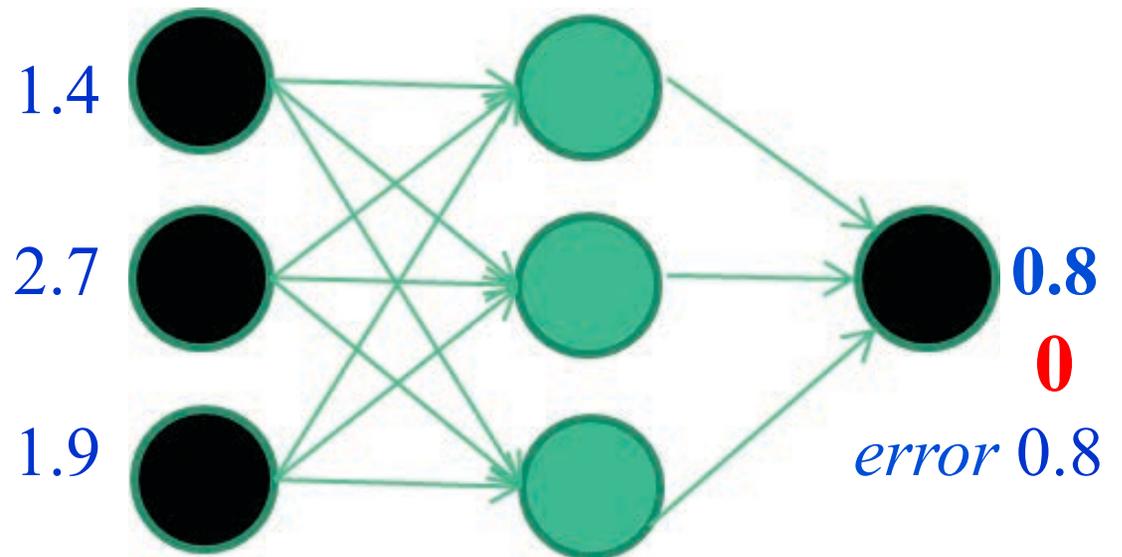
Feed it through to get output



Training data

| <i>Fields</i> | <i>class</i> |
|---------------|--------------|
| 1.4 2.7 1.9 | 0 |
| 3.8 3.4 3.2 | 0 |
| 6.4 2.8 1.7 | 1 |
| 4.1 0.1 0.2 | 0 |
| etc ... | |

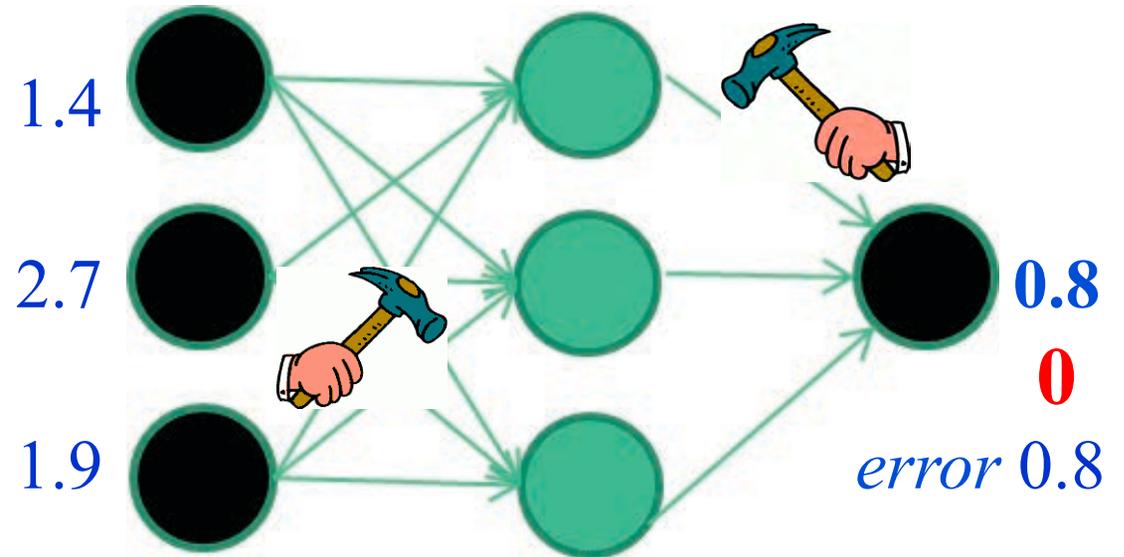
Compare with target output



Training data

| <i>Fields</i> | <i>class</i> |
|---------------|--------------|
| 1.4 2.7 1.9 | 0 |
| 3.8 3.4 3.2 | 0 |
| 6.4 2.8 1.7 | 1 |
| 4.1 0.1 0.2 | 0 |
| etc ... | |

Adjust weights based on error



Training data

Fields *class*

1.4 2.7 1.9 0

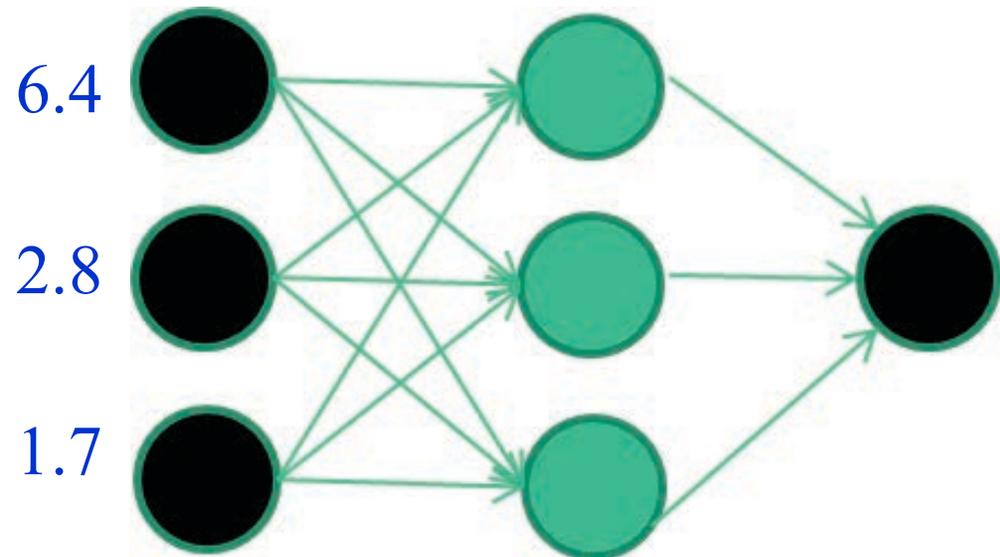
3.8 3.4 3.2 0

6.4 2.8 1.7 1

4.1 0.1 0.2 0

etc ...

Present a training pattern



Training data

Fields *class*

1.4 2.7 1.9 0

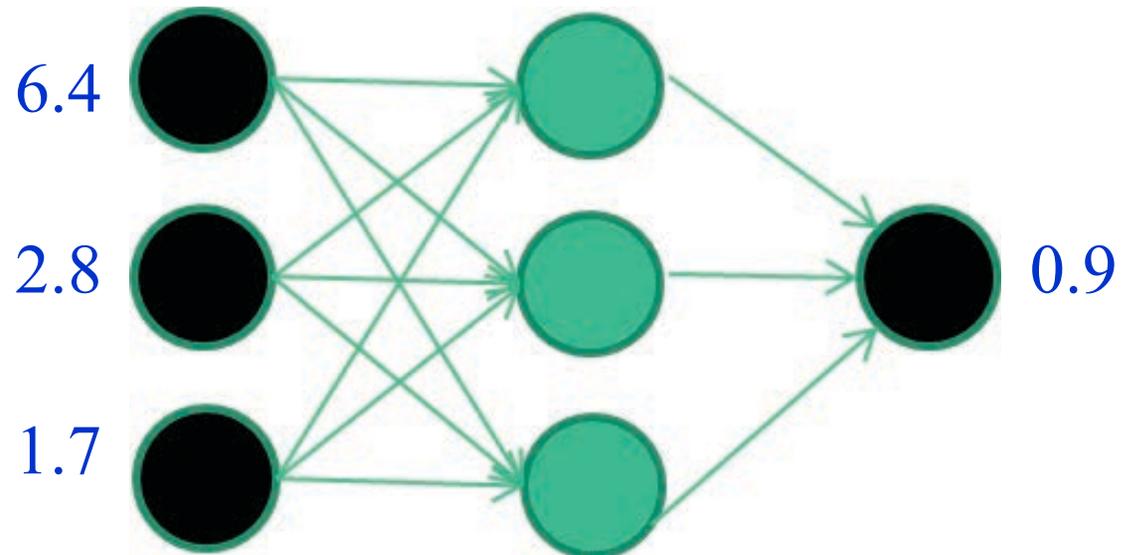
3.8 3.4 3.2 0

6.4 2.8 1.7 1

4.1 0.1 0.2 0

etc ...

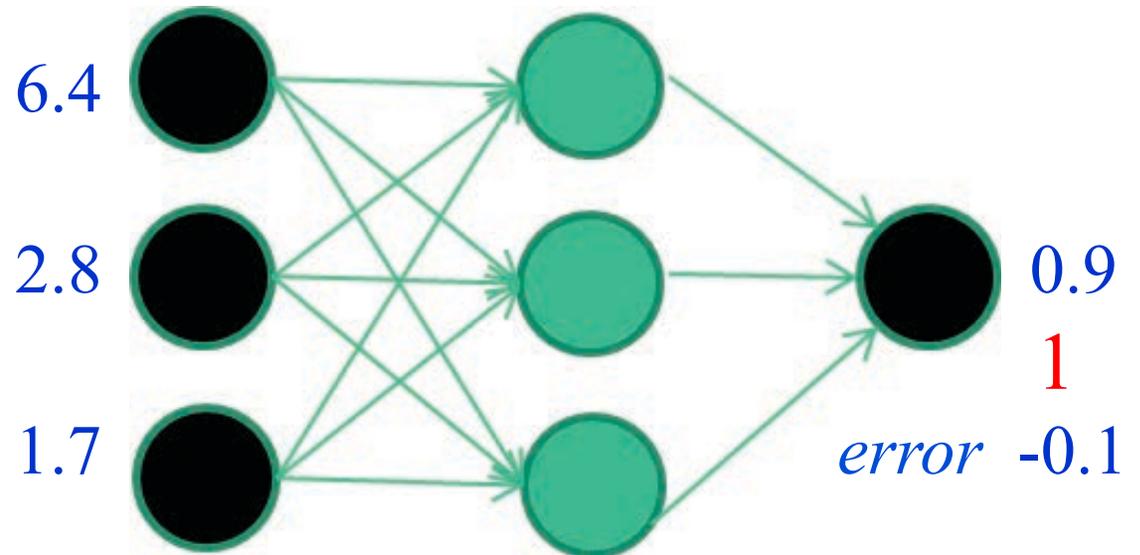
Feed it through to get output



Training data

| <i>Fields</i> | <i>class</i> |
|---------------|--------------|
| 1.4 2.7 1.9 | 0 |
| 3.8 3.4 3.2 | 0 |
| 6.4 2.8 1.7 | 1 |
| 4.1 0.1 0.2 | 0 |
| etc ... | |

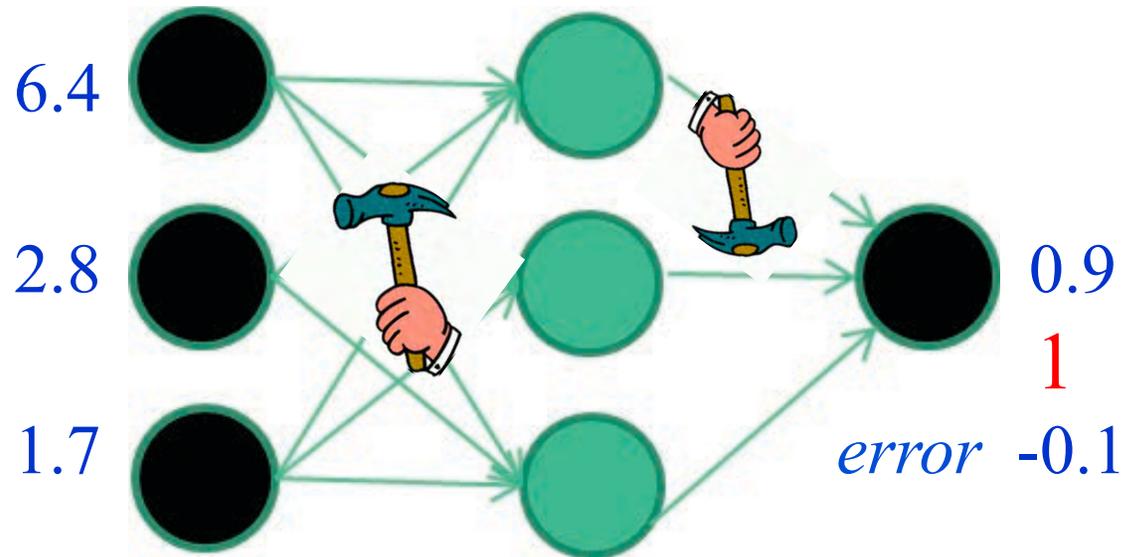
Compare with target output



Training data

| <i>Fields</i> | <i>class</i> |
|---------------|--------------|
| 1.4 2.7 1.9 | 0 |
| 3.8 3.4 3.2 | 0 |
| 6.4 2.8 1.7 | 1 |
| 4.1 0.1 0.2 | 0 |
| etc ... | |

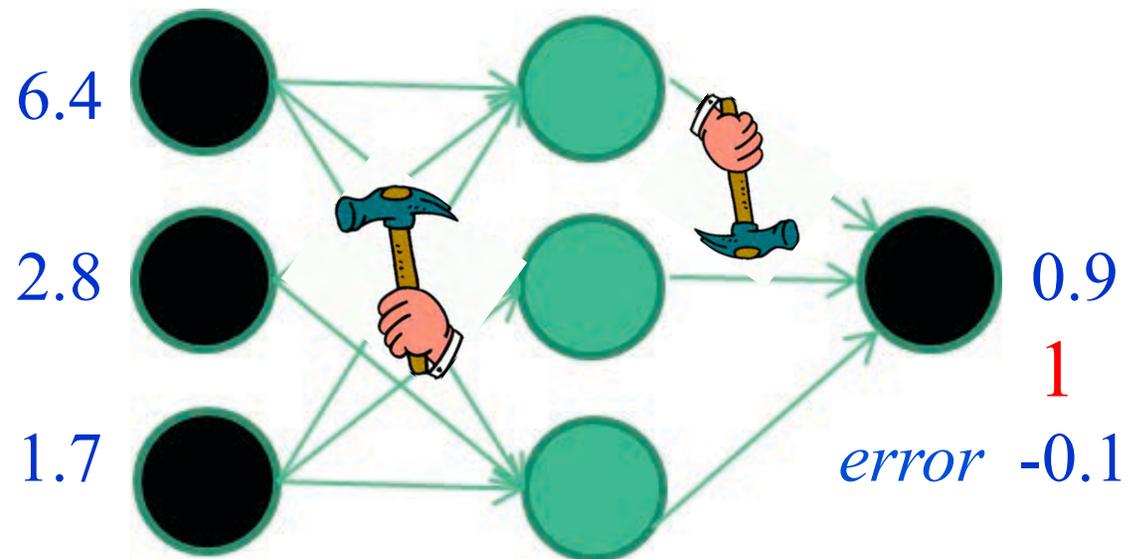
Adjust weights based on error



Training data

| <i>Fields</i> | <i>class</i> |
|---------------|--------------|
| 1.4 2.7 1.9 | 0 |
| 3.8 3.4 3.2 | 0 |
| 6.4 2.8 1.7 | 1 |
| 4.1 0.1 0.2 | 0 |
| etc ... | |

And so on

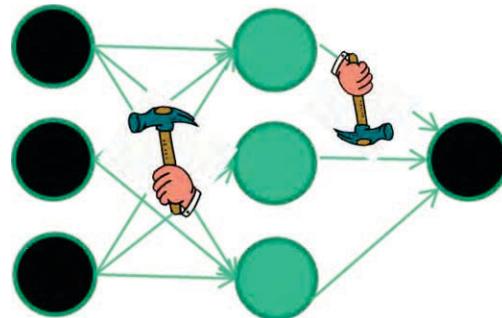


Repeat this thousands, maybe millions of times – each time taking a random training instance, and making slight weight adjustments

Algorithms for weight adjustment are designed to make changes that will reduce the error

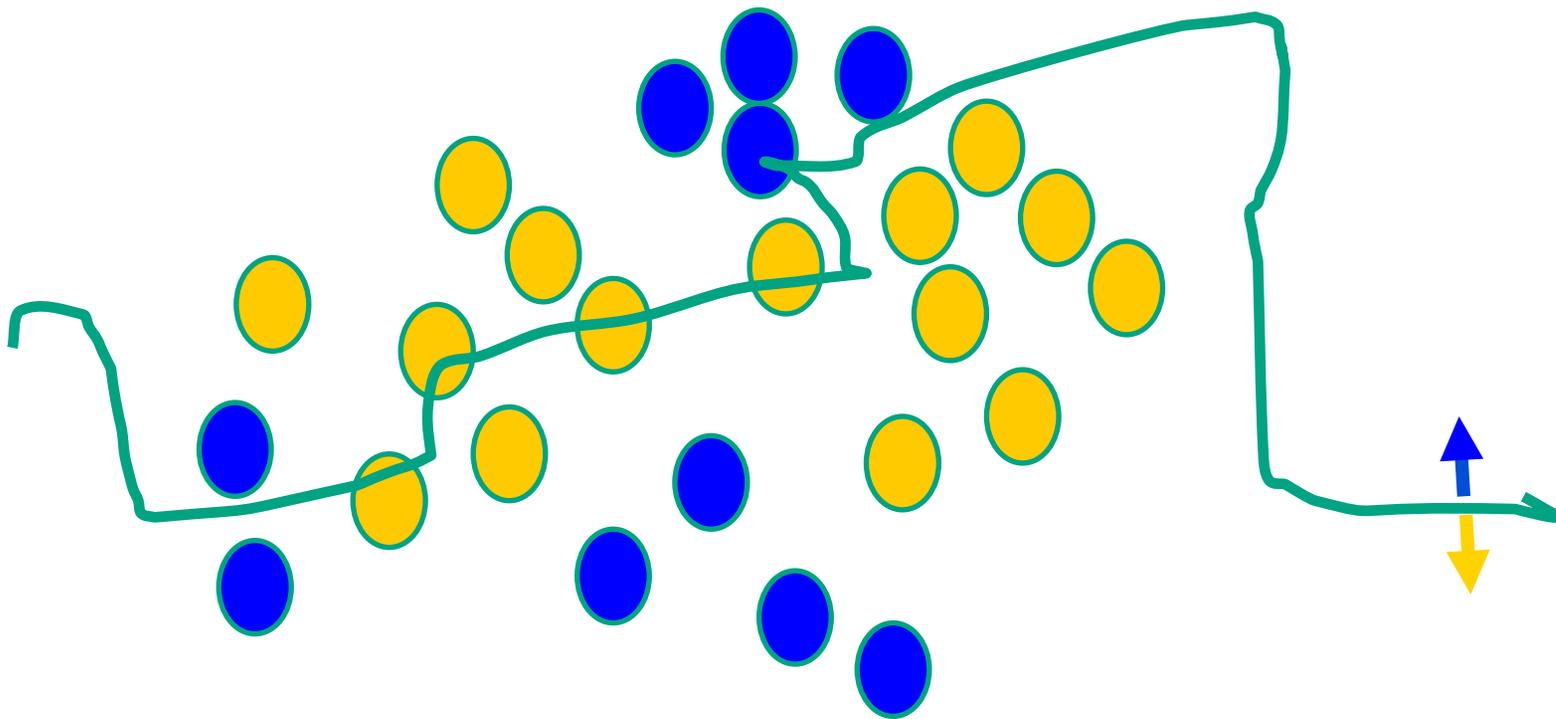
The Main Points to Remember

- weight-learning algorithms for NNs are “simple”
- they work by making thousands and thousands of tiny adjustments, each making the network do better at the most recent pattern, but perhaps a little worse on many others
- but, by “luck”, eventually this tends to be good enough to learn effective classifiers for many real applications



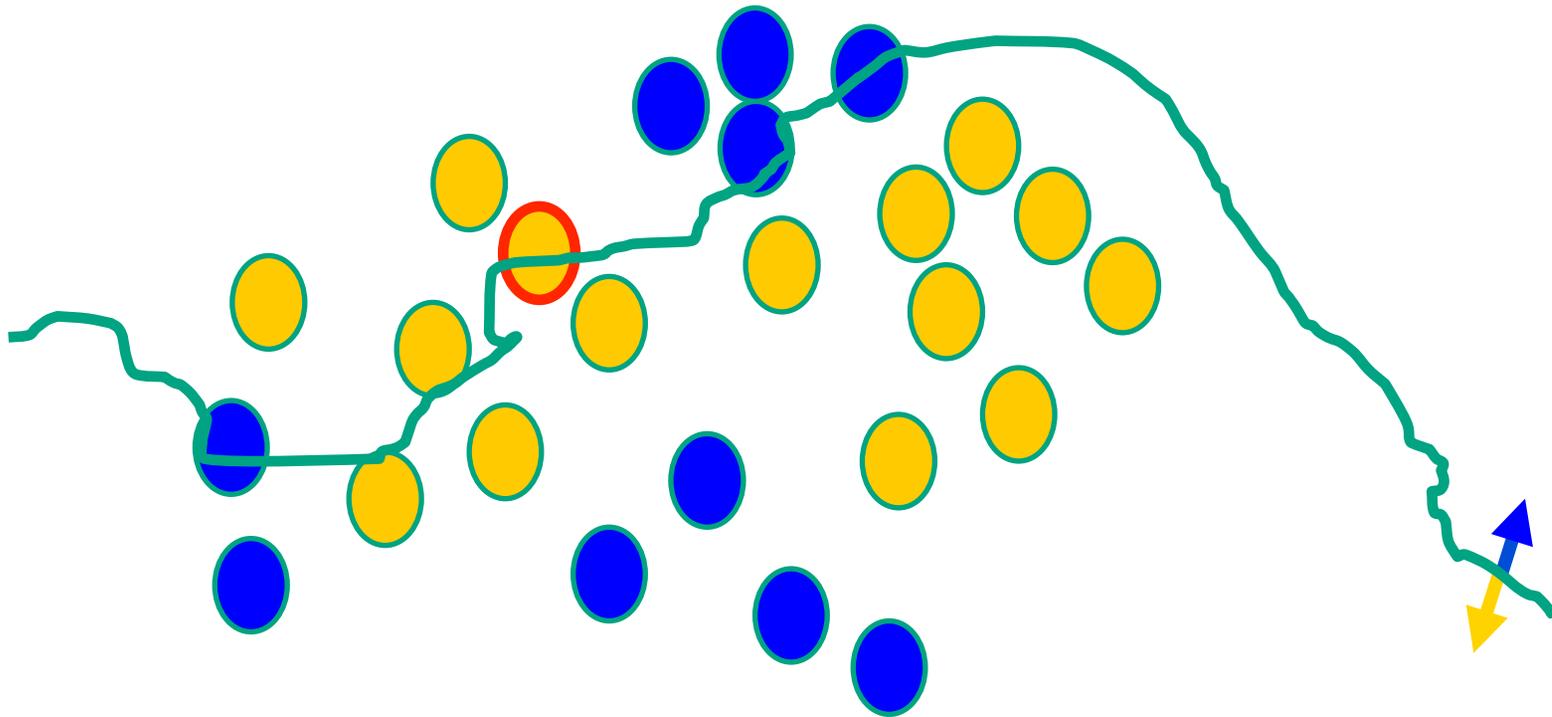
The Decision Boundary Perspective

Initial random weights



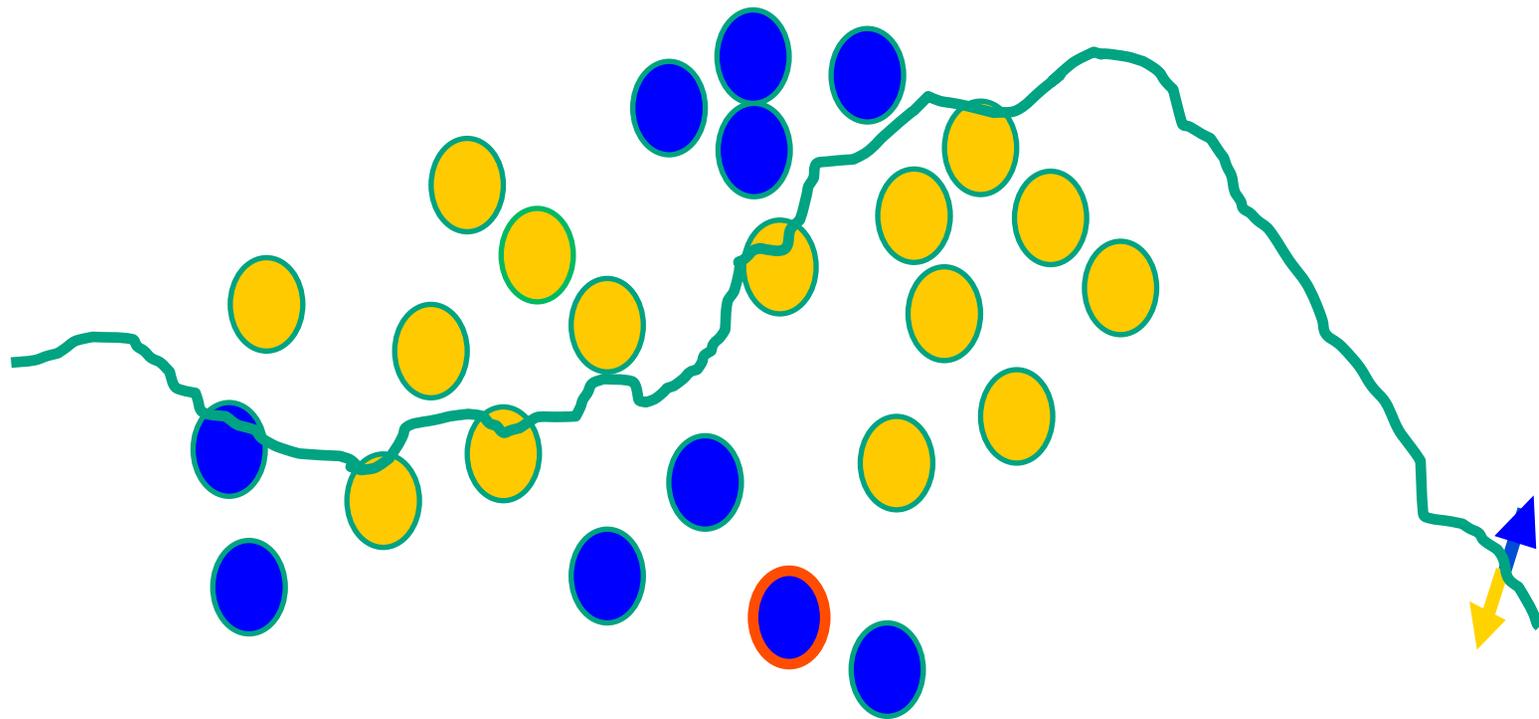
The Decision Boundary Perspective

Present a training instance / adjust the weights



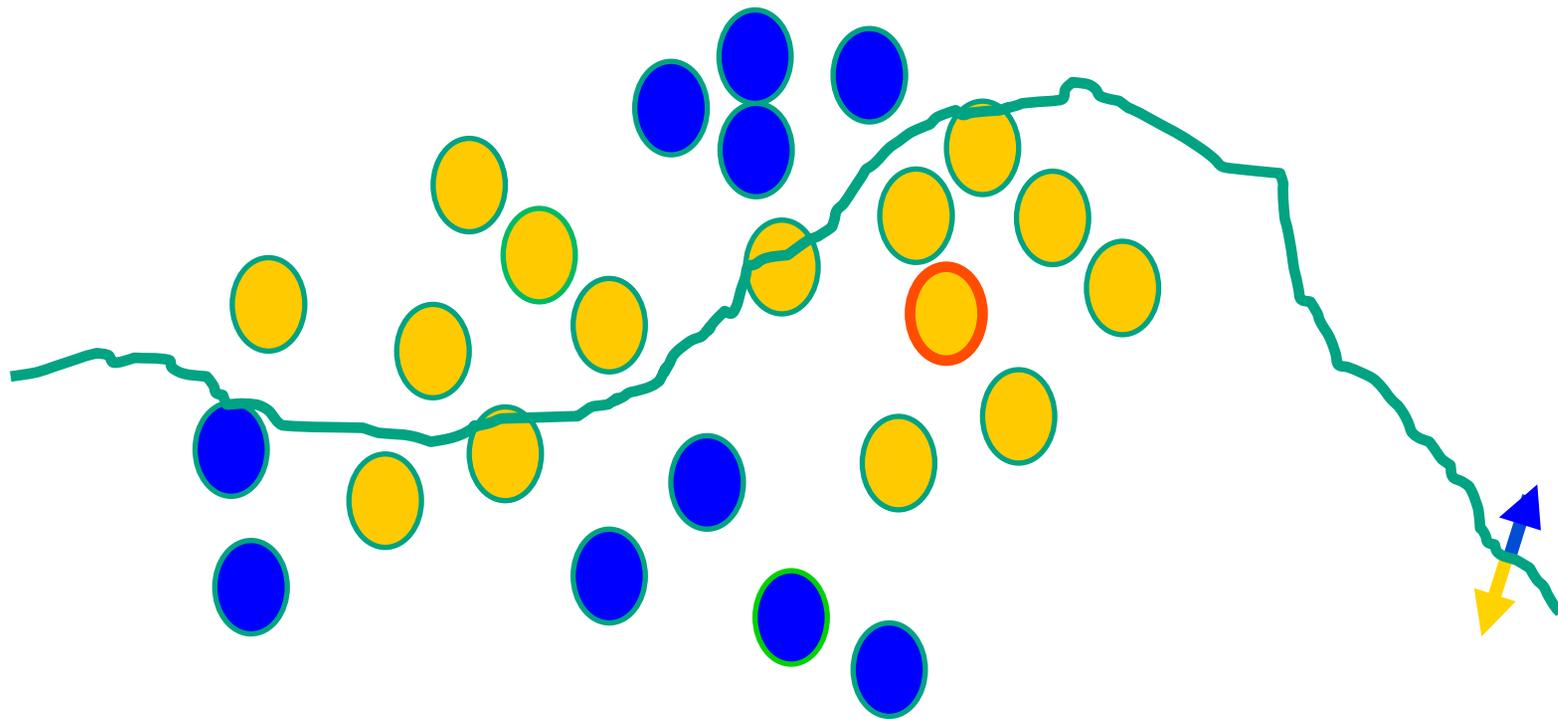
The Decision Boundary Perspective

Present a training instance / adjust the weights



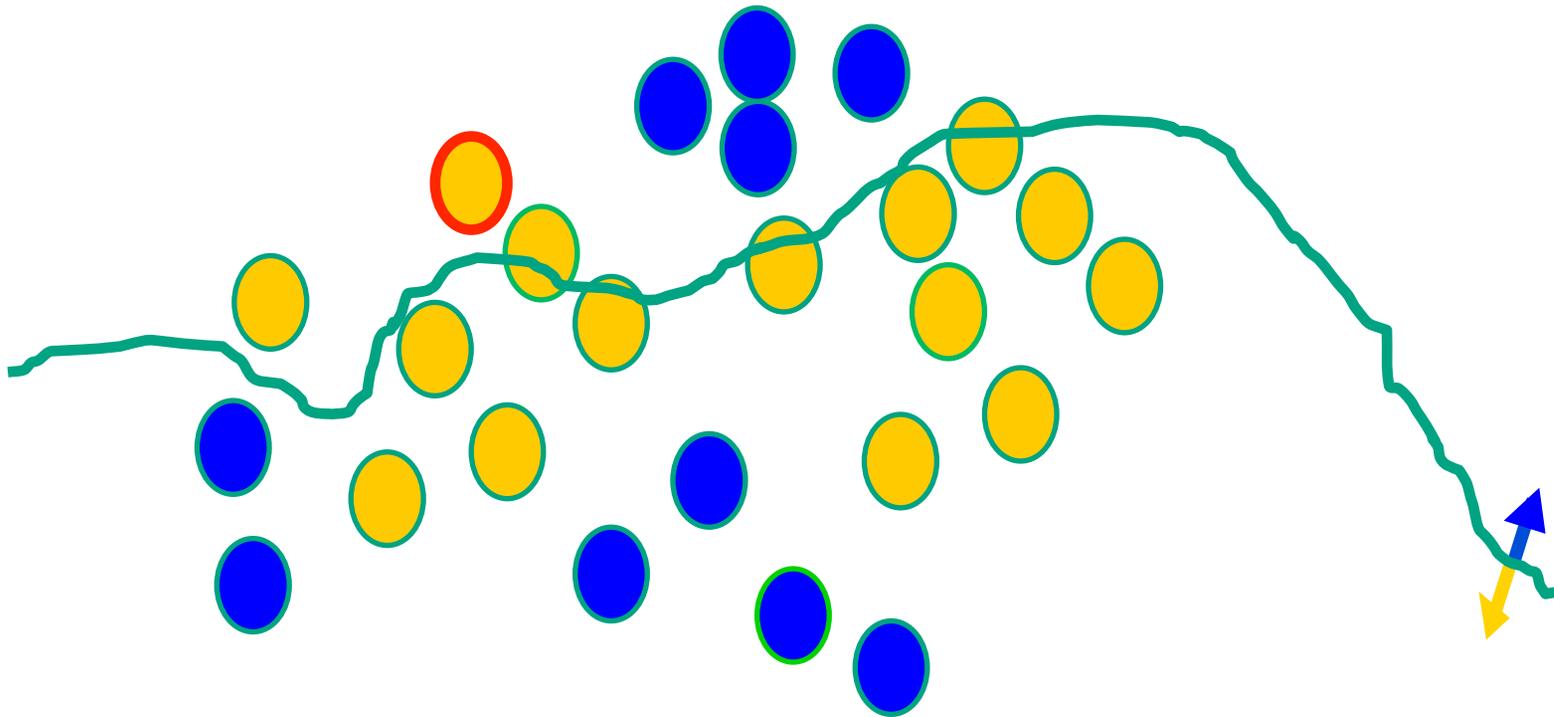
The Decision Boundary Perspective

Present a training instance / adjust the weights



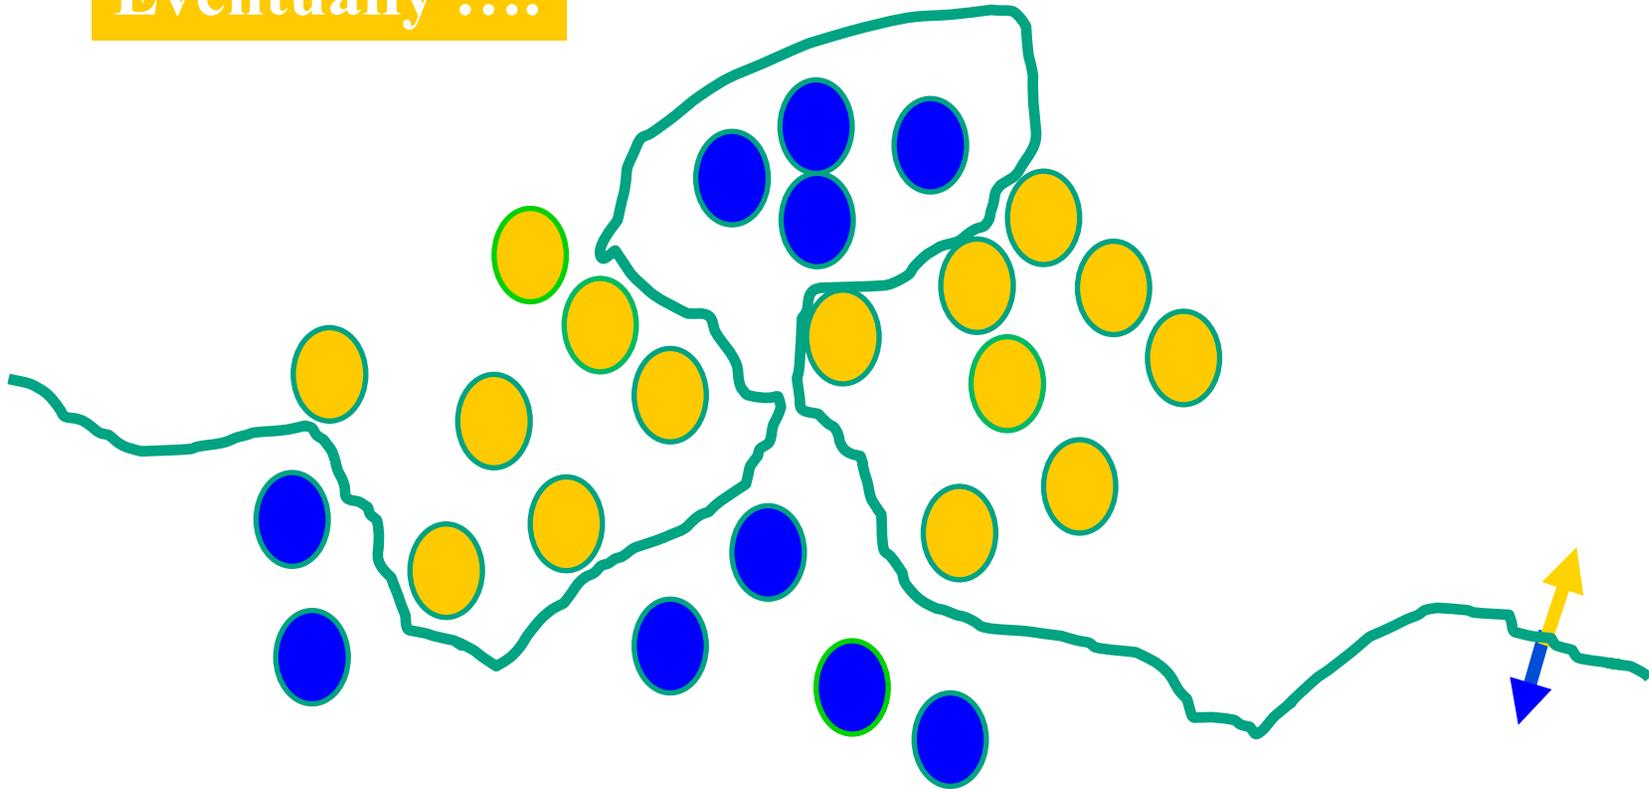
The Decision Boundary Perspective

Present a training instance / adjust the weights



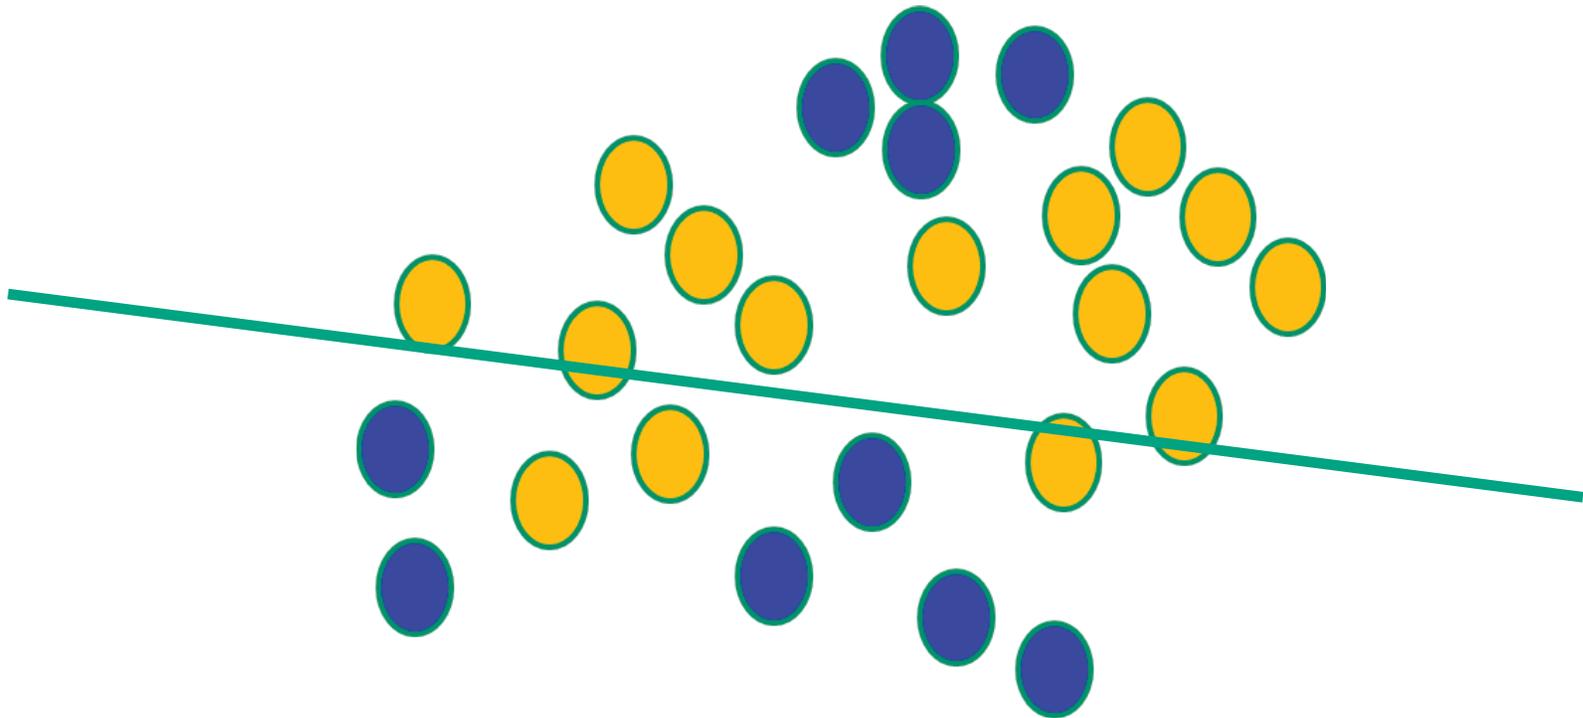
The Decision Boundary Perspective

Eventually

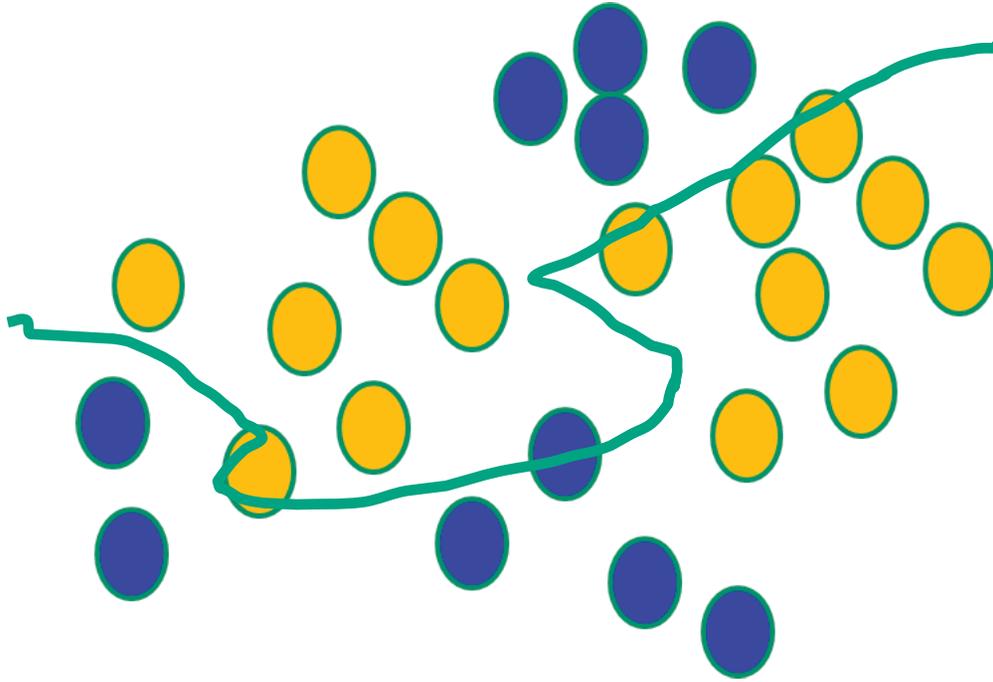


Linear Activation Functions

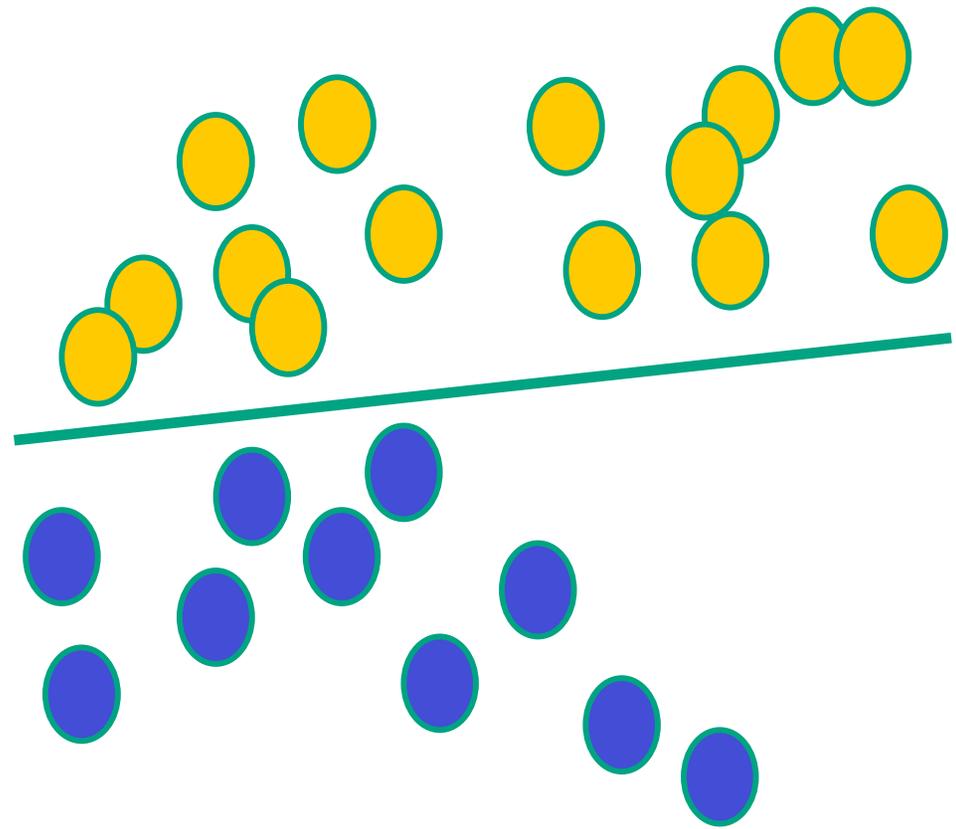
If $f(x)$ is linear, the NN can *only* draw straight decision boundaries (even if there are many layers of units)



NNs use nonlinear $f(x)$ so they can draw complex boundaries, but keep the data unchanged



SVMs only draw straight lines, but they transform the data first in a way that makes that OK



Limitations of Neural Networks

Random initialization + densely connected networks lead to:

- High cost
 - Each neuron in the neural network can be considered as a logistic regression.
 - Training the entire neural network is to train all the interconnected logistic regressions.
- Difficult to train as the number of hidden layers increases
 - Recall that logistic regression is trained by gradient descent.
 - In backpropagation, gradient is progressively getting more dilute. That is, below top layers, the correction signal δn is minimal.
- Stuck in local optima
 - The objective function of the neural network is usually not convex.
 - The random initialization does not guarantee starting from the proximity of global optima.
- Solution
 - Deep Learning/Learning multiple levels of representation

What exactly is Deep Learning ?

Why is it generally better than other methods on image, speech and certain other types of data?

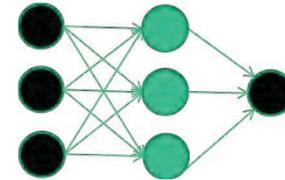
The short answer:

“Deep Learning” means using a **neural network** with **several layers of nodes** between input and output

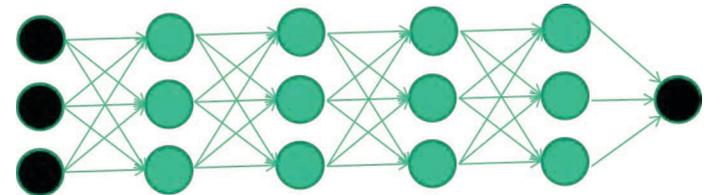
The series of layers between input & output do **feature identification and processing in a series of stages**, just as our brains seem to.

Multi-layer neural networks have been around for about 25 years... So, what's actually new?

We have always had good algorithms for learning the weights in networks with 1 hidden layer

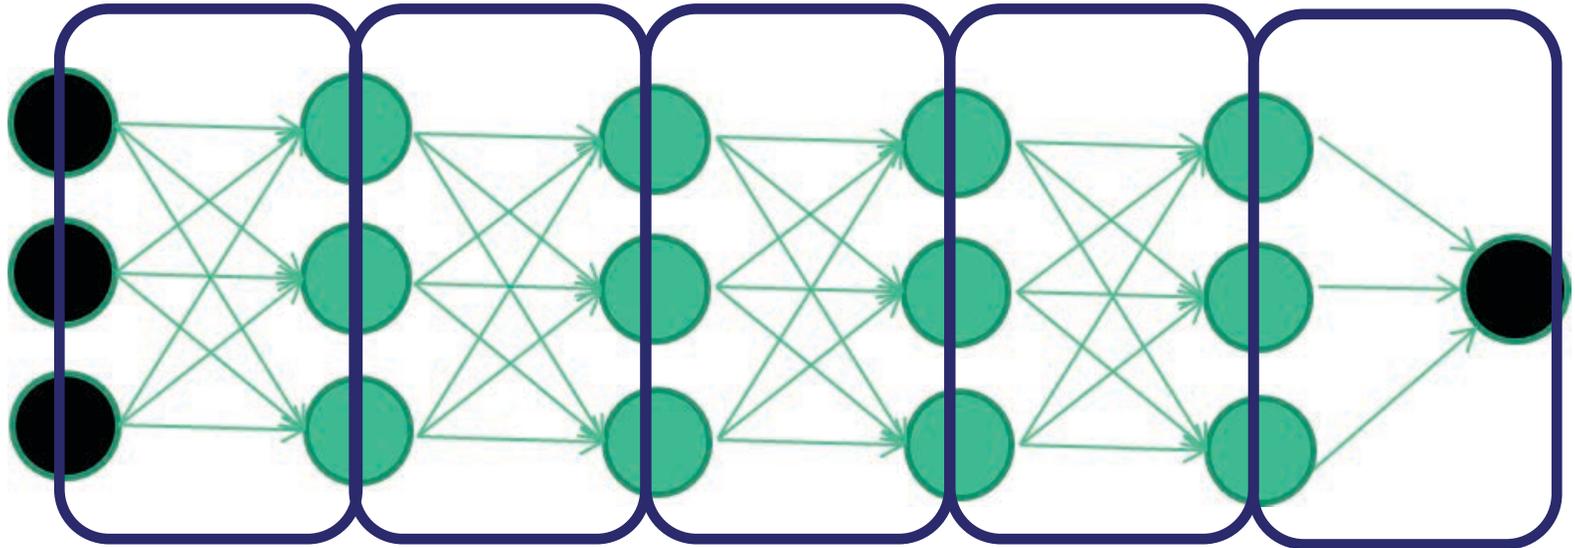


But these algorithms are not good at learning the weights for networks with more hidden layers



What's new is: algorithms for training many-layer networks

How to Train a Multi-Layer Network



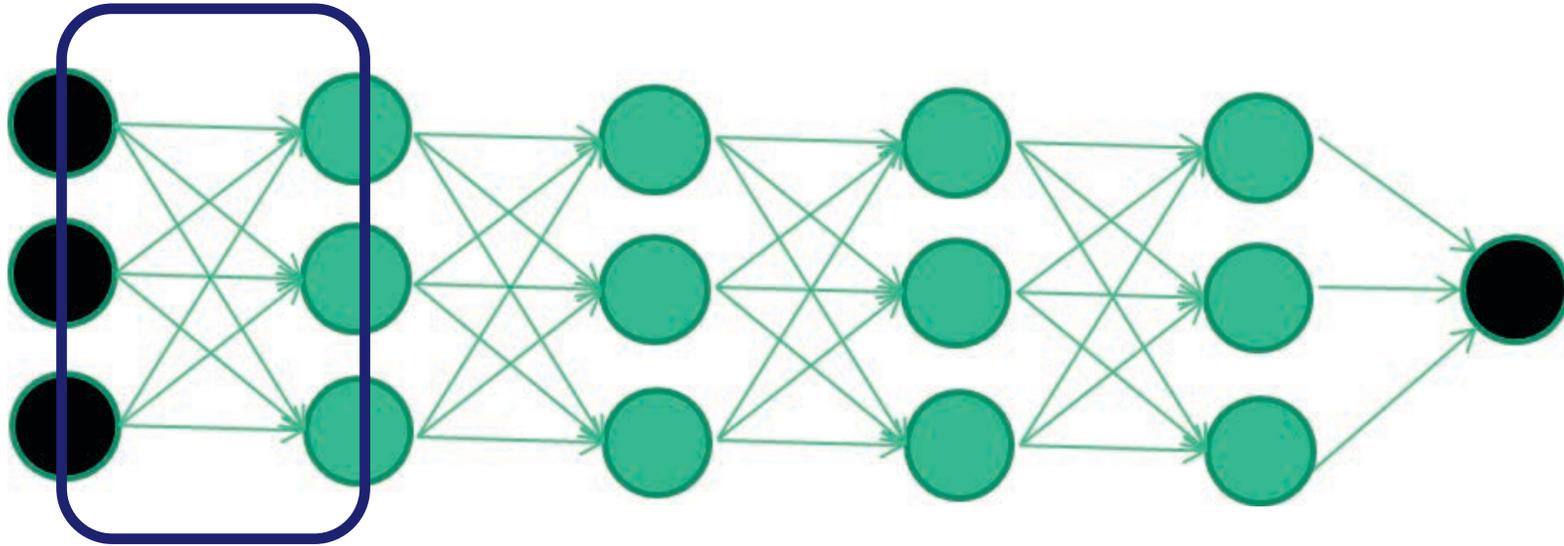
Train **this** layer first

then **this** layer

then **this** layer

then **this** layer

finally **this** layer



*EACH of the (non-output) layers is trained to be an **auto-encoder**.*

Basically, it is forced to learn good features that describe what comes from the previous layer.

Networks for Deep Learning

- ***Deep Belief Networks*** and ***Autoencoders*** employs layer-wise unsupervised learning to initialize each layer and capture multiple levels of representation simultaneously.
 - Hinton, G. E, Osindero, S., and Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18:1527-1554.
 - Bengio, Y., Lamblin, P., Popovici, P., Larochelle, H. (2007). Greedy Layer-Wise Training of Deep Networks, *Advances in Neural Information Processing Systems* 19
- ***Convolutional Neural Network*** organizes neurons based on animal's visual cortex system, which allows for learning patterns at both local level and global level.
 - Y. LeCun, L. Bottou, Y. Bengio and P. Haffner: Gradient-Based Learning Applied to Document Recognition, *Proceedings of the IEEE*, 86(11):2278-2324, November 1998



Yann LeCun

(60, born in Paris,
now lives in NYC)

LeNet image recognition
inventor of backpropagation
methods for training, and of
convolutional neural nets
current director of Artificial
Intelligence at Facebook

Deep Belief Networks

- A *deep belief network* (DBN) is a probabilistic, generative model made up of multiple layers of hidden units.
 - A composition of simple learning modules that make up each layer
- A DBN can be used to generatively pre-train a DNN by using the learned DBN weights as the initial DNN weights.
 - Back-propagation or other discriminative algorithms can then be applied for fine-tuning of these weights.
- Advantages:
 - Particularly helpful when limited training data are available
 - These pre-trained weights are closer to the optimal weights than are randomly chosen initial weights.

Convolutional Neural Networks

- Convolutional Neural Networks are inspired by mammalian visual cortex.
 - The visual cortex contains a complex arrangement of cells, which are sensitive to small sub-regions of the visual field, called a receptive field. These cells act as local filters over the input space and are well-suited to exploit the strong spatially local correlation present in natural images.
 - Two basic cell types:
 - Simple cells respond maximally to specific edge-like patterns within their receptive field.
 - Complex cells have larger receptive fields and are locally invariant to the exact position of the pattern.

“Godfathers of AI” Awarded 2018 Turing Prize



Yann LeCun
(Facebook &
New York Univ.)



Geoff Hinton
(Google &
Univ. Toronto)



Yoshua Bengio
(Element AI &
Univ. Montreal)

Convolutional Neural Network for Image Classification



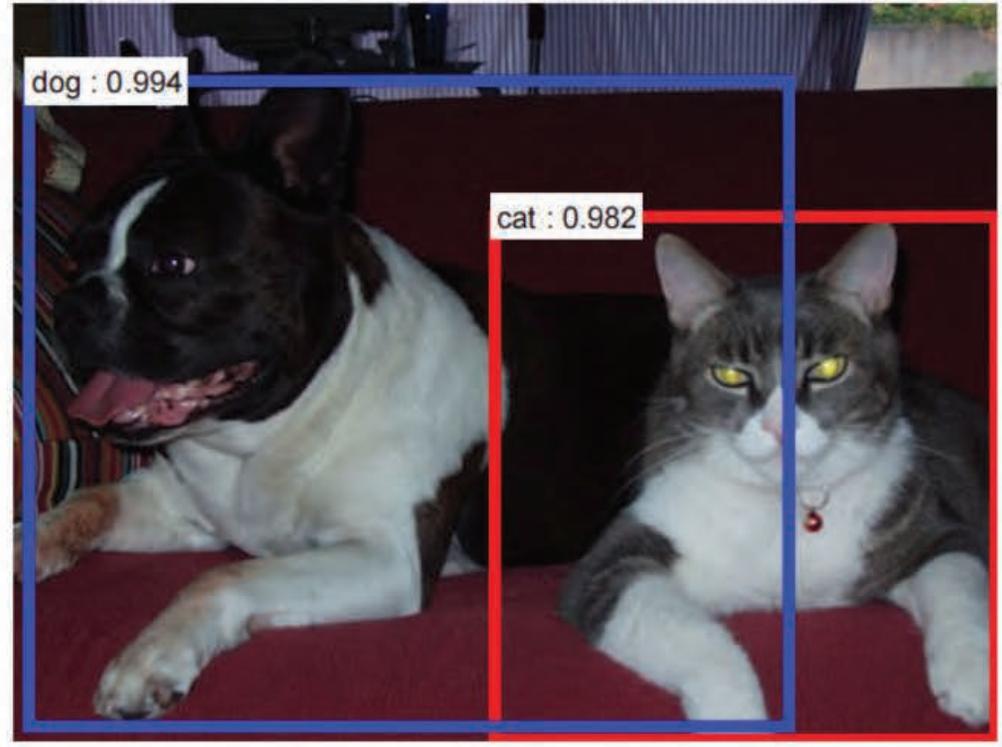
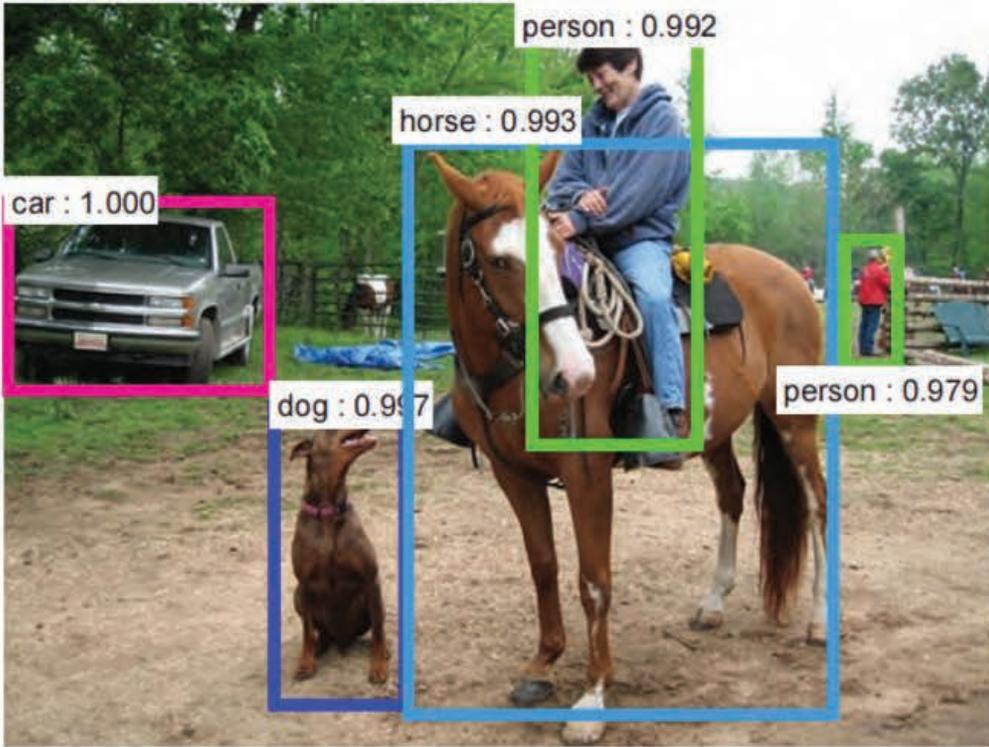
night bridge city suspension bridge river

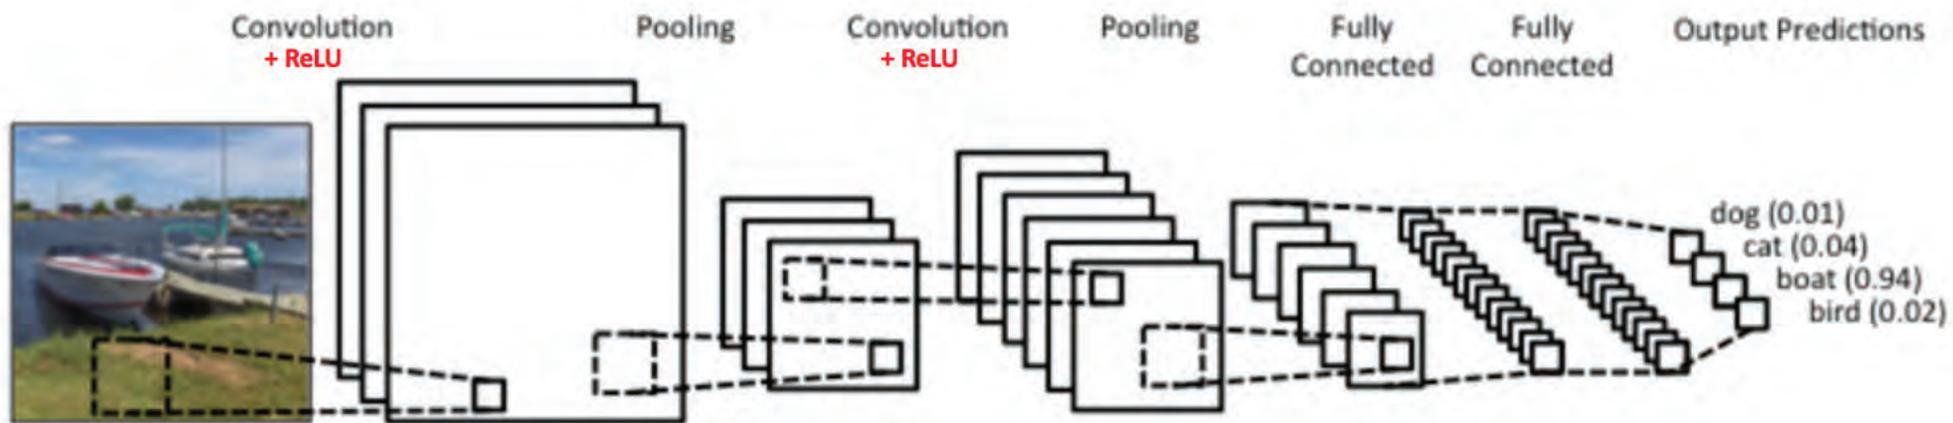


train subway railroad railway station transportation



competition tennis athlete stadium ball many spectators





| | | | | |
|---|---|---|---|---|
| 1 | 1 | 1 | 0 | 0 |
| 0 | 1 | 1 | 1 | 0 |
| 0 | 0 | 1 | 1 | 1 |
| 0 | 0 | 1 | 1 | 0 |
| 0 | 1 | 1 | 0 | 0 |

Image

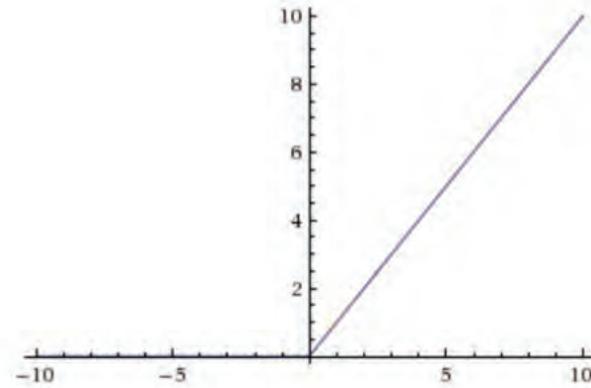
| | | |
|---|---|---|
| 1 | 0 | 1 |
| 0 | 1 | 0 |
| 1 | 0 | 1 |

Filter

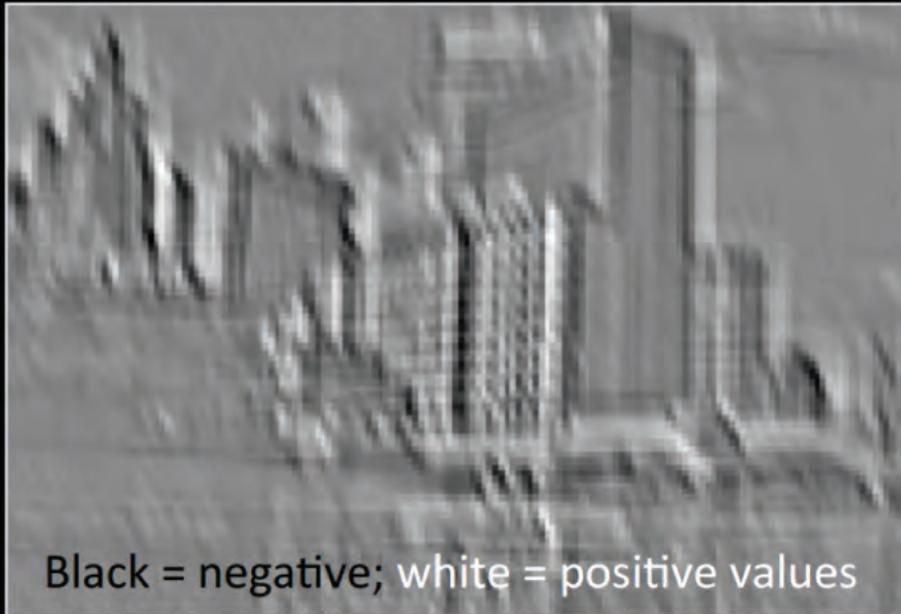
| Operation | Filter | Convolved Image |
|---|--|---|
| Identity | $\begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ |  |
| Edge detection | $\begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ -1 & 0 & 1 \end{bmatrix}$ |  |
| | $\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}$ |  |
| | $\begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}$ |  |
| Sharpen | $\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$ |  |
| Box blur (normalized) | $\frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$ |  |
| Gaussian blur (approximation) | $\frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$ |  |

The ReLU (Rectified Linear Unit) Operation

Output = Max(zero, Input)



Input Feature Map

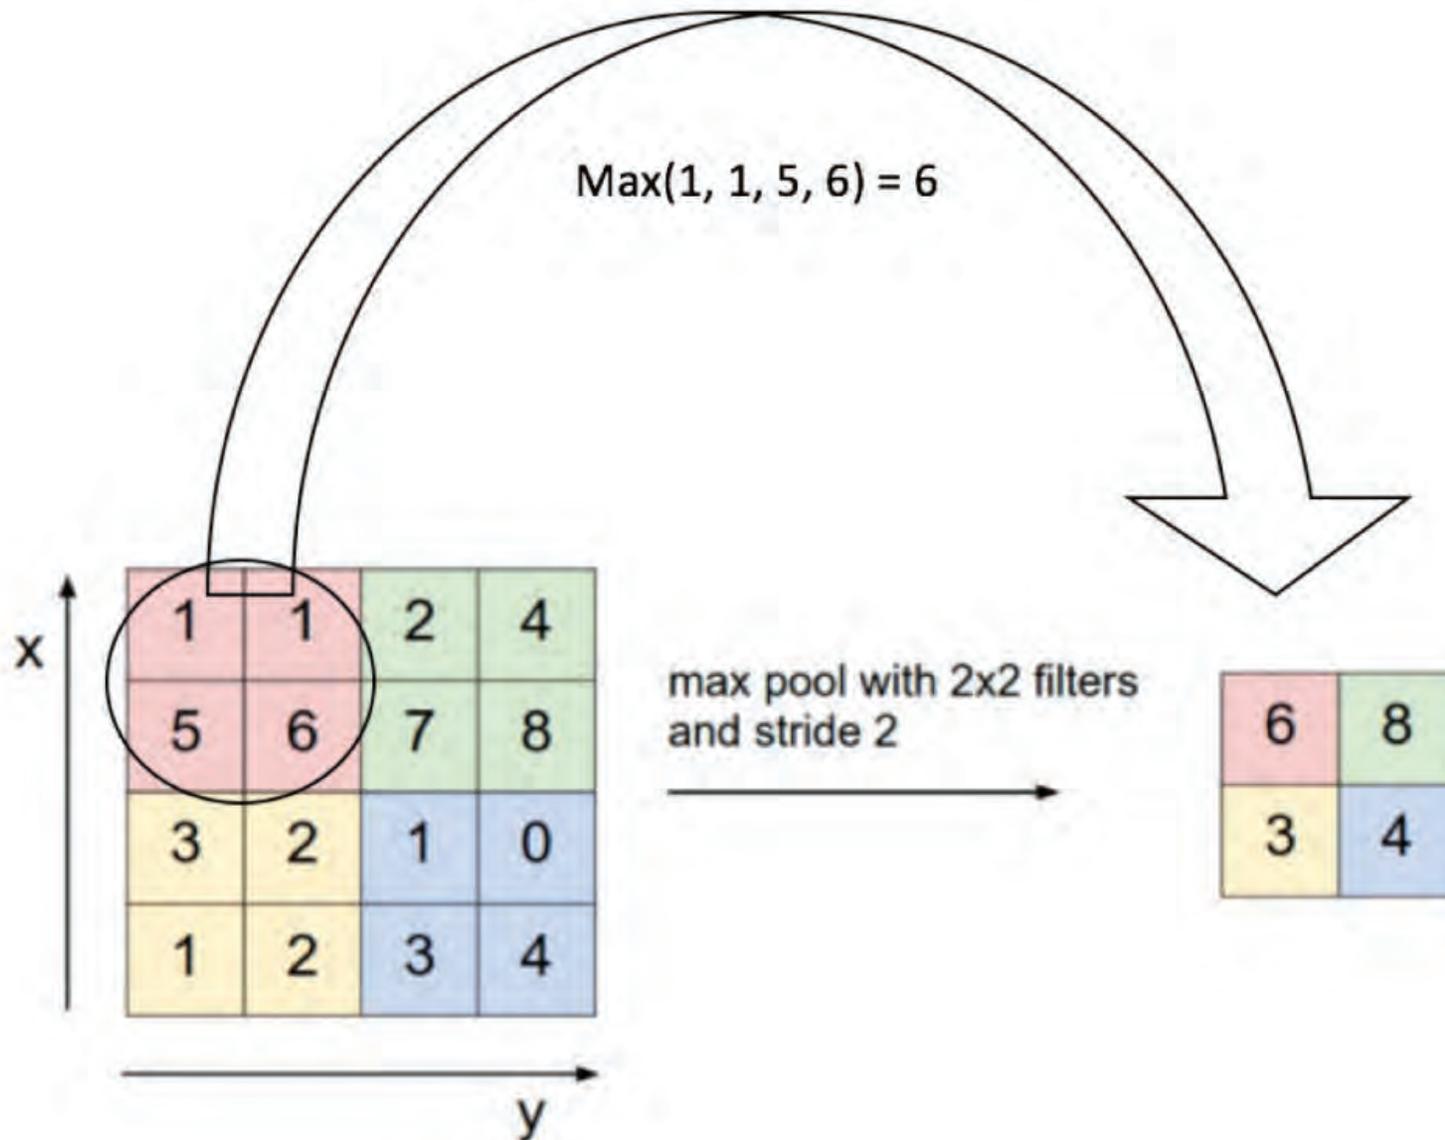


ReLU
→

Rectified Feature Map

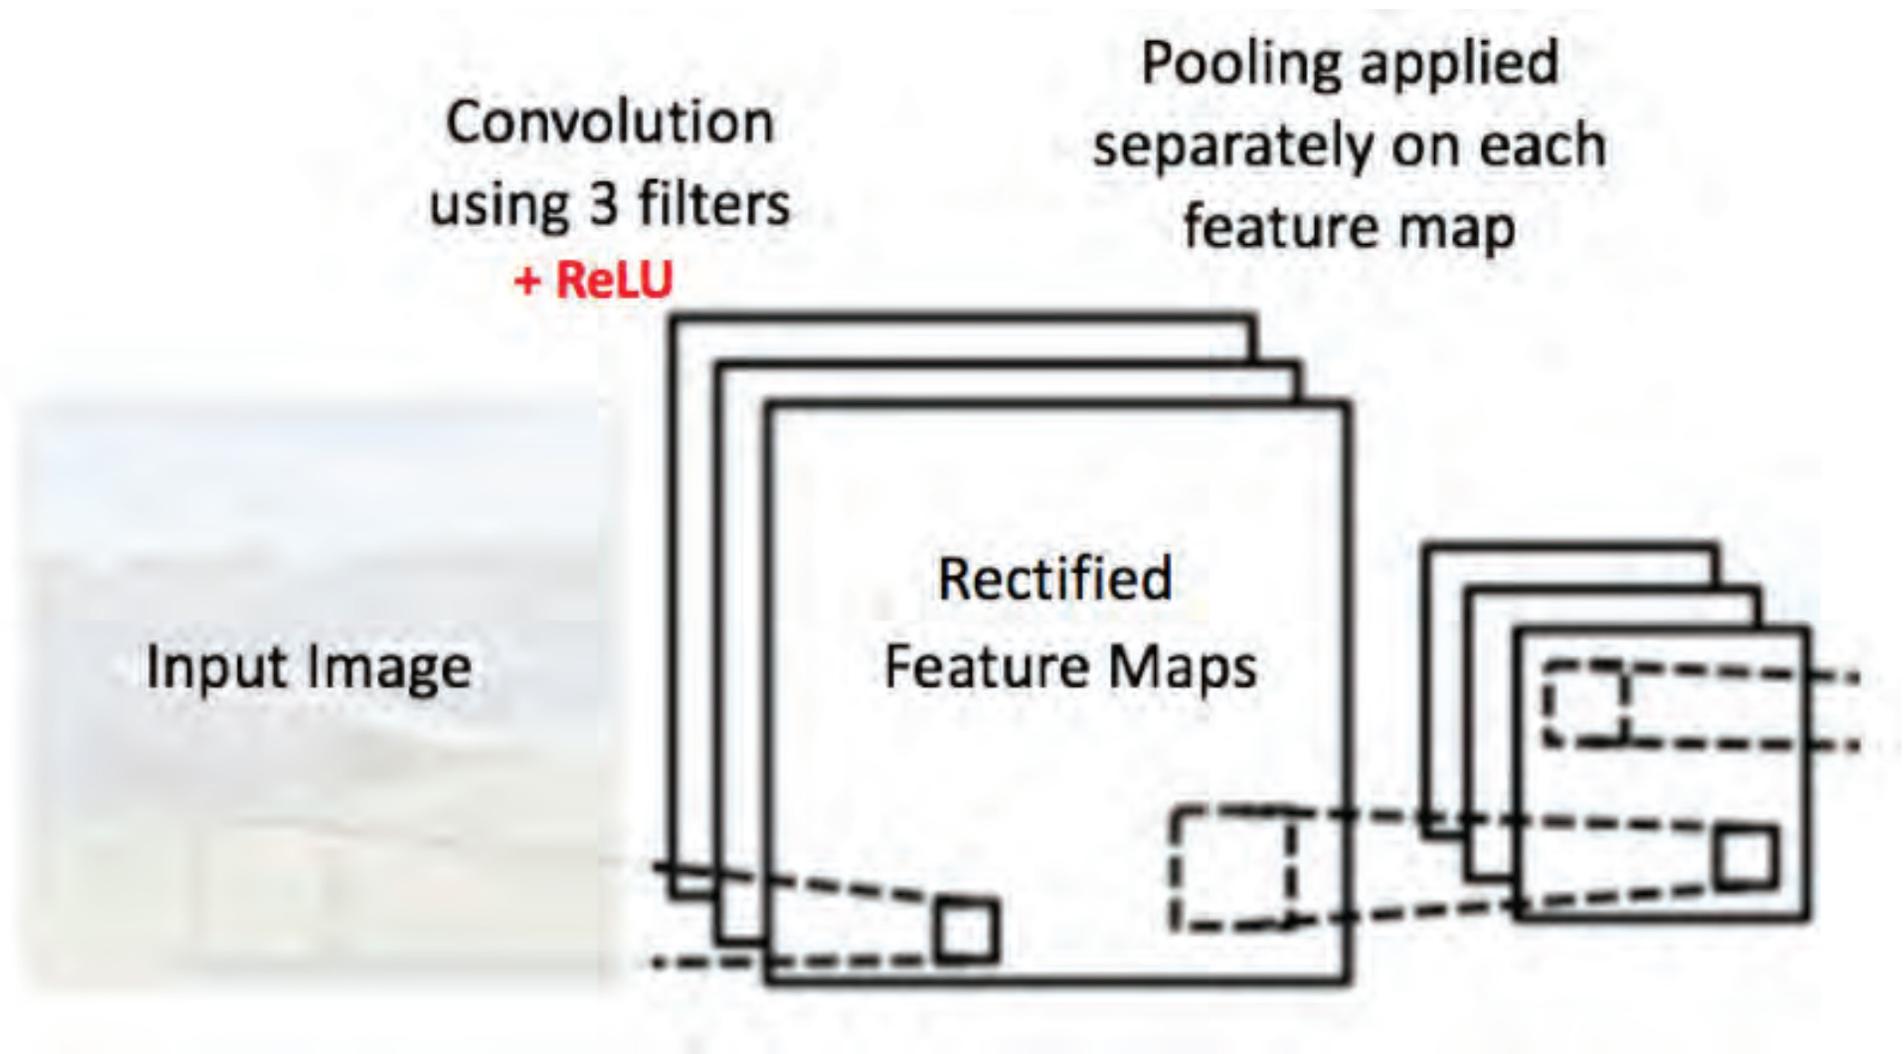


The Max Pooling Operation

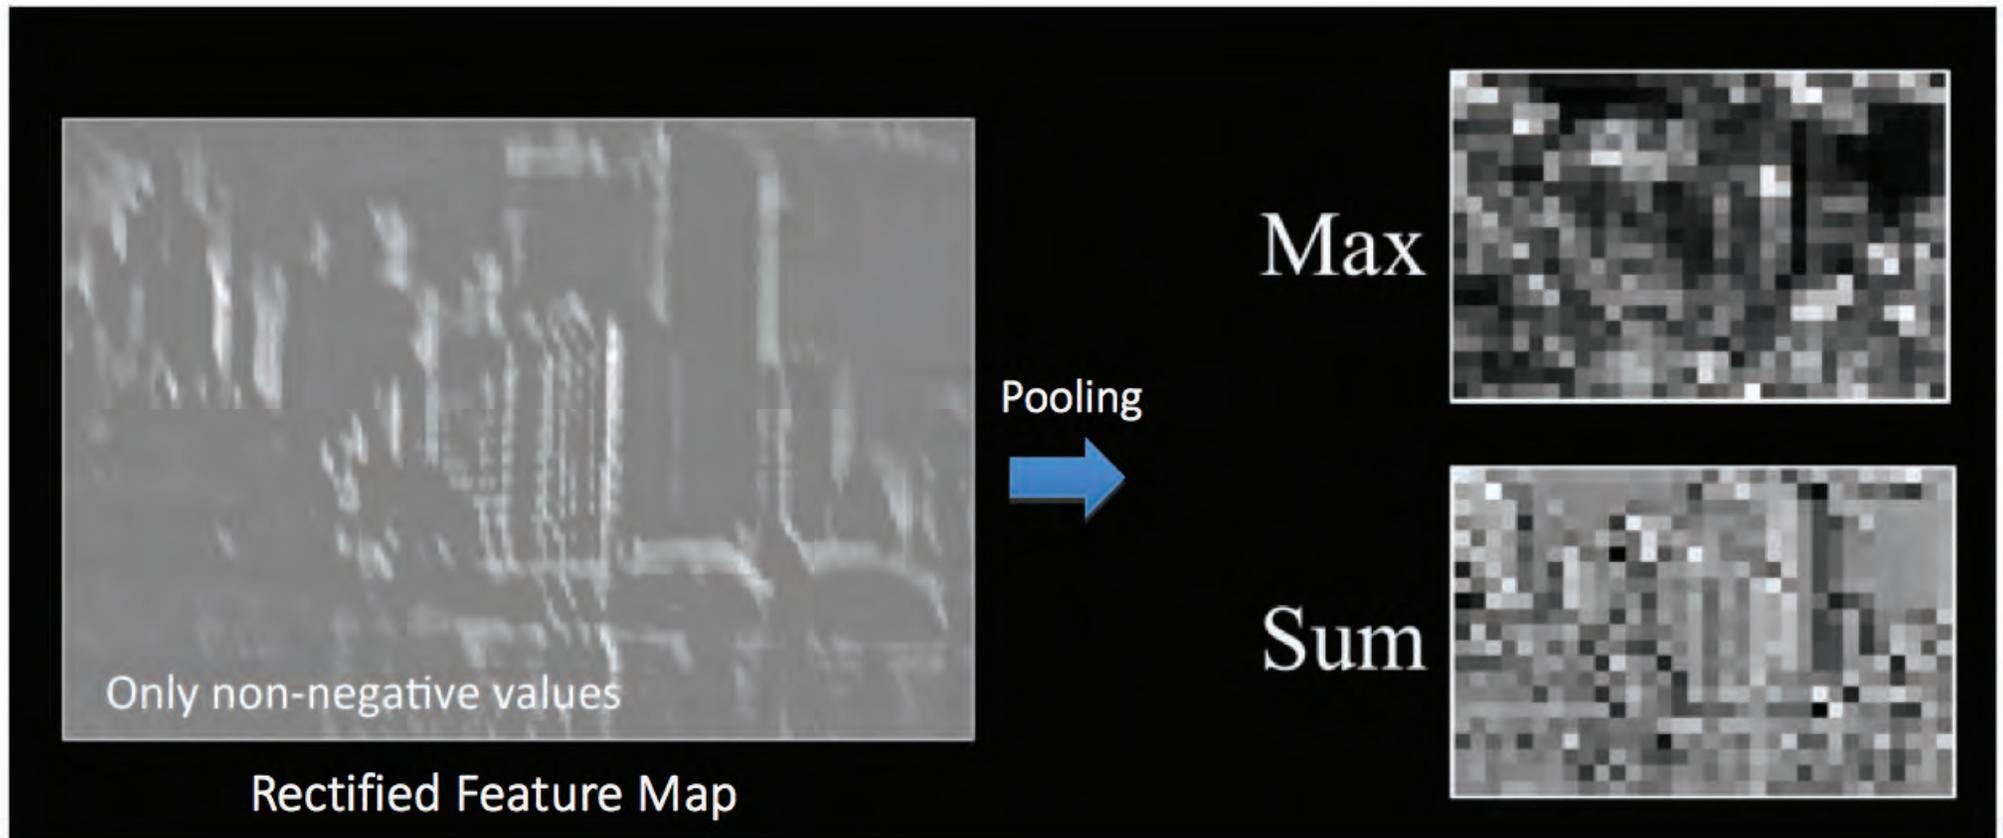


Rectified Feature Map

Pooling Applied to Rectified Feature Maps



Pooling Applied to Rectified Feature Maps



Training of a Convolutional Neural Net

Step 1: Initialize all filters and parameters / weights with random values.

Step 2: The network takes a training image as input, goes through the forward propagation step (convolution, ReLU and pooling operations along with forward propagation in the Fully Connected layer) and finds the output probabilities for each class.

Lets say the output probabilities for the boat image above are [0.2, 0.4, 0.1, 0.3]

Since weights are randomly assigned for the first training example, output probabilities are also random.

Step 3: Calculate the total error at the output layer (summation over all 4 classes)

$$\text{Total Error} = \sum \frac{1}{2} (\text{target probability} - \text{output probability})^2$$

Step 4: Use Backpropagation to calculate the gradients of the error with respect to all weights in the network and use gradient descent to update all filter values / weights and parameter values to minimize the output error.

The weights are adjusted in proportion to their contribution to the total error.

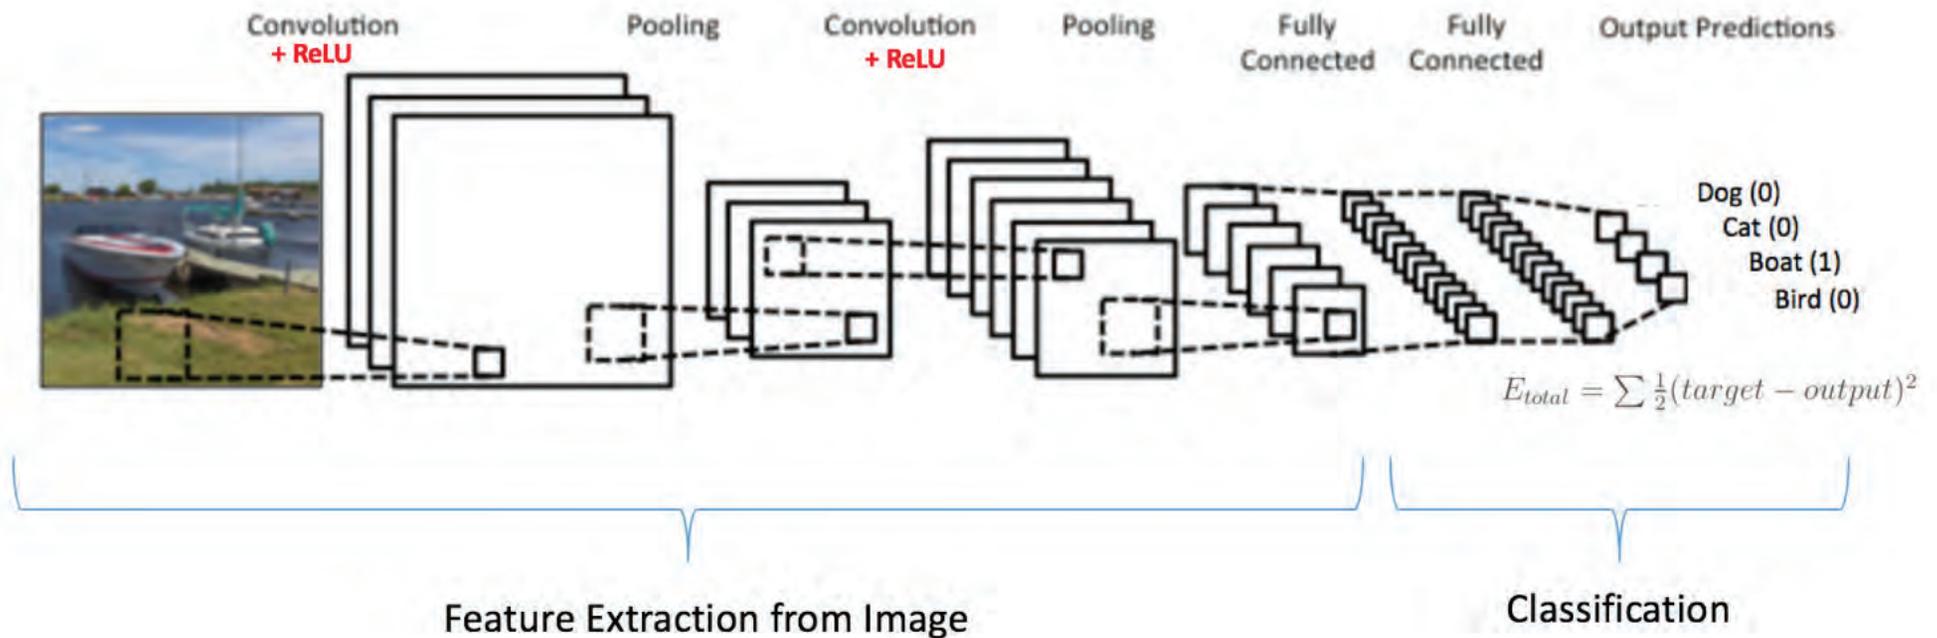
When the same image is input again, output probabilities might now be [0.1, 0.1, 0.7, 0.1], which is closer to the target vector [0, 0, 1, 0].

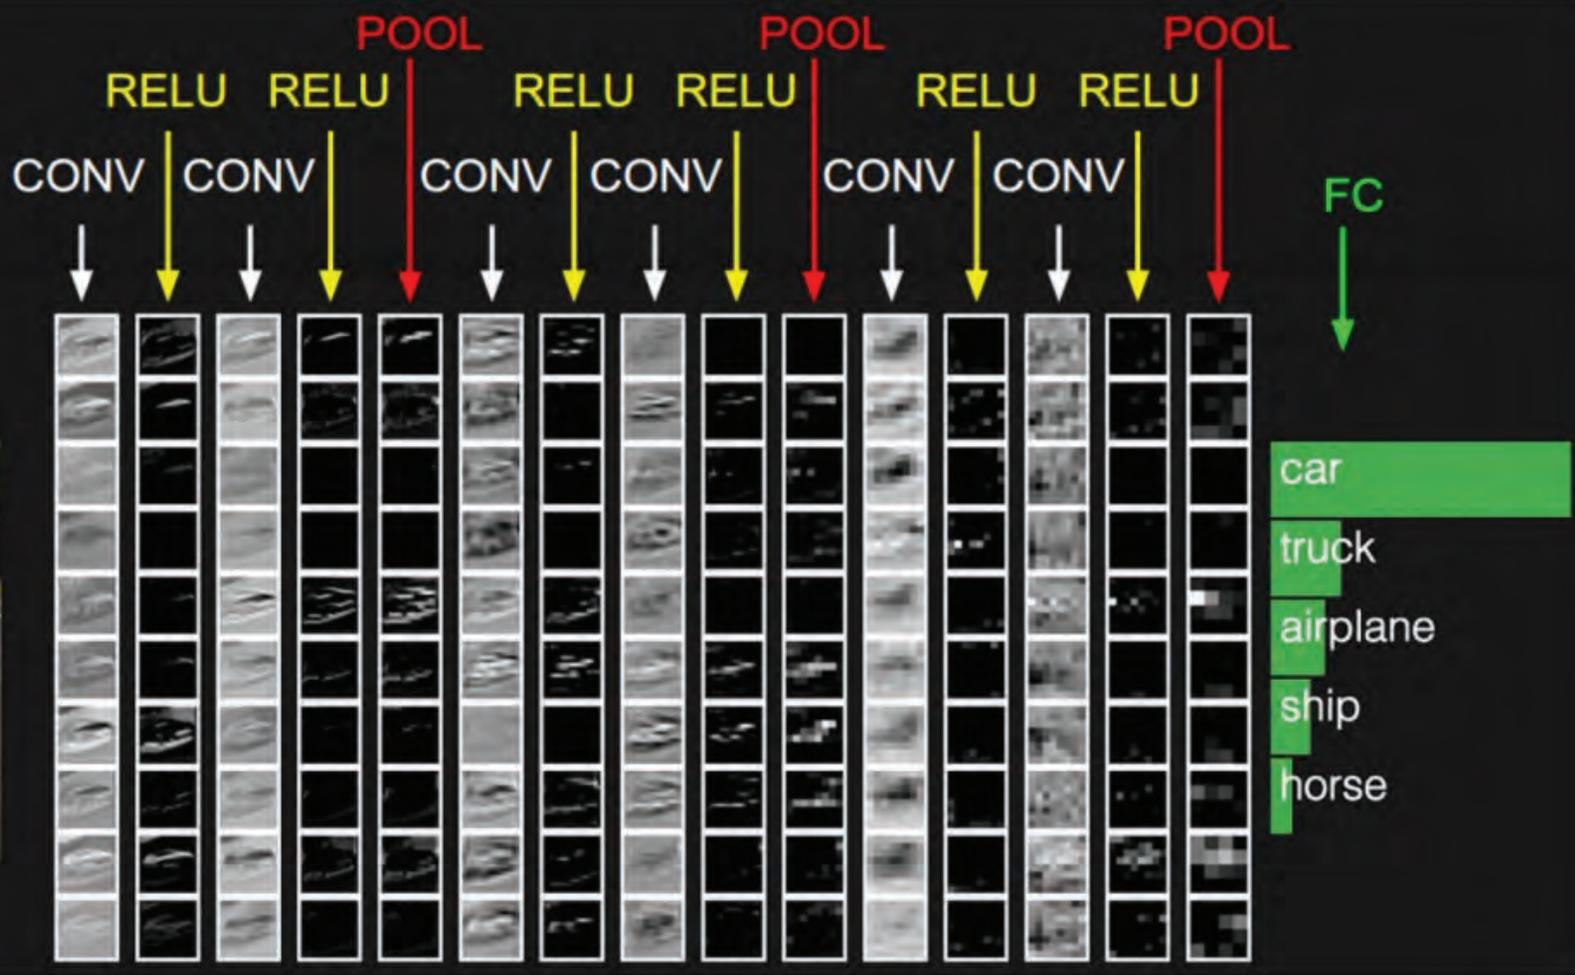
This means that the network has learnt to classify this particular image correctly by adjusting its weights / filters such that the output error is reduced.

Parameters like number of filters, filter sizes, architecture of the network etc. have all been fixed before Step 1 and do not change during training process – only the values of the filter matrix and connection weights get updated.

Step 5: Repeat steps 2-4 with all images in the training set.

Convolutional Neural Nets: Putting It All Together





0123456789

Output Layer

FC Layer 2

FC Layer 1

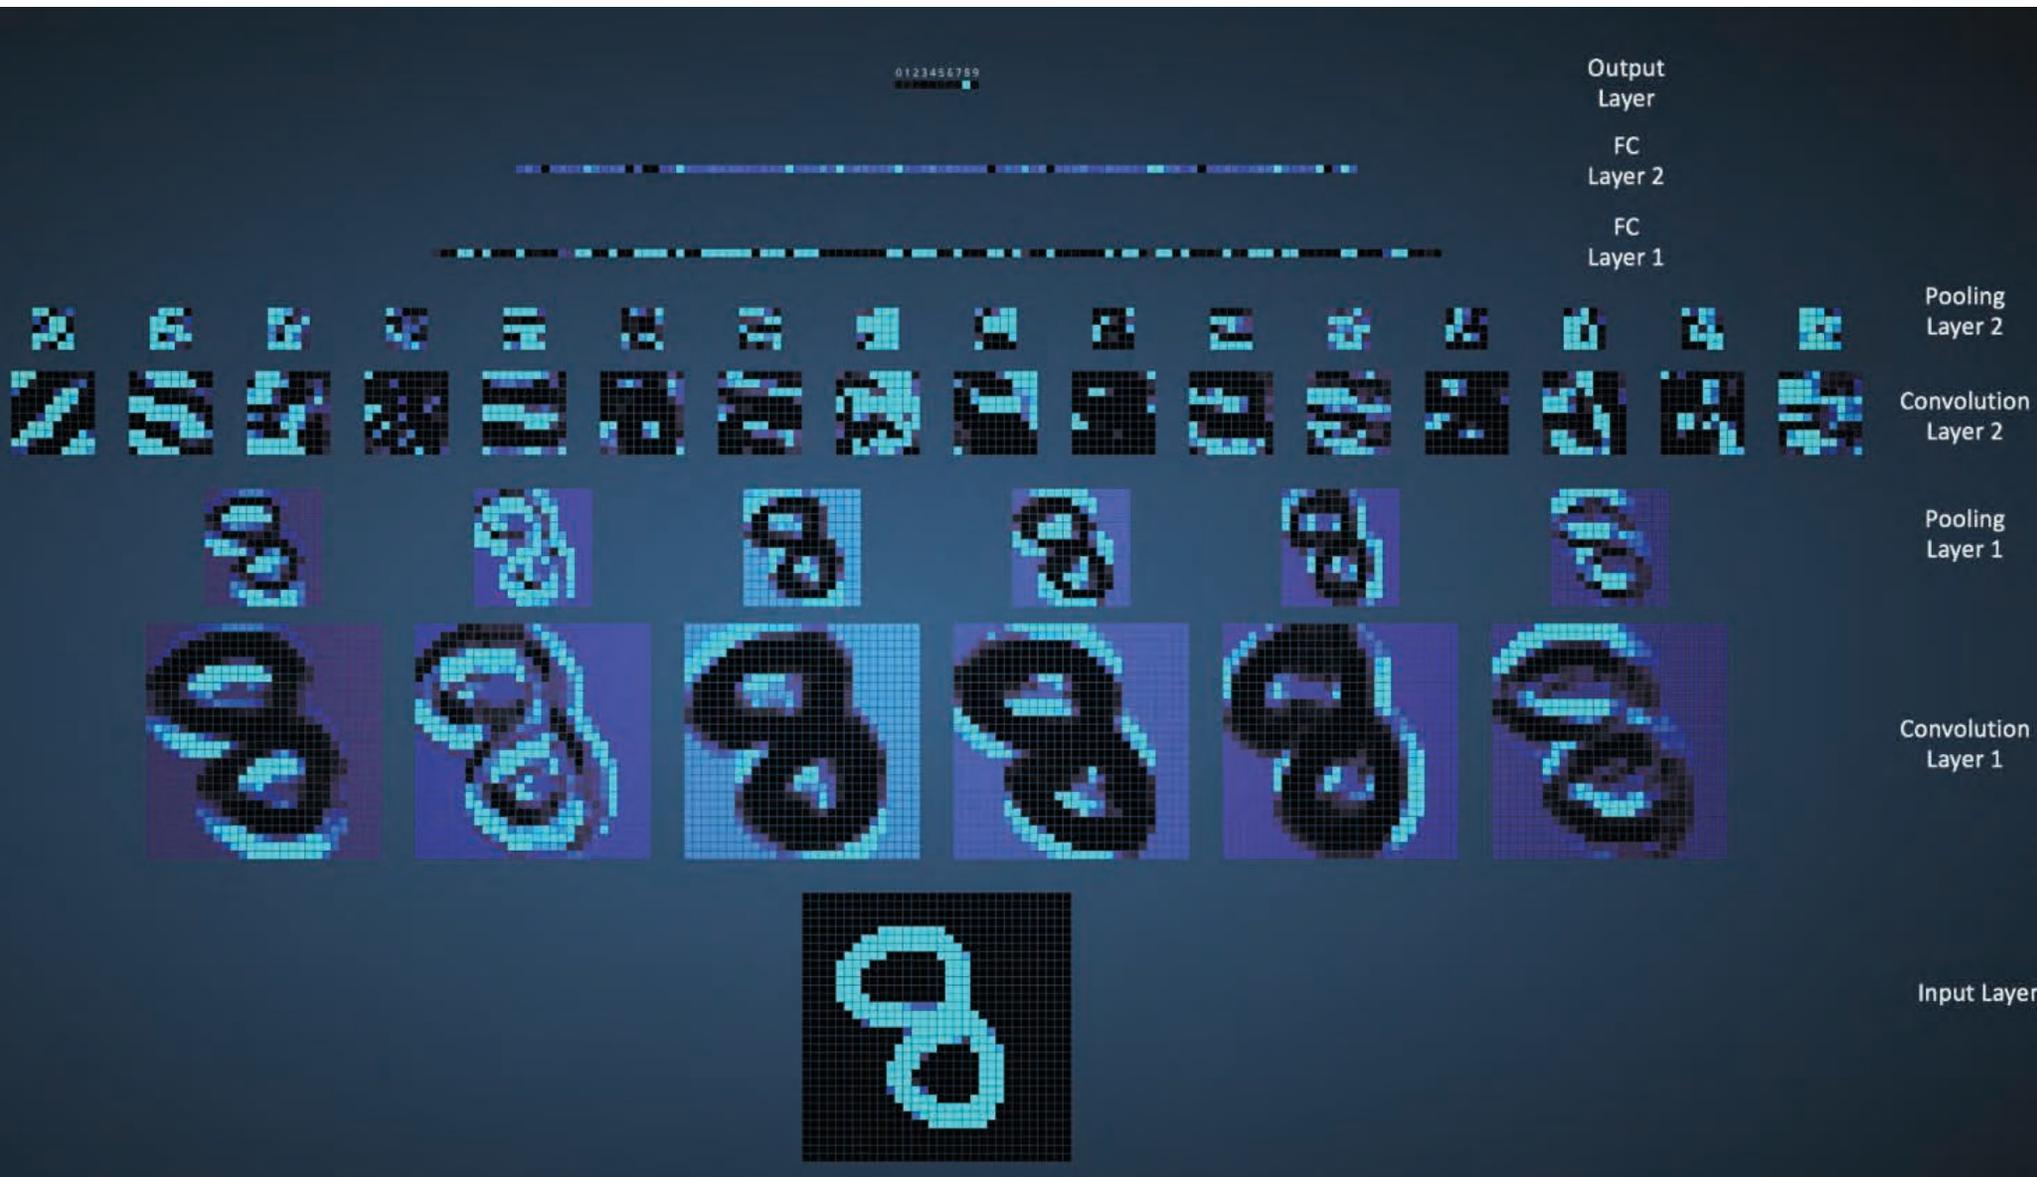
Pooling Layer 2

Convolution Layer 2

Pooling Layer 1

Convolution Layer 1

Input Layer



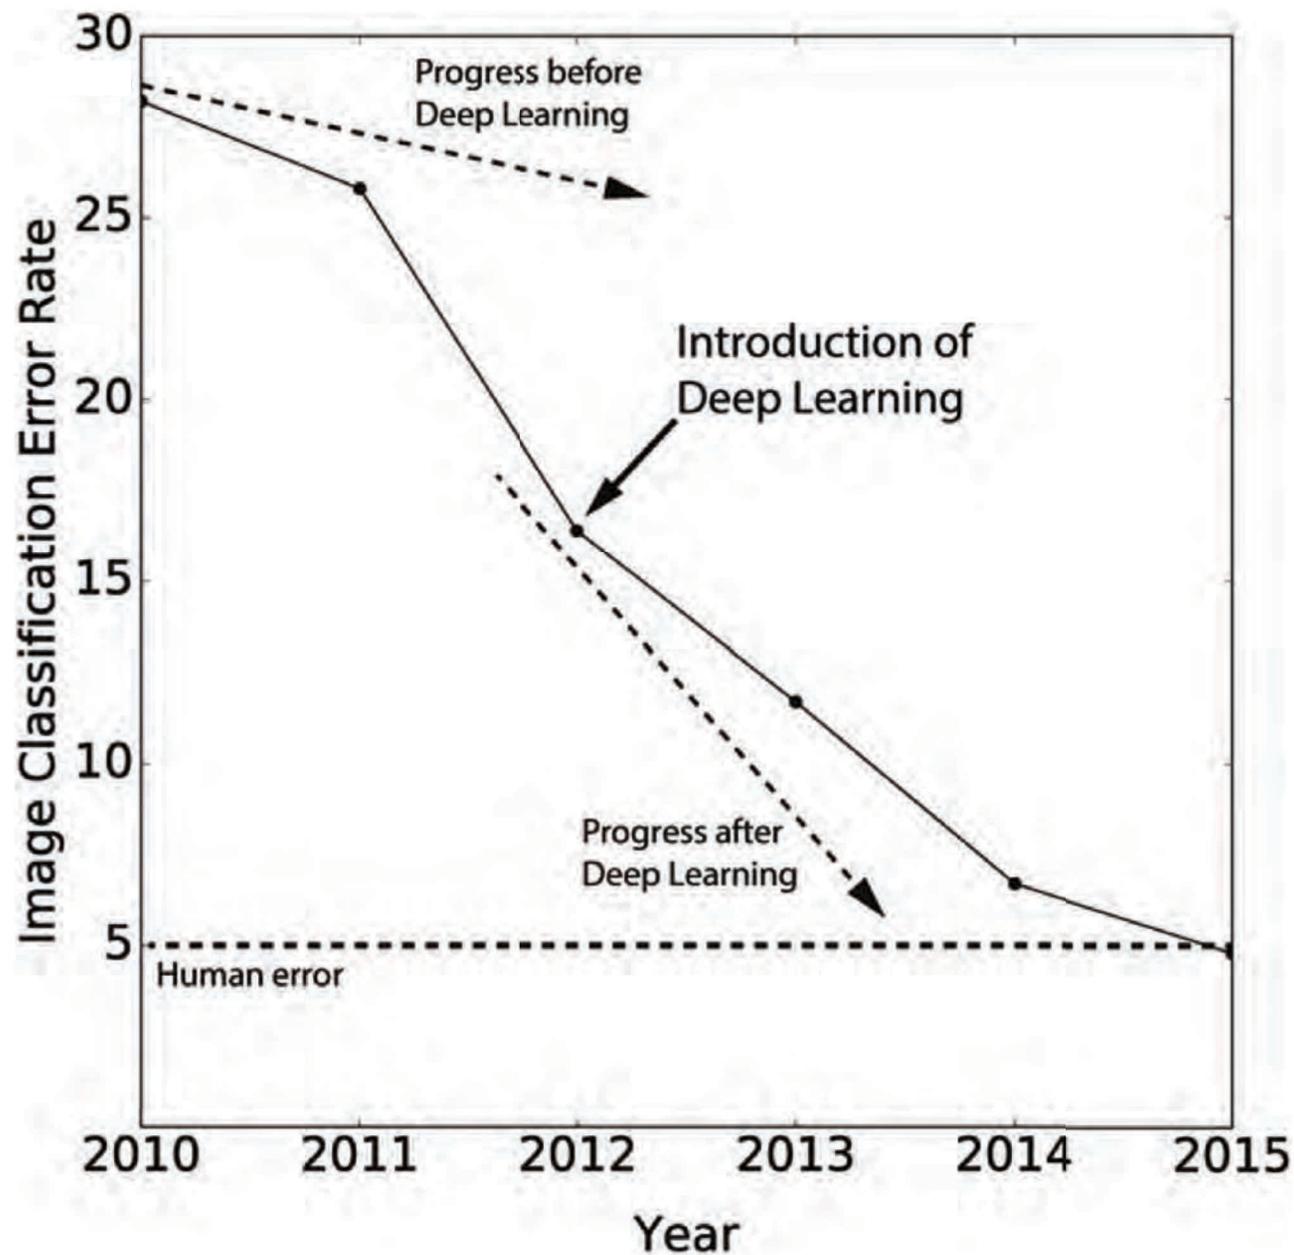


Figure 4. Historical error rate of the best performing image classification algorithms in the annual ImageNet competition.^[37] Established models of computer vision stagnated at 25-30%. The introduction of deep learning in 2012 led to a significant improvement to ~15%, and human level accuracy (~5%) for image classification was achieved by 2015.



TensorFlow is an open source library for machine learning tasks developed by Google and first released in November 2015.

It is a “second generation” system for machine learning, based on deep learning neural networks.

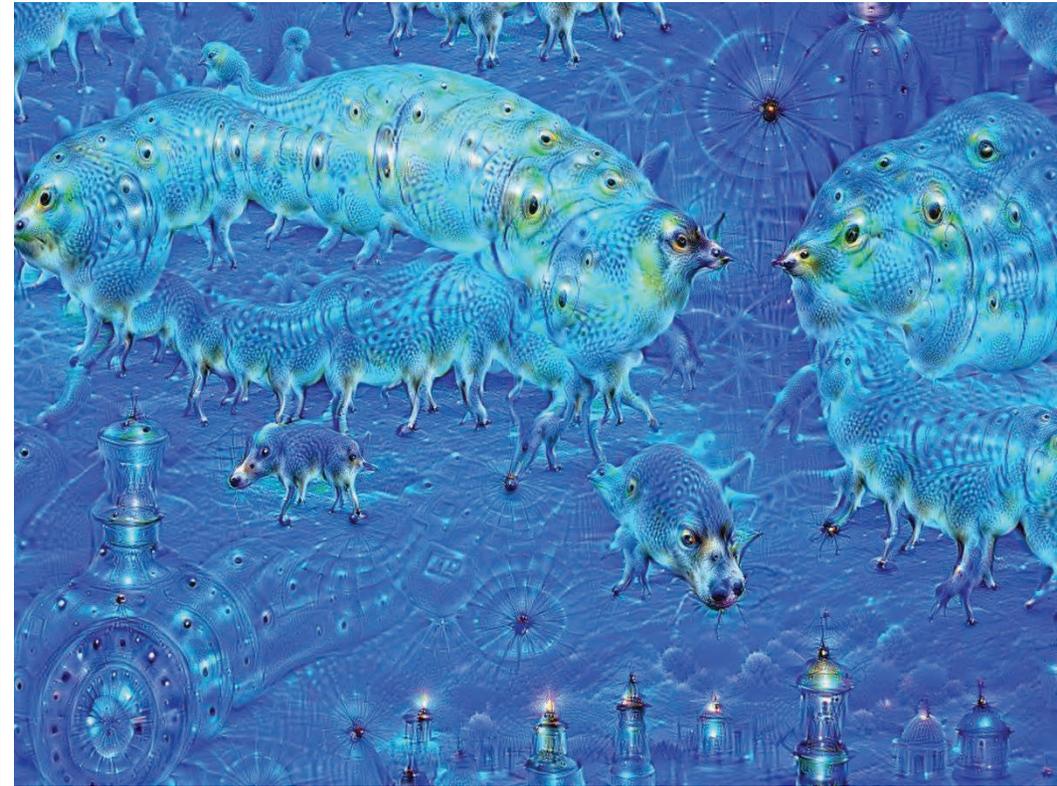
RackBrain now handles a large number of Google searches, and is powered by TensorFlow.

TensorFlow calculations are generally expressed as stateful dataflow graphs. Nodes in the graph represent mathematical operations, while edges are multidimensional arrays ("tensors") communicated between them. The name, TensorFlow, derives from the operations which neural networks perform on the arrays themselves.

DeepDream - Convolutional Neural Network

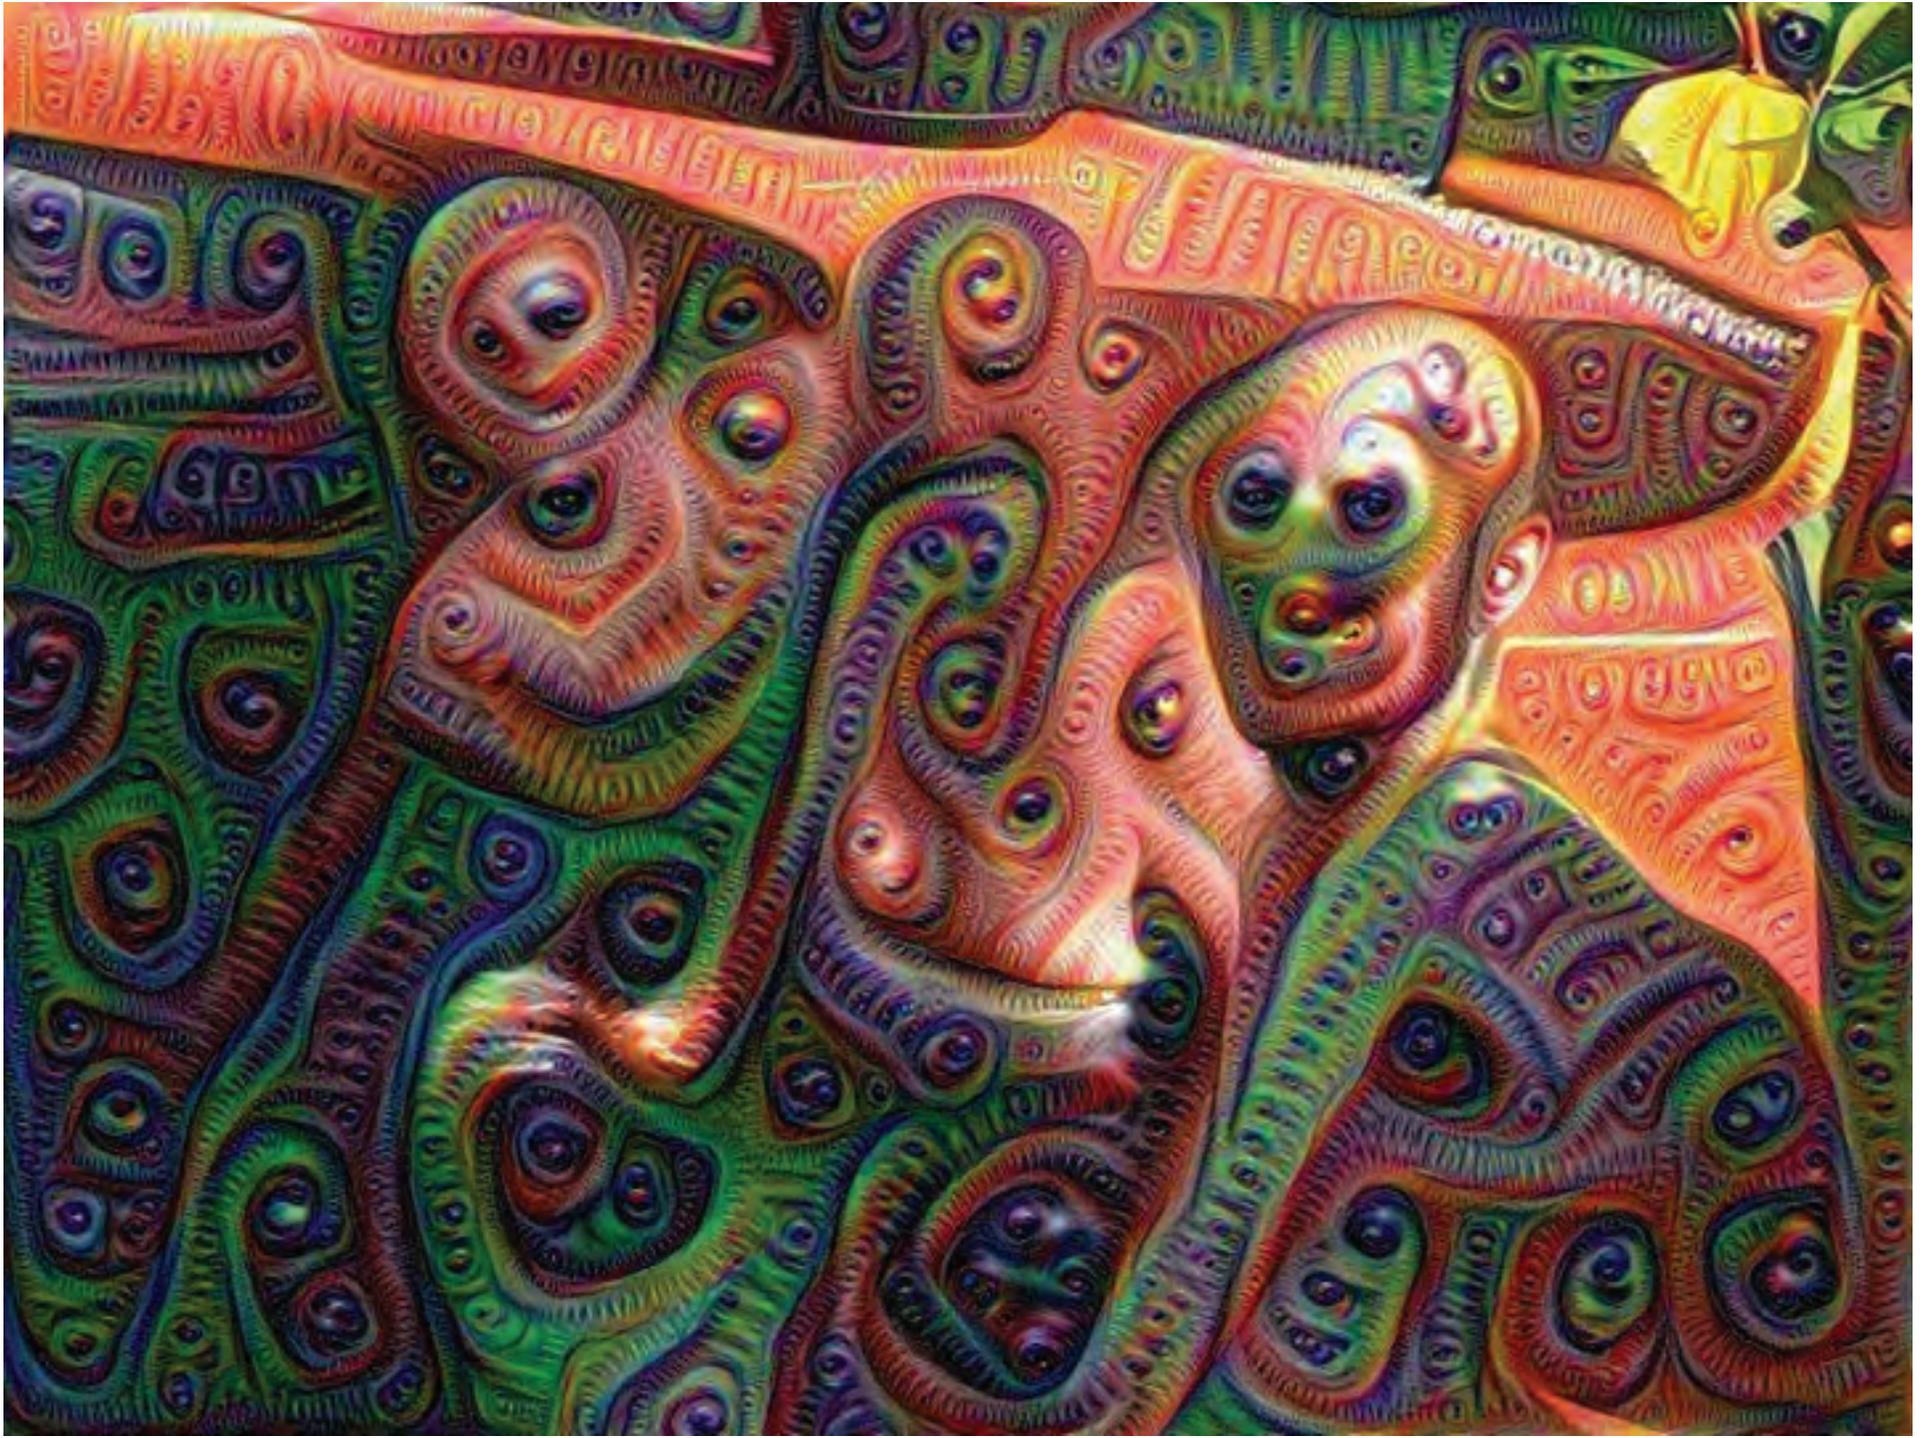


Original Image



**After 10 Iterations
of DeepDream**

Three Men in a Pool (DeepDream)



nature

THE INTERNATIONAL WEEKLY JOURNAL OF SCIENCE

At last — a computer program that
can beat a champion Go player **PAGE 484**

ALL SYSTEMS GO

CONSERVATION

SONGBIRDS À LA CARTE

Illegal harvest of millions
of Mediterranean birds

PAGE 452

RESEARCH ETHICS

SAFEGUARD TRANSPARENCY

Don't let openness backfire
on individuals

PAGE 459

POPULAR SCIENCE

WHEN GENES GOT 'SELFISH'

Dawkins's calling
card forty years on

PAGE 462

NATURE.COM/NATURE

28 January 2016 £10

Vol. 529, No. 7587



In *Nature*, 27 January 2016

- “DeepMind’s program AlphaGo beat Fan Hui, the European Go champion, five times out of five in tournament conditions...”
- “AlphaGo was not preprogrammed to play Go: rather, it learned using a general-purpose algorithm that allowed it to interpret the game’s patterns.”
- “...AlphaGo program applied **deep learning** in neural networks (convolutional NN) — brain-inspired programs in which connections between layers of simulated neurons are strengthened through examples and experience.”

WHEN A USER TAKES A PHOTO,
THE APP SHOULD CHECK WHETHER
THEY'RE IN A NATIONAL PARK...

SURE, EASY GIS LOOKUP.
GIMME A FEW HOURS.

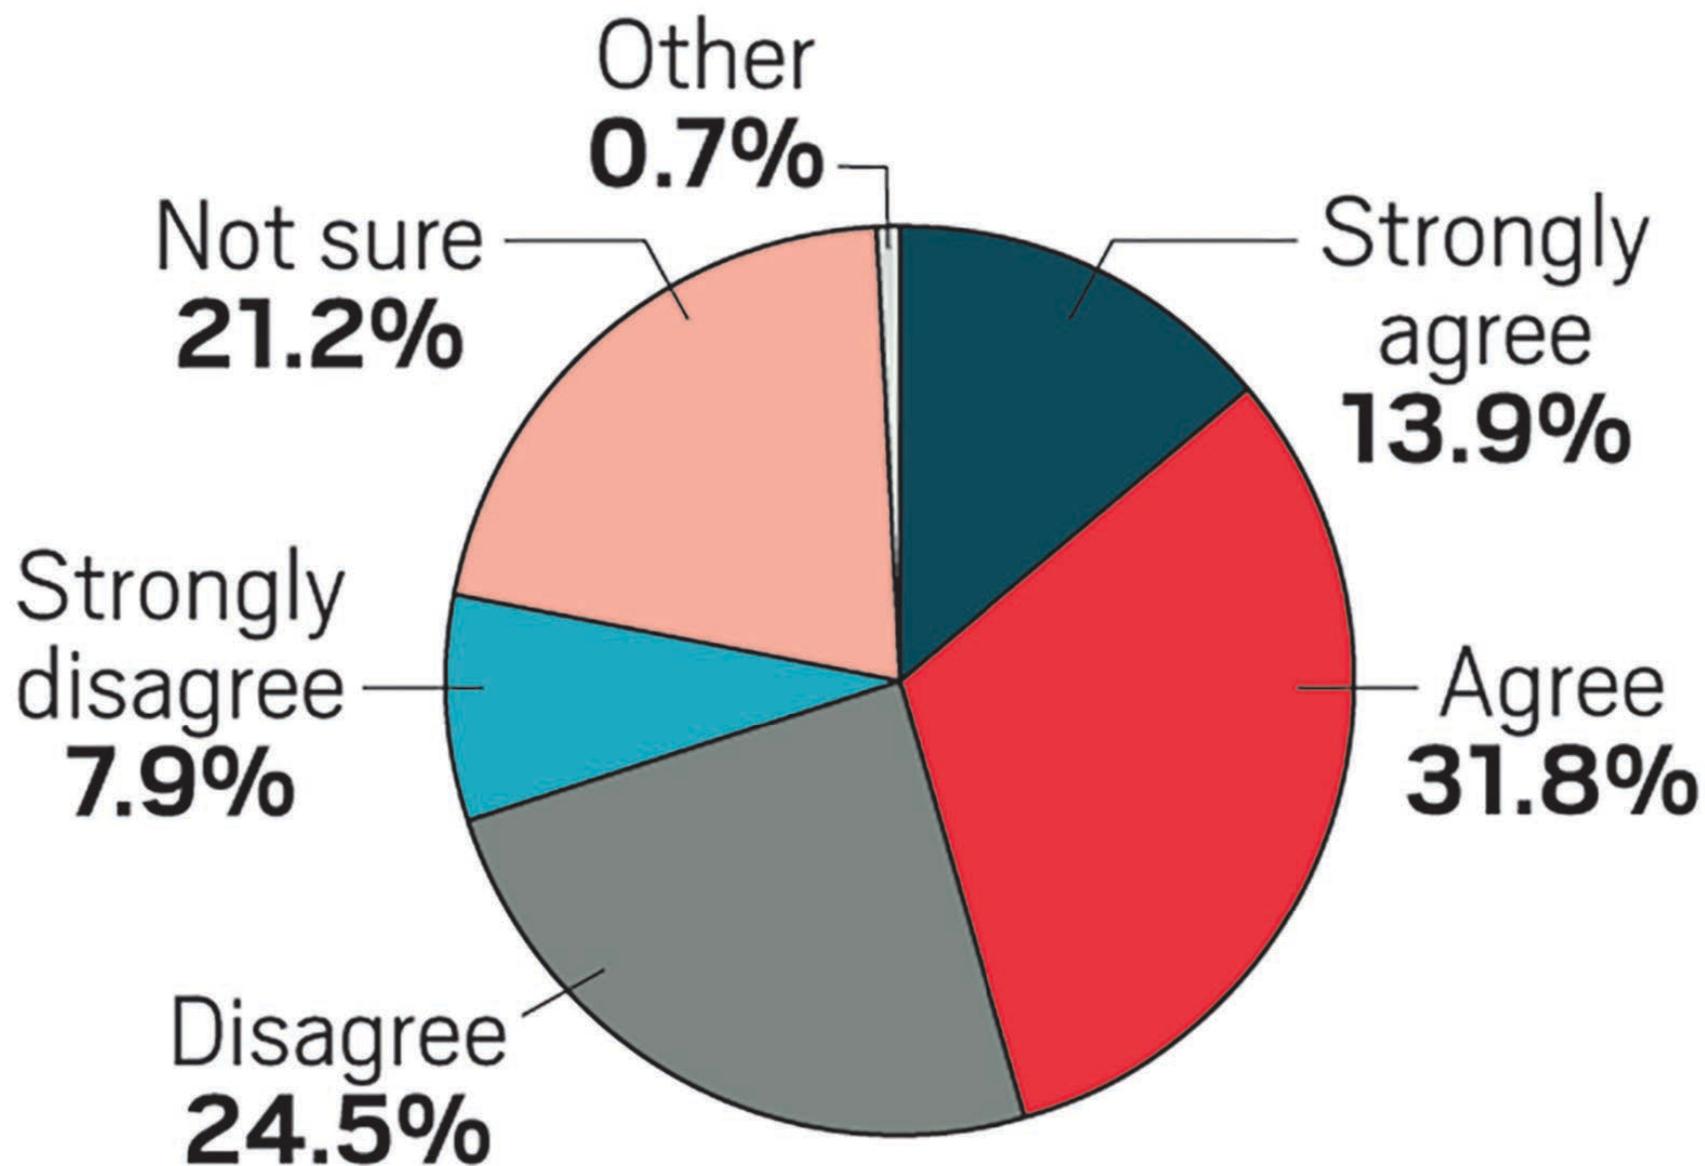
... AND CHECK WHETHER
THE PHOTO IS OF A BIRD.

I'LL NEED A RESEARCH
TEAM AND FIVE YEARS.



IN CS, IT CAN BE HARD TO EXPLAIN
THE DIFFERENCE BETWEEN THE EASY
AND THE VIRTUALLY IMPOSSIBLE.

Do you think machine learning is overhyped?



Chemical & Engineering News, August 2018

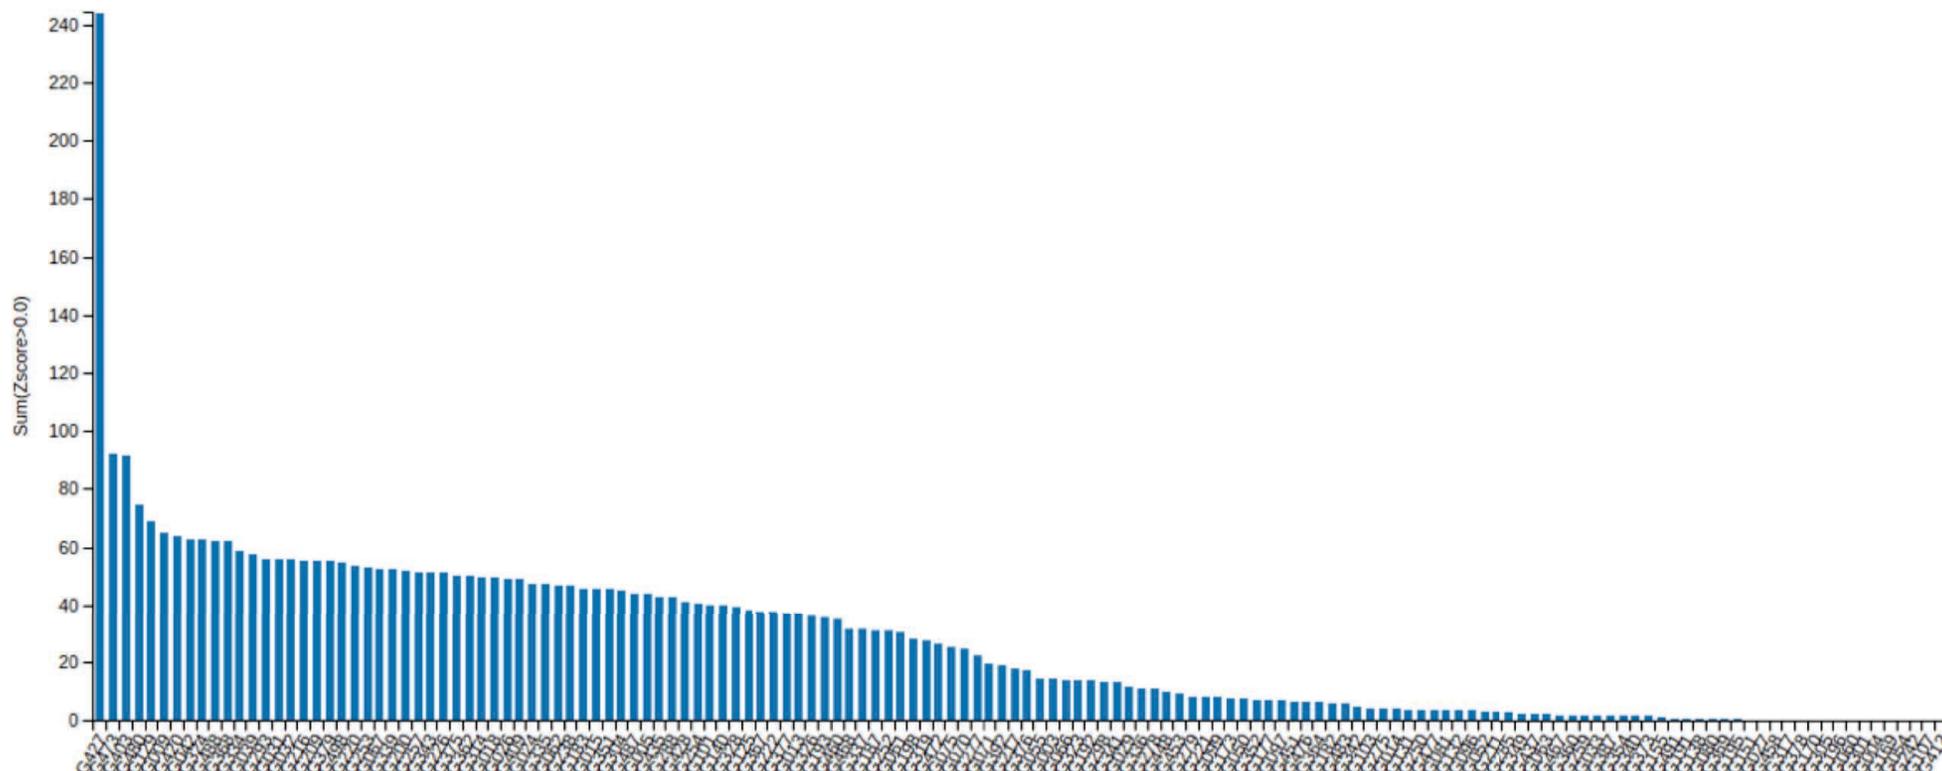
Despite all this promise—or perceived promise—one thing that machine learning isn't is magic. “Let's be realistic,” says **George Dahl**, a computer scientist at Google. “Machine learning is nonlinear regression,” a simple type of statistical analysis in which collected data are “fit” with model parameters. Dahl won a Merck & Co. machine-learning competition while a graduate student in **Geoffrey Hinton**'s group at the University of Toronto.

“Since machine learning simply is interpolation within big data sets, it will remain difficult or impossible to use in areas where it is hard to generate large sets of reliable data, because extrapolation by machine learning can and will produce wildly wrong answers.”—**Bernd Hartke**, professor, University of Kiel

Deep neural networks' ability to learn even from very complex data makes them especially attractive in pharmaceutical chemistry. Abraham Heifets is cofounder and CEO of **Atomwise**, which makes deep-neural-network-based software for predicting binding affinities of drug candidates to targets in the body. He calls the introduction of deep neural networks a fundamental change in machine learning.

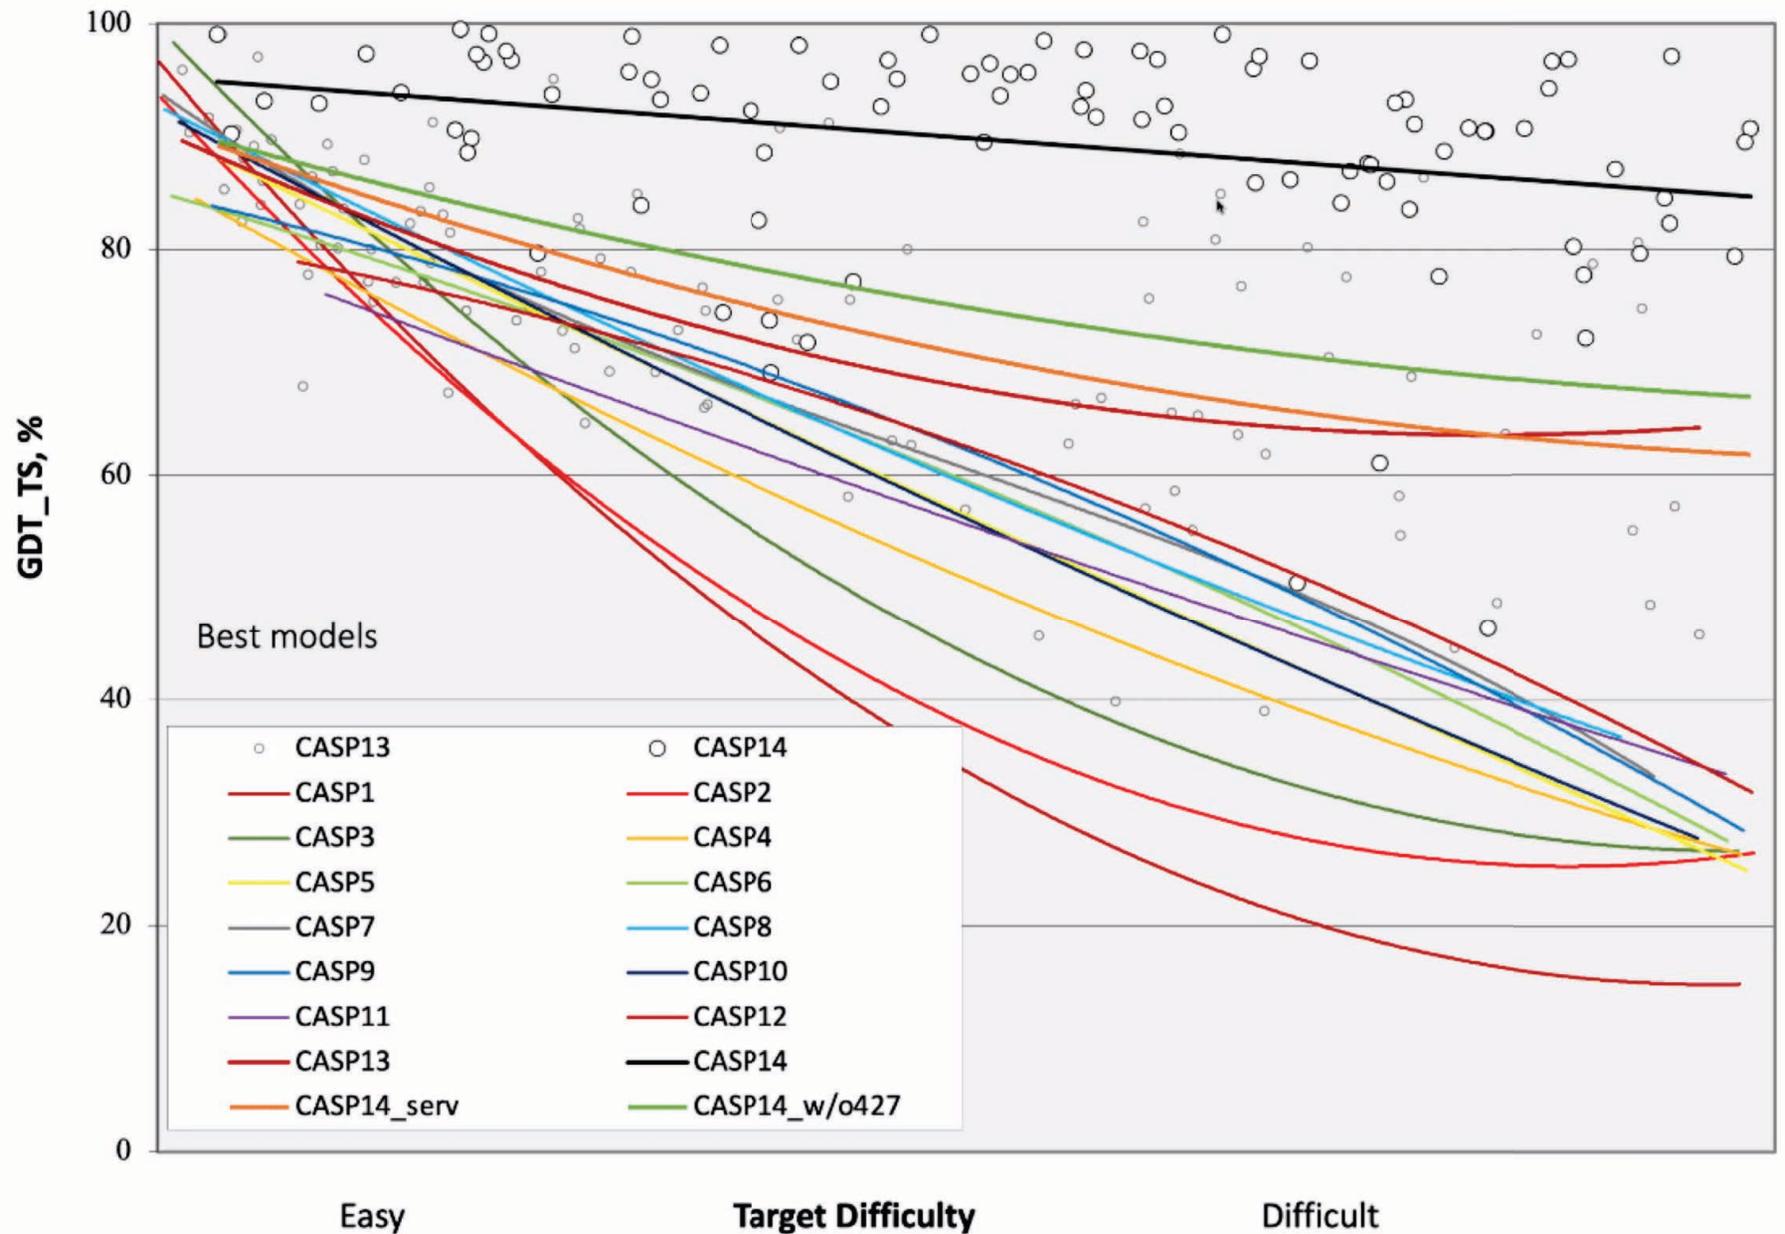
Some have accused Atomwise of overhyping machine learning's capabilities. A **2015 TechCrunch article** quoted the firm's cofounder Alexander Levy as saying the Atomwise software allowed him to predict a cure for measles from his living room.

Z-Score Sums for Groups in CASP14



mean Z-score for AlphaFold 2 = 2.5
for “hard” targets = 3.8 (“IQ” > 160)

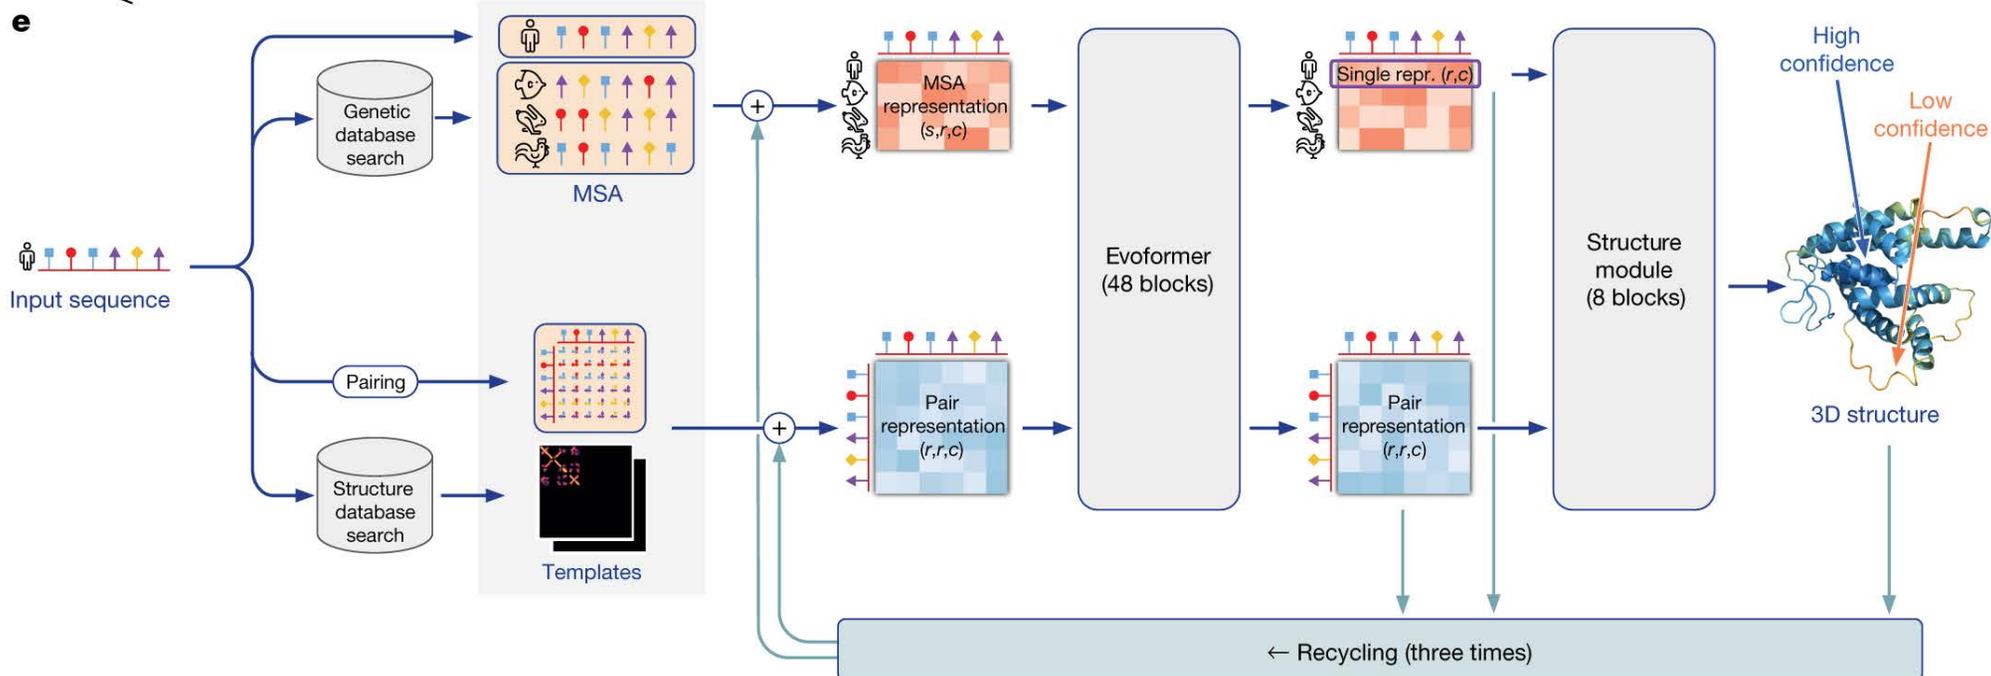
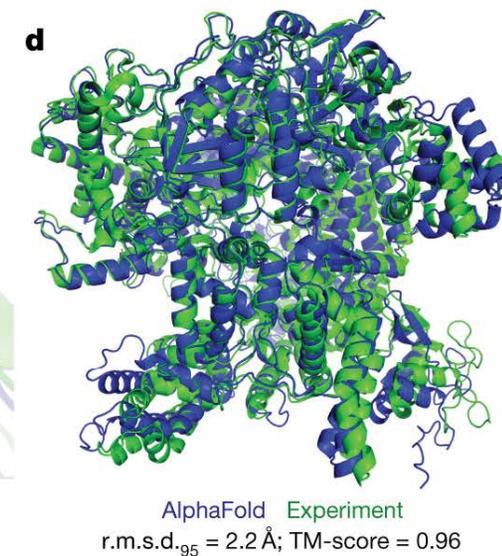
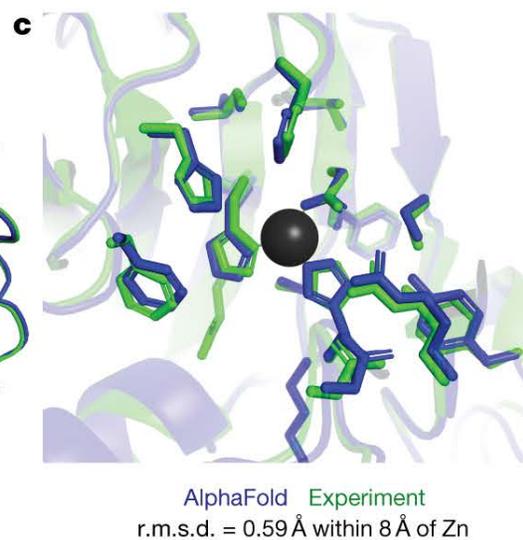
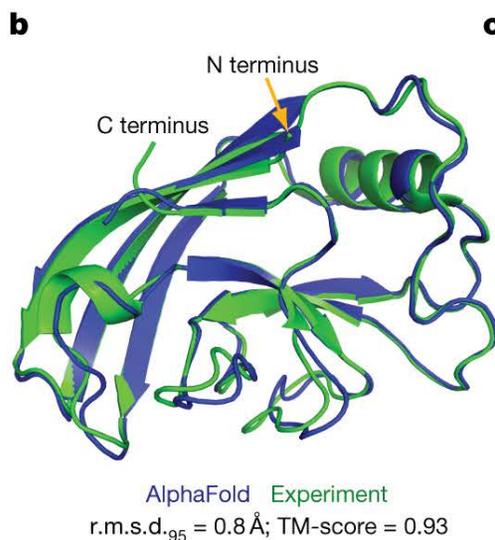
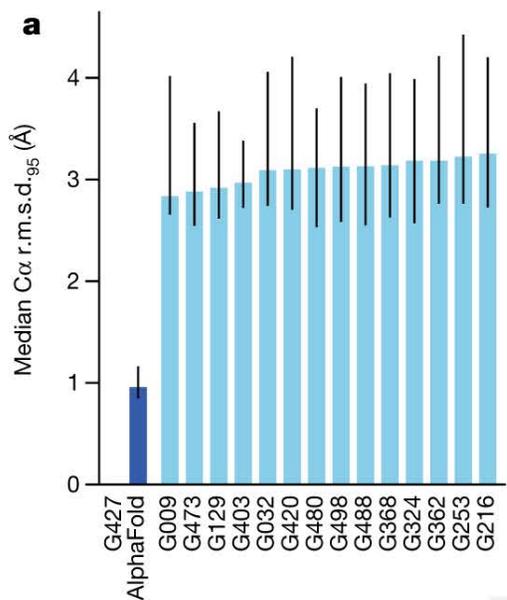
Historical CASP Results



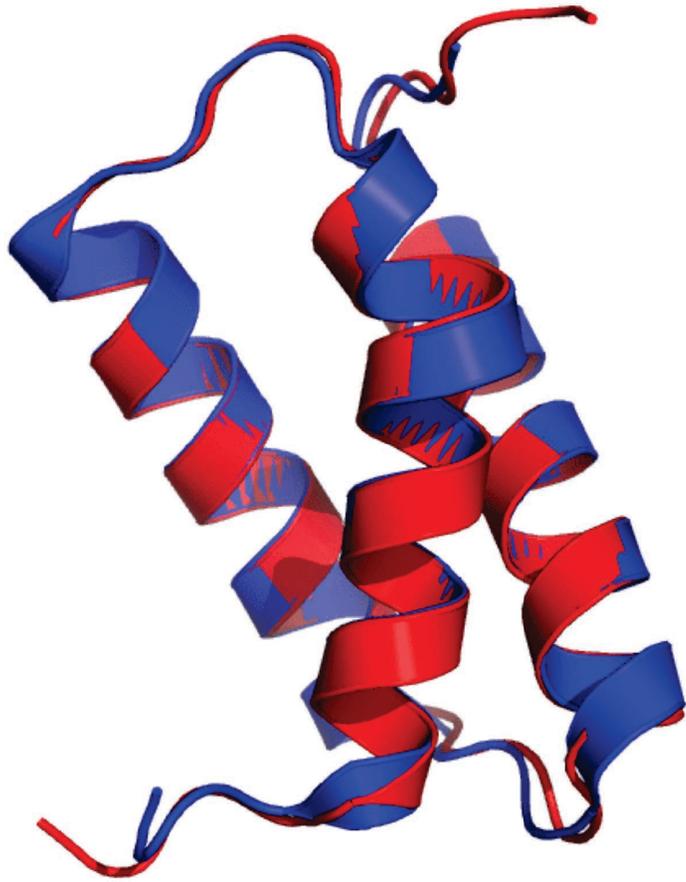
GDT=60 Overall Fold Correct

GDT=80 Sidechains Correct

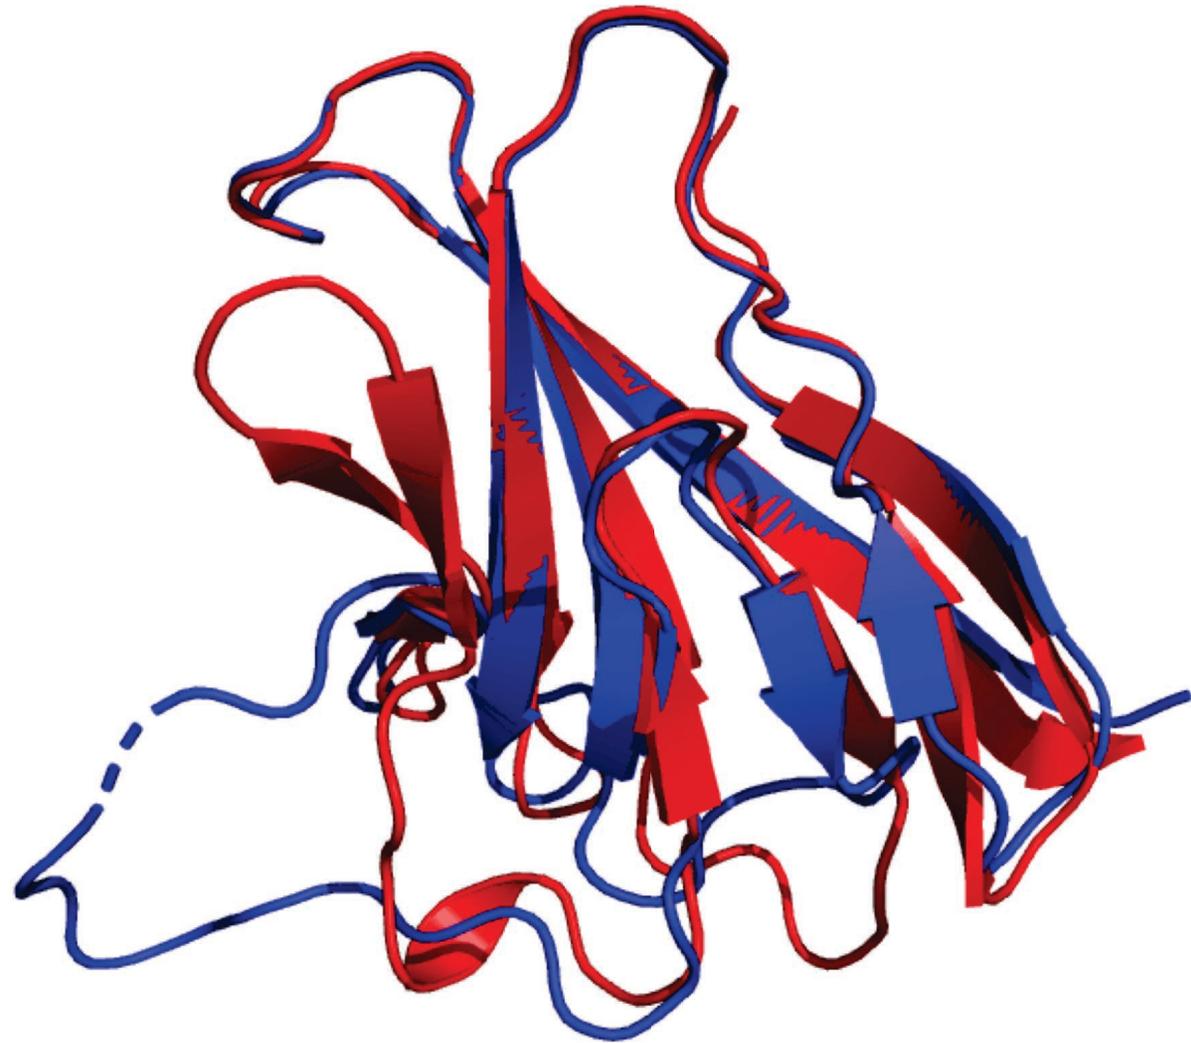
AlphaFold 2 Produces Highly Accurate Structures



Sample CASP14 Results from AlphaFold 2

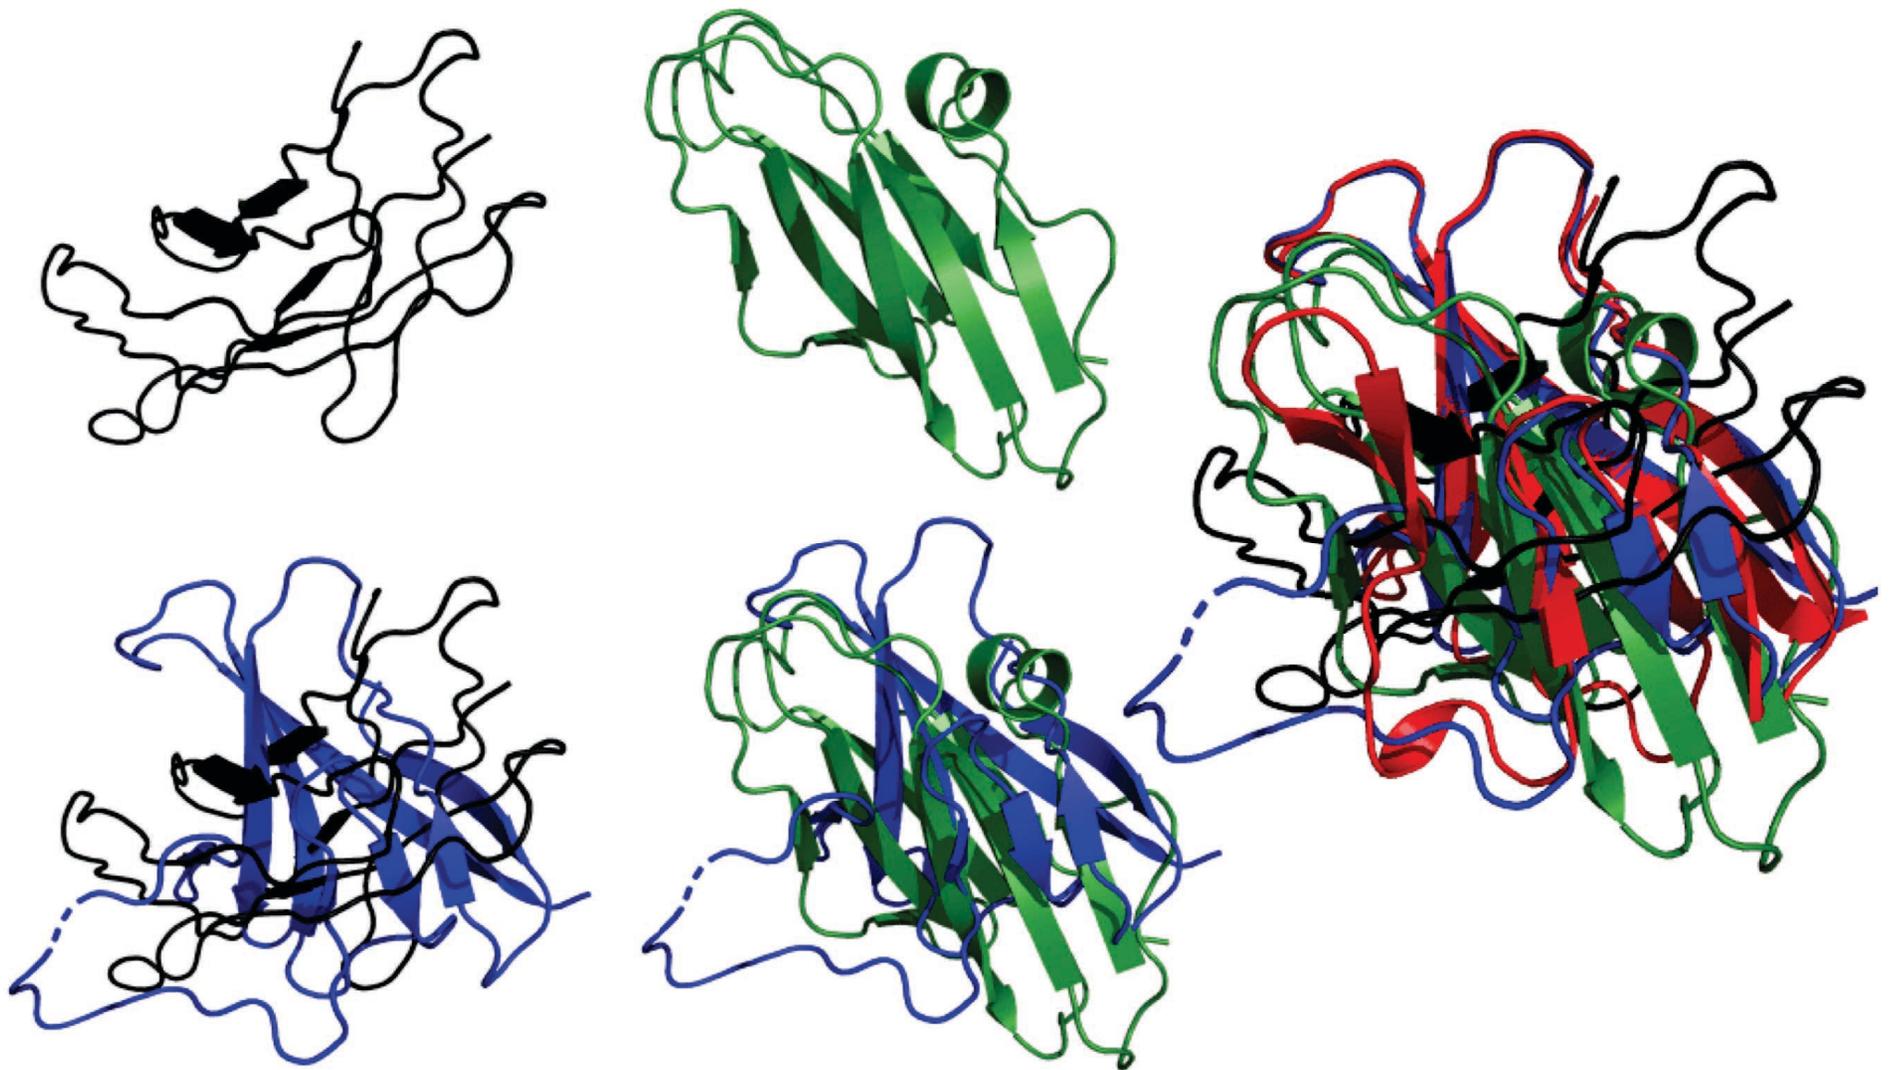


Bacteriophage T4 SPACKLE
(total RMSD of 0.48 Ang)



SARS-CoV-2 ORF8
("poor" result for AlphaFold2)

Comparison for SARS-CoV-2 ORF8



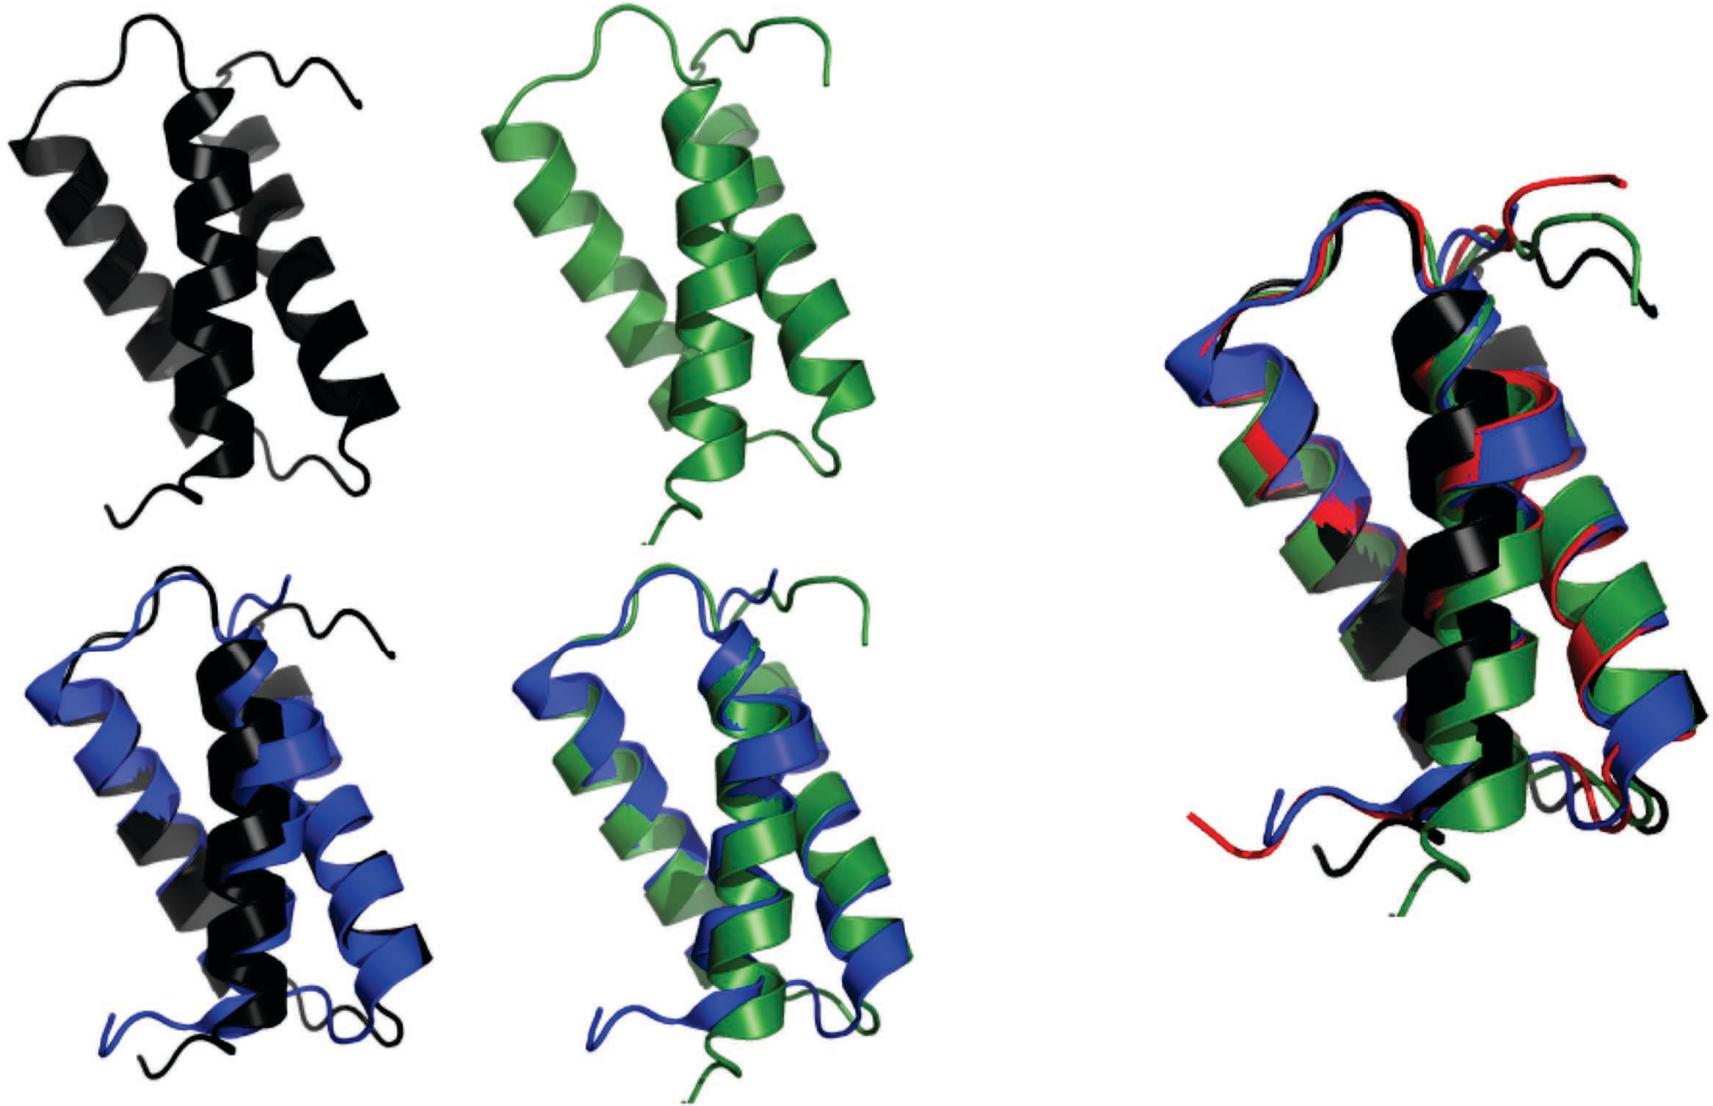
● *Expt*

● *Baker*

● *AlphaFold 2*

● *Zhang*

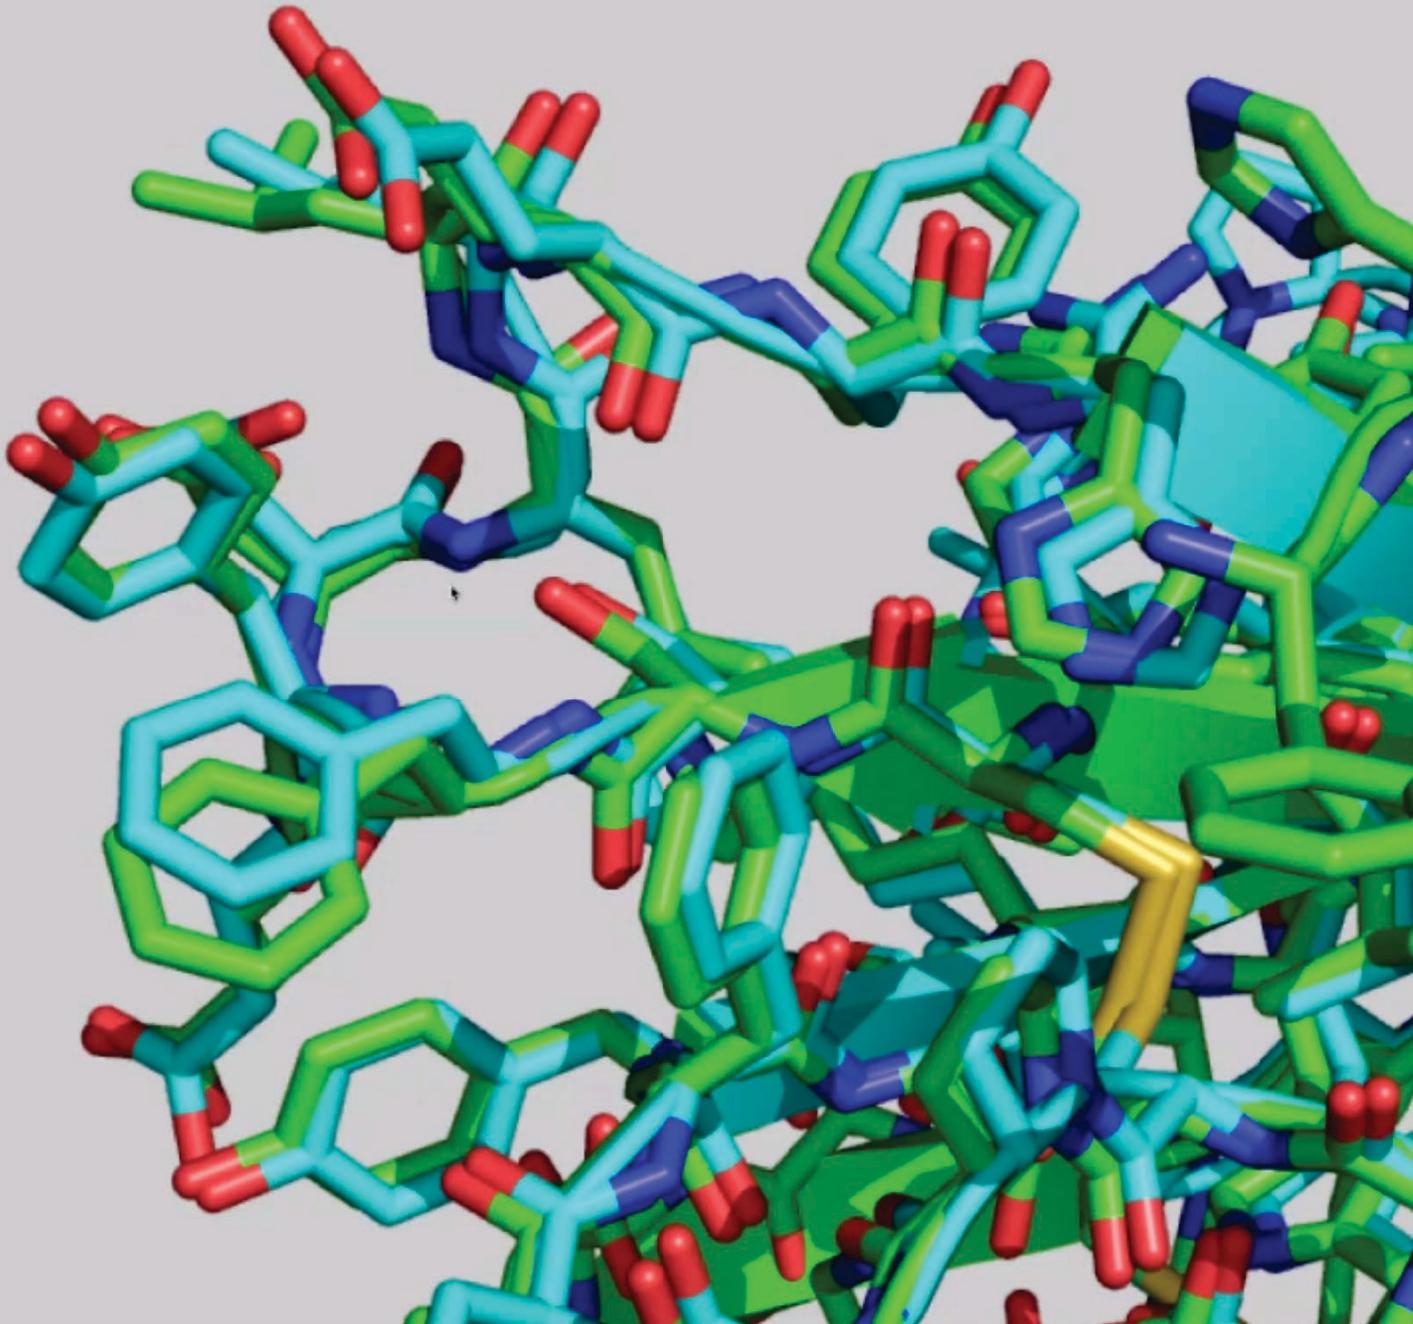
Comparison for Bacteriophage T4 SPACKLE



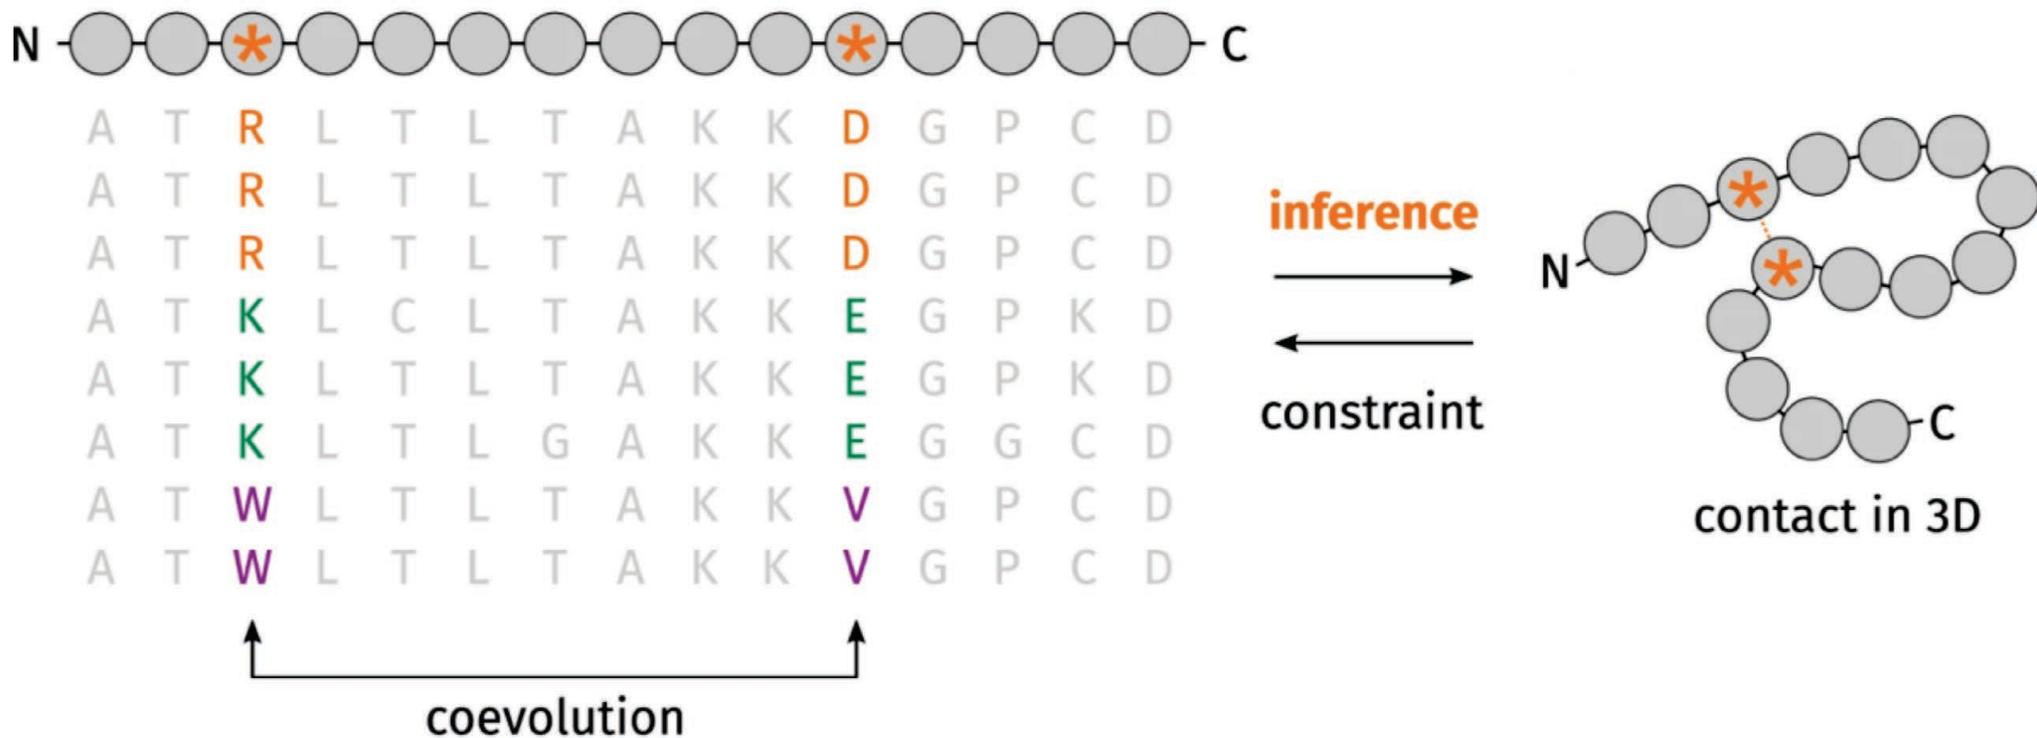
● *Expt*
● *Baker*

● *AlphaFold 2*
● *Zhang*

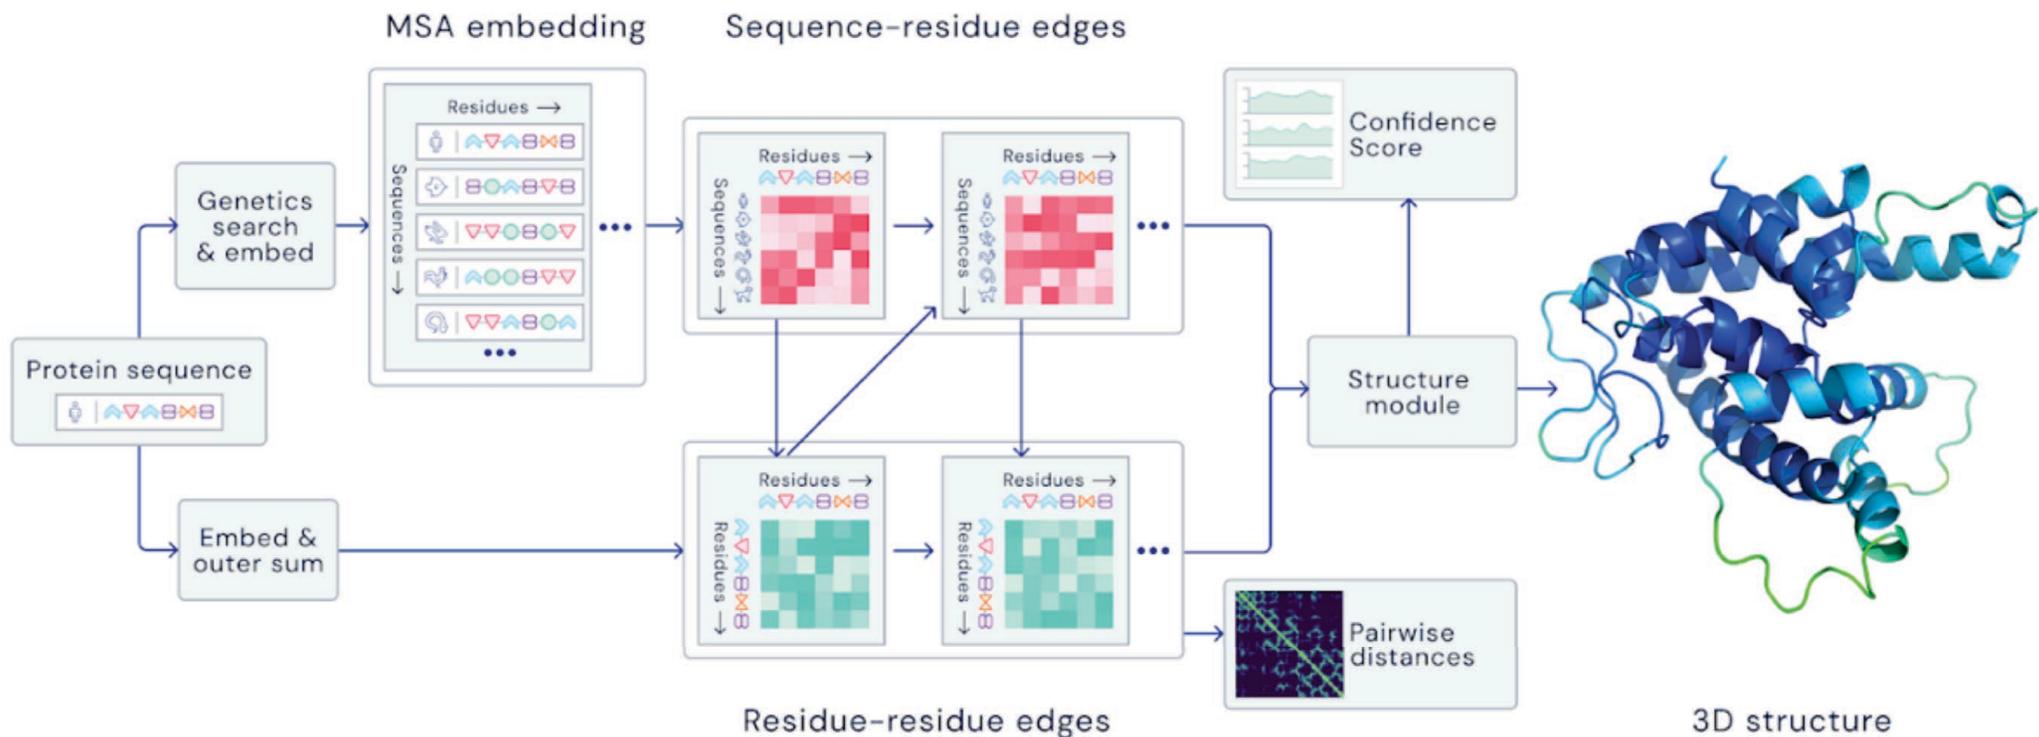
AlphaFold 2 Side Chain Accuracy



Sequence Covariance is a Major Data Source Used Effectively by AlphaFold 2



Neural Net Design for AlphaFold 2



the predictive process proceeds in iterative fashion by ***“passing information back and forth between the MSA and the internal representation of the protein”***

Architectural Details of AlphaFold 2

