How Neural Networks See the World

Building Intuition for a Broader AI Discussion

Book was great!

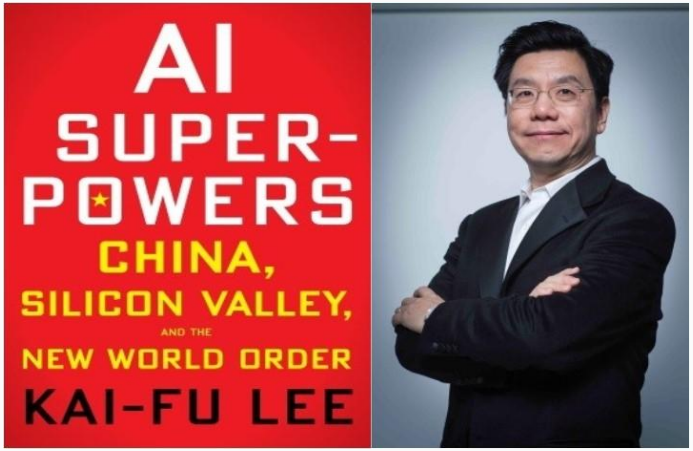

Broader AI conversation.

But how does all this “AI” stuff actually work?

What is deep learning?

How do neural networks learn?

Engineering Perception

I offer a developer perspective

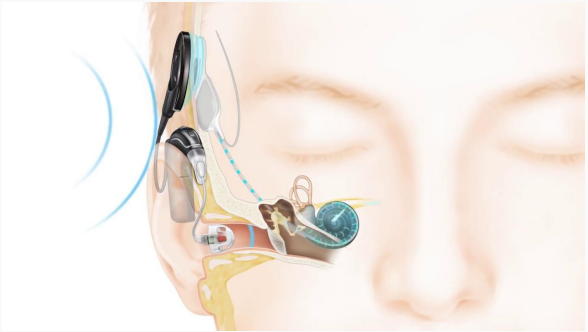

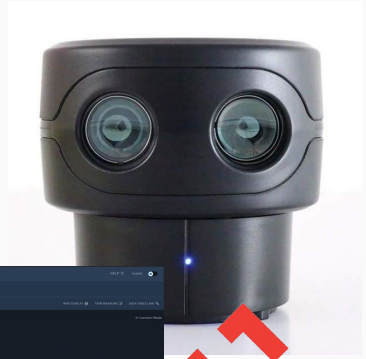

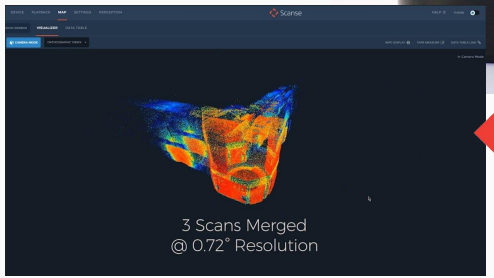

RIPCORD

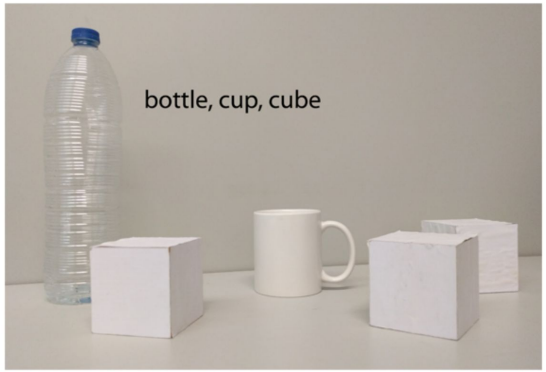

(a) Image classification

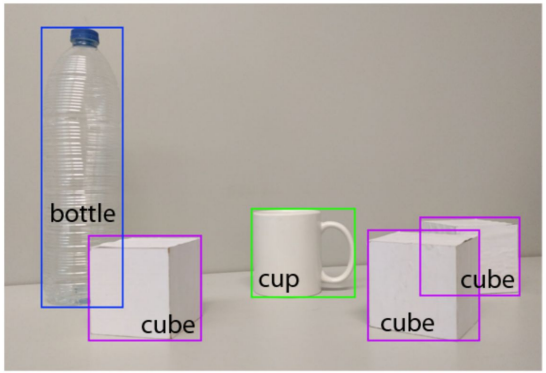

(b) Object localization

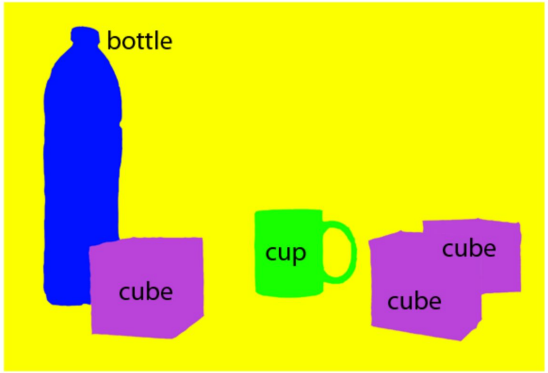

(c) Semantic segmentation

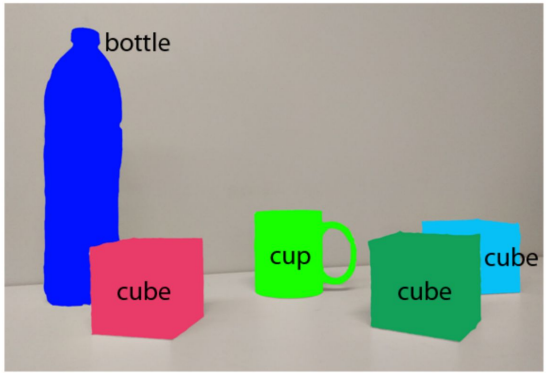

(d) Instance segmentation

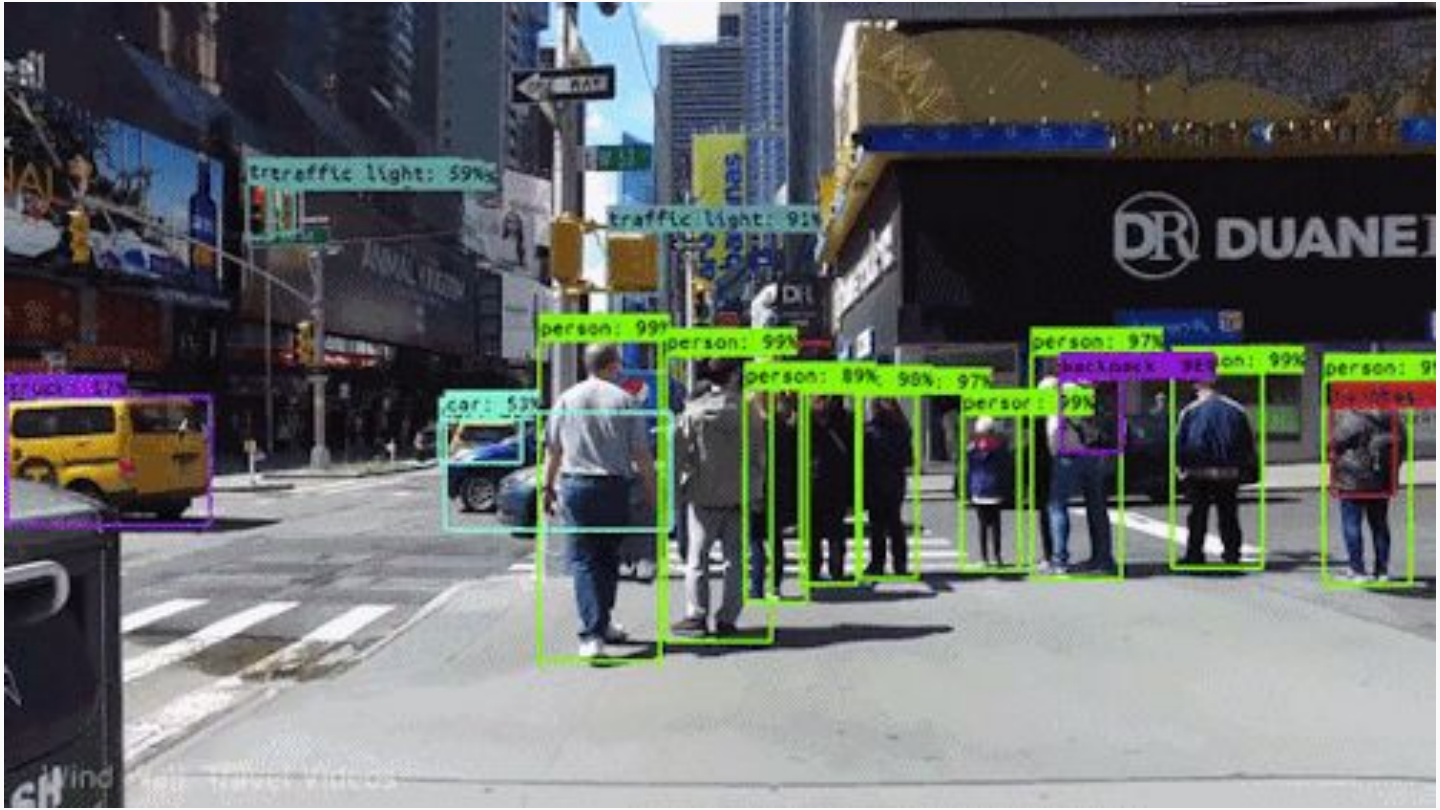
traffic light: 59%
traffic light: 91%
person: 99%
person: 99%
person: 89%
person: 97%
car: 53%

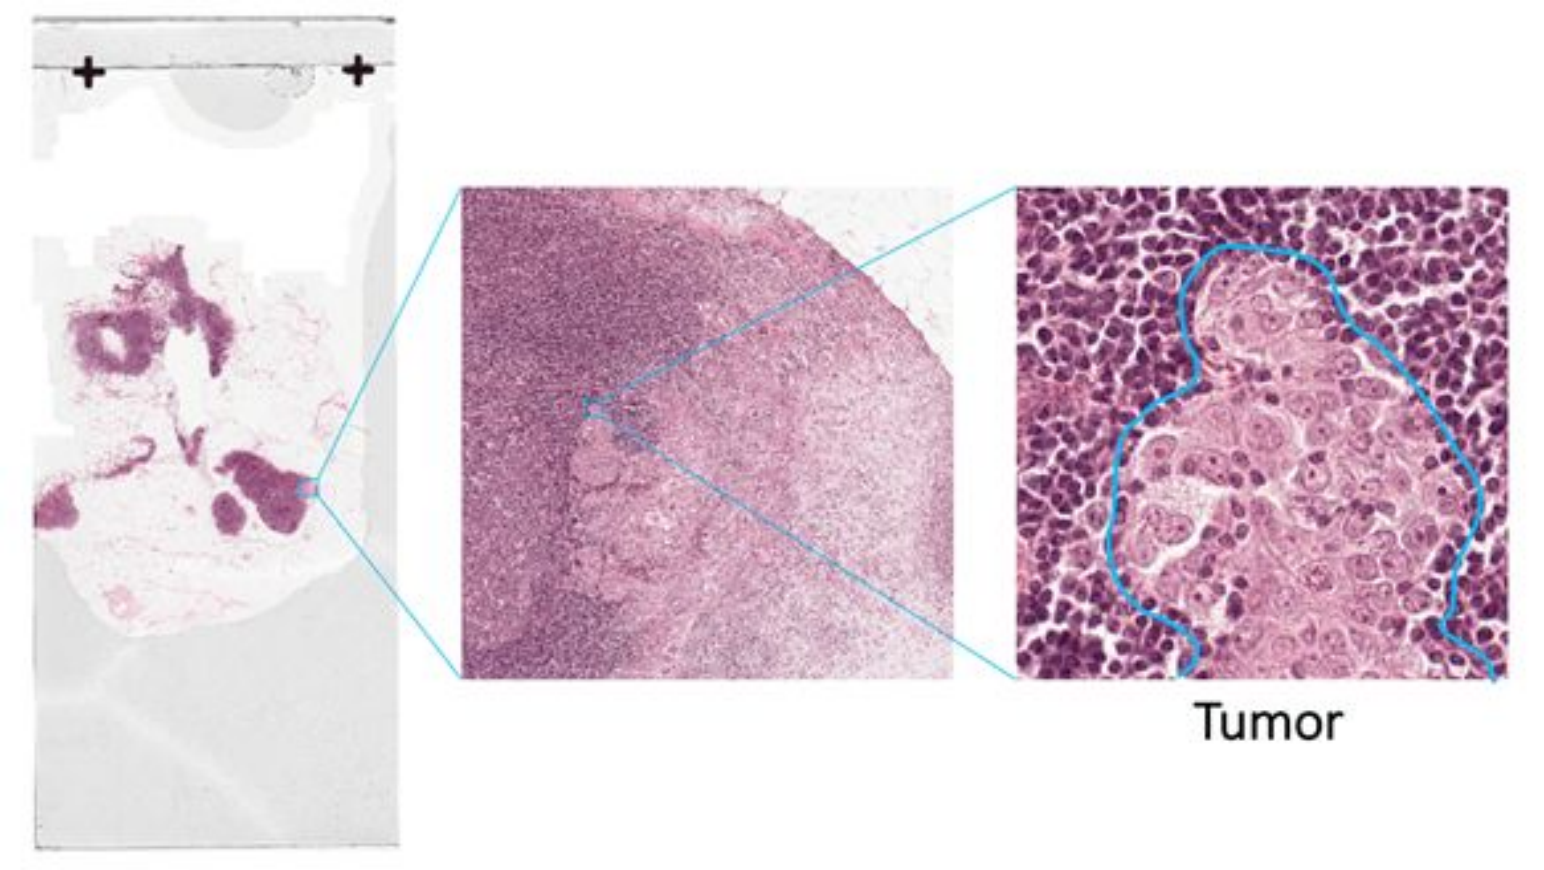
Tumor

How do we (humans) see the world?

Our brains get **excited**

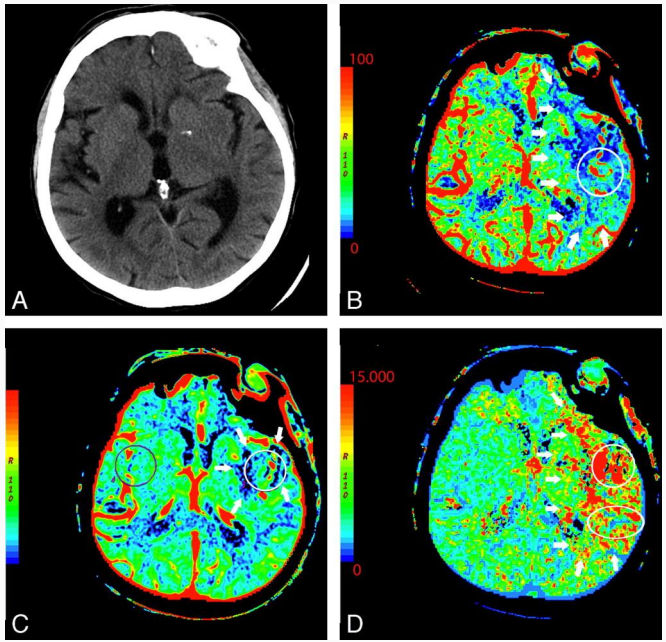

Our brains get **excited**

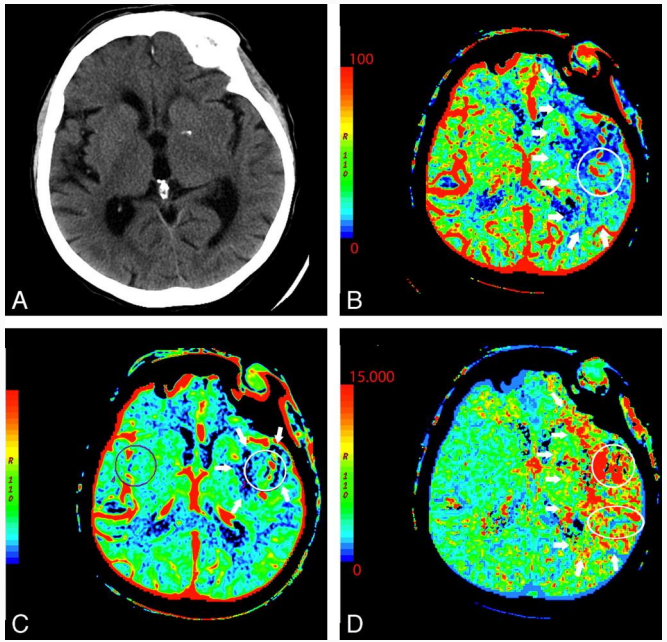

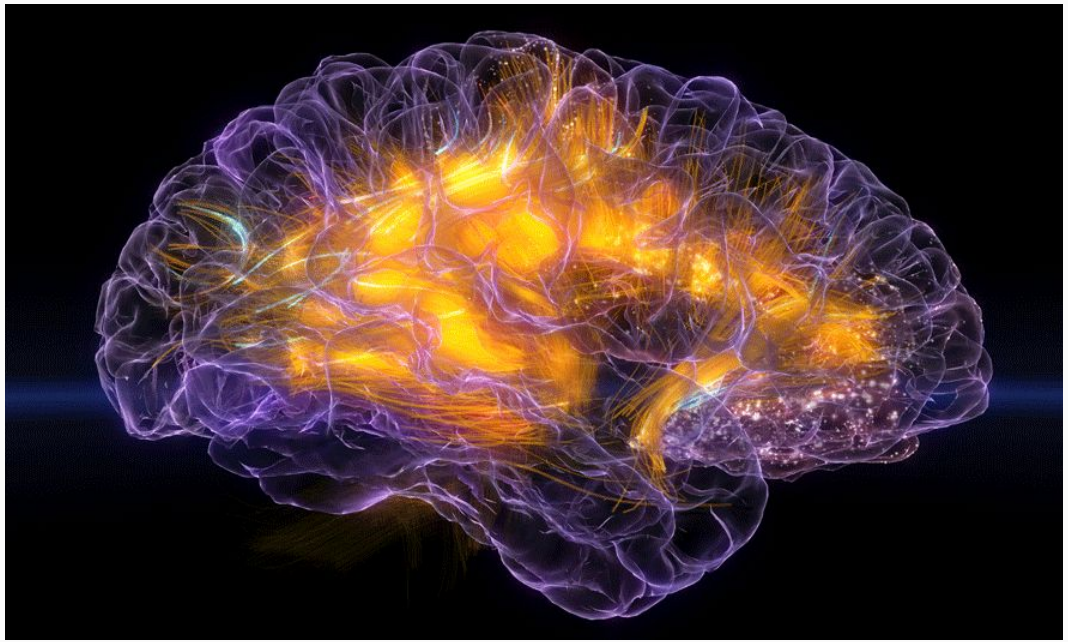

Experience leverages **different pathways**

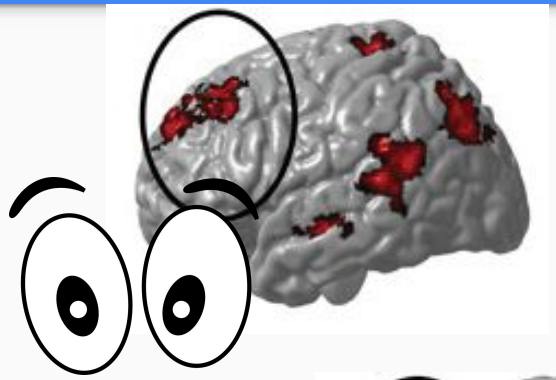

Average

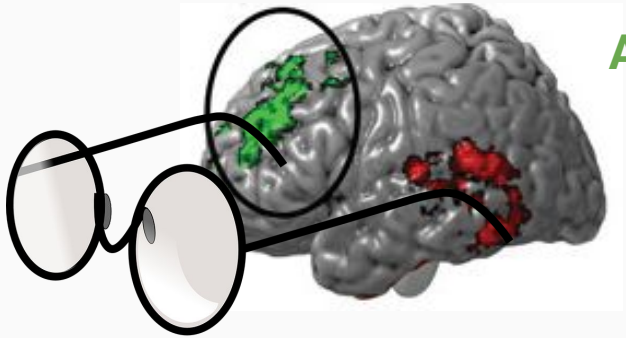

Attorney

Question:

What does it mean to **understand** something?

Neural **pathways** are prepared and ready.

What does it mean to **learn** something?

Reinforce, recruit, adapt, or create new neural pathways.

Question:

What does it mean to **understand** something?

You are prepared to interpret **new examples**.

What does it mean to **learn** something?

Studying or observing more examples, in preparation.

In deep learning,
neural networks **study existing data**
in order to interpret new data.

Classification is a simple way to interpret

- Managing variability
- Recognizing patterns
- This vs. That
- Concise communication about complex groups

Task:

Describe this simple data

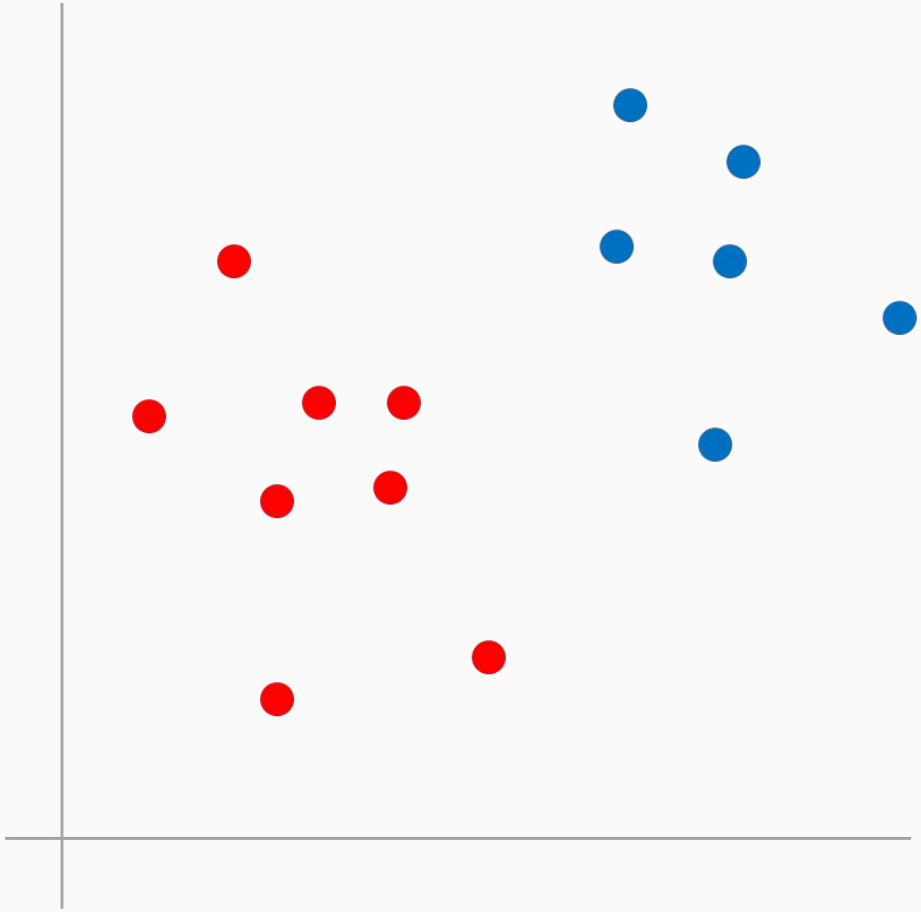

Rules are concise

RULE #1:

Points ABOVE the line are BLUE

RULE #2:

Points BELOW the line are RED

Leverage math to formally describe the rule:

$y = mx + b$

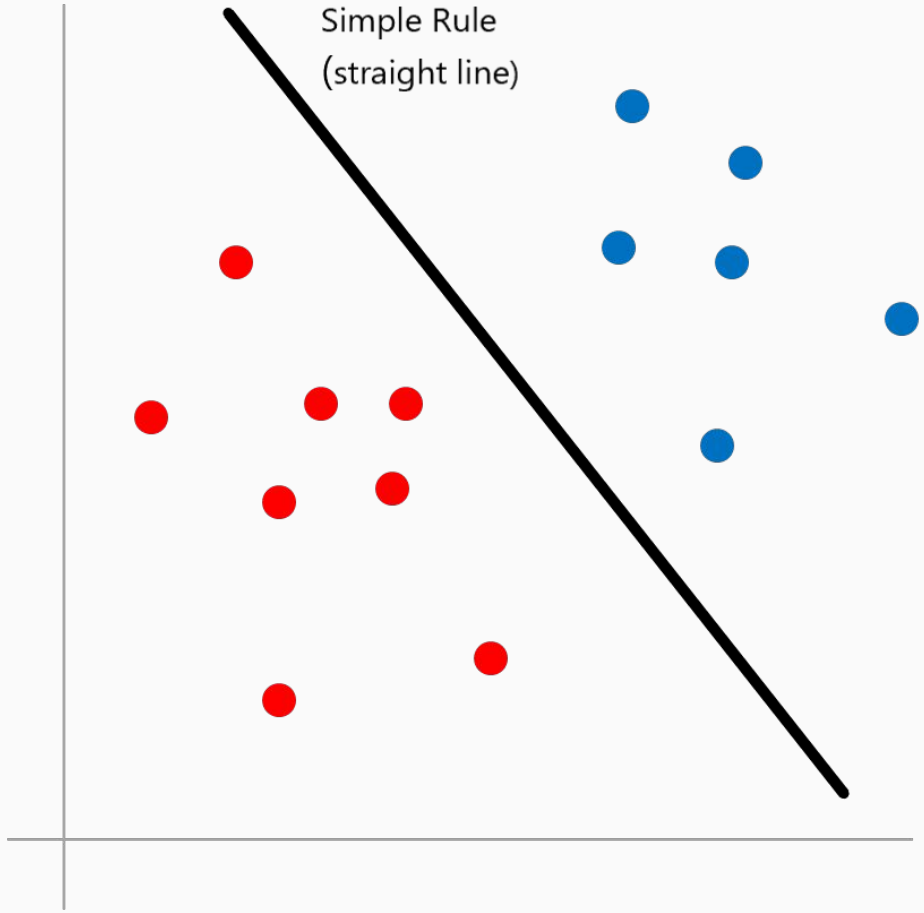

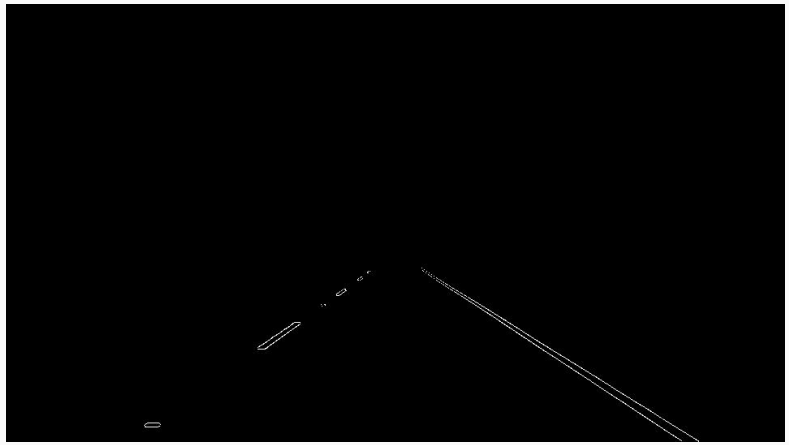

Rule: Lane Lines are WHITE

Rule: Road is below horizon

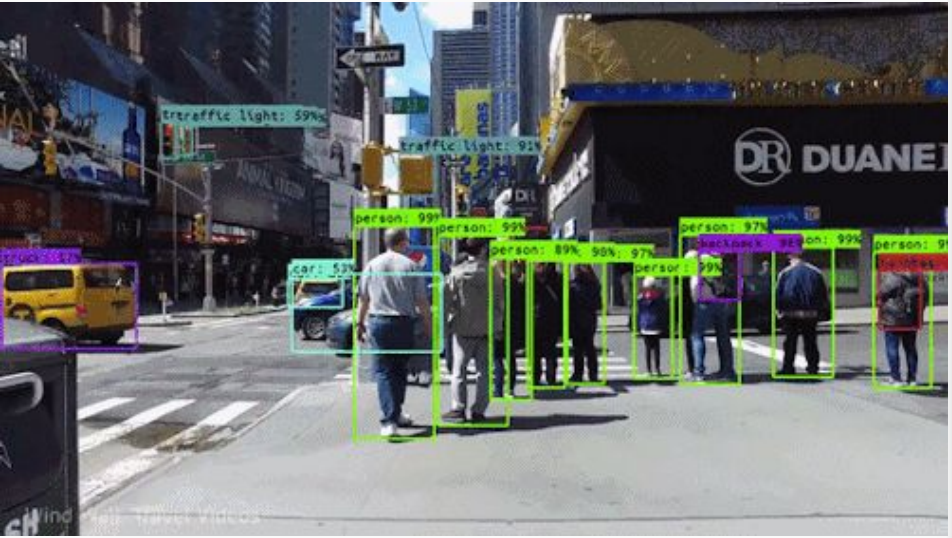

But real data is messy, and writing complex rules is hard.

Deep Learning makes complex discrimination possible.

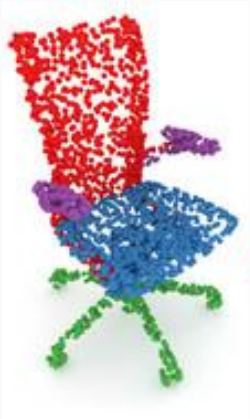

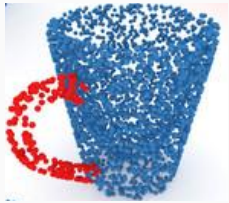

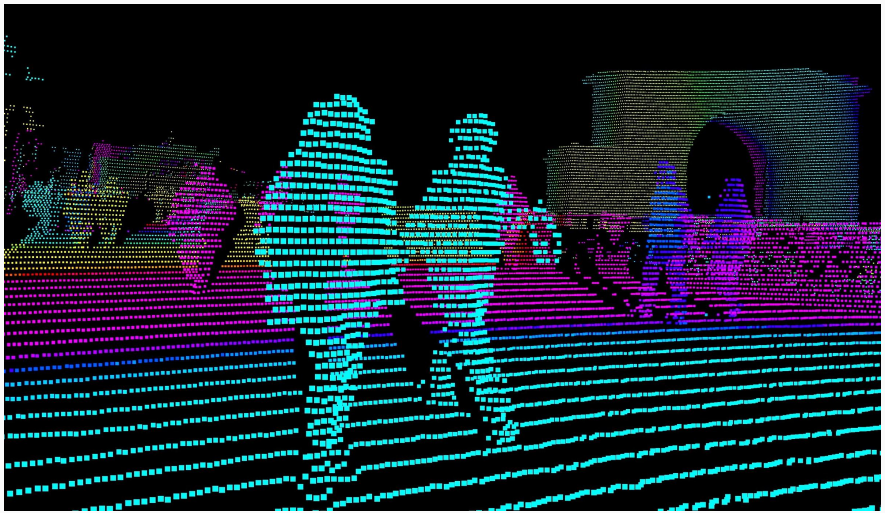

But real data is messy, and writing complex rules is hard.

Neural networks **appear** to learn complex rules.

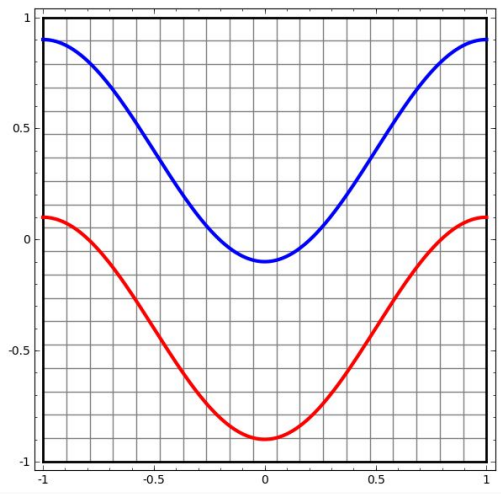

Complex Data

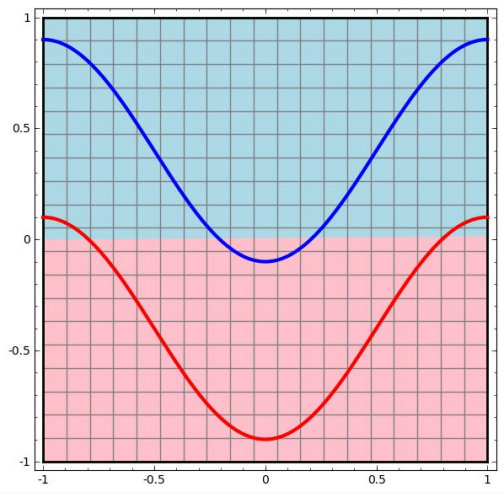

Simple Rule
(inadequate)

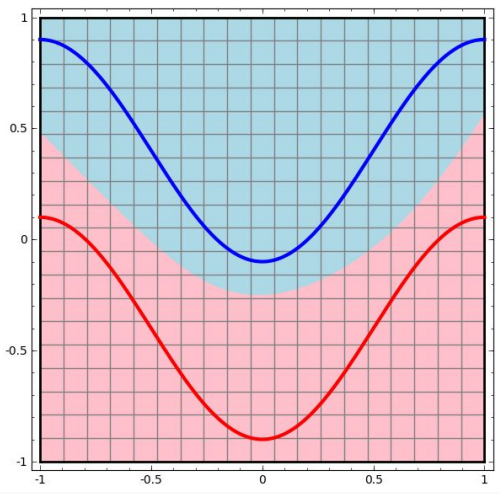

Complex Rule

Actually:

Neural networks apply simple rules in a complex way.

To apply simple rules, neural networks distort space.

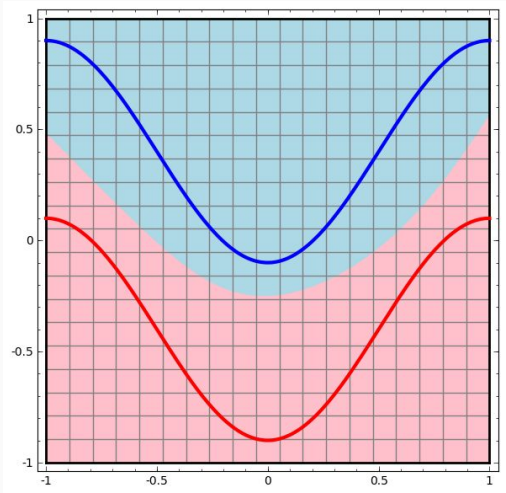

Complex Rule

=

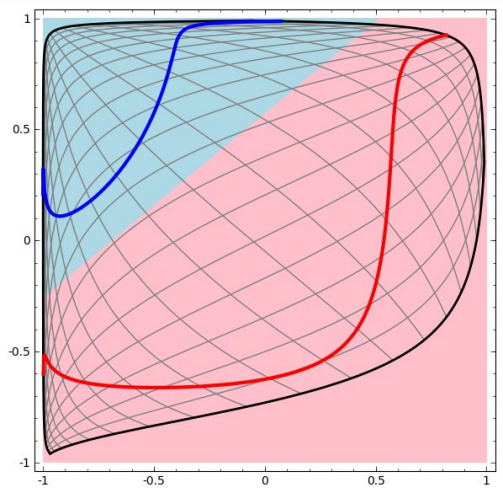

Simple Rule

Change Your Perspective

In order to understand data, neural networks distort how it is presented.

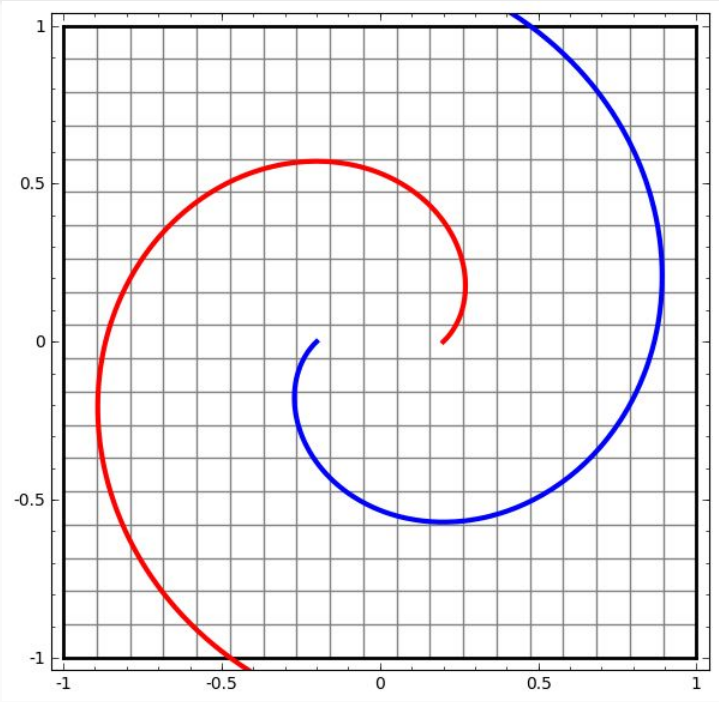

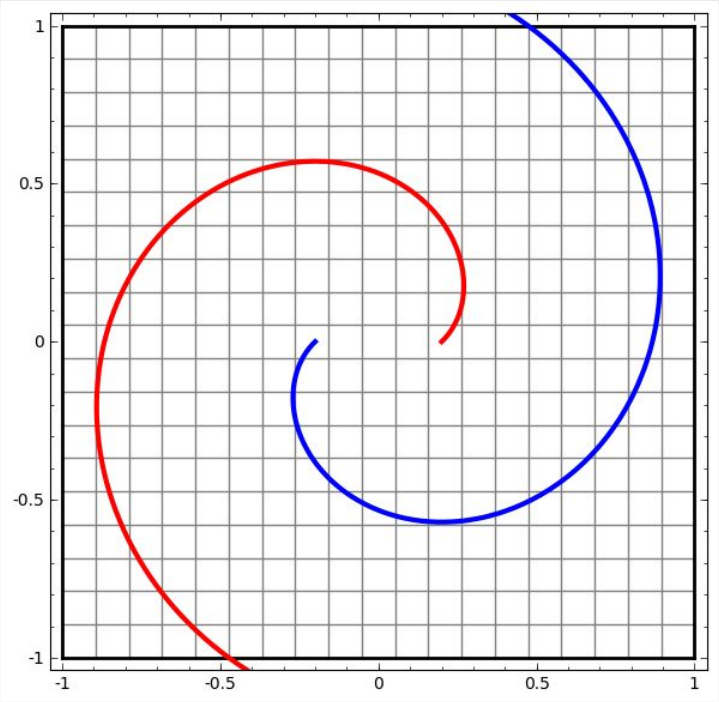

Takeaway:

An experienced neural network is prepared to interpret new examples.

AI landscape is filled with jargon.

Just know **learning** is fundamentally unique.

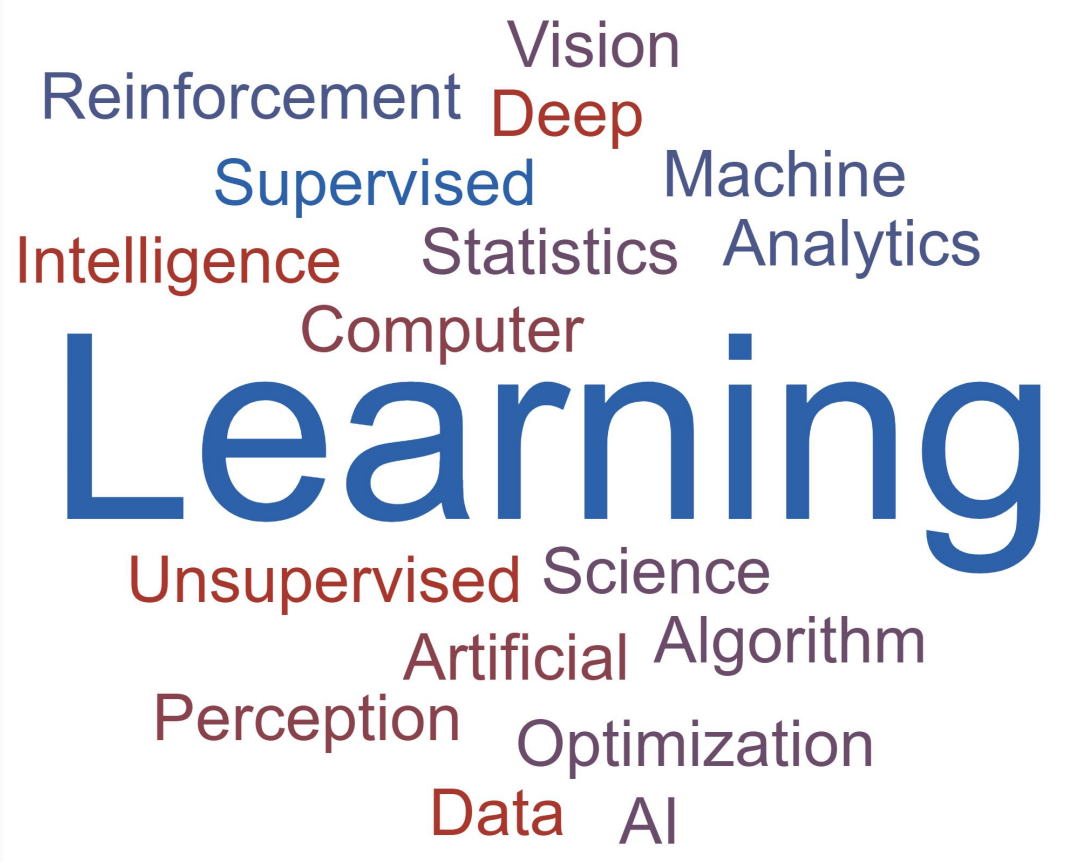

Unlike traditional problem solving, learning is an indirect solution.

With **ALGORITHMS**, behavior is designed.

Algorithms are:

- Written instructions
- Rule driven
- Provable
- Guaranteed Behavior

With **Learning**, environment and curriculum are designed. Behavior emerges indirectly.

Unlike traditional problem solving, learning is an indirect solution.

With **ALGORITHMS**, behavior is designed.

Algorithms are:

- Written
- Rule driven
- Provable
- Guaranteed Behavior

Stop Defining Rules

With **Learning**, environment and curriculum are designed. Behavior emerges indirectly.

Start Learning Rules

Neural Networks are like students, that learn in familiar terms.

Neural networks study data.

Often under supervision that demands memorization.

Developers are teachers, creating classroom environment to facilitate learning.

Question:

What is a network?

Connections are the building blocks of networks

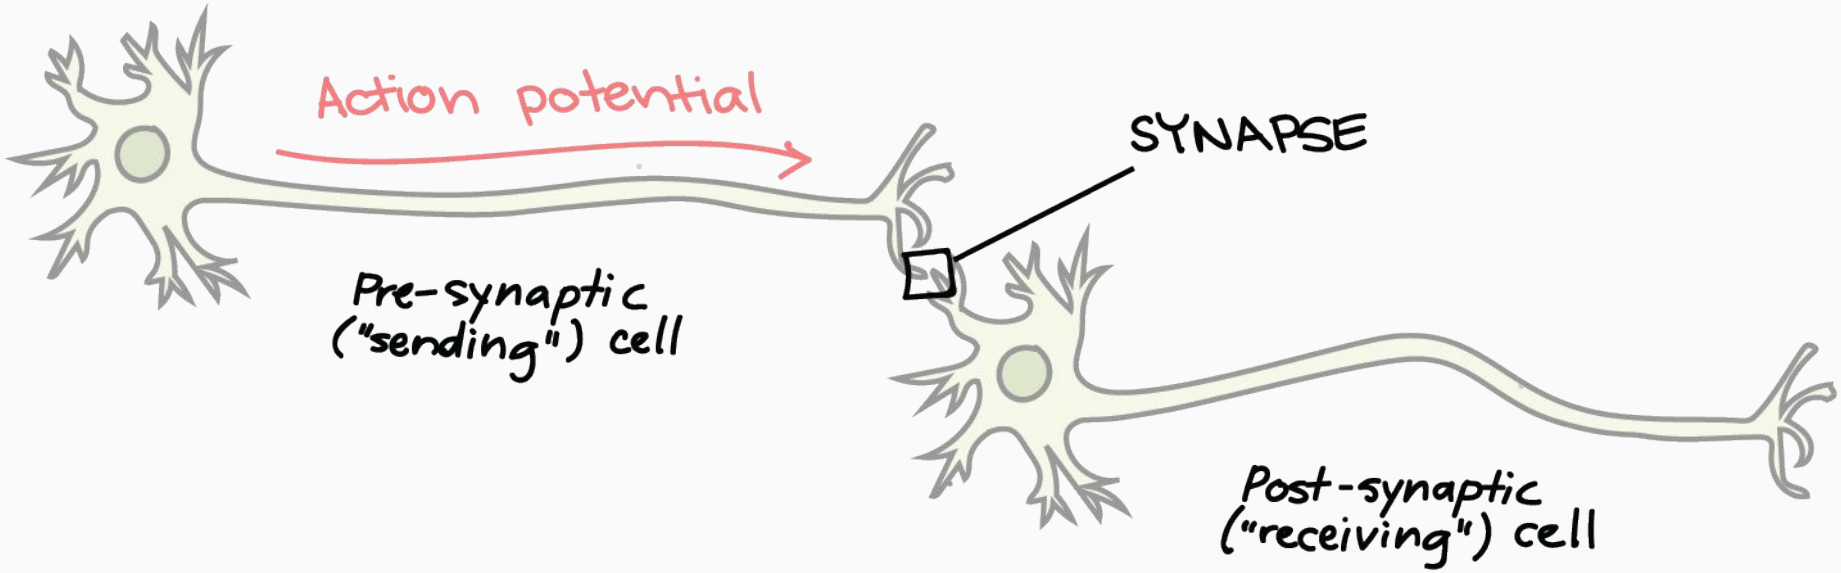

Networks are comprised of many tangled connections

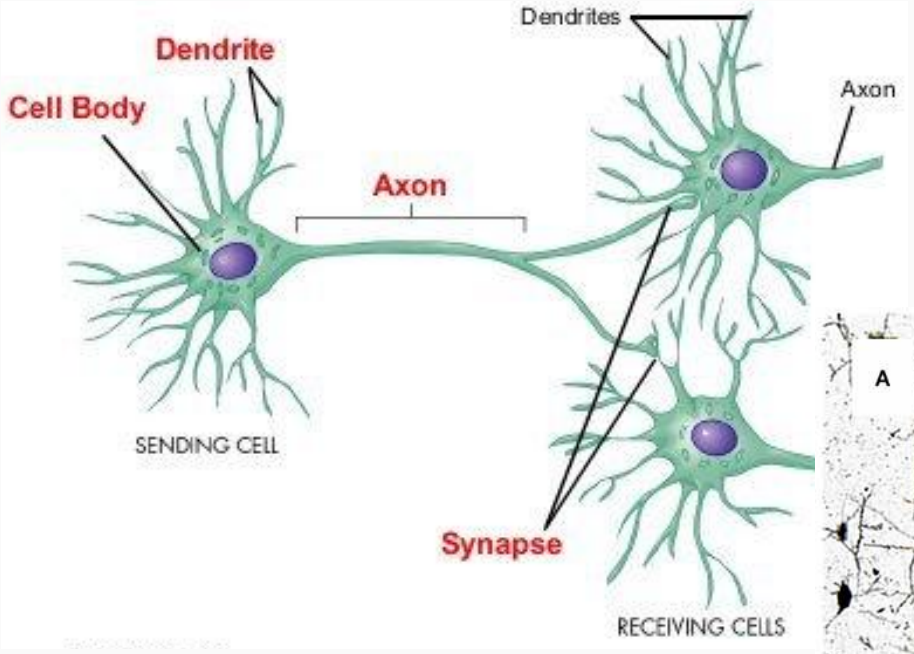

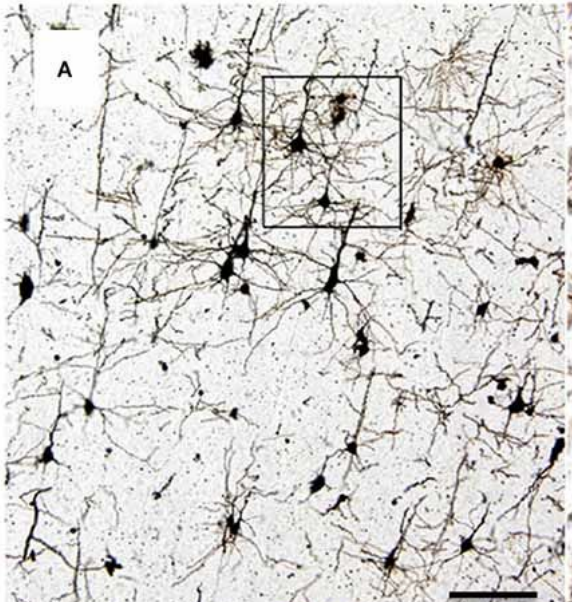

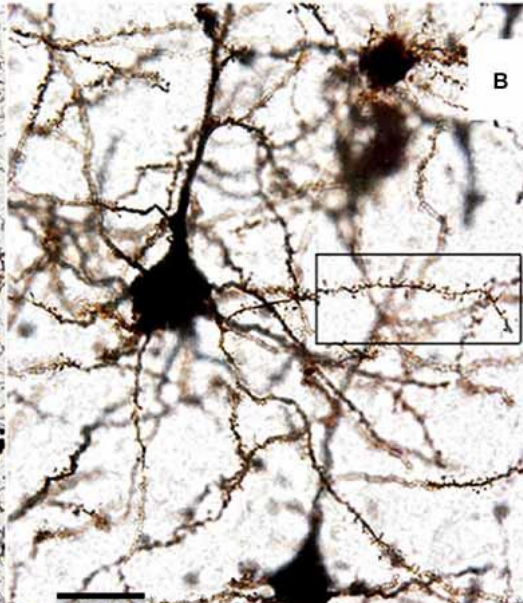

From the outside, networks have **inputs** and **outputs**

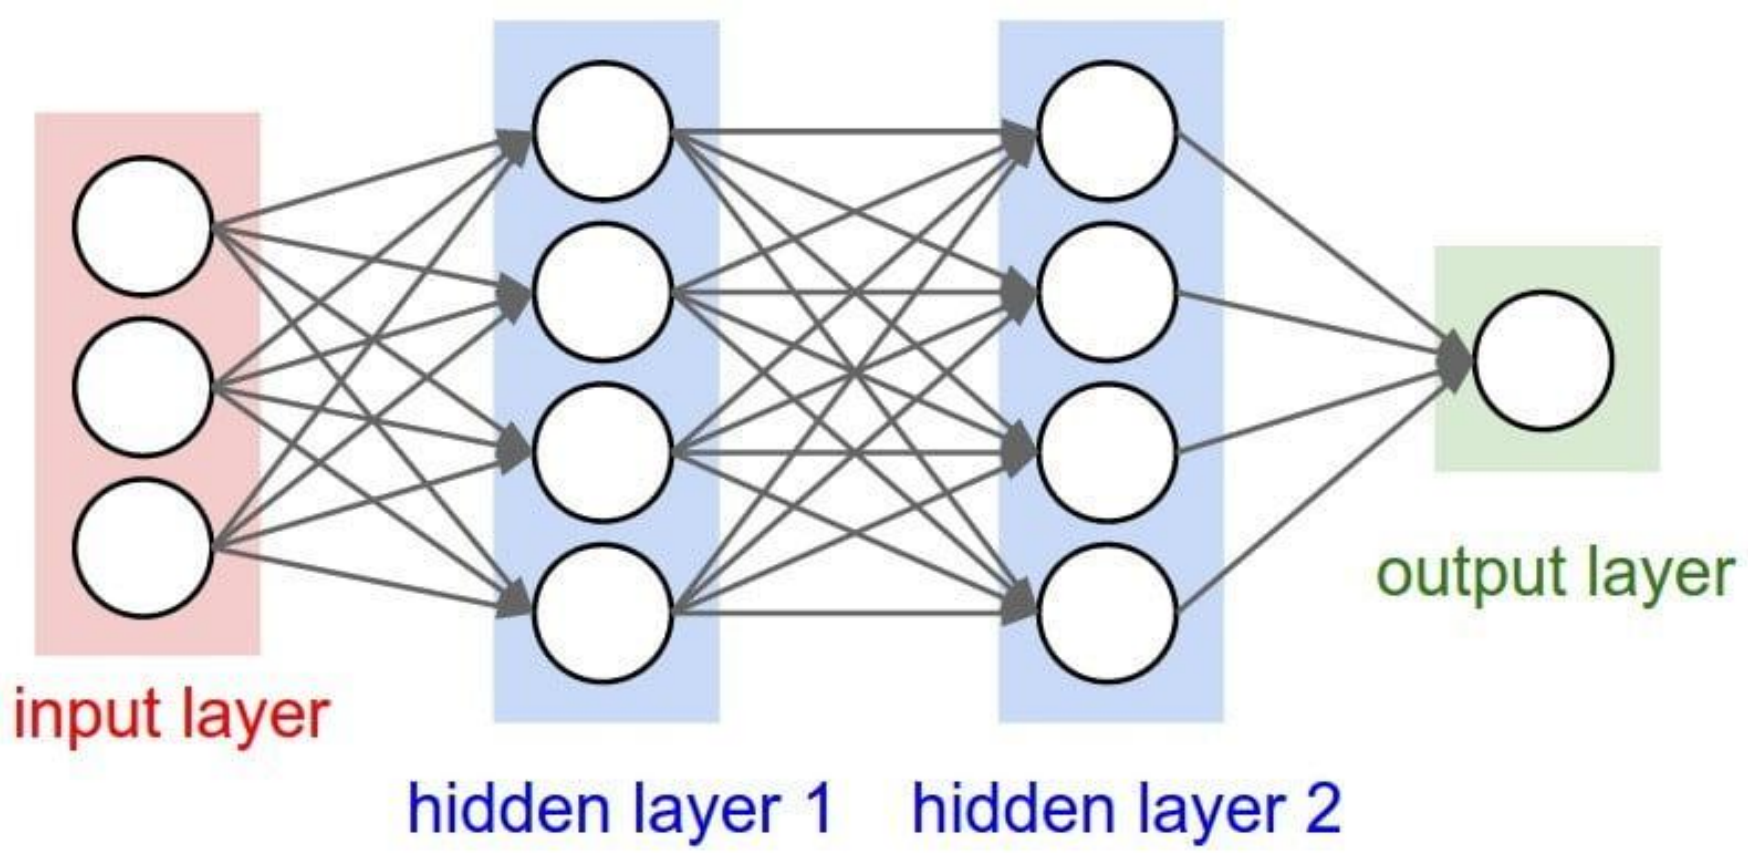

You can design
network **architectures**
for pattern recognition

Seeing the bigger picture
requires abstraction and distance.

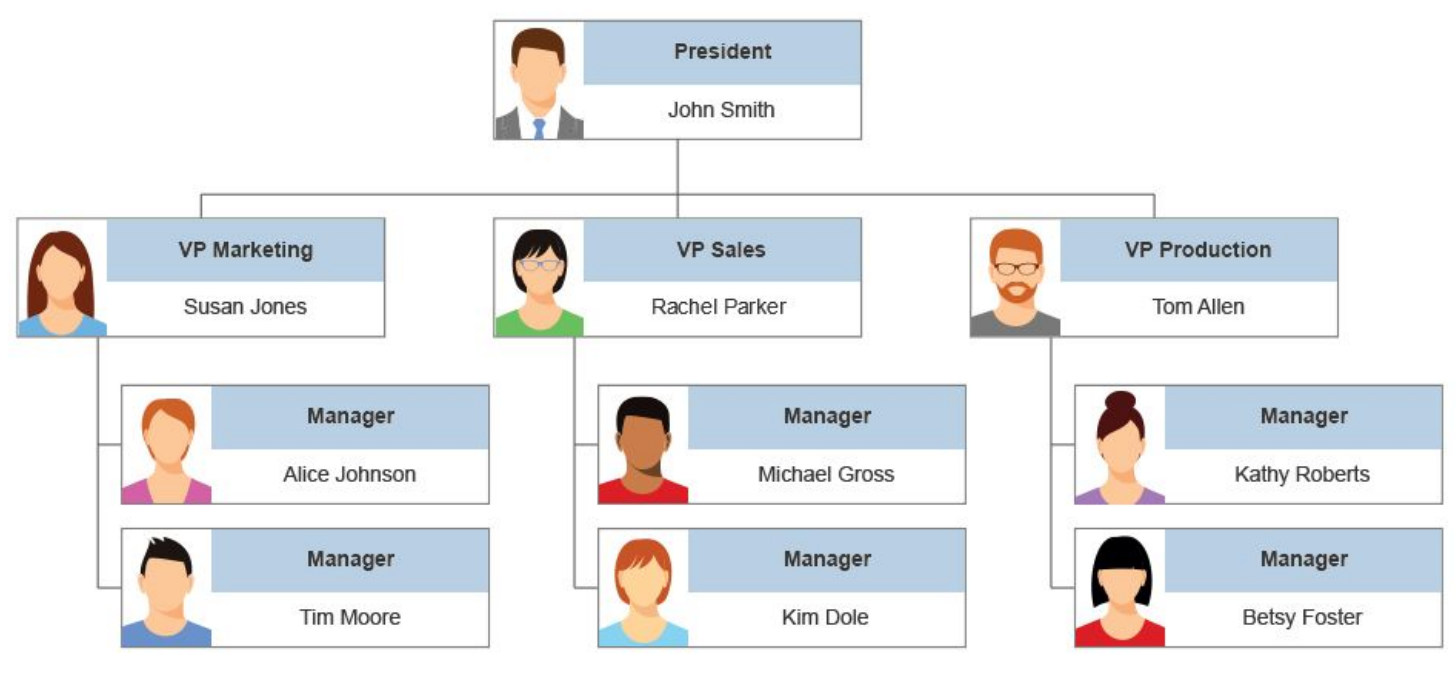

Sensing pyramids abstract details

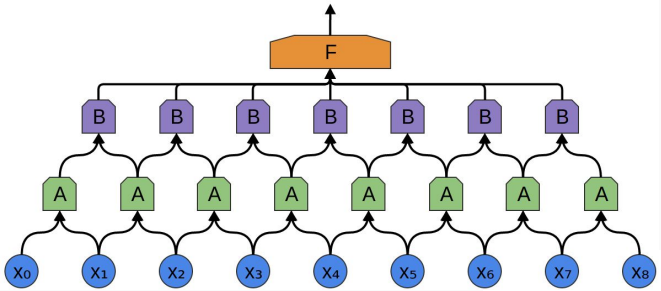

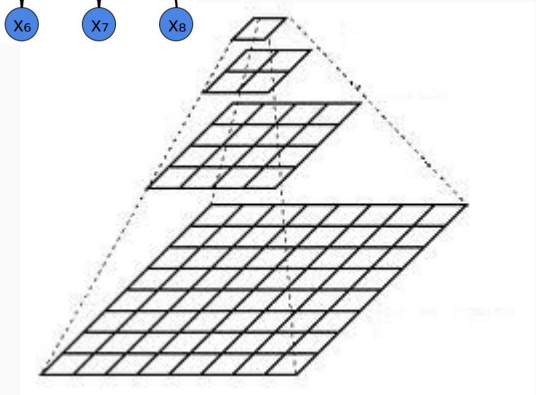

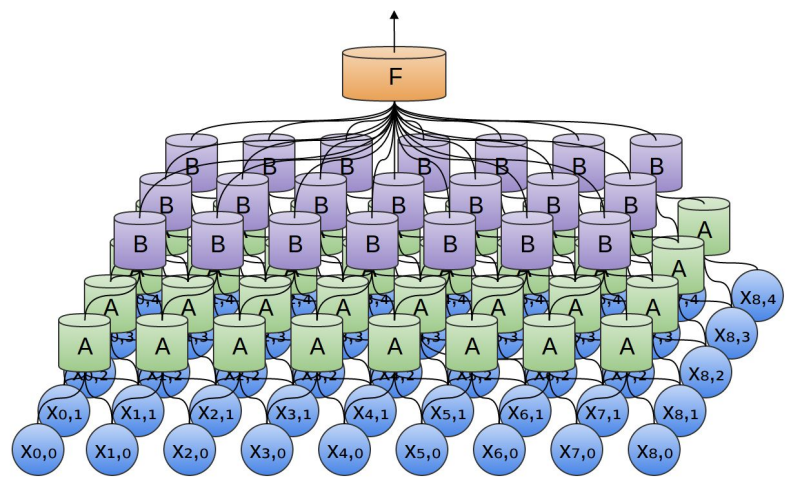

Example: Classify handwritten digits

0
1
2
3
4
5
6
7
8
9

0
1
2
3
4
5
6
7
8
9

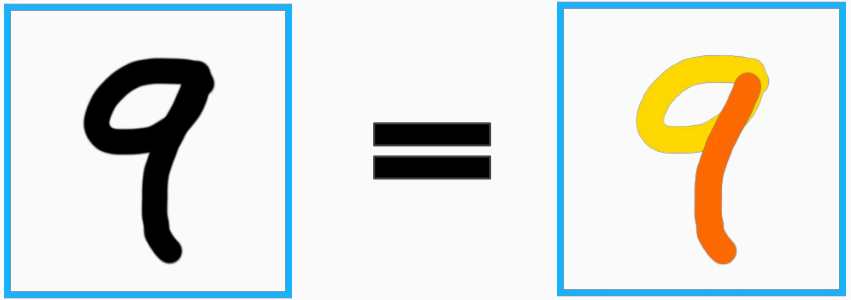

Nested Patterns

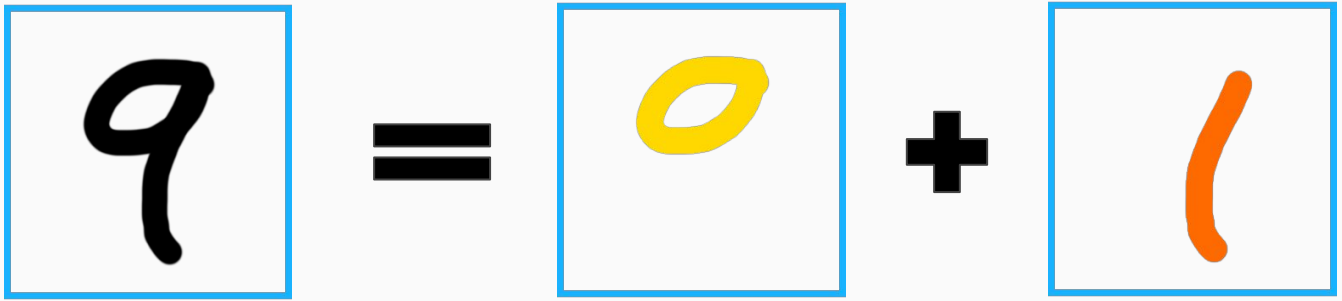

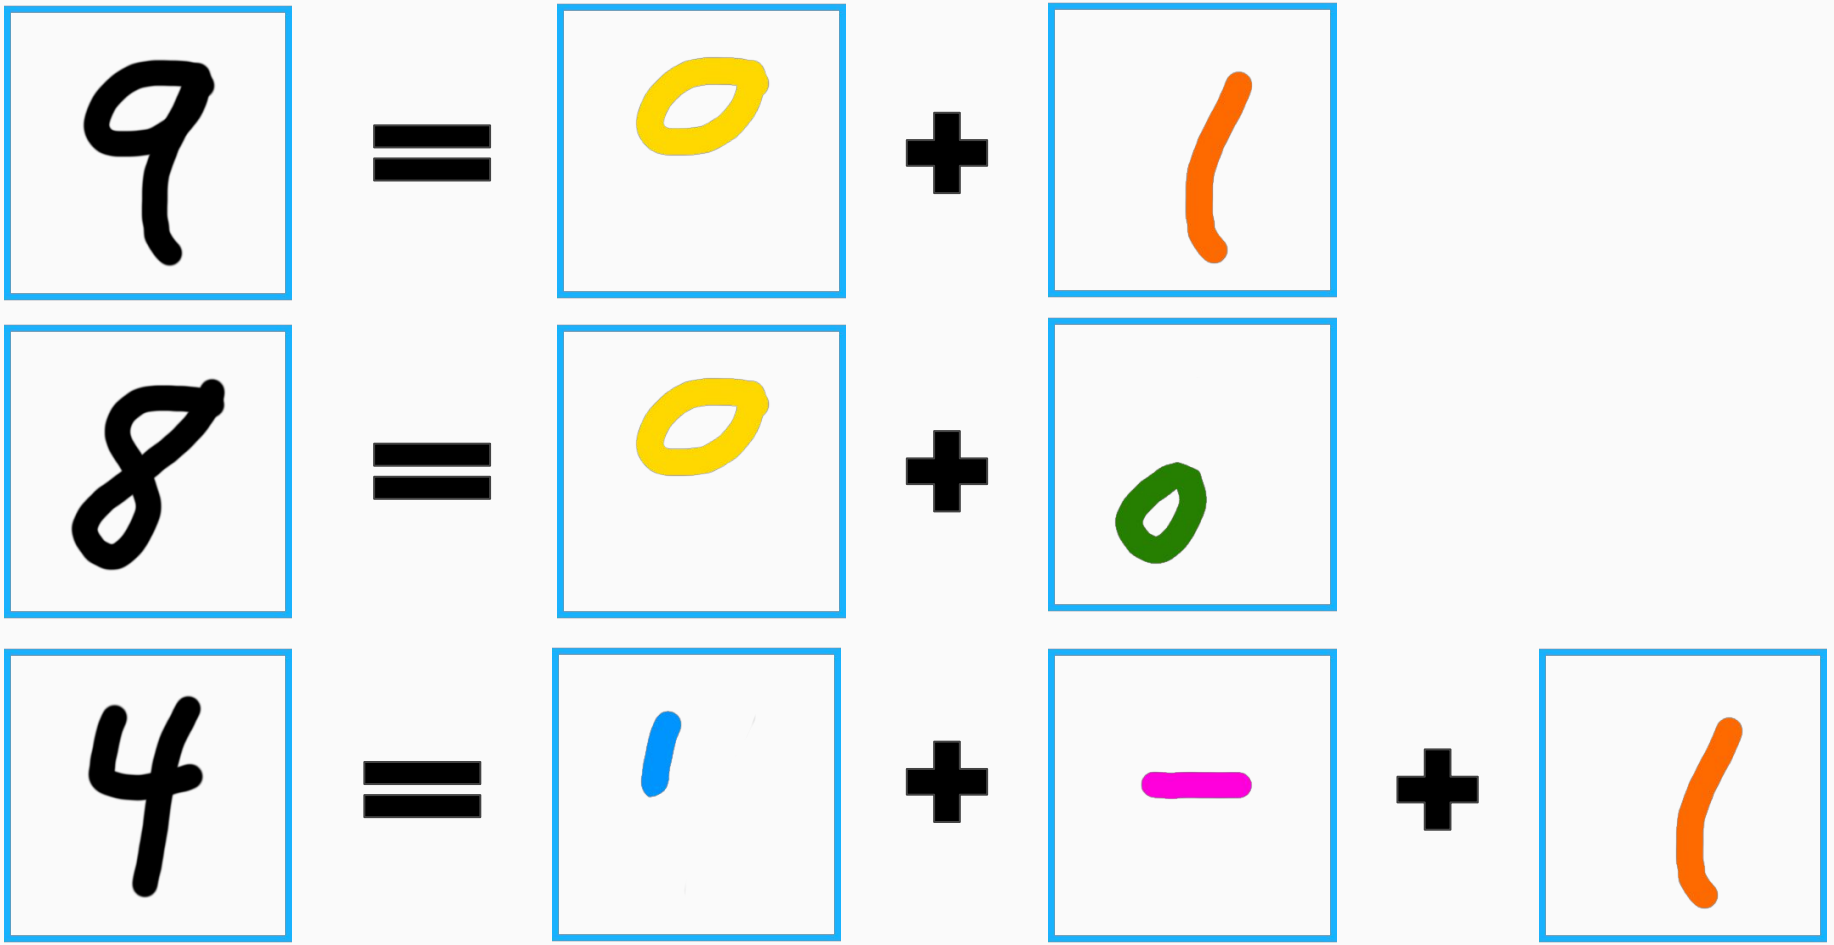
9 = 0 + 1
8 = 0 + 0
4 = 1 + − + 1

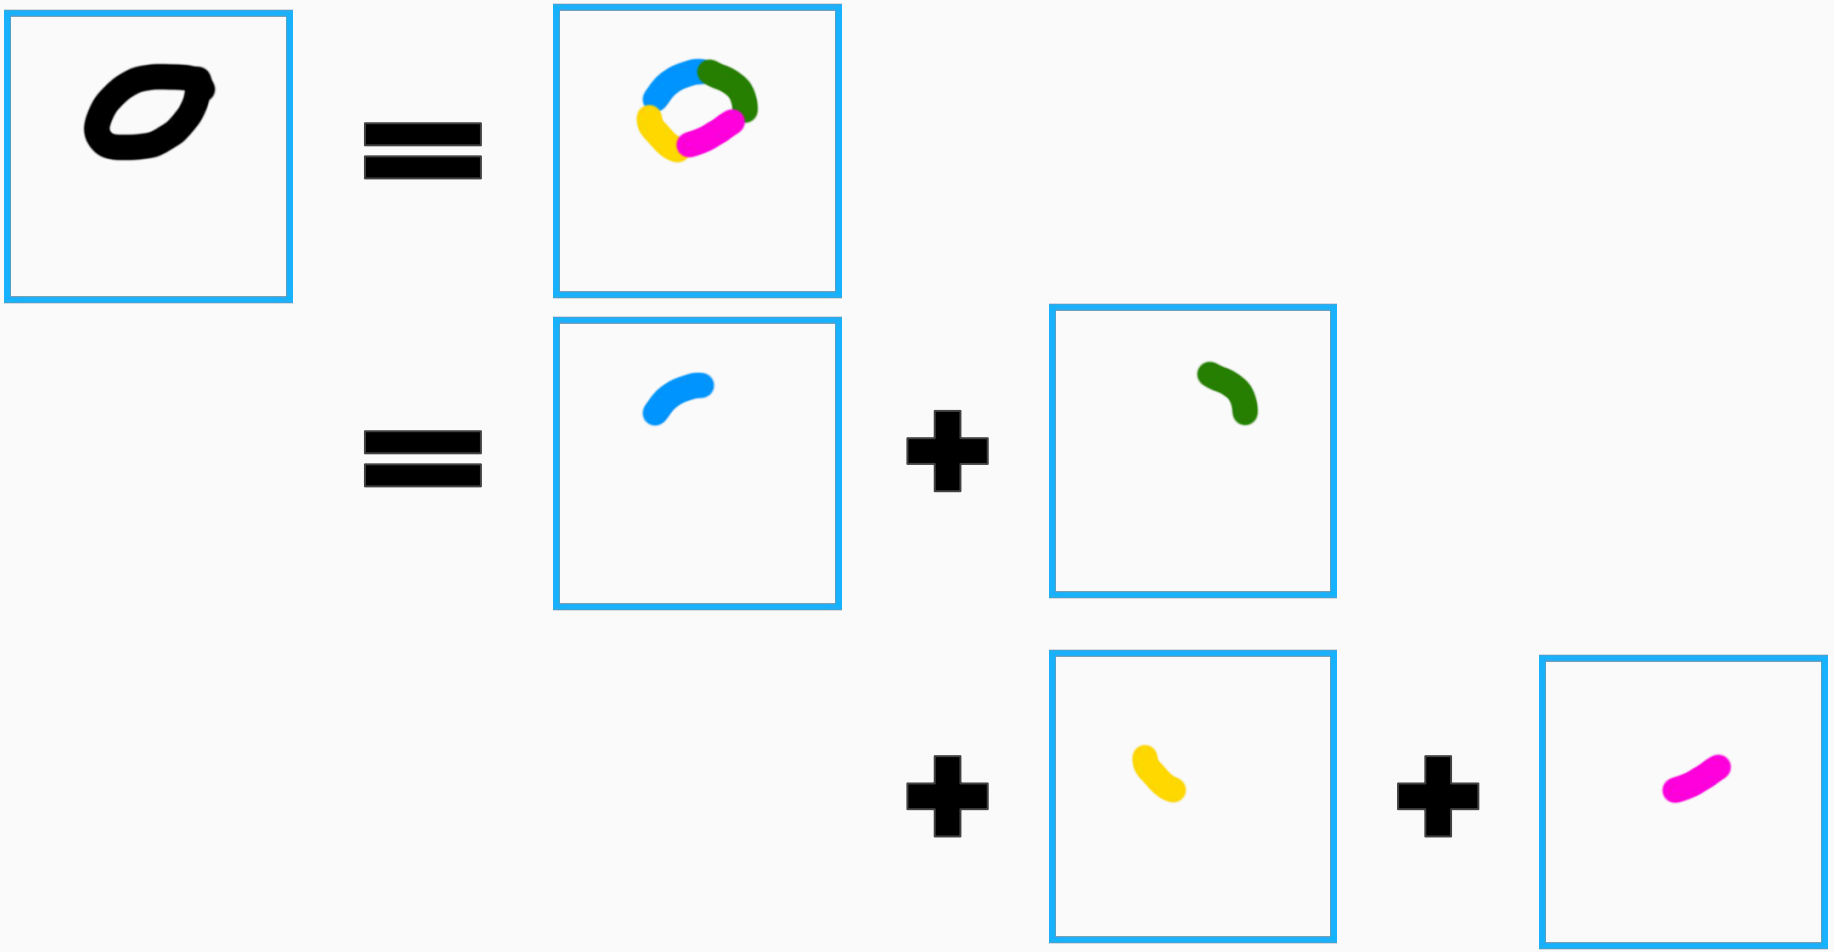

Nested Patterns

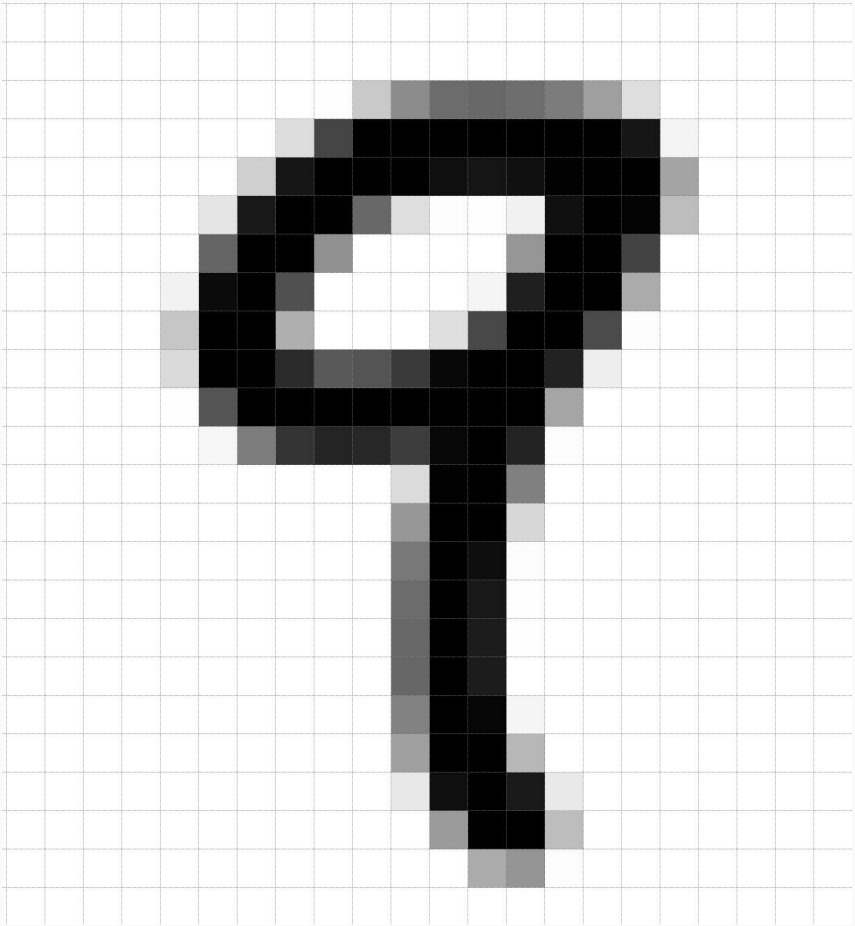

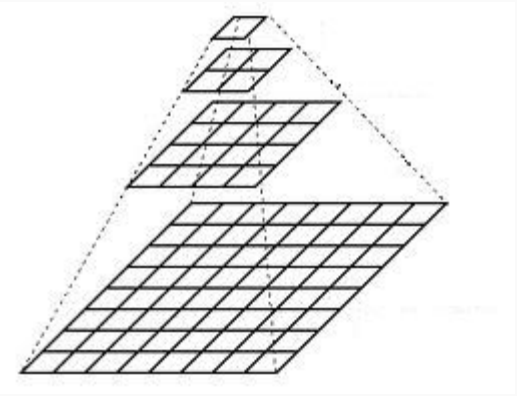

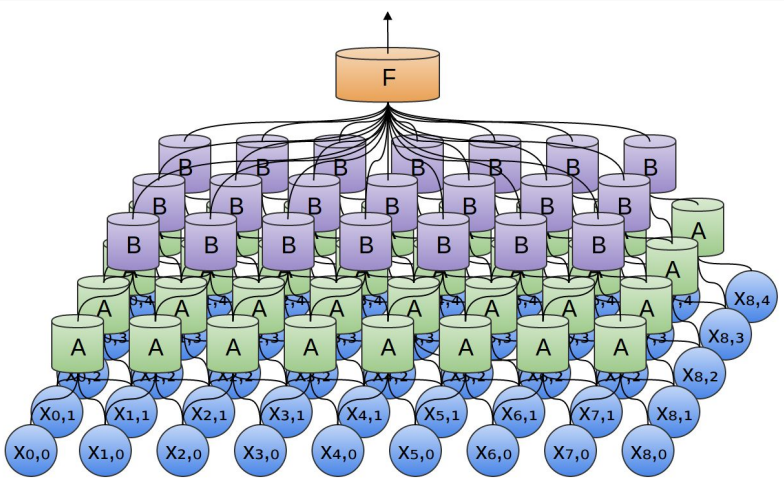

We limit detail at a certain level (pixels)

Human perception is **discrete**

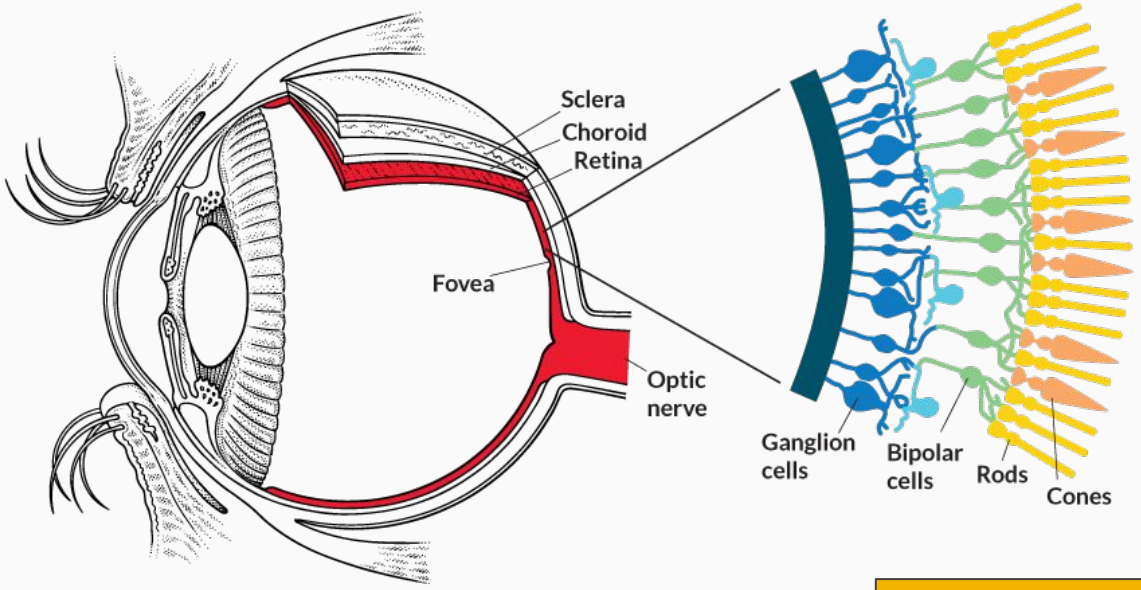

Light

At one end, a network must **connect to data**

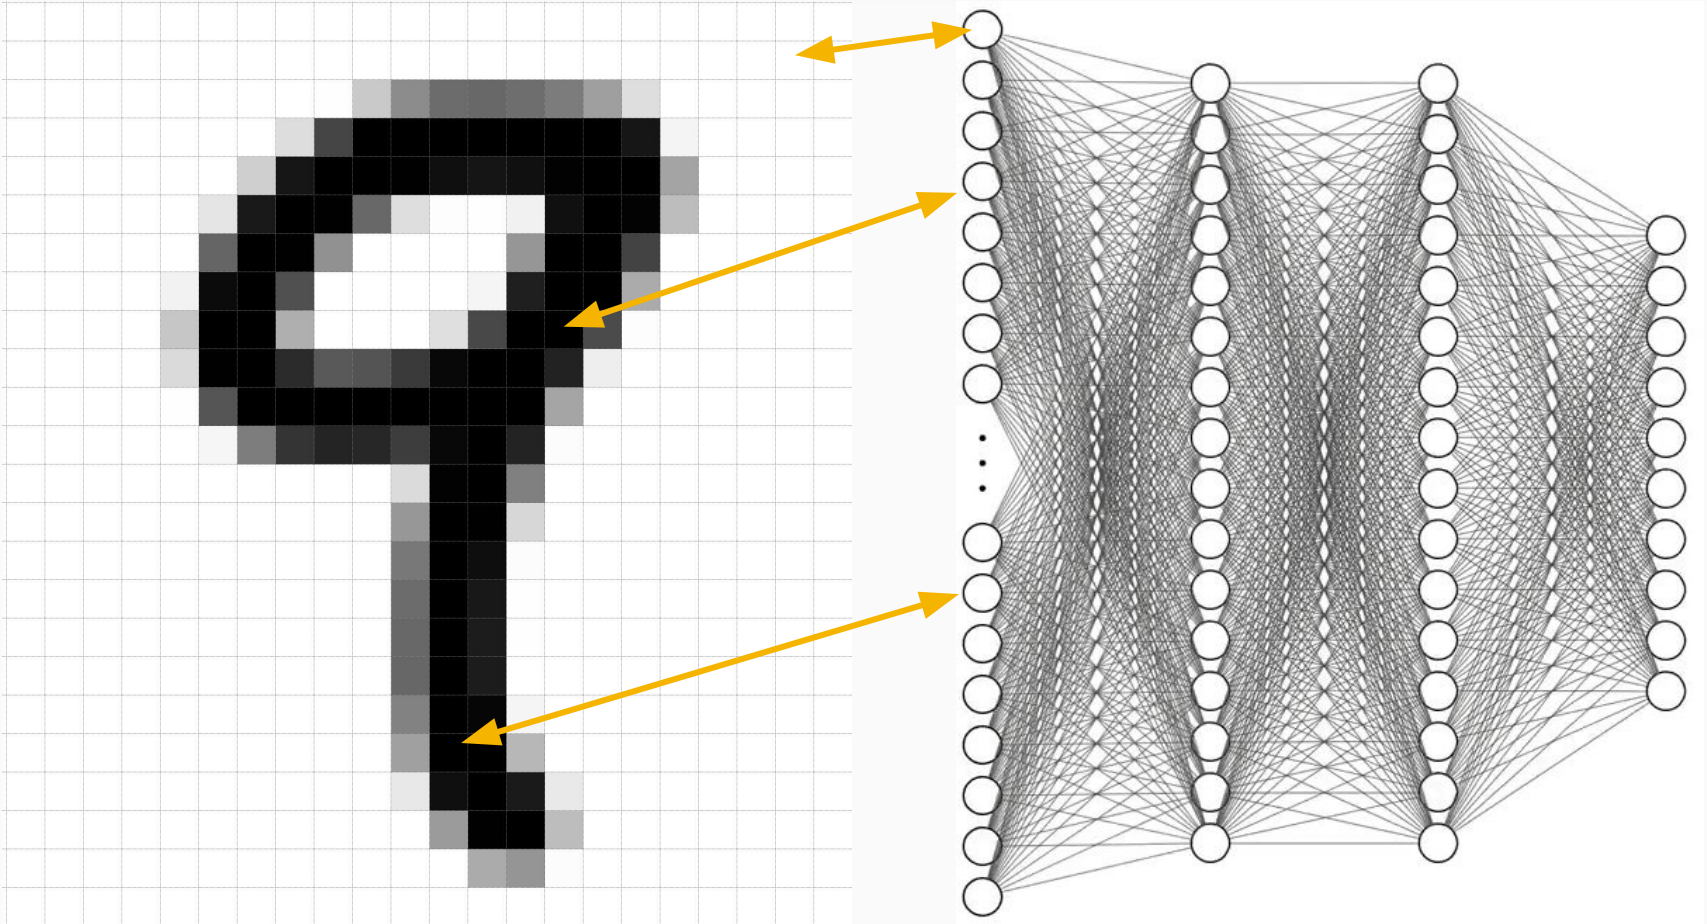

Different inputs will **excite** different pathways in the network

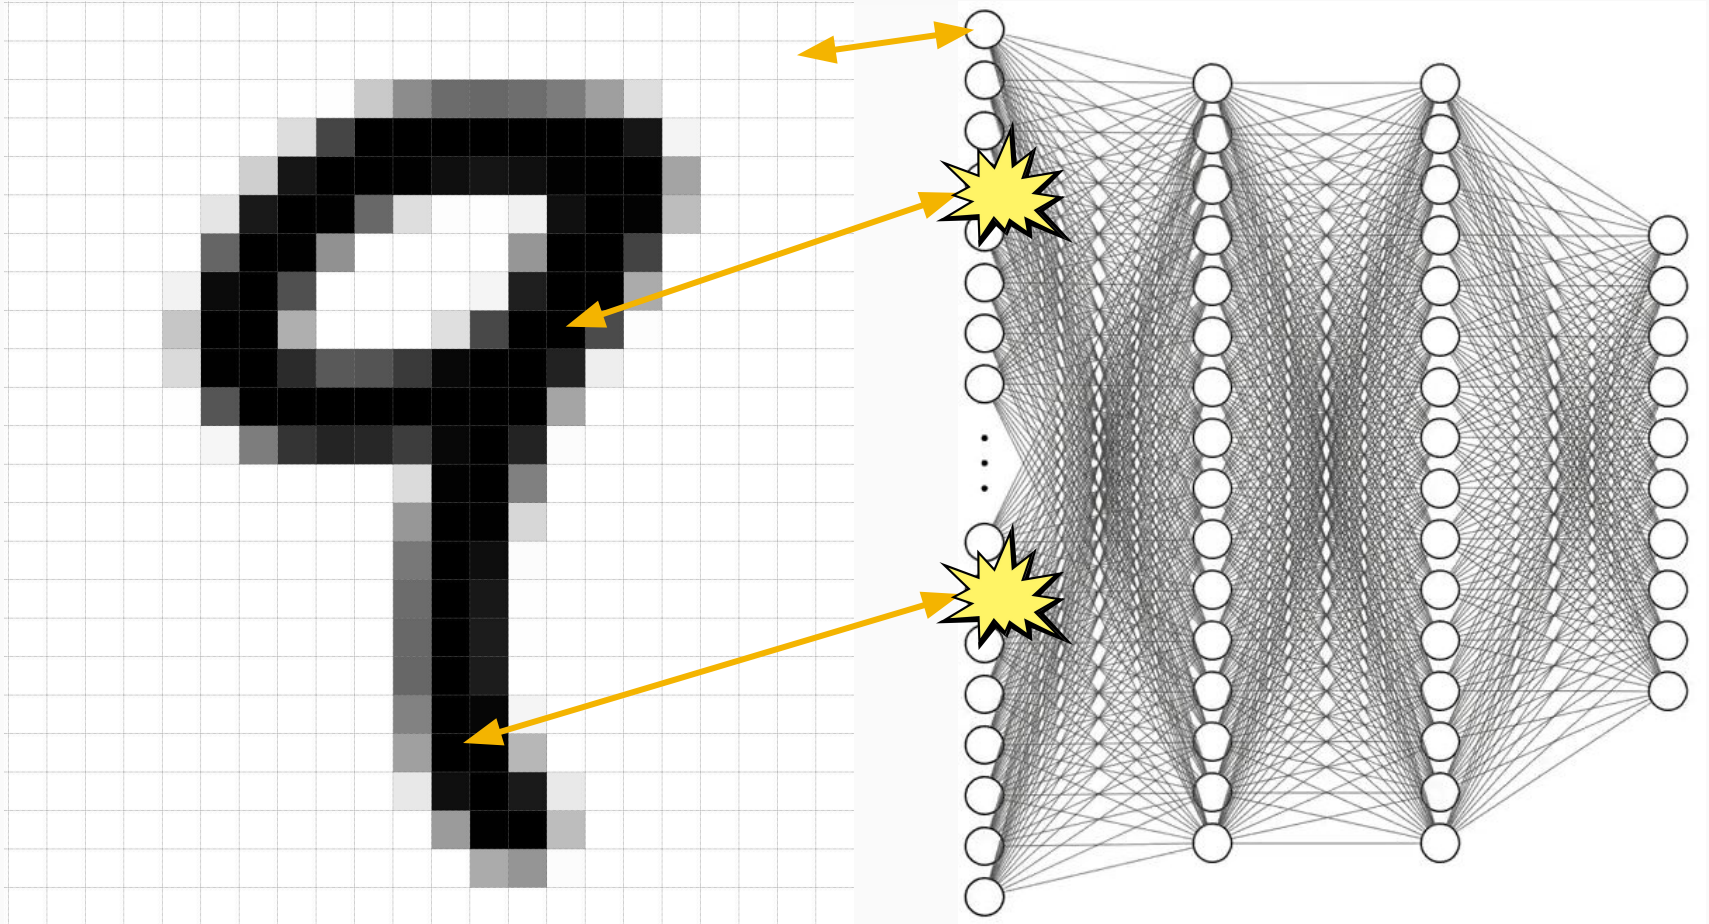

Different inputs will **excite** different pathways in the network

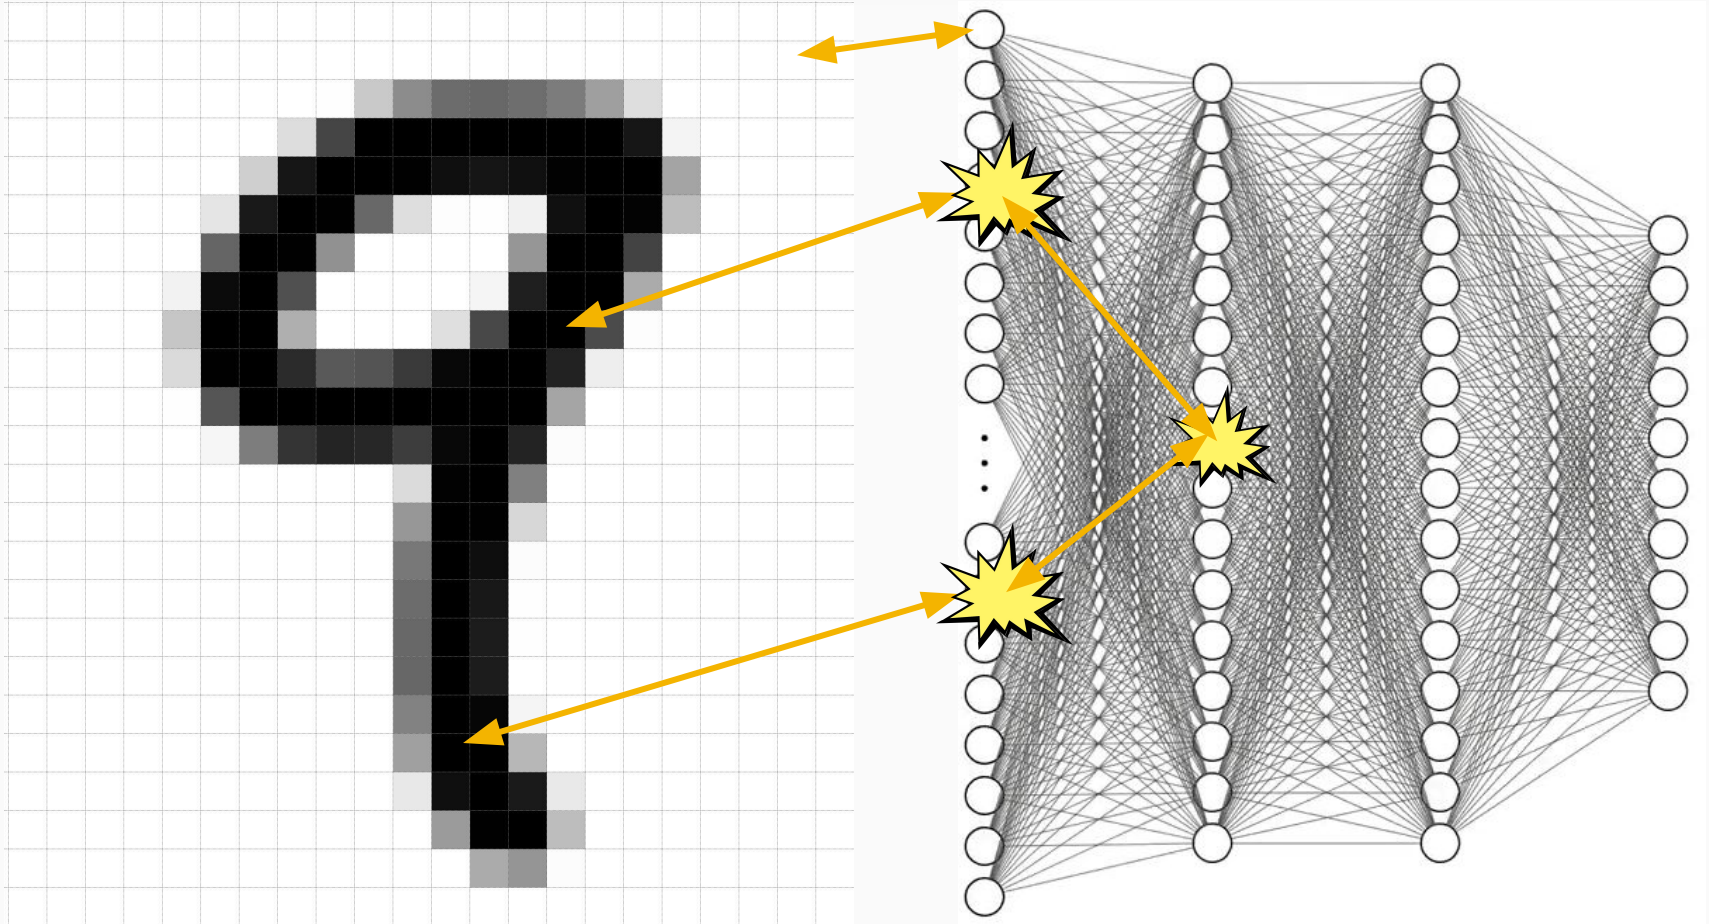

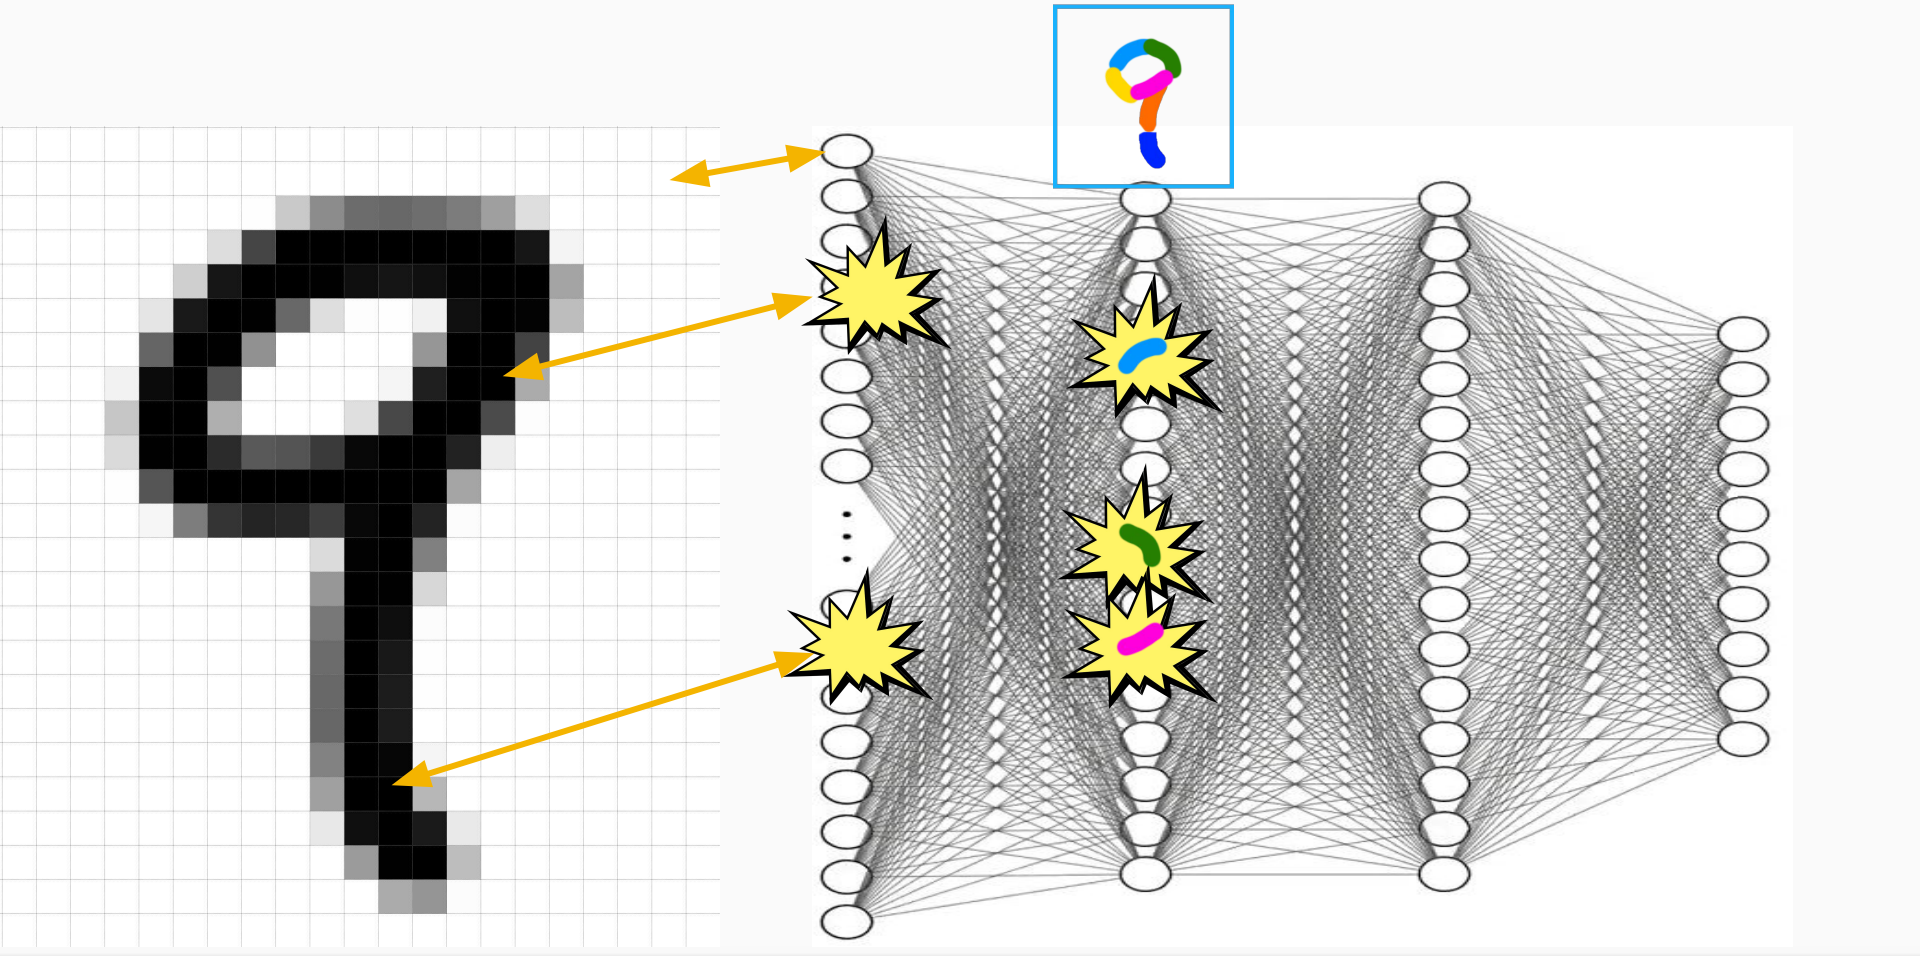

Networks consider granular details, in order to recognize **larger patterns.**

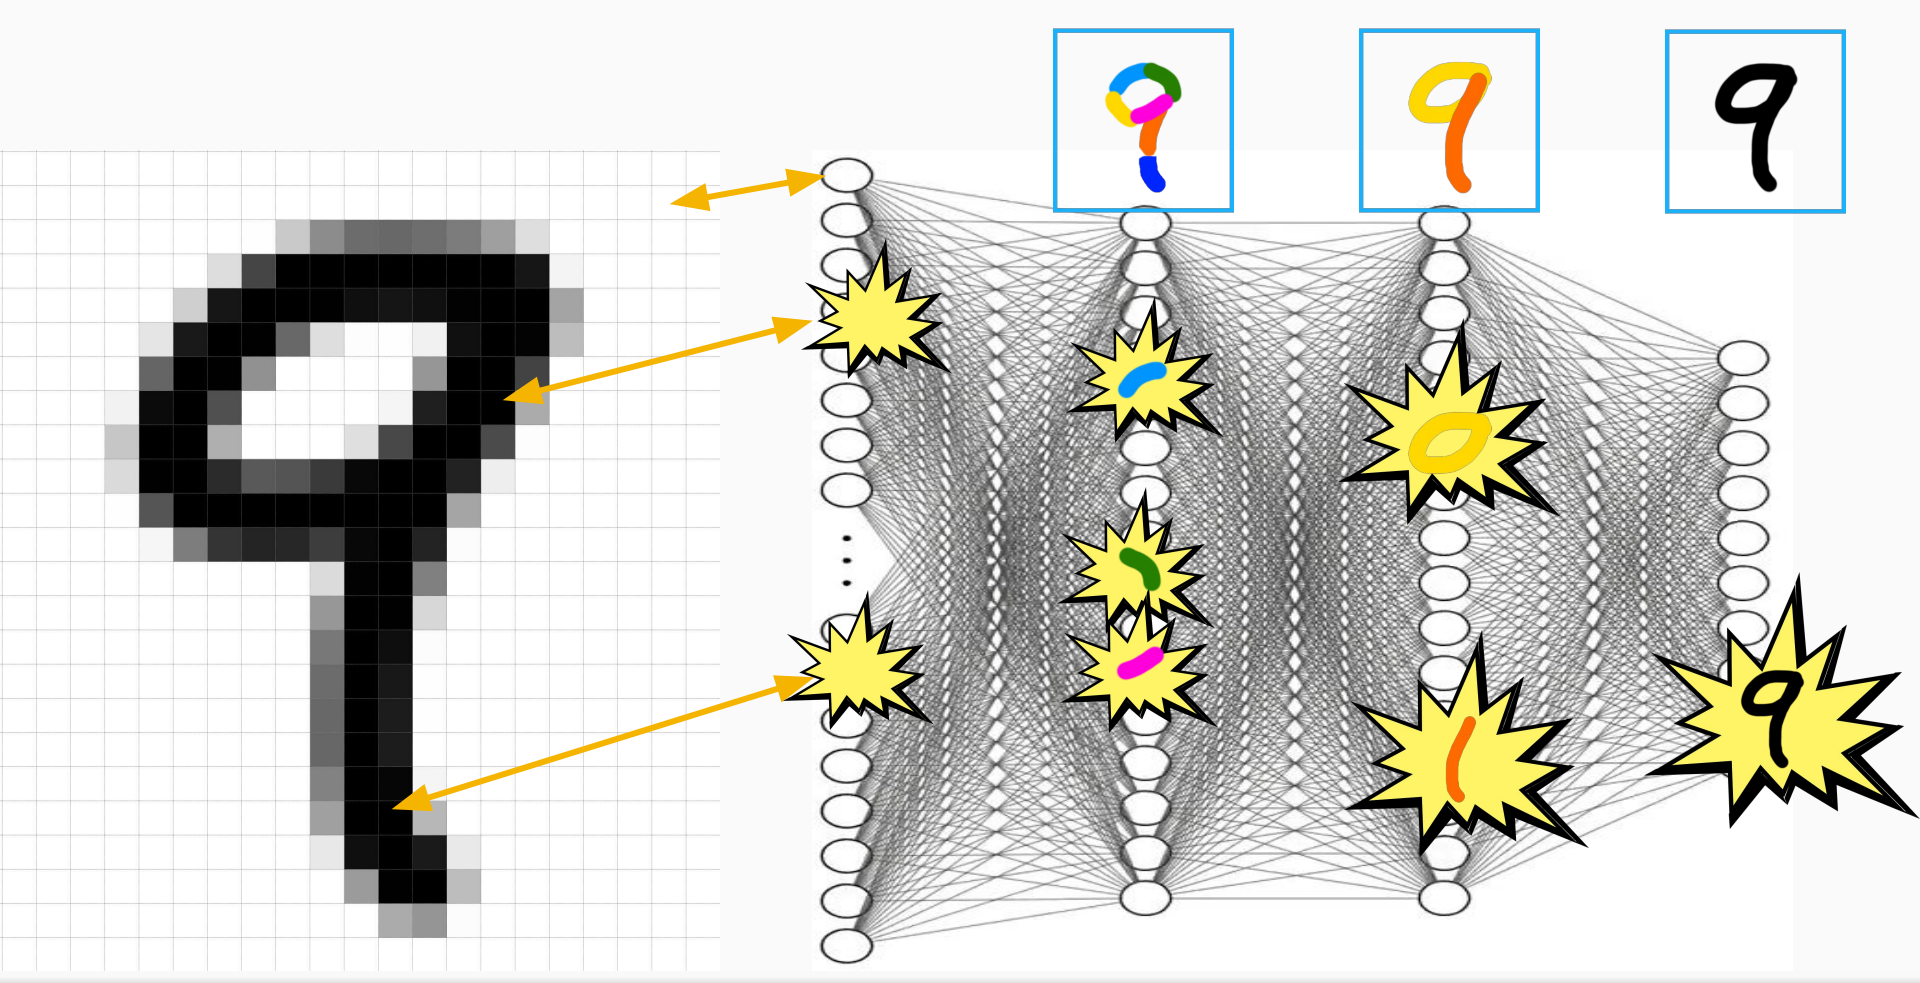

Networks consider granular details, in order to recognize **larger patterns.**

Neural networks see the world in **nested patterns**

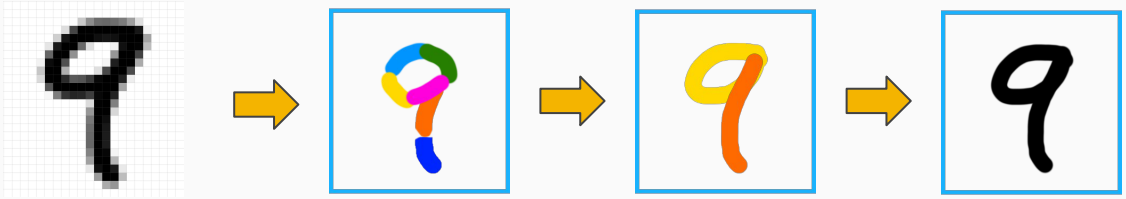

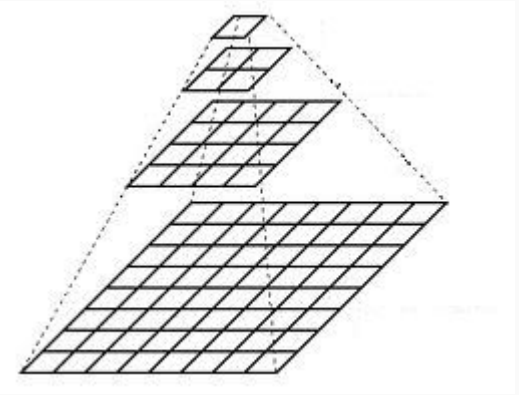

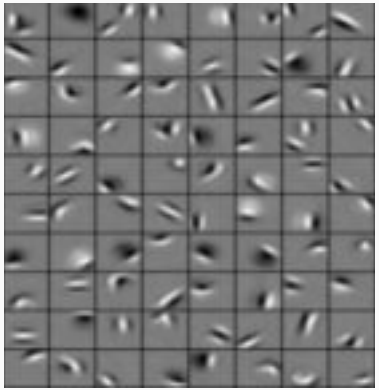

Nested patterns
are everywhere.

Neural Networks “see” the patterns in data like an image...

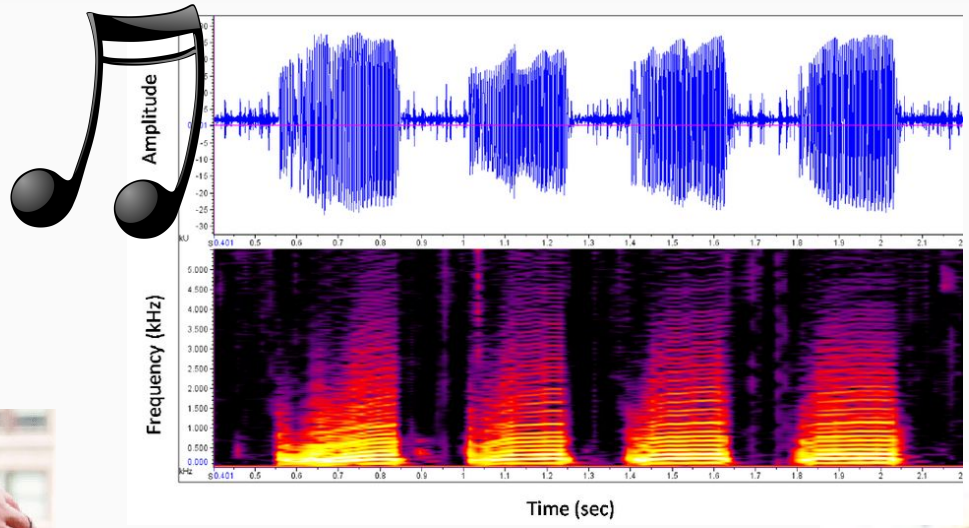

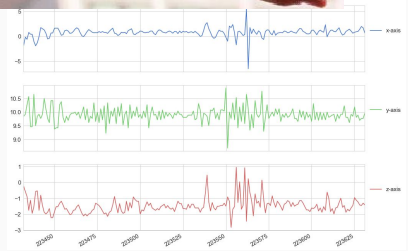

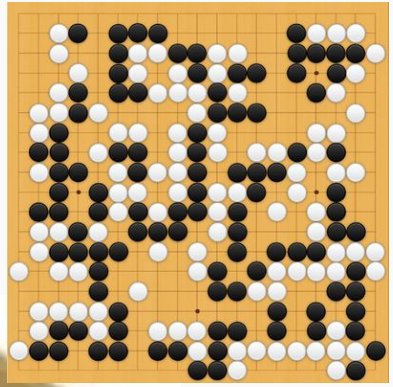

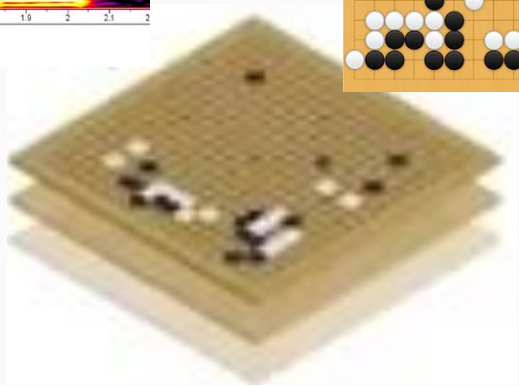

Question:

How does a neural network know which pathways are important?

Learning a rule is like repeatedly failing an exam...

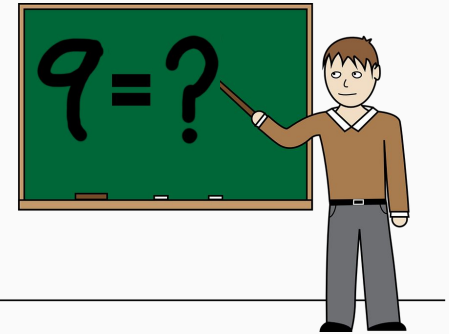

6

8

9

This is a slow, data hungry evolution

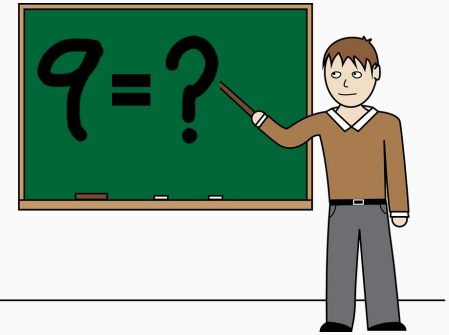

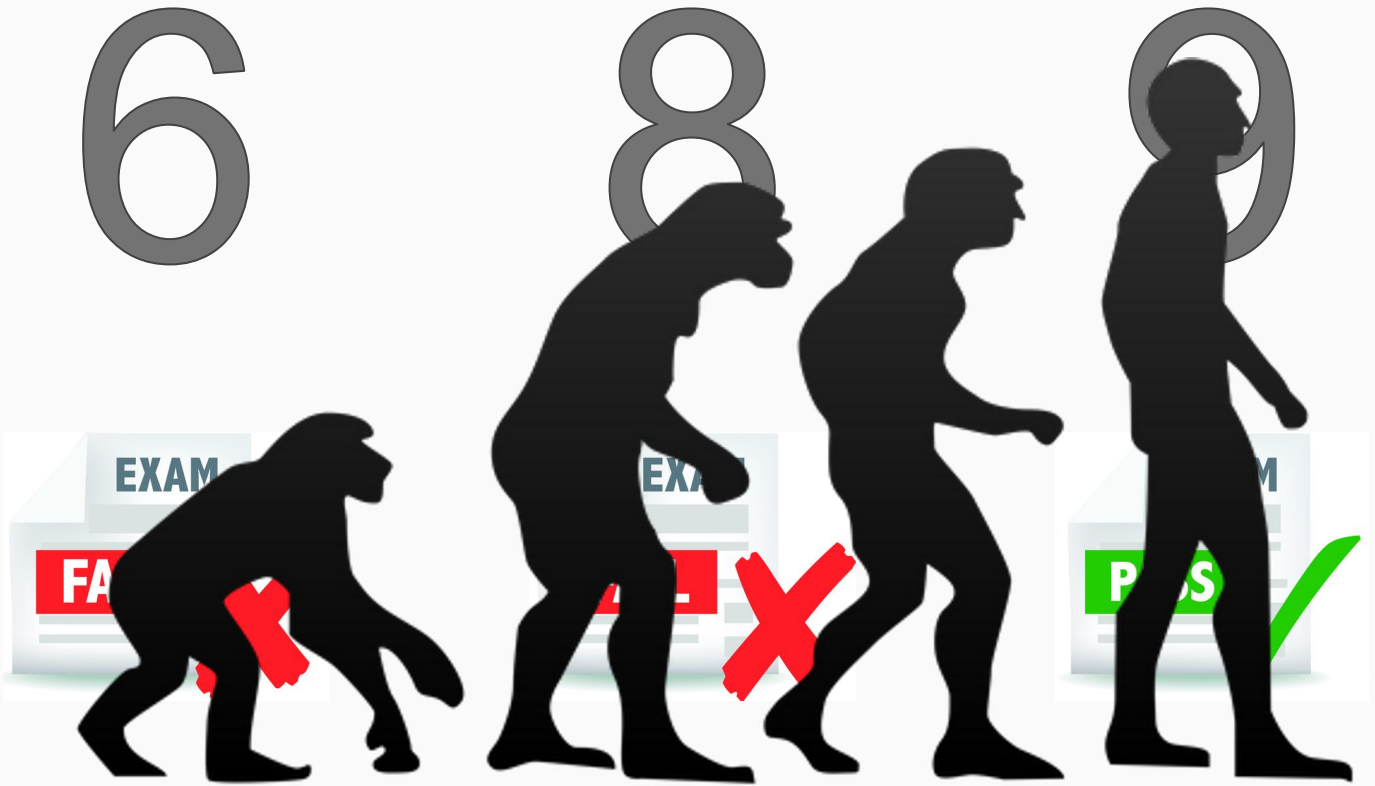

Learning by memorization requires many examples.

Learning by memorization requires many examples.

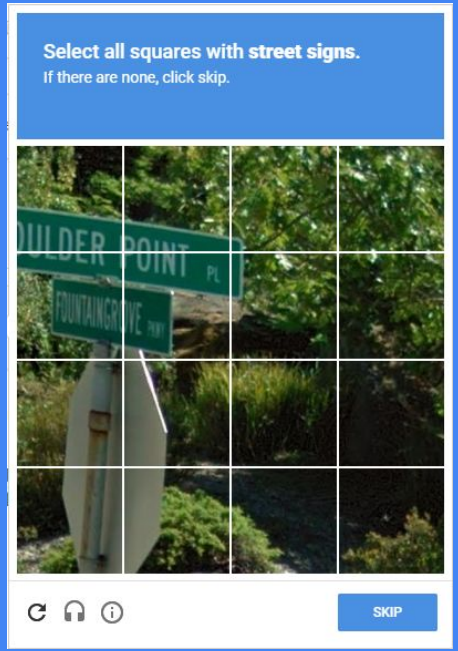

Current state of deep learning

- STUDY data to interpret new examples
- Build representations that exploit simple rules
- Recognize nested patterns
- Abstract details in complex inputs, present concise outputs

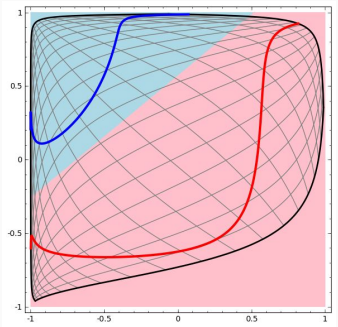

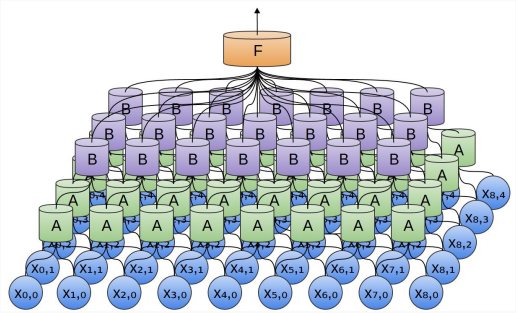

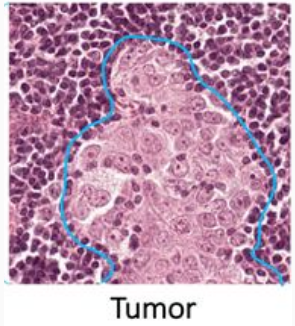

Reinforcement Learning will move beyond interpretation to strategy

Reinforcement learning involves exploration

- Learning by “doing”
- Goal driven
- Attempting actions and evaluating outcomes
  - Simulated or real
- Data is created, NOT provided
  - Removes need for existing data

Reinforcement learning involves exploration

- Learning by “doing”
- Goal driven
- Attempting actions and evaluating outcomes
  - Simulated or real
- Data is created, NOT provided
  - Removes need for existing data

Explore safely in simulation, Transfer knowledge to reality

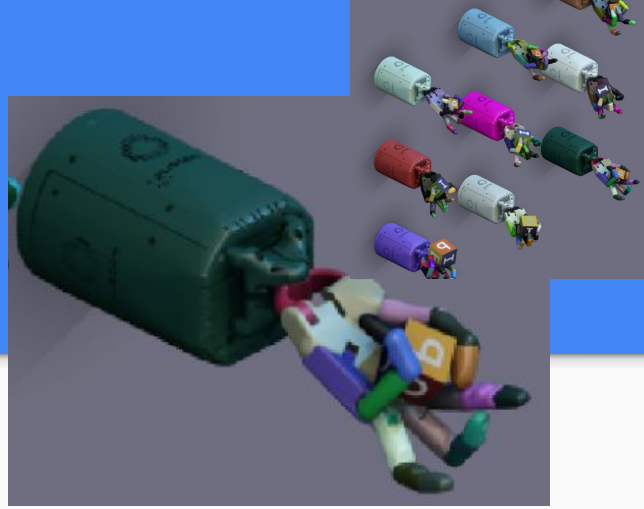

FINGER PIVOTING

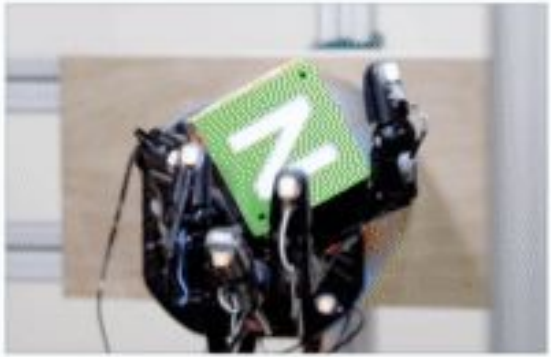

SLIDING

FINGER GAITING

Exploration leads to **novel strategy**, (compared to learning by example)

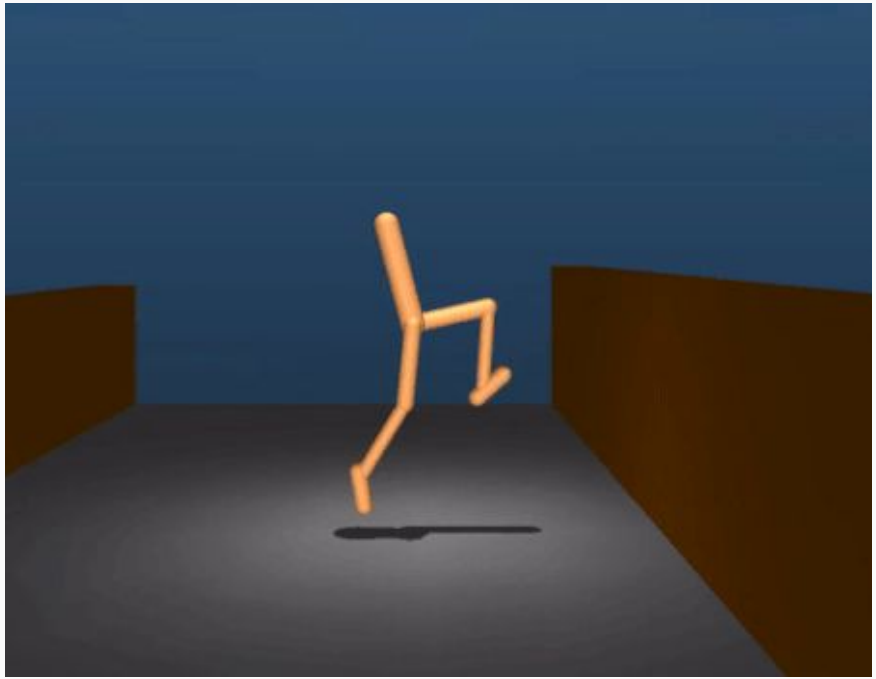

Self-play avoids learning plateau, ignoring limitations of teacher

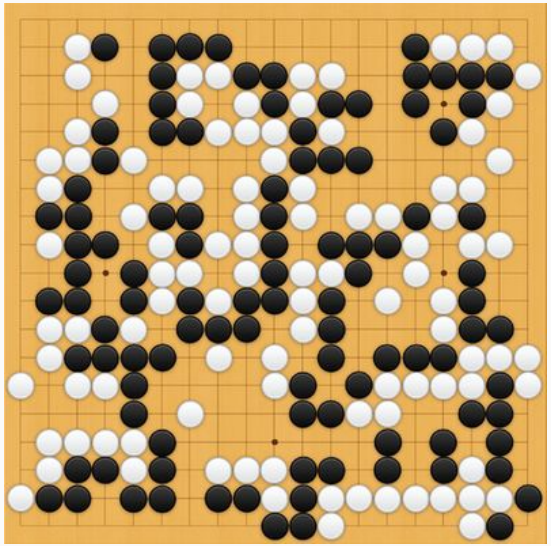

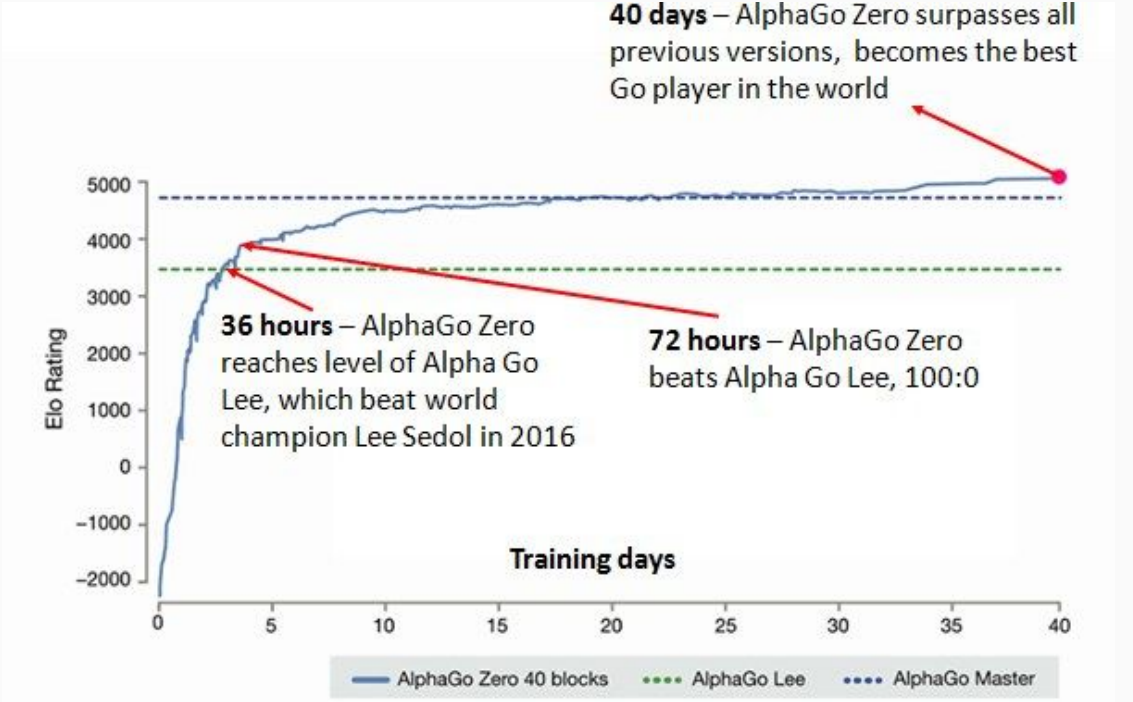

Cooperative strategy is difficult to define, but can be reinforced

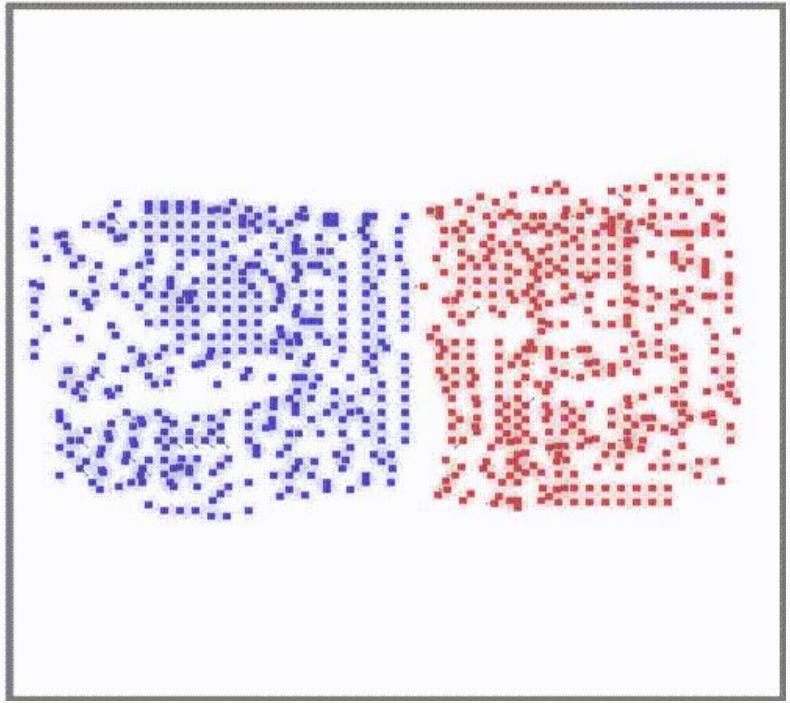

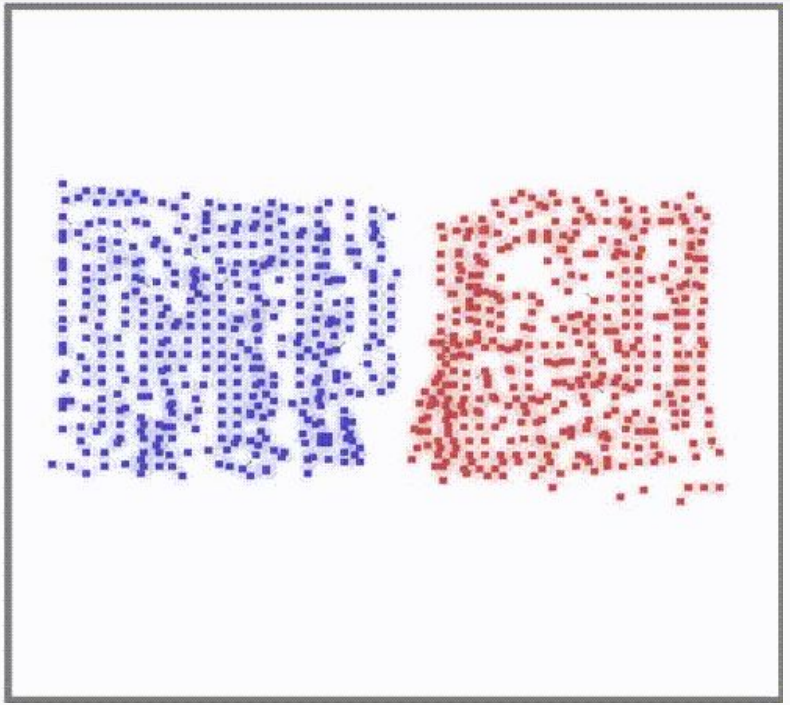

Deep Learning

Sensing systems

Reinforcement Learning

Decision making systems

Deep learning will **improve our senses**

Identify things we can't see

- Anomalies, Problems and Targets
  - Tumors, leads, customers,
- Needs and Opportunities
  - Which student is struggling?
  - Which patient needs medical attention or therapy?

Summarize complex details, and simply point us to the problem.

Reinforcement Learning will **optimize**, and **solve problems**

We'll provide high level goals, or desired outcomes

- Pick up the cup
- Reduce congestion
- Increase profit
- Maximize yield
- Retain user
- **Improve health**

Neural Networks
Demystified

Thank you!

Contact: david@questionablyartificial.com

Bonus:
(if there is time)

Training Neural Networks is a tiny part of a Cyclical, Integrated Process

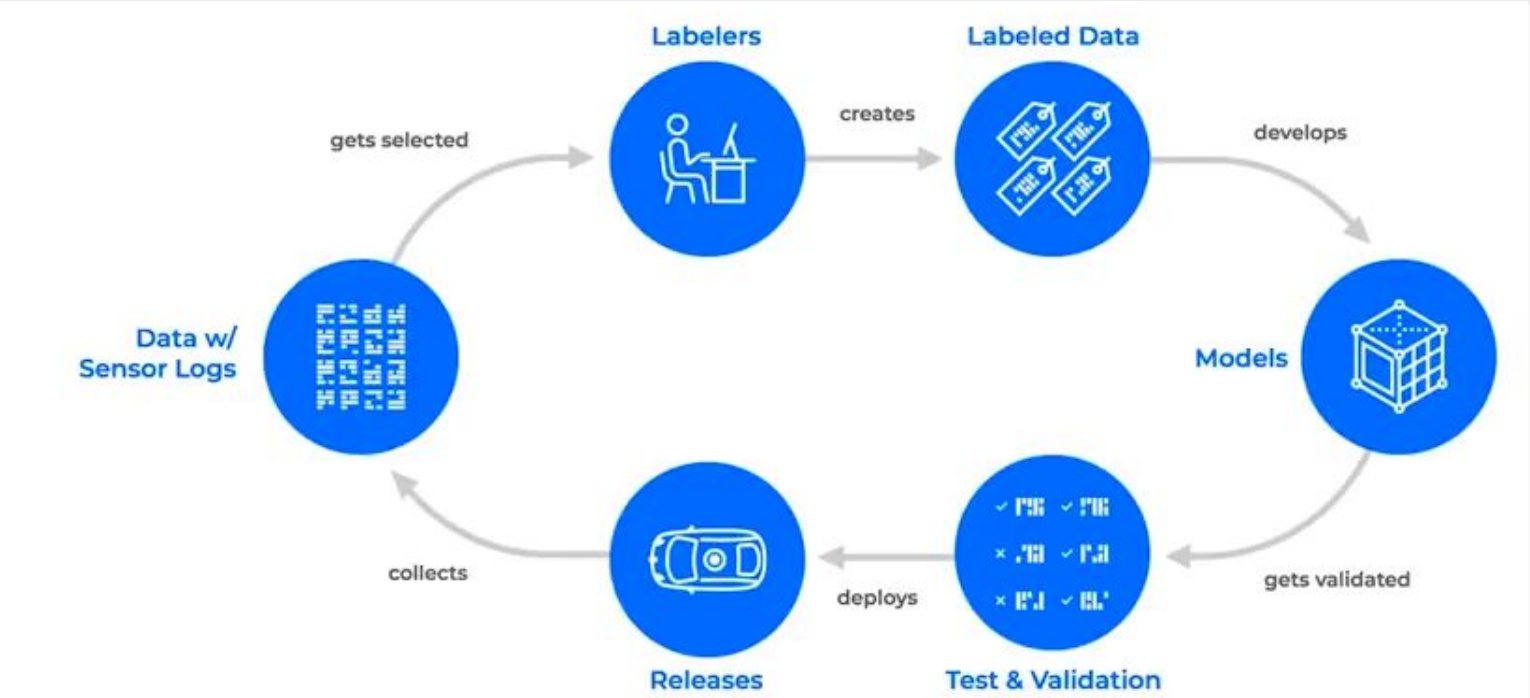

AI is the cognitive component, but the grander problem is Sensing Infrastructure.

Cognition cannot be developed without a sensing infrastructure.

Development infrastructure is crucial

Labeling

Simulation

Gyms

Visualization

Dev Tools

Networking

AI development requires diverse specialists, more than AI engineers.

- Developer tools
- Low level hardware optimization
- Edge compute and sensing
- Sensor fusion
- Sensor manufacturers
- Data storage, organization, and retrieval.

Translation

Complex Input

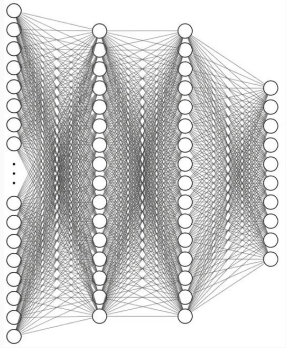

Concise Output

Translation

Complex Input

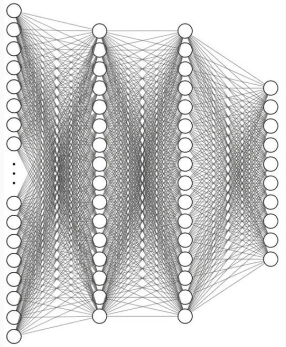

Concise Format

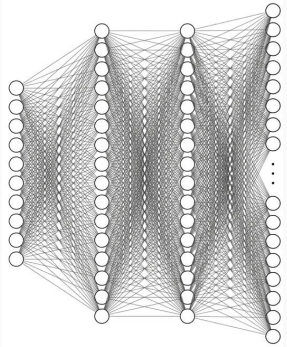

Complex Output

Translation

Complex Input

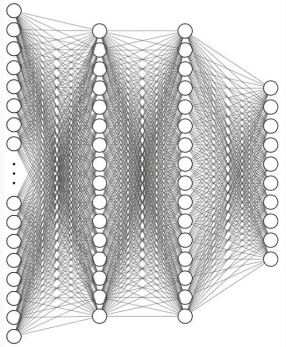

Translation

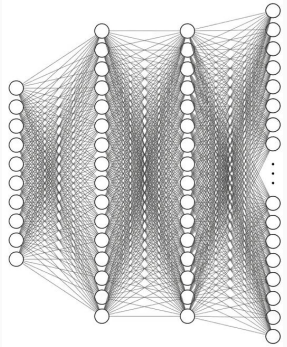

Complex Output

Translation

Complex Input

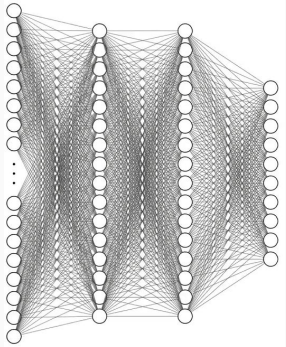

hello

Translation

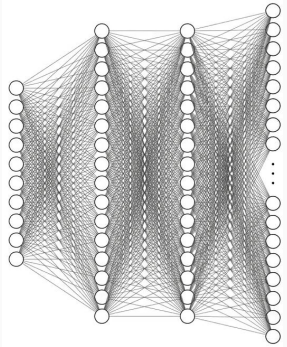

Complex Output

你好

Neural networks can learn to translate, even when we don’t don’t know how

Complex Input

Translation

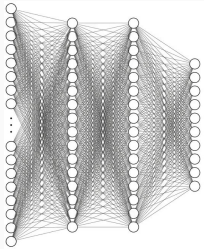

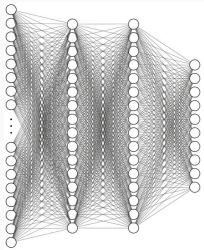

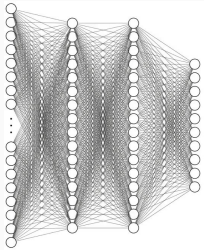

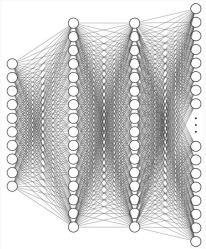

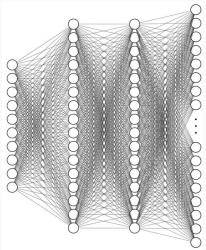

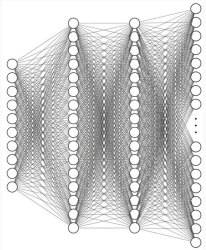

Complex Output

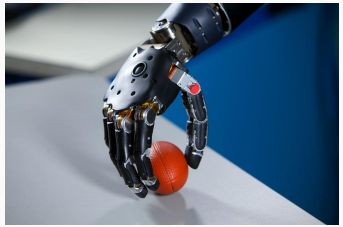

Translation

Complex Input

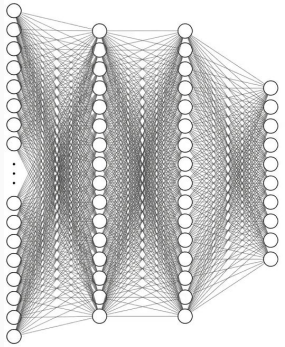

Translation

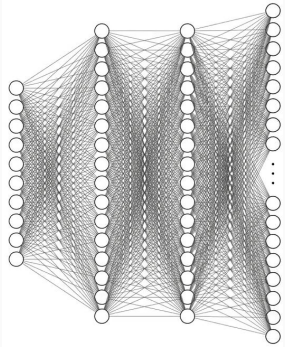

Complex Output

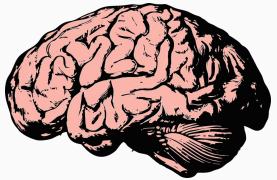

Translation

Complex Input

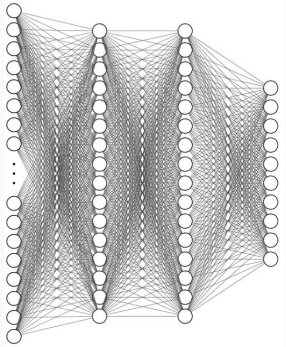

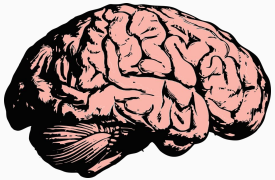

Translation

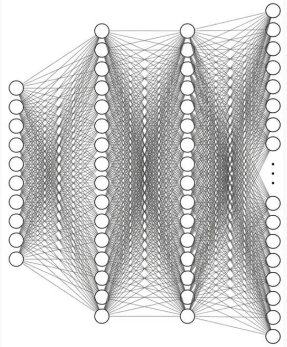

Complex Output

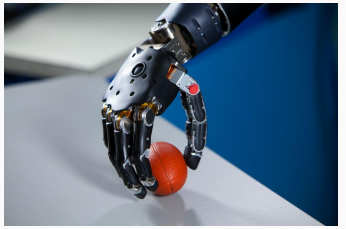

Translation

Complex Input

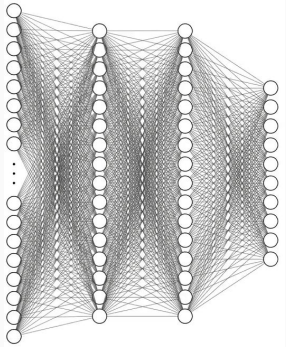

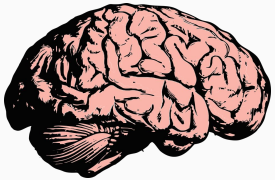

Translation

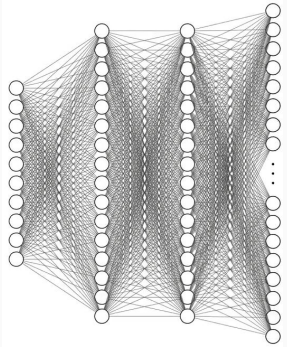

Complex Output

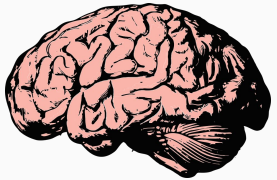

Human perception is **discrete**

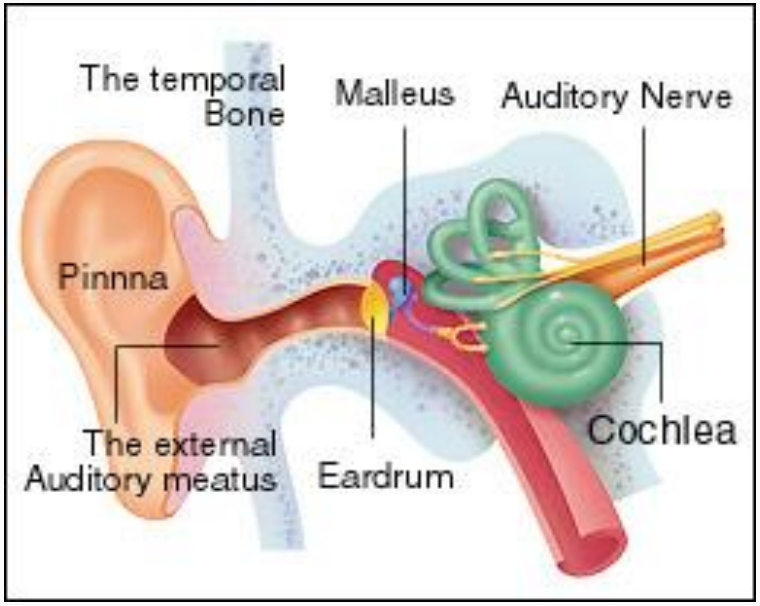

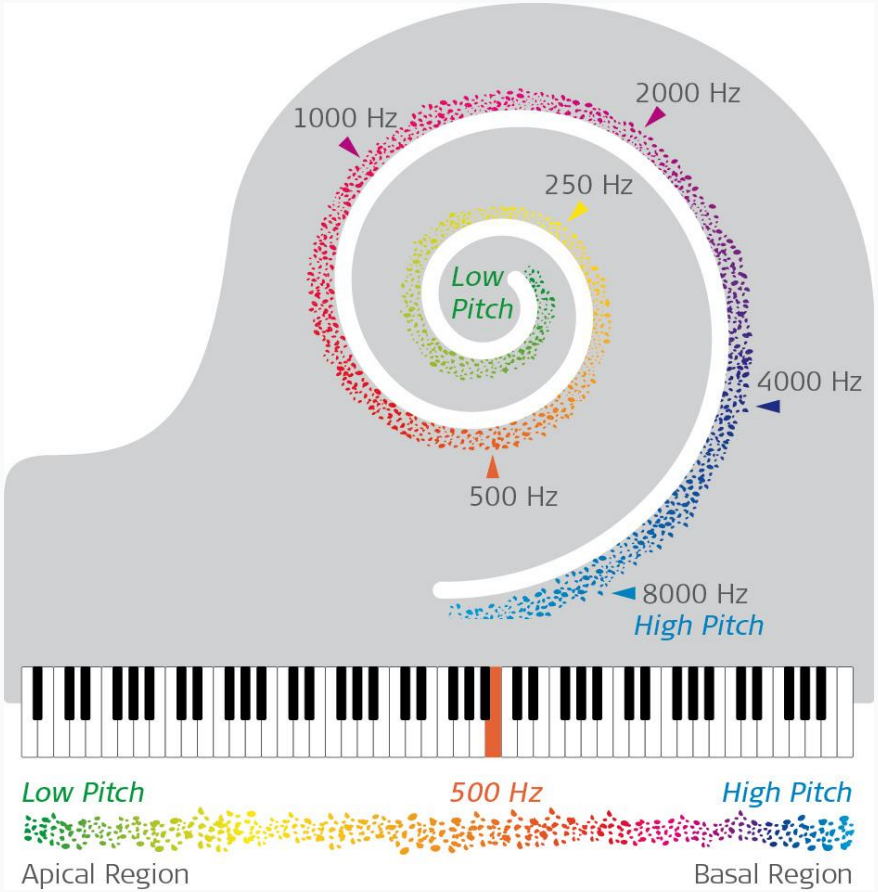

References

Some imagery pulled from various sources, most notably:

- OpenAI
- colah.github.io
- Drago Anguelov (Waymo) MIT Self-Driving Car Diagram
- 3Blue1Brown