The Brain and the Modern AI

Drastic **differences** and curious **similarities**

Ilya Kuzovkin

UNIVERSITY OF TARTU 1632

OFFWORLD

A supercomputer that works like the human brain has just been turned on

Manchester University just switched on the world's largest neuromorphic supercomputer

INSIGHTS

Before you know it!

It won't be long before AI mimics a human brain to do more, with less – just like we can.

Science News Technology Space NASA Artificial Intelligence Weird Science Science of Sci-Fi

MIT Scientists Design a Computer That Works Just Like a Human Brain

Researchers developed algorithms that mimic the human brain (and tl results don't suck)

MARKETS BUSINESS INVESTING TECH POLITICS CNBC TV

MIT researchers develop new chip design to take us closer to computers that work like human brains

PUBLISHED MON, OCT 8 2018 • 8:26 AM EDT | UPDATED WED, OCT 10 2018 • 10:29 AM EDT

IBM is teaching AI to behave more like the human brain

Soon machines will be able to pay atte

SPACE PHYSICS CHEMISTRY BIOLOGY ENERGY & ENVIRONMENT

Intel Developing an AI Chip That Acts Like a Human Brain

A supercomputer that works like the human brain has just been turned on

Manchester University just switched on the world's largest neuromorphic supercomputer

INSIGHTS

Before you know it!

It won't be long before AI mimics a human brain to do more, with less – just like we can.

Science News Technology Space NASA Artificial Intelligence Weird Science Science of Sci-Fi

MIT Scientists Design a Computer That Works Just Like a Human Brain

Researchers developed algorithms that mimic the human brain (and t… results don't suck)

MARKETS BUSINESS INVESTING TECH POLITICS CNBC TV

MIT researchers develop new chip design to take us closer to computers that work like human brains

IBM is teaching AI to behave more like the human brain

Soon machines will be able to pay atte…

SPACE PHYSICS CHEMISTRY BIOLOGY ENERGY & ENVIRONMENT

Intel Developing an AI Chip That Acts Like a Human Brain

TOP 5 THIS MONTH MEDIUM TIPS SUBMIT

Do neural networks really work like neurons?

WIRED BUSINESS CULTURE GEA…

LEE SIMMONS

Why Artificial Intelligence Is Not Like Your Brain–

Contrary to belief, artificial intellige… rocket scientist.

future tense

The Wrong Cognitive Measuring Stick

Why it's a mistake to compare A.I. with human inte…

geeks

Home › SCIENCE › Stop comparing human intelligence and AI – they are completely different

Stop comparing human intelligence and AI – they are completely different

TechTalks HOME BLOG TIPS & TRICKS WHAT IS INTERVIEWS

Home › Blog › There's a huge difference between AI and human intelligence—so let's stop comparing...

Blog

There's a huge difference between AI and human intelligence—so let's stop comparing them

The empty bra…

Your brain does not process information, retrie… knowledge or store memories. In short: your br… is not a computer

Why Artificial Intelligence Will Never Compete With the Human Brain


A supercomputer that works like the human brain has just been turned on
Manchester University just switched on the world's largest neuromorphic supercomputer
INSIGHTS
Before you know it!
It won't be long before AI mimics a human brain to do more, with less – just like we can.
MIT Scientists Design a Computer That Works Just Like a Human Brain
Researchers developed algorithms that mimic the human brain (and the results don't suck)
MIT researchers develop new chip design to take us closer to computers that work like human brains
IBM is teaching AI to behave more like the human brain
SPACE PHYSICS CHEMISTRY BIOLOGY ENERGY & ENVIRONMENT
Intel Developing an AI Chip That Acts Like a Human Brain
WIRED
Do neural networks really work like neurons?
Why Artificial Intelligence Is Not Like Your Brain
Contrary to belief, artificial intellige rocket scientist.
future tense
The Wrong Cognitive Measuring Stick
Why it's a mistake to compare A.I. with human int
Stop comparing human intelligence and AI – they are completely different
TechTalks HOME BLOG TIPS & TRICKS WHAT IS INTERVIEWS
There's a huge difference between AI and human intelligence—so let's stop comparing them
The empty bra
Your brain does not process information, retri knowledge or store memories. In short: your br is not a computer
Why Artificial Intelligence Will Never Compete With the Human Brain
TOP 5 THIS MONTH MEDIUM TIPS SUBMIT
How neuroscience enables better Artificial Intelligence design
MENU nature International journal of science
OUTLOOK · 24 JULY 2019 · CLARIFICATION 21 AUGUST 2019
How AI and neuroscience drive each other forwards
Three Invaluable Ways AI and Neuroscience Are Driving Each Other Forward
BLOG POST RESEARCH
AI and Neuroscience: A virtuous circle
Neuron
REVIEW | VOLUME
Neuroscience-Inspired Artificial Intelligence
Demis Hassabis • Dharshan Kumaran • Christopher Summerfield • Matthew Botvinick
convergence of neuroscience and artificial lligence
Science
Neuroscience and artificial intelligence can help improve each other
July 9, 2019 9.23pm AEST
SHARE
PERSPECTIVE | NEUROSCIENCE
Using neuroscience to develop artificial intelligence

We need to structure this conversation

Suggestion

David Marr's

Three levels of analysis

We need to structure this conversation

Suggestion

David Marr's

Three levels of analysis

Goal of the computation

What is the purpose of computation?

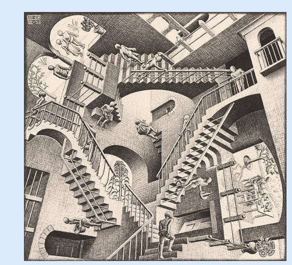

We need to structure this conversation

Suggestion

David Marr's

Three levels of analysis

Goal of the computation

What is the purpose of computation?

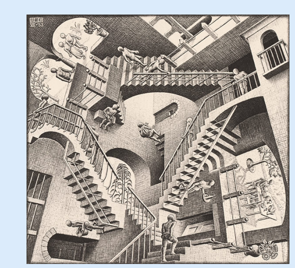

Algorithm and representation

What representations does the system use?

What processes are in use to manipulate representations?

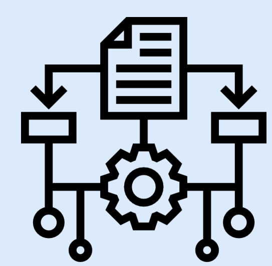

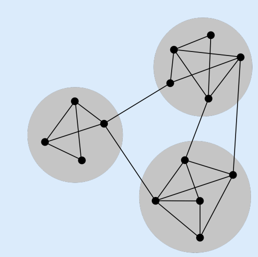

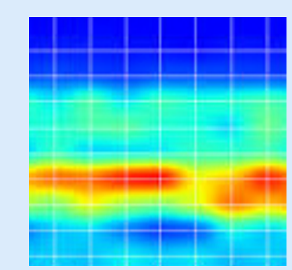

We need to structure this conversation

Suggestion

David Marr's

Three levels of analysis

Goal of the computation

What is the purpose of computation?



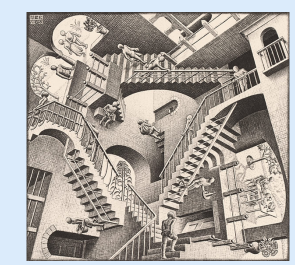

Algorithm and representation

What representations does the system use?

What processes are in use to manipulate representations?

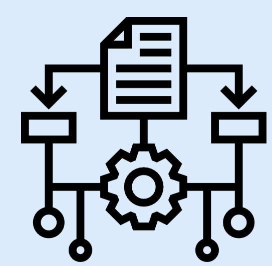
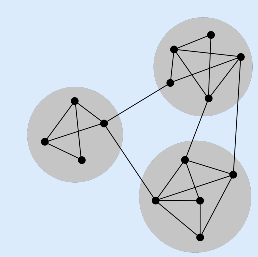
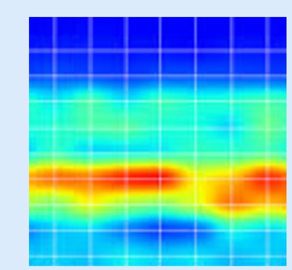

Implementation

How is the system physically realized?

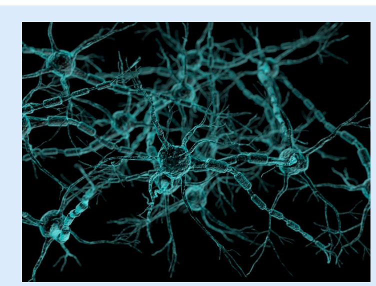
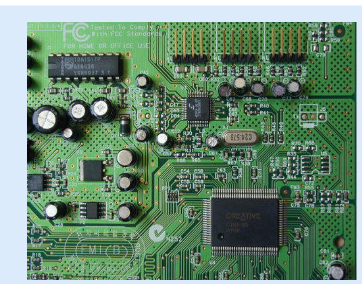

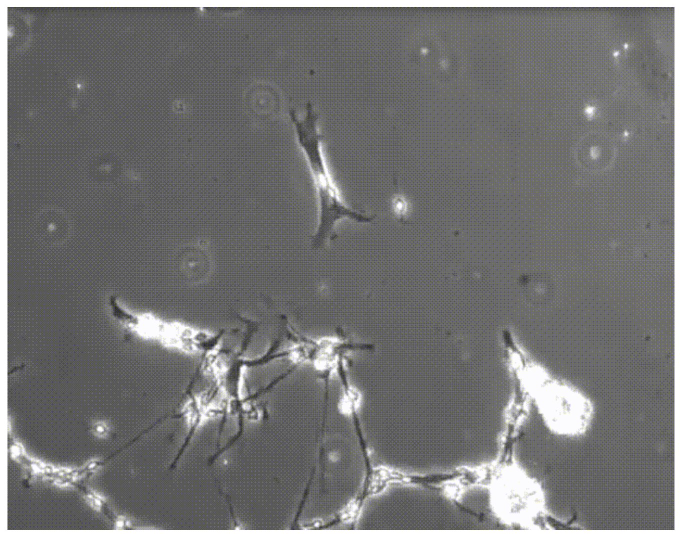

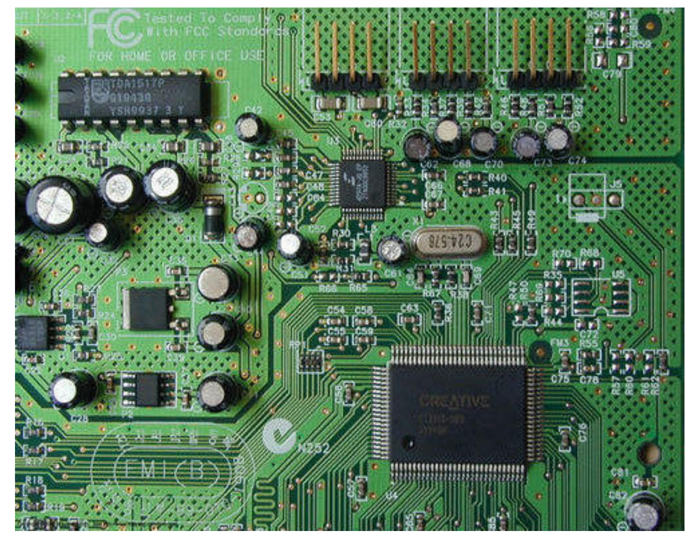

Implementation

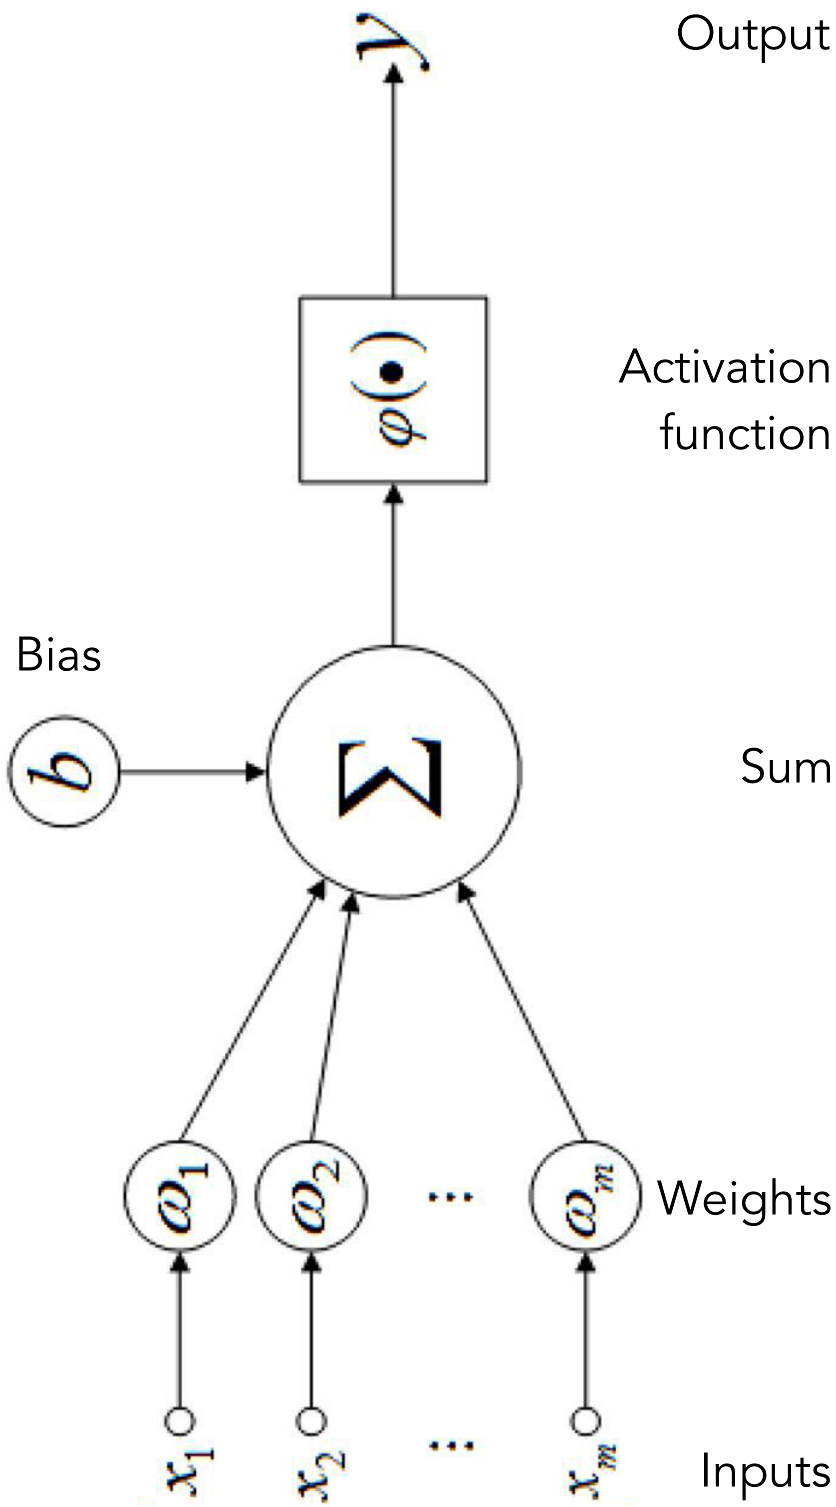

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595 https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

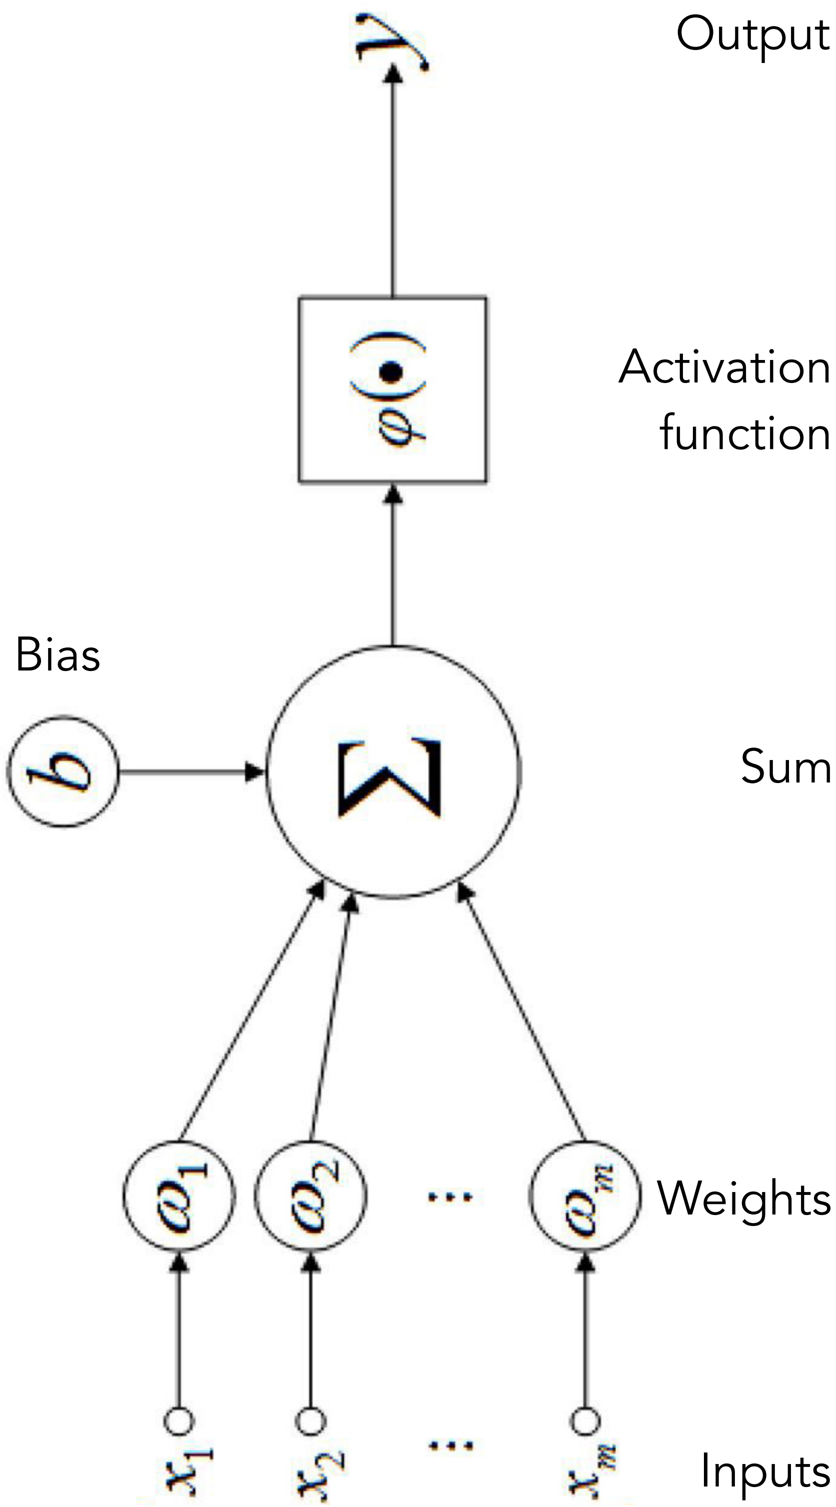

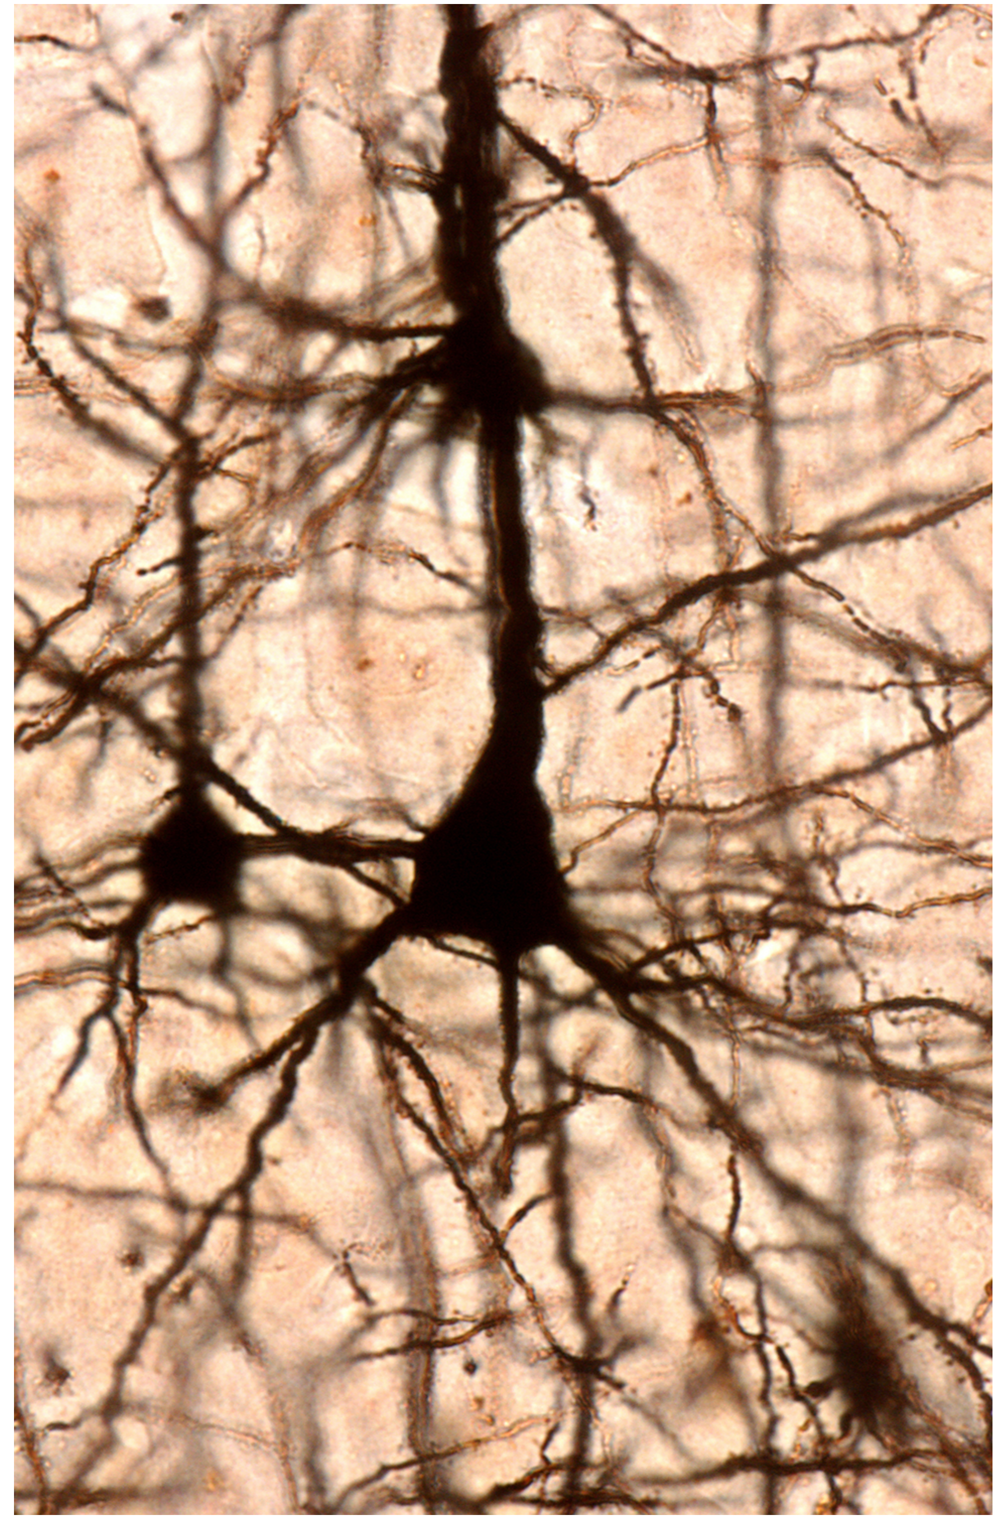

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595 https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

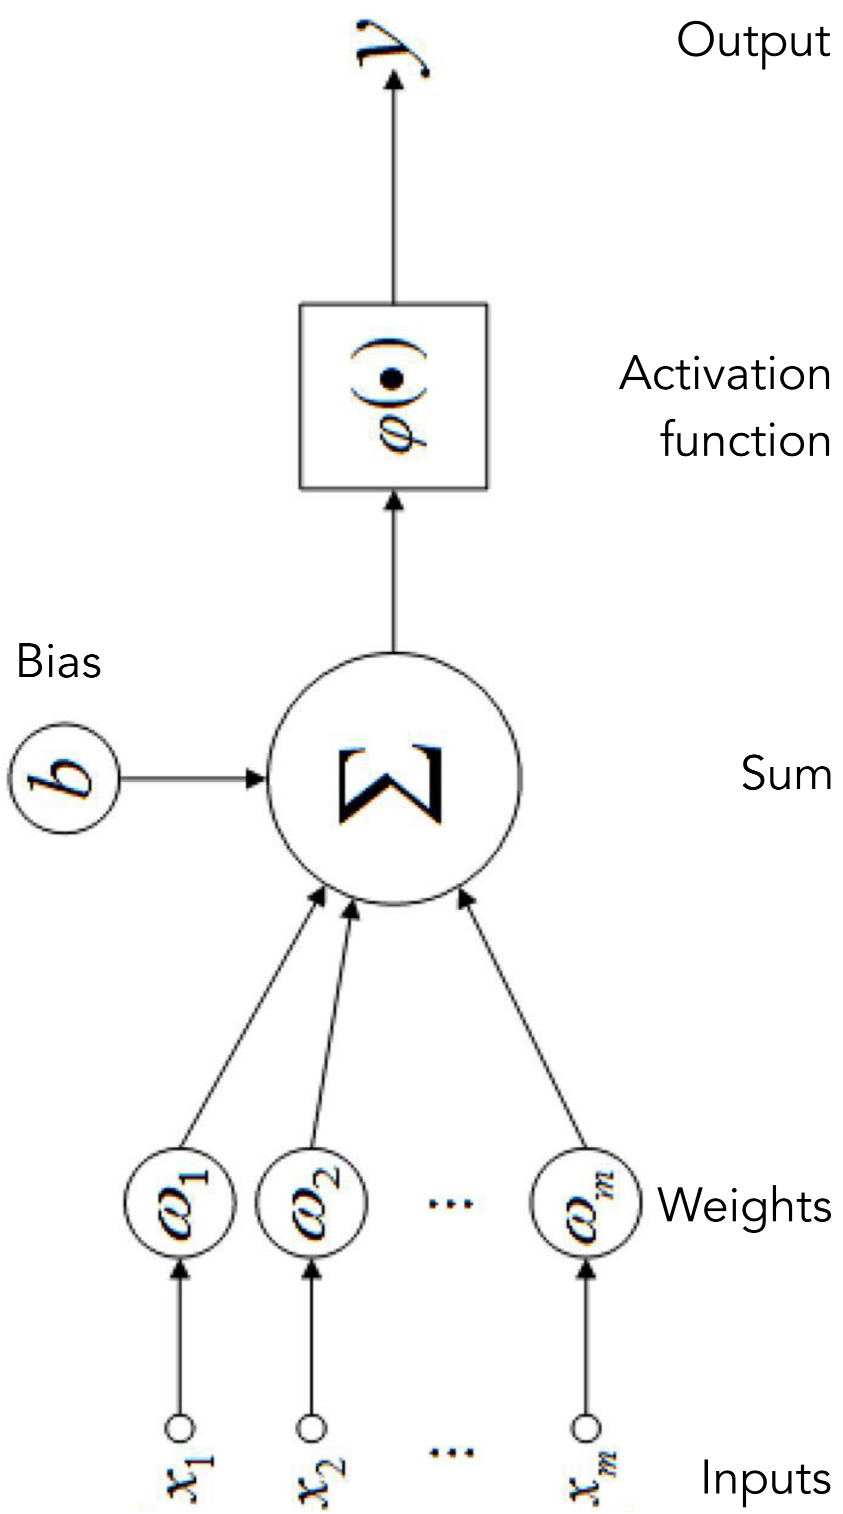

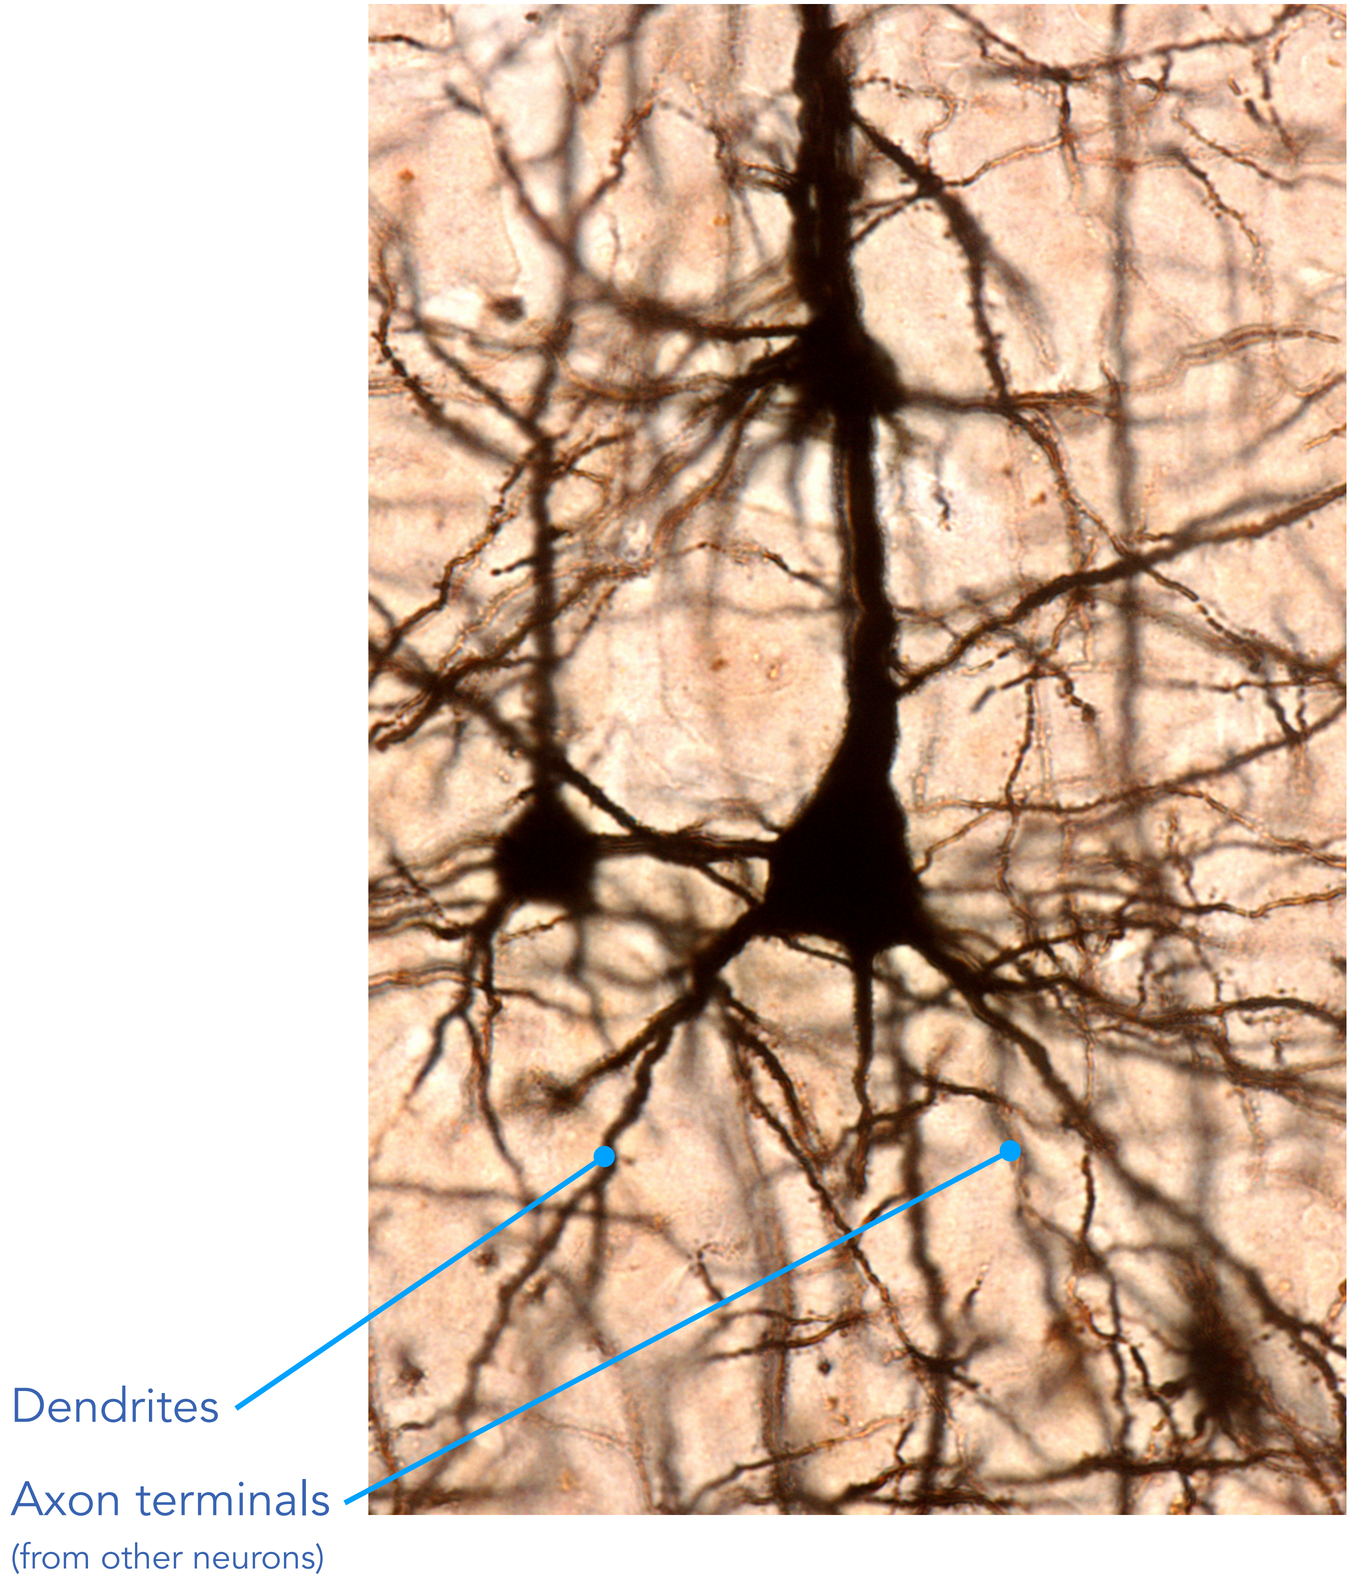

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595 https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

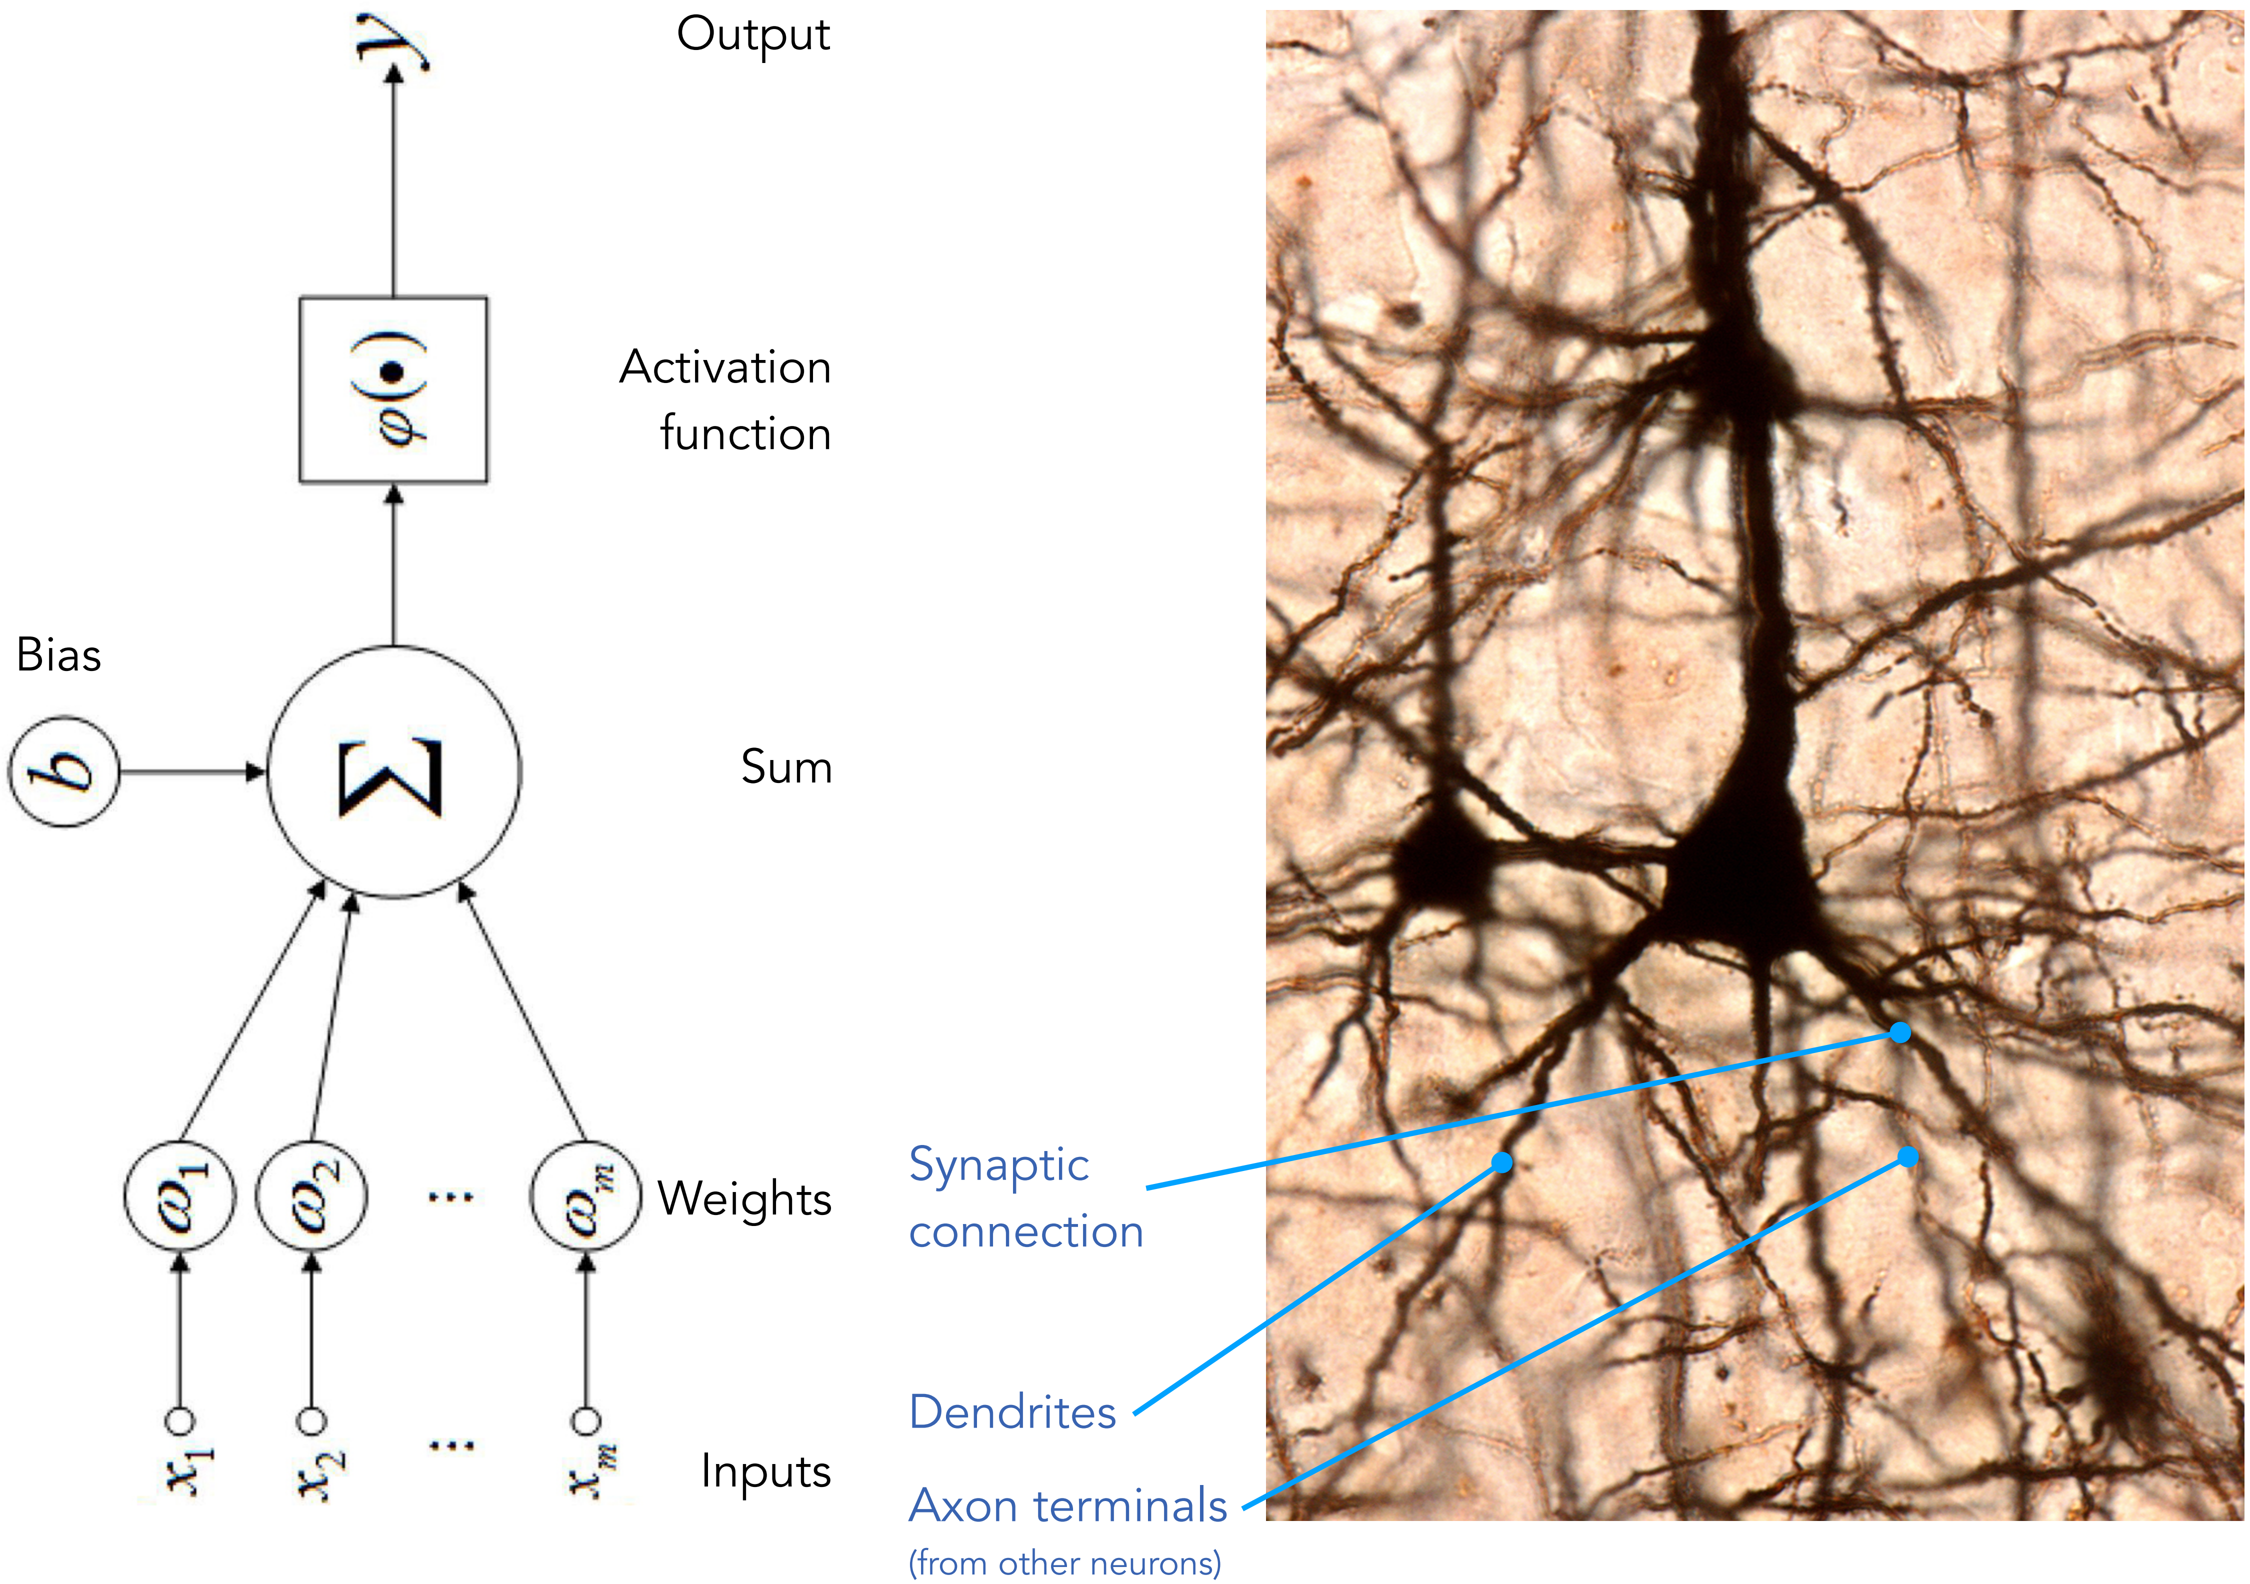

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595 https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

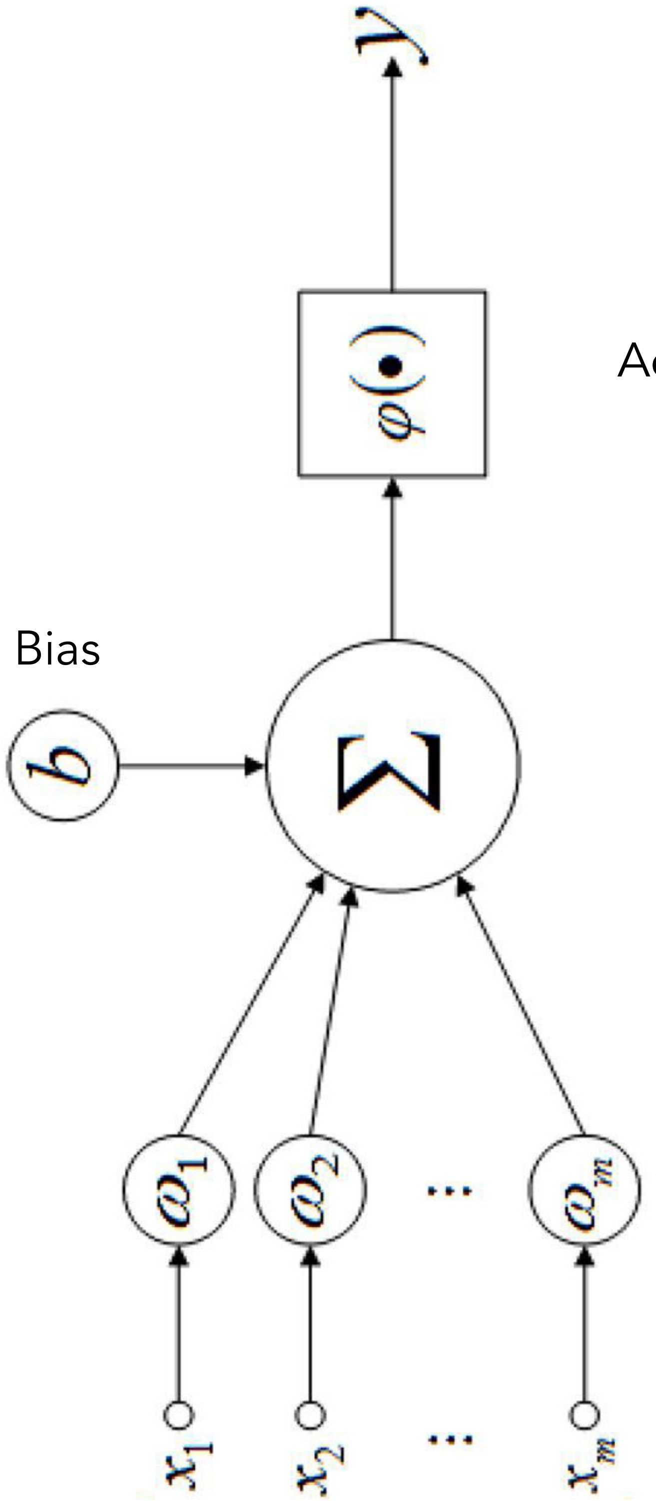

Output

Activation
function

Bias

Sum

Weights

Inputs

Cell
depolarization

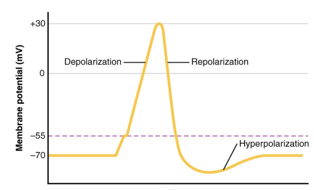

Synaptic
connection

Dendrites

Axon terminals
(from other neurons)

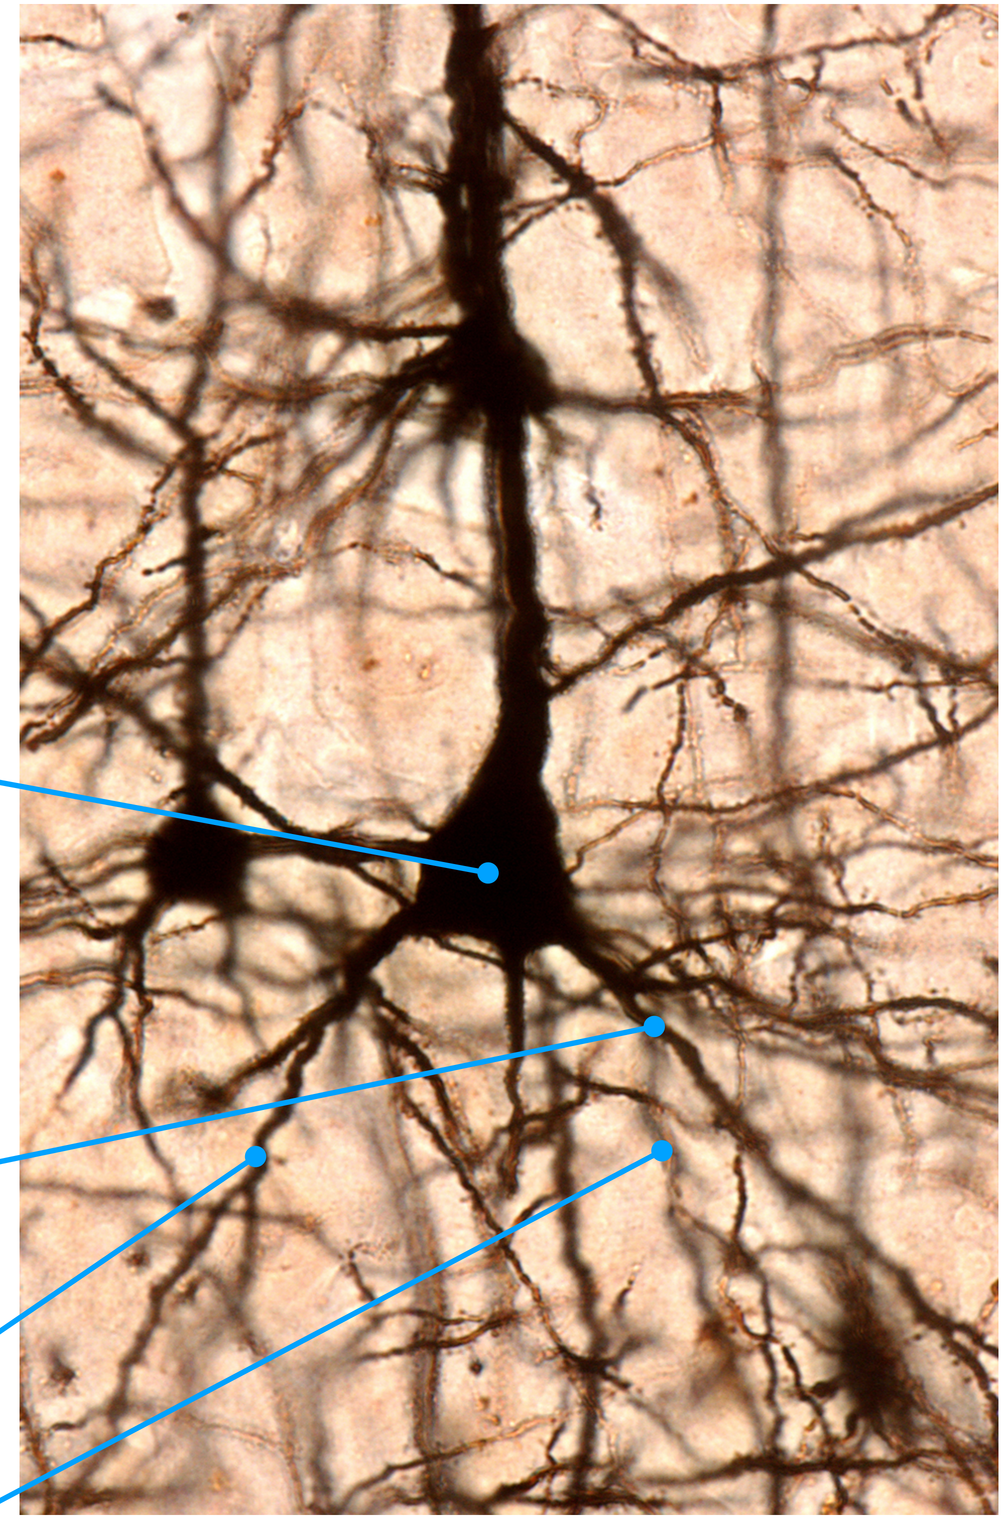

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595 https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

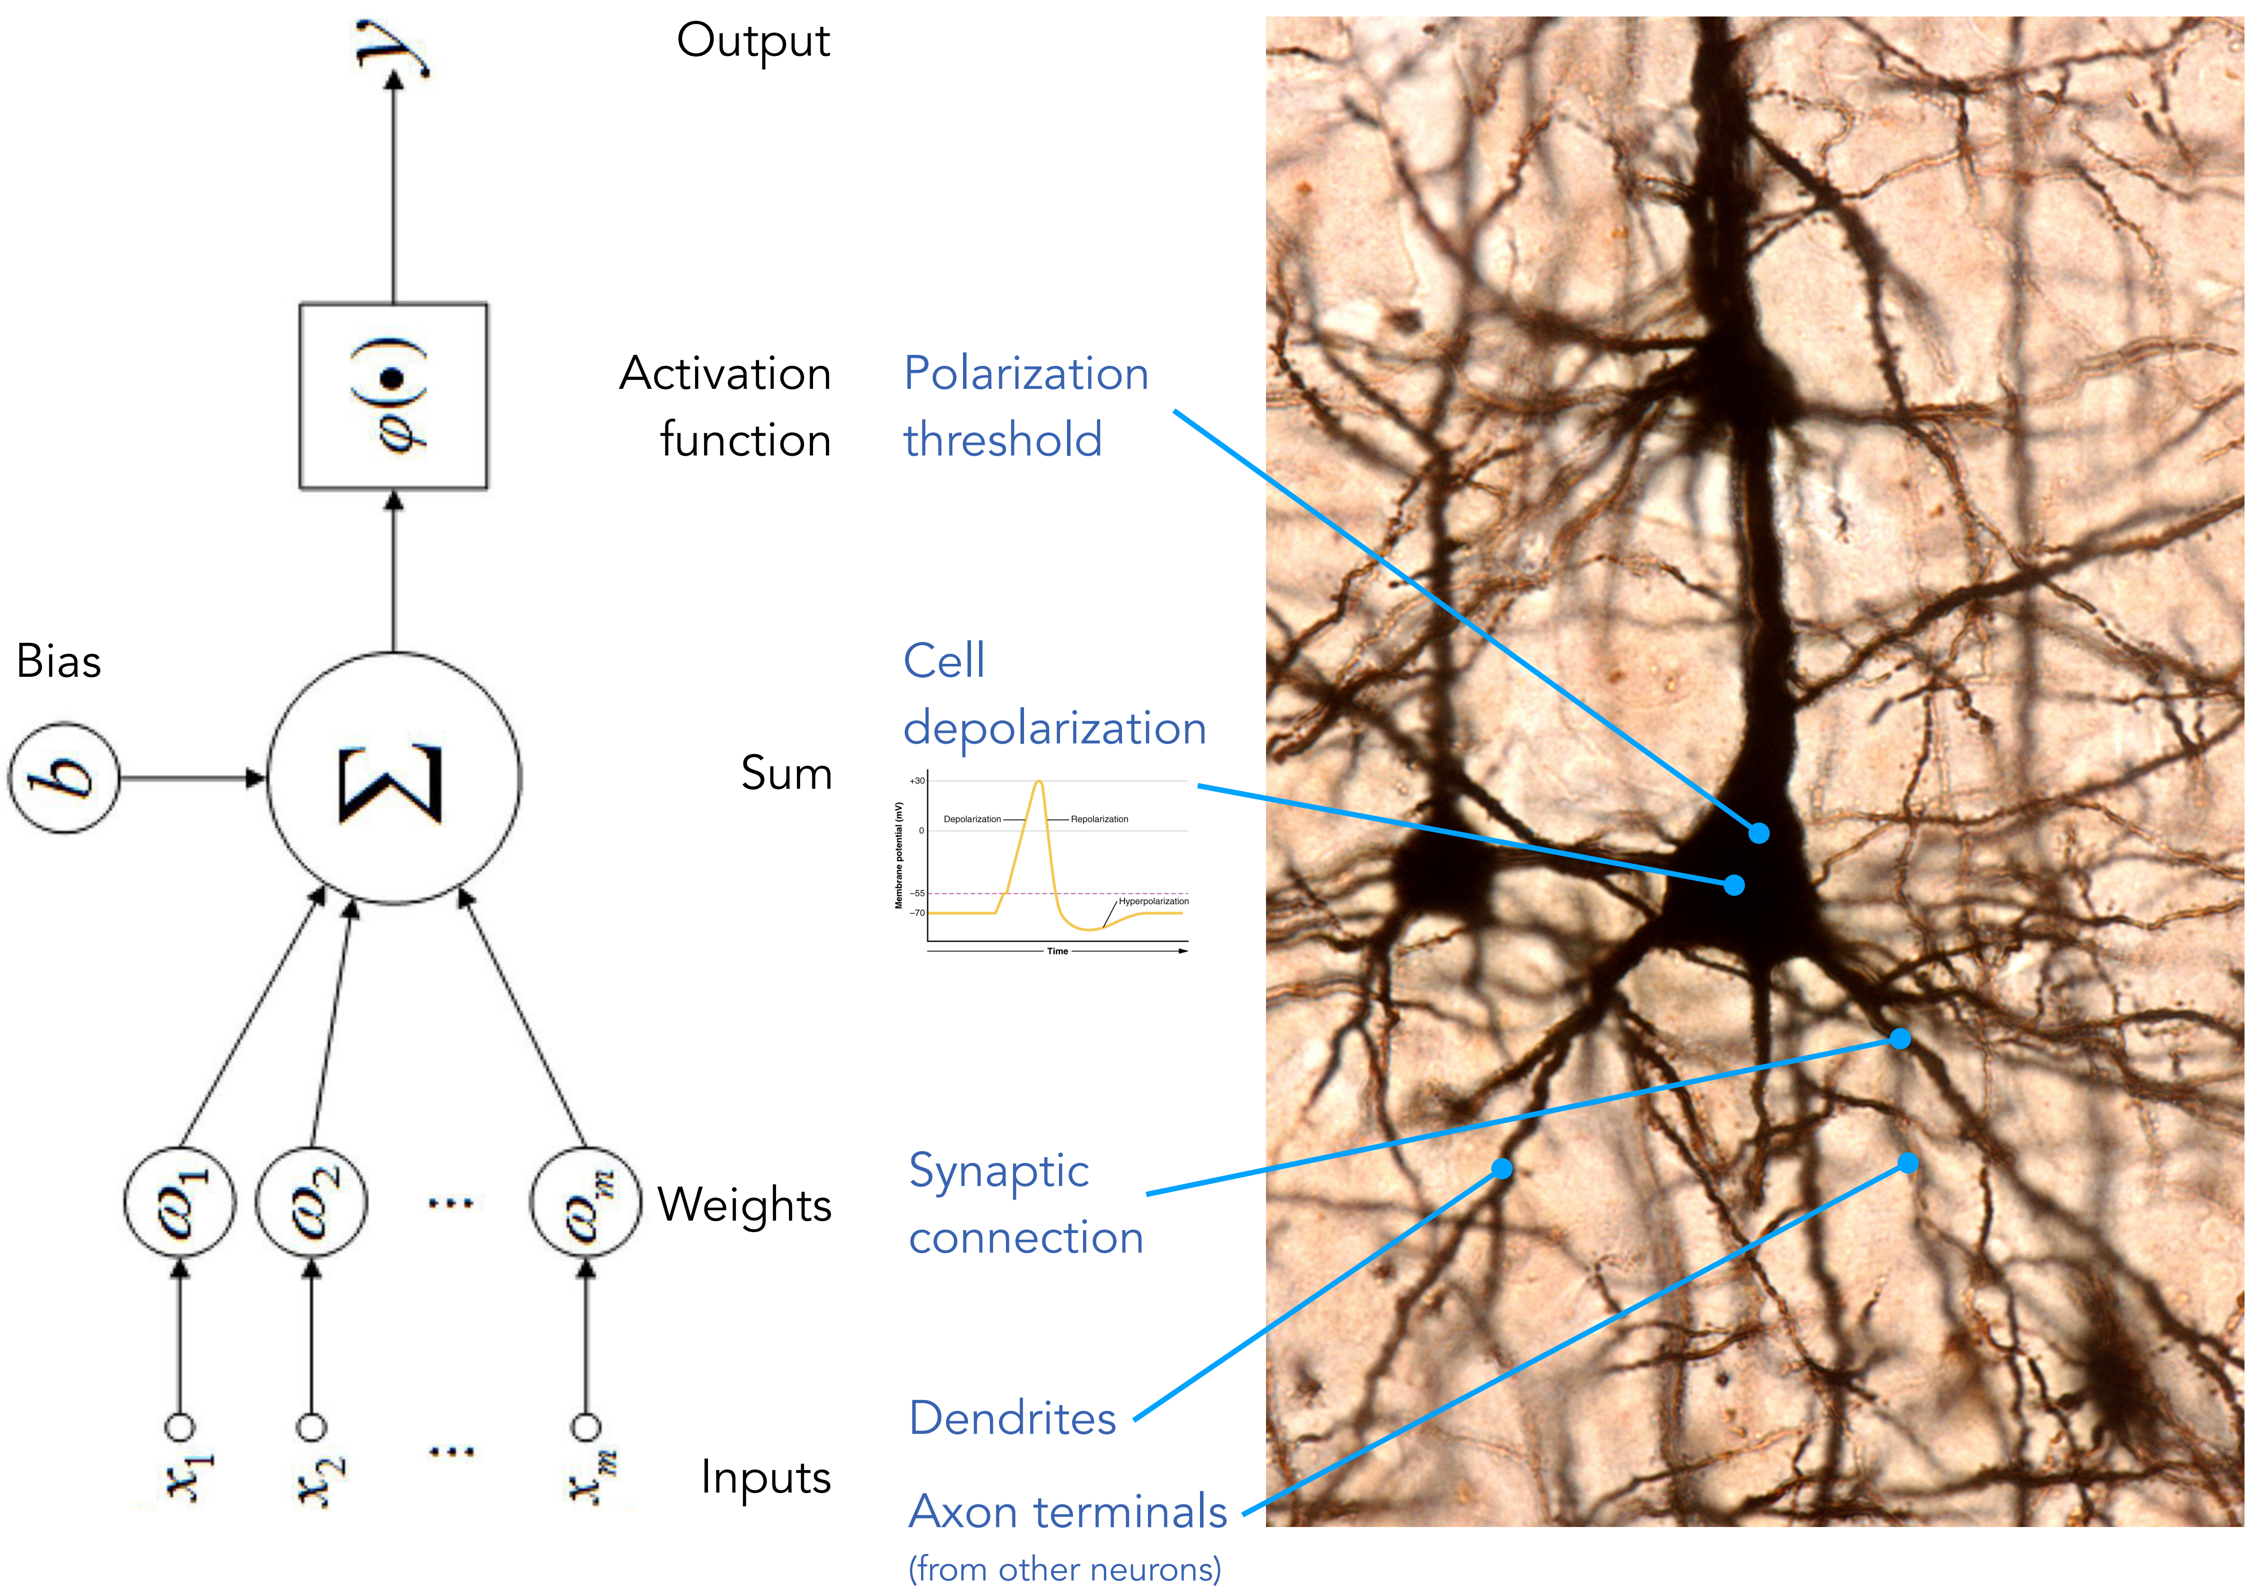

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595 https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

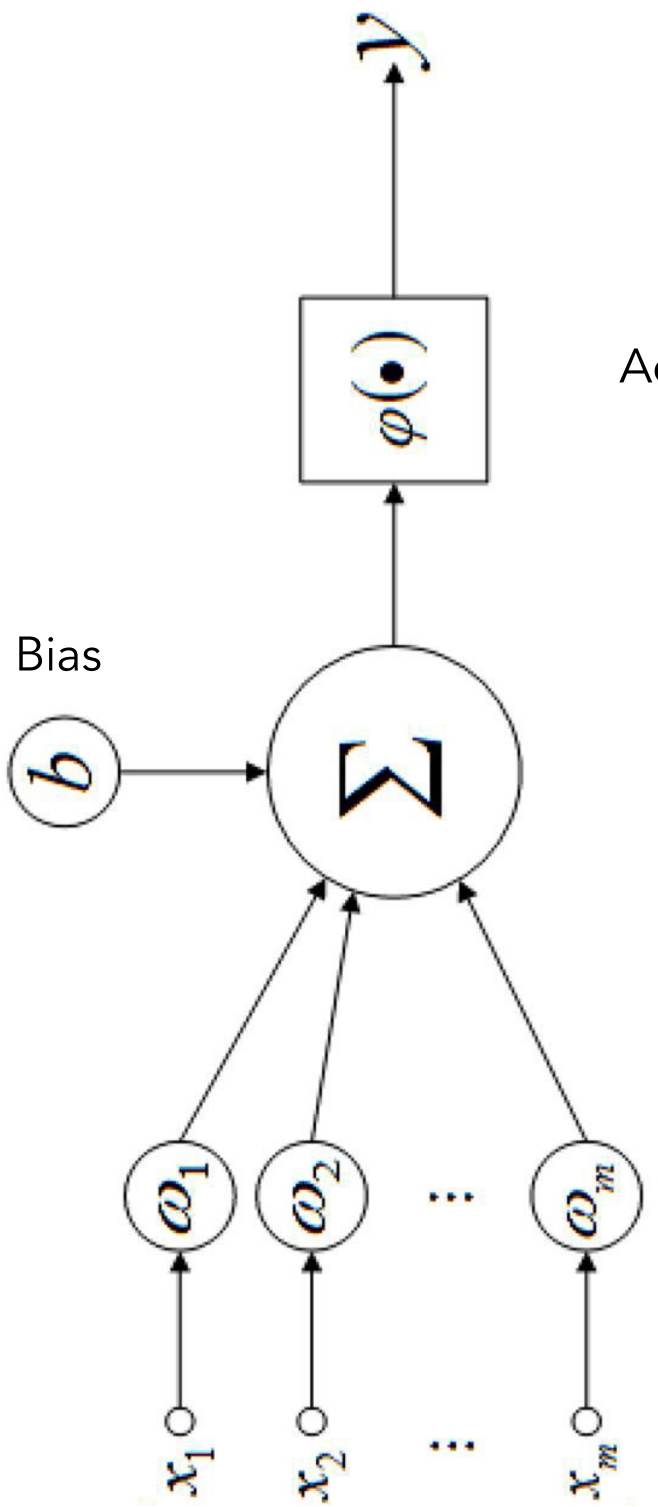

Output

Action potential (aka Spike)

Activation function

Polarization threshold

Sum

Cell depolarization

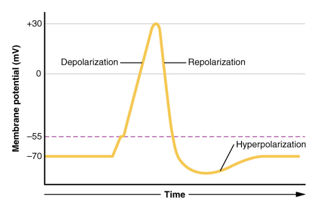

Weights

Synaptic connection

Dendrites

Inputs

Axon terminals
(from other neurons)

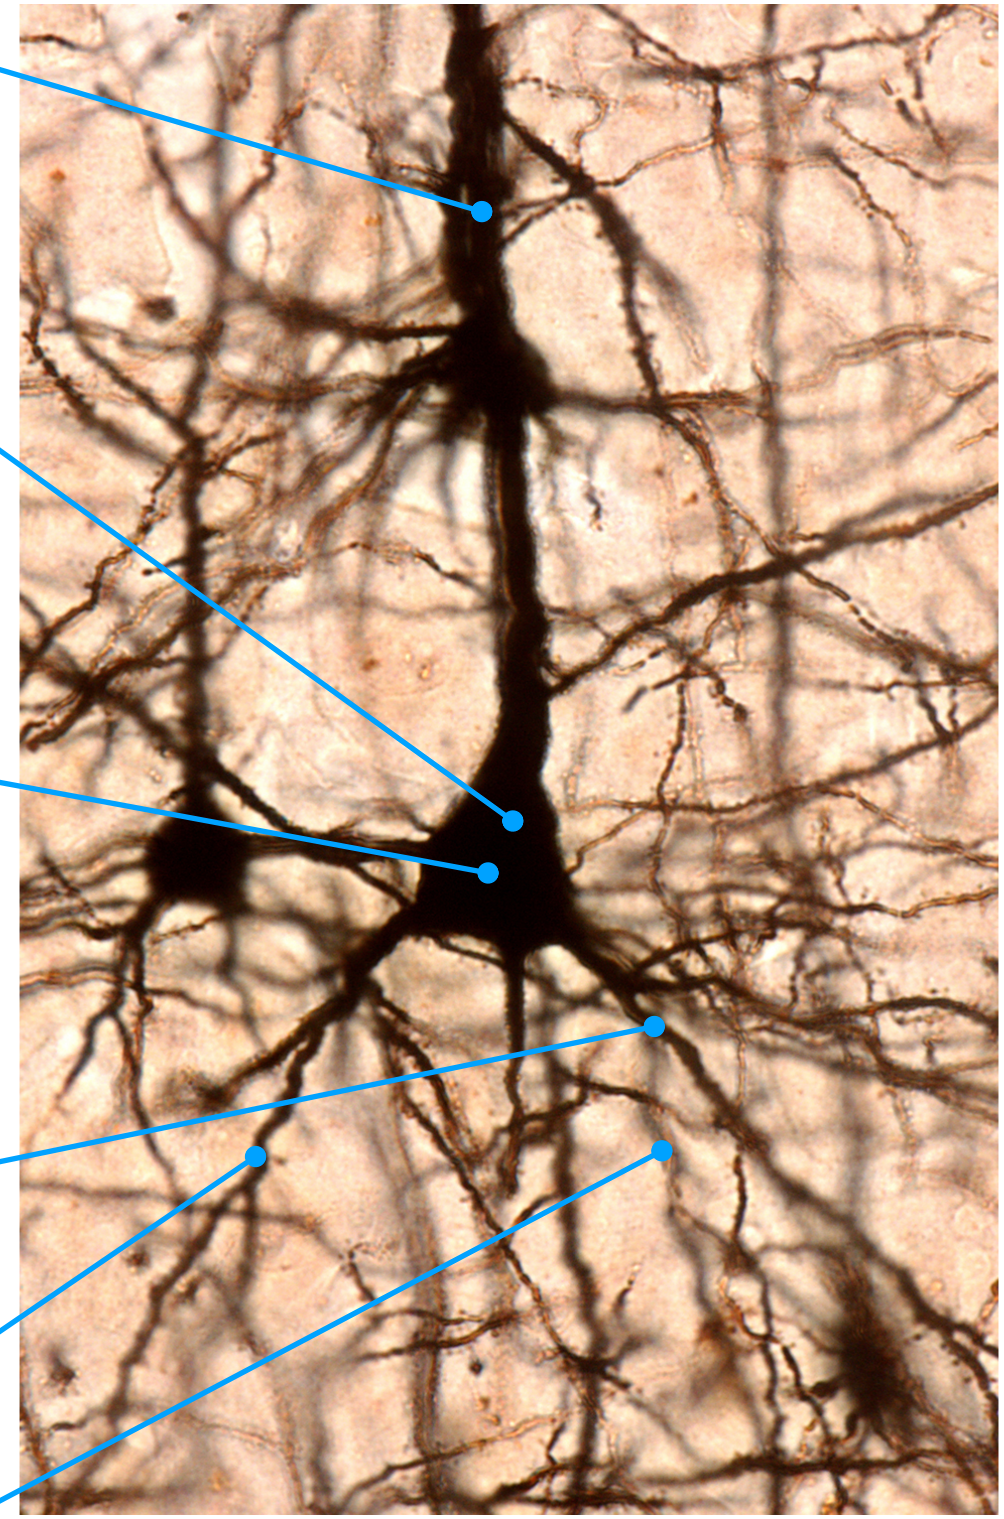

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595
https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e
https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

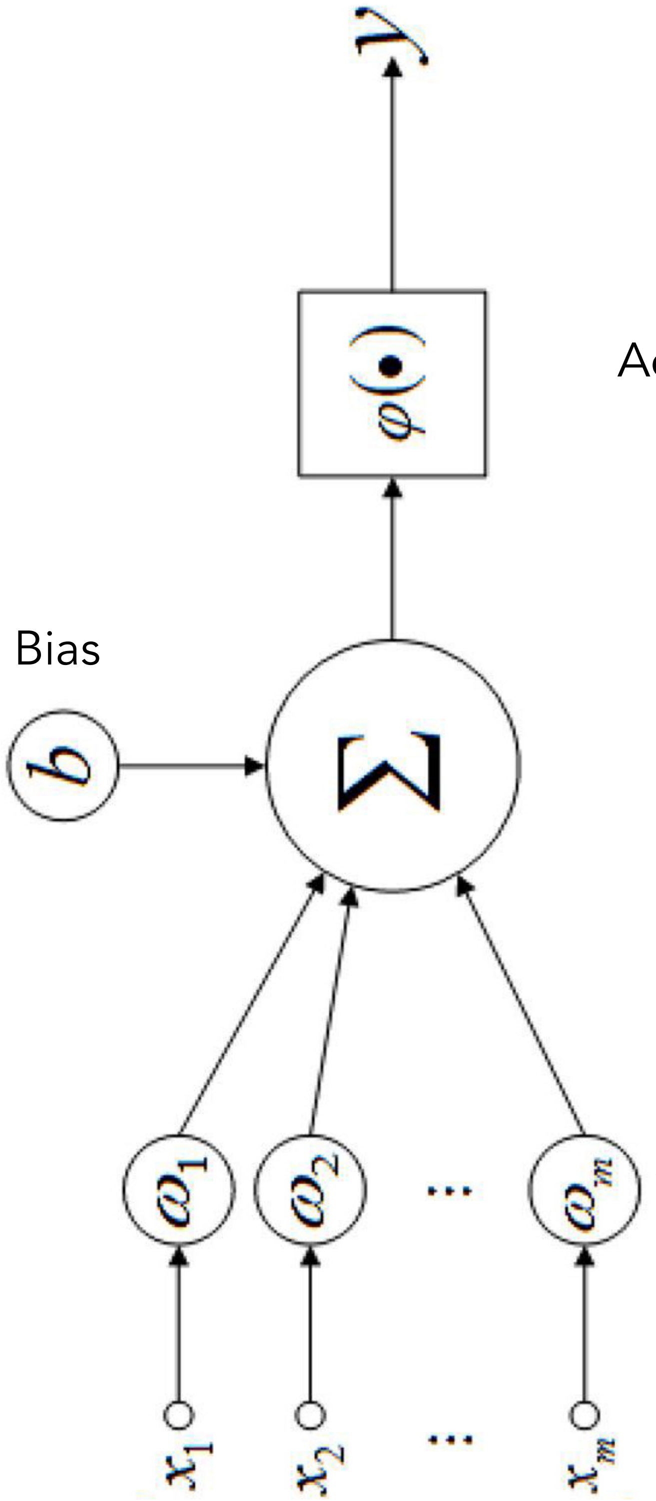

Output — Action potential (aka Spike)

Activation function — Polarization threshold

Bias

Sum — Cell depolarization

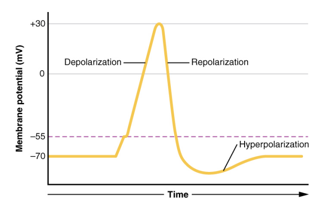

Weights — Synaptic connection

Dendrites

Inputs — Axon terminals (from other neurons)

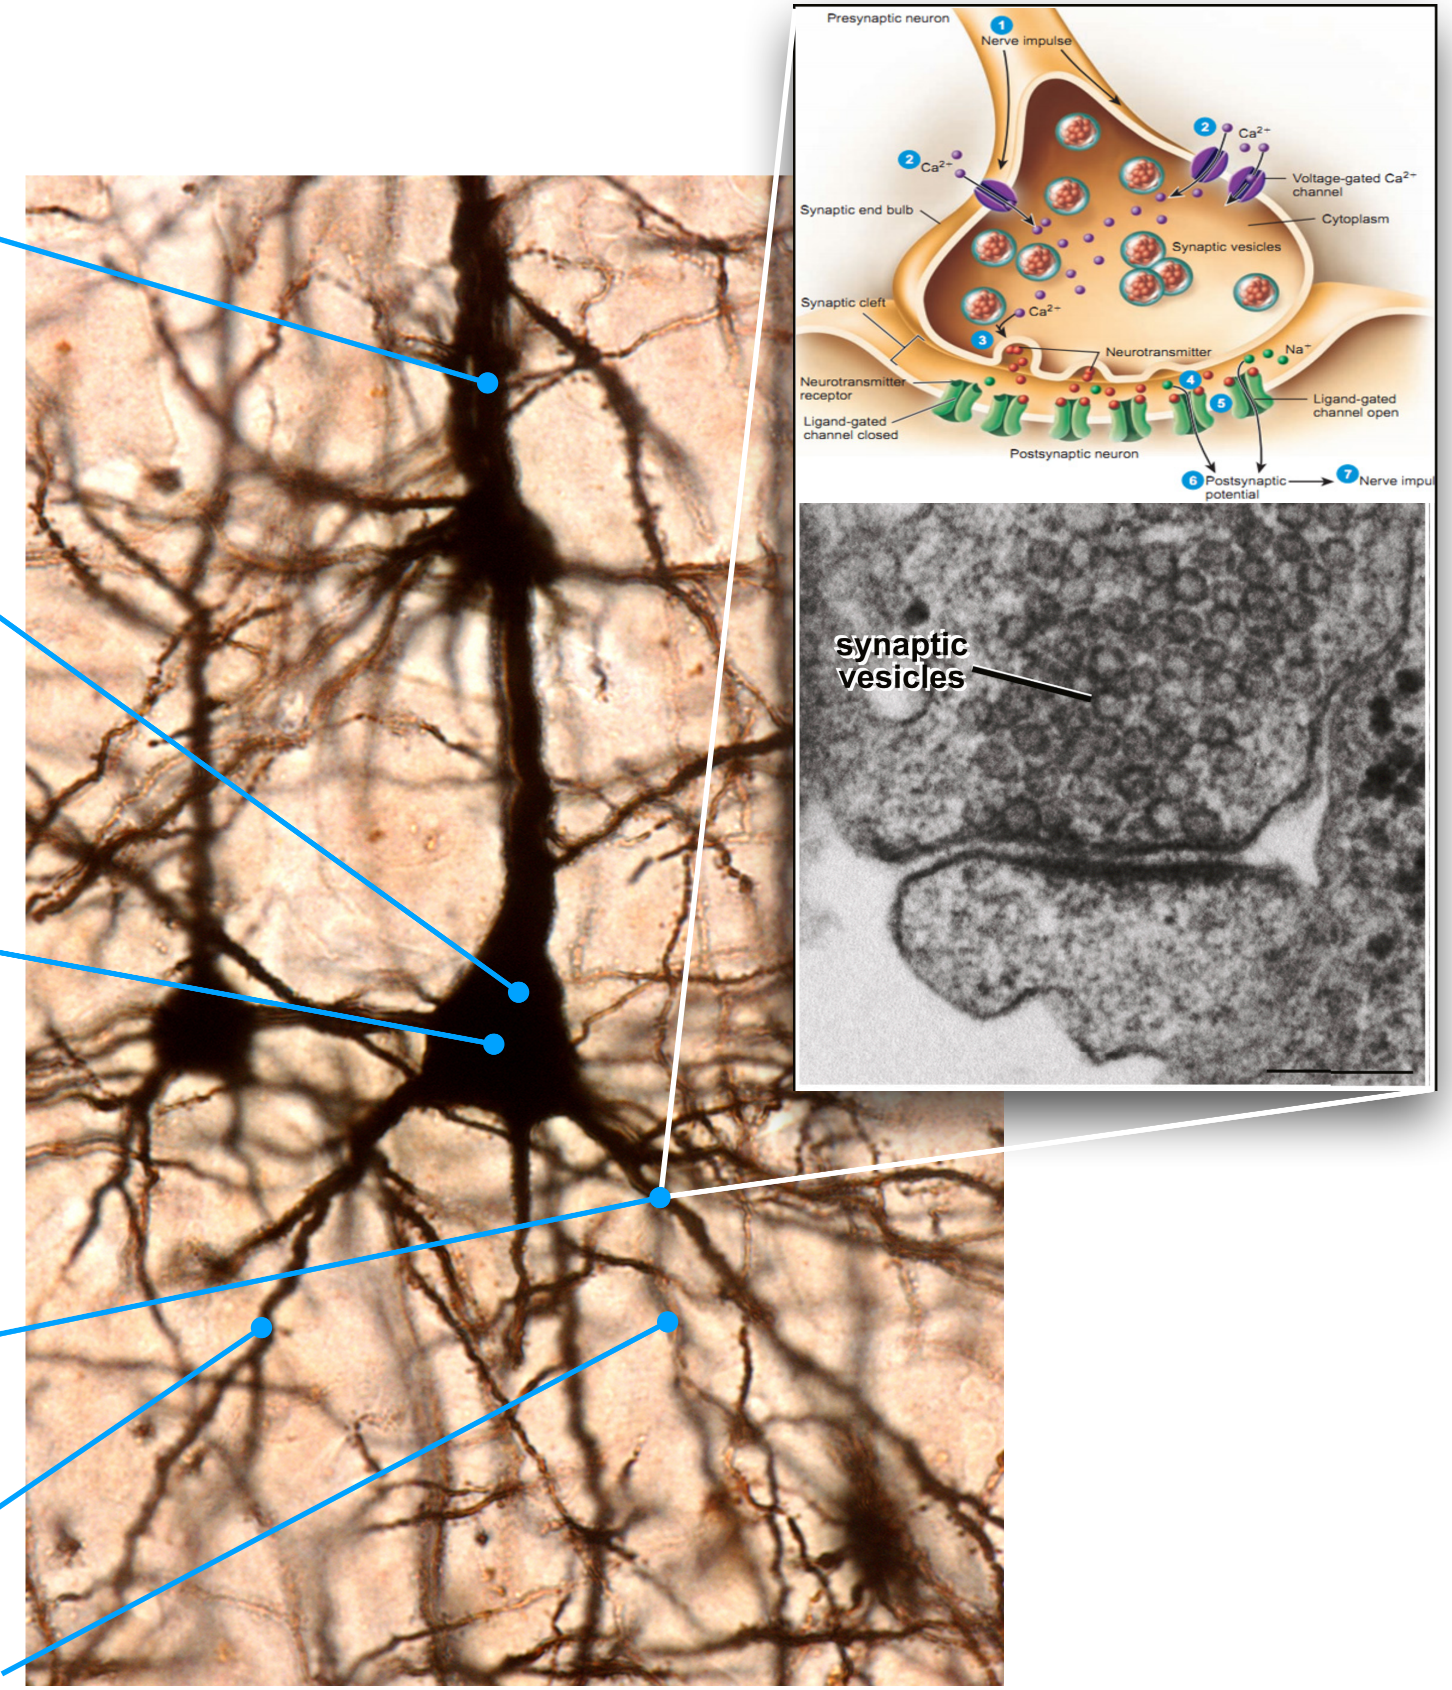

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595
https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e
https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

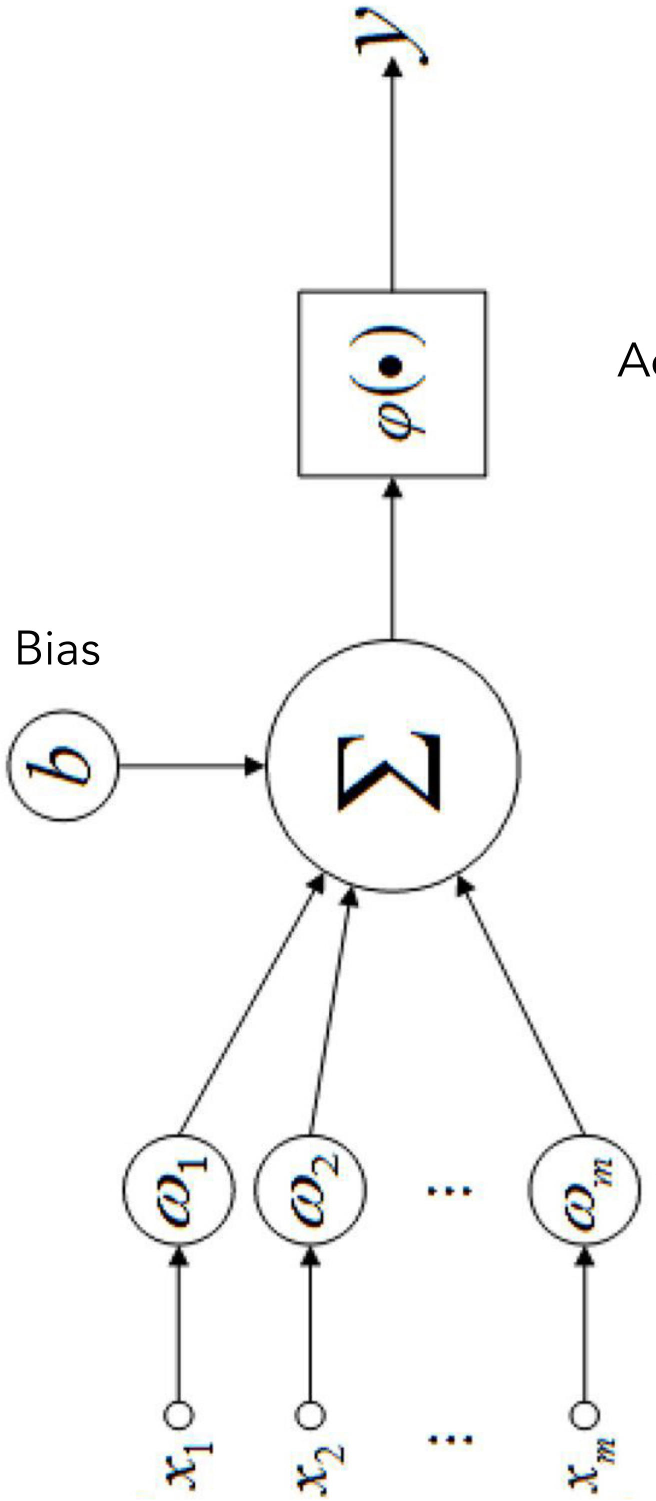

Output | Action potential (aka Spike)

Activation function | Polarization threshold

Sum | Cell depolarization

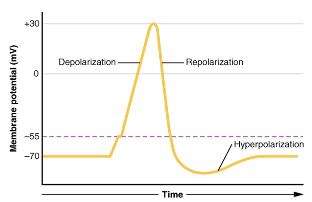

Weights | Synaptic connection

Dendrites

Inputs | Axon terminals (from other neurons)

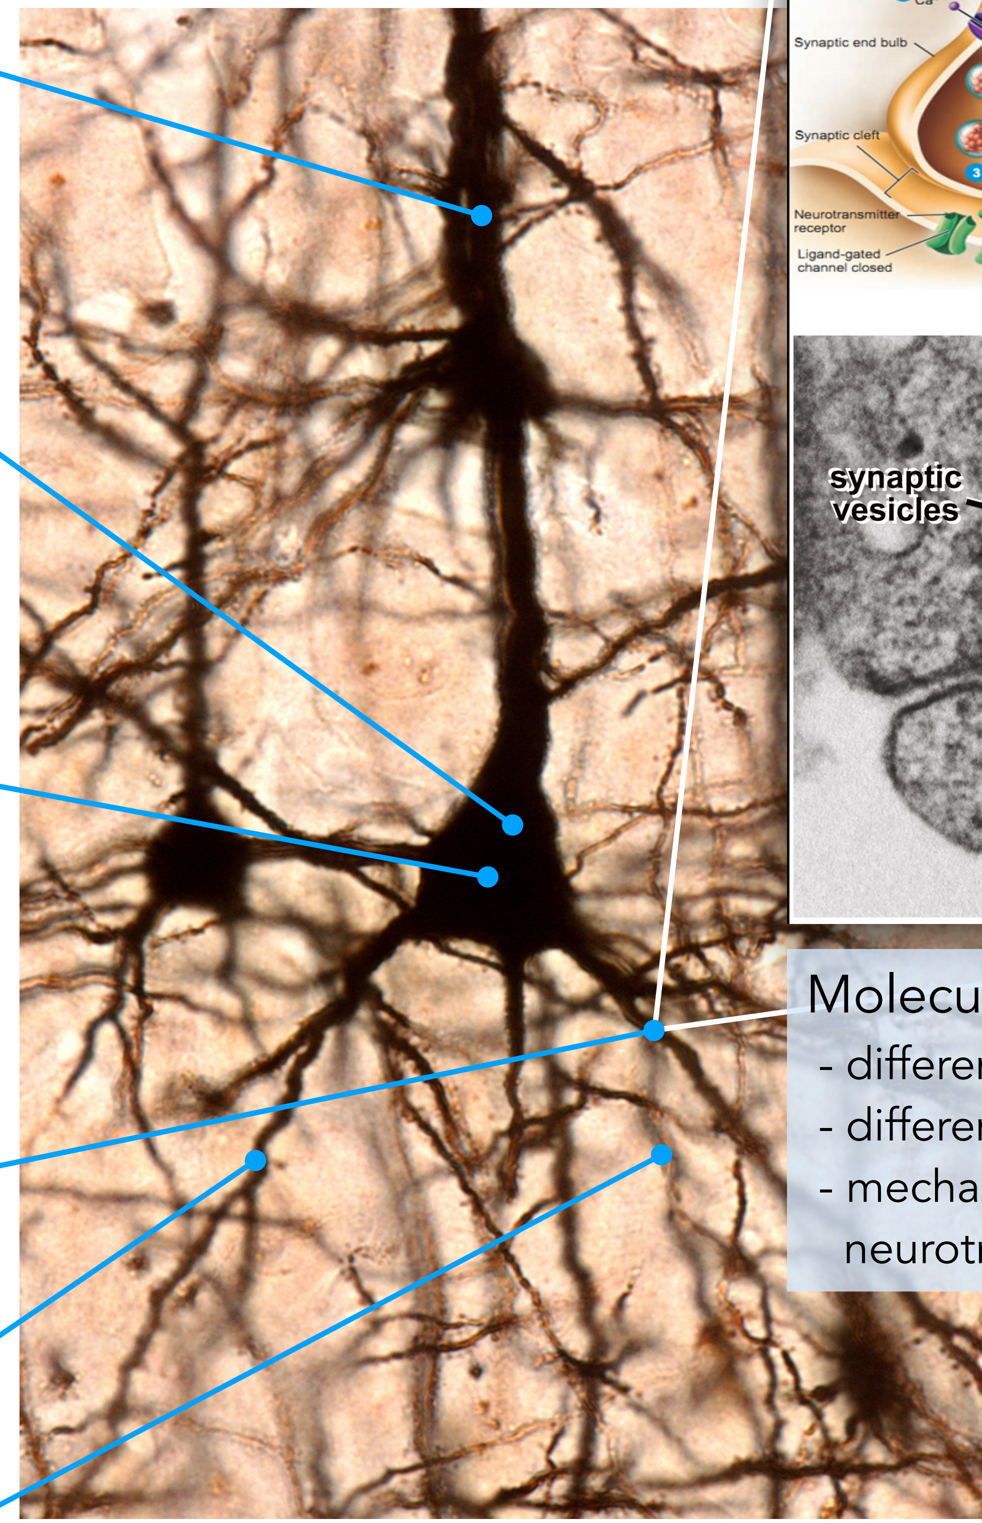

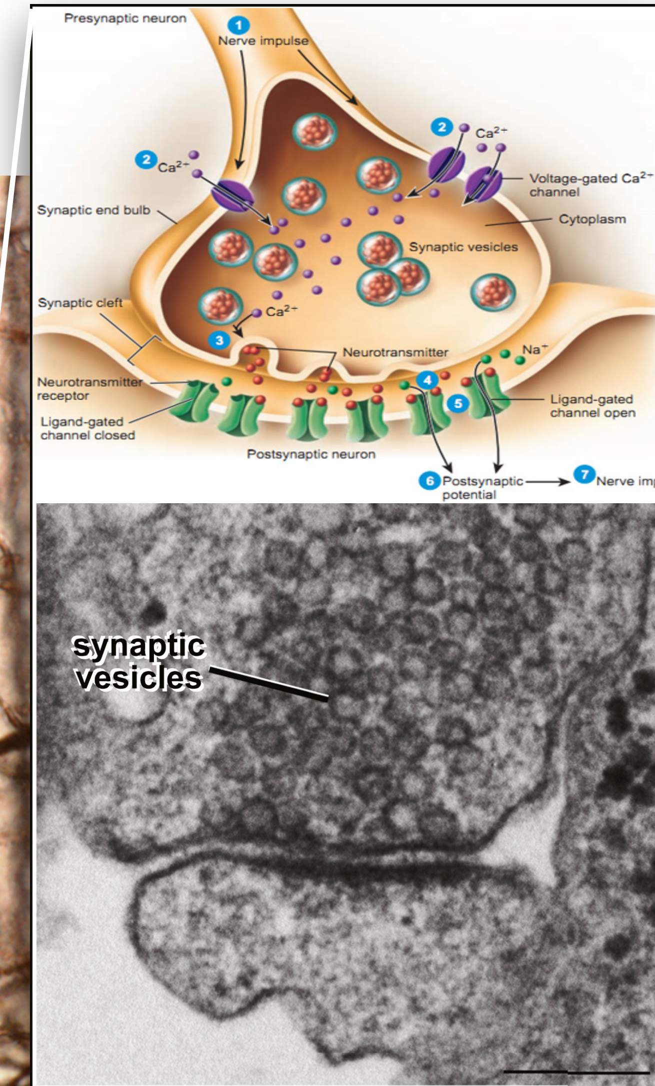

Molecular neurobiology
- different types of channels
- different receptors
- mechanism of neurotransmitter release

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595
https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e
https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

Bias

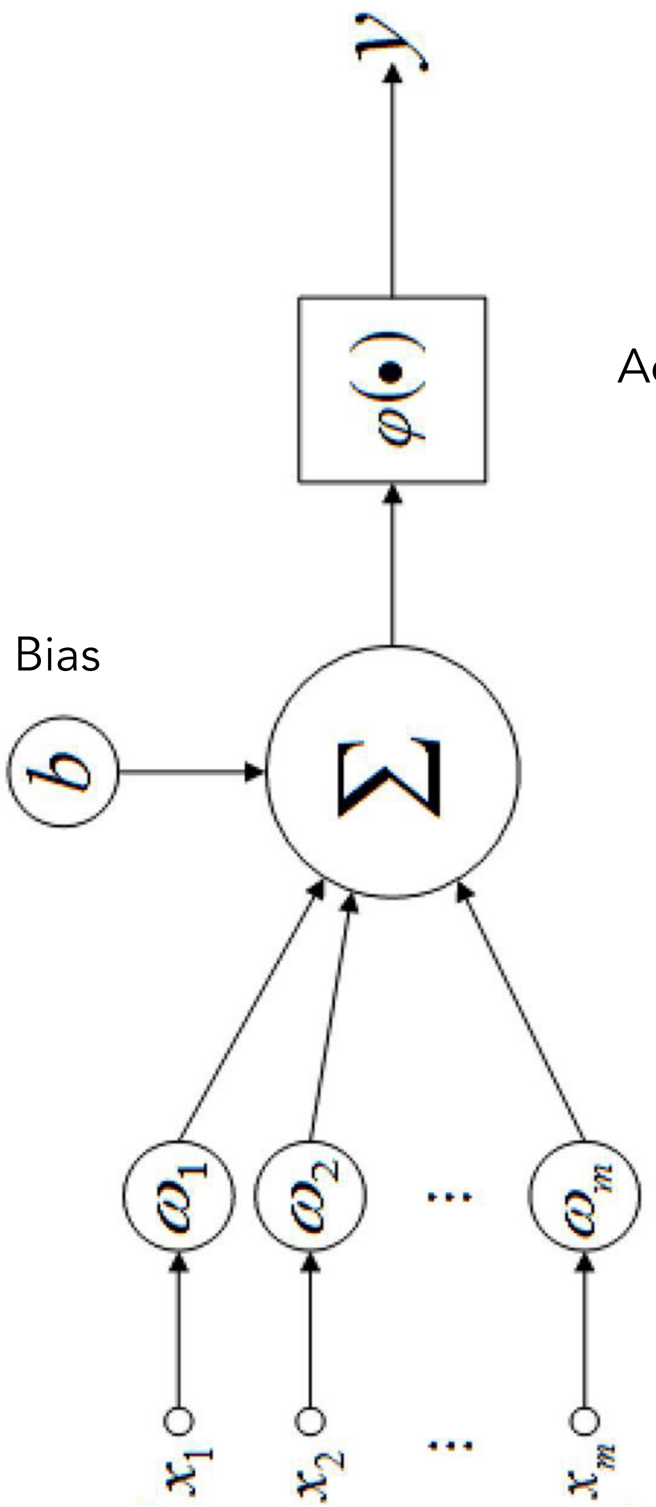

Output — Action potential (aka Spike)

Activation function — Polarization threshold

Sum — Cell depolarization

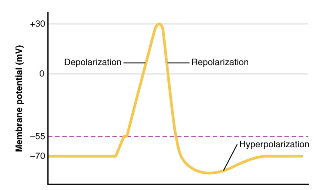

Weights — Synaptic connection

Inputs — Dendrites

Axon terminals (from other neurons)

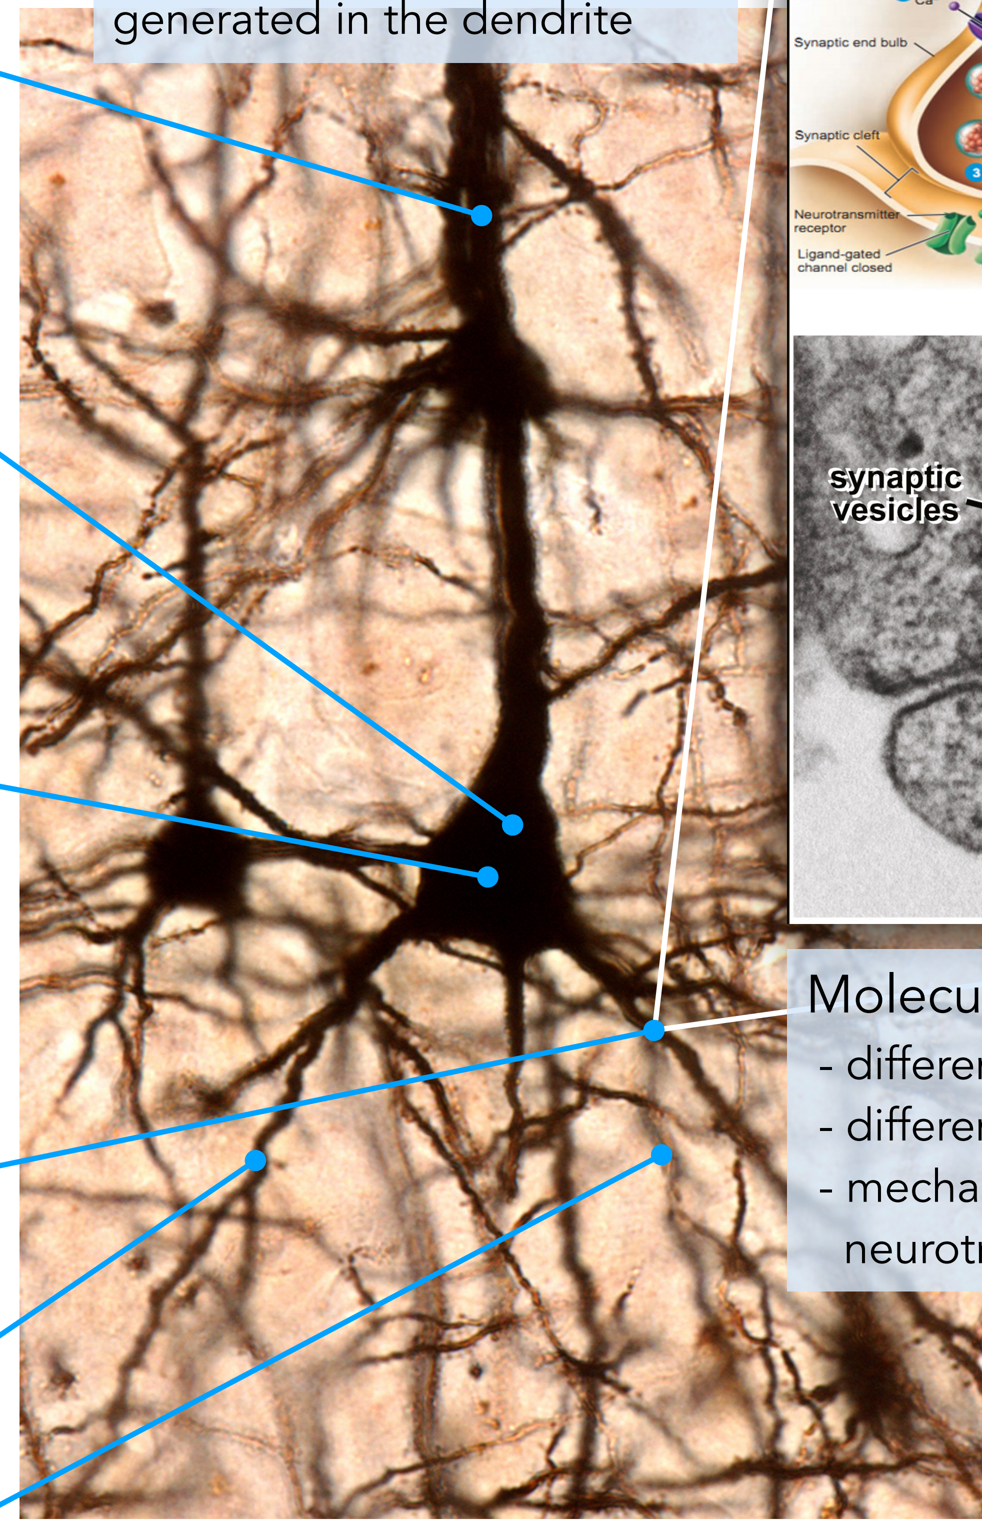

Dendritic spike
action potential can be generated in the dendrite

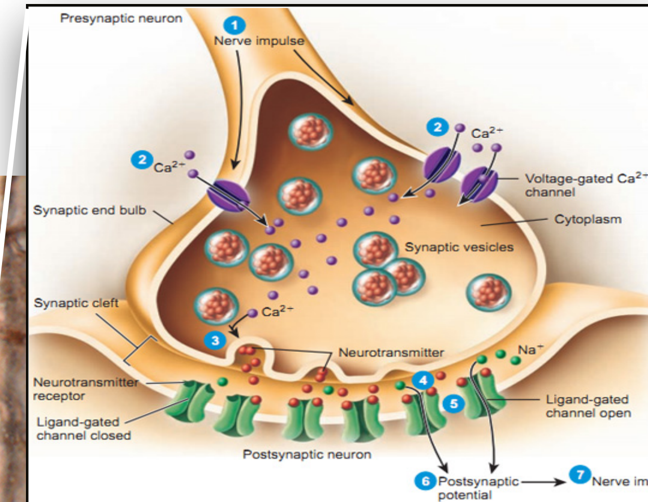

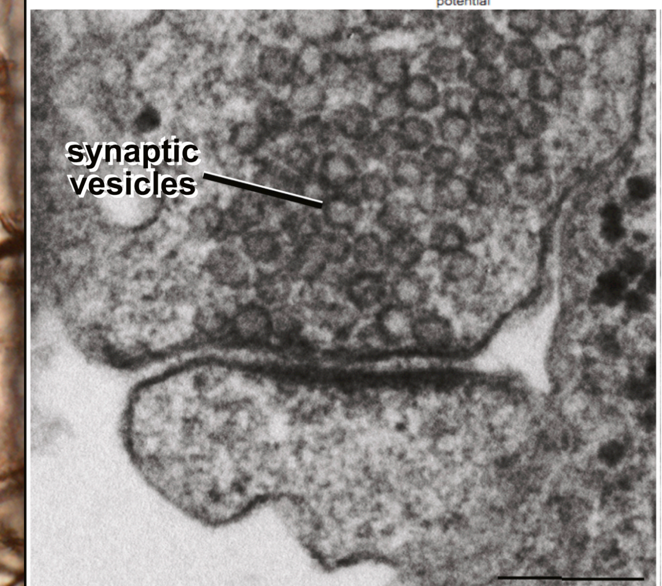

Molecular neurobiology
- different types of channels
- different receptors
- mechanism of neurotransmitter release

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595 https://antranik.org/synaptic-transmission-by-somatic-motorneurons/

https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

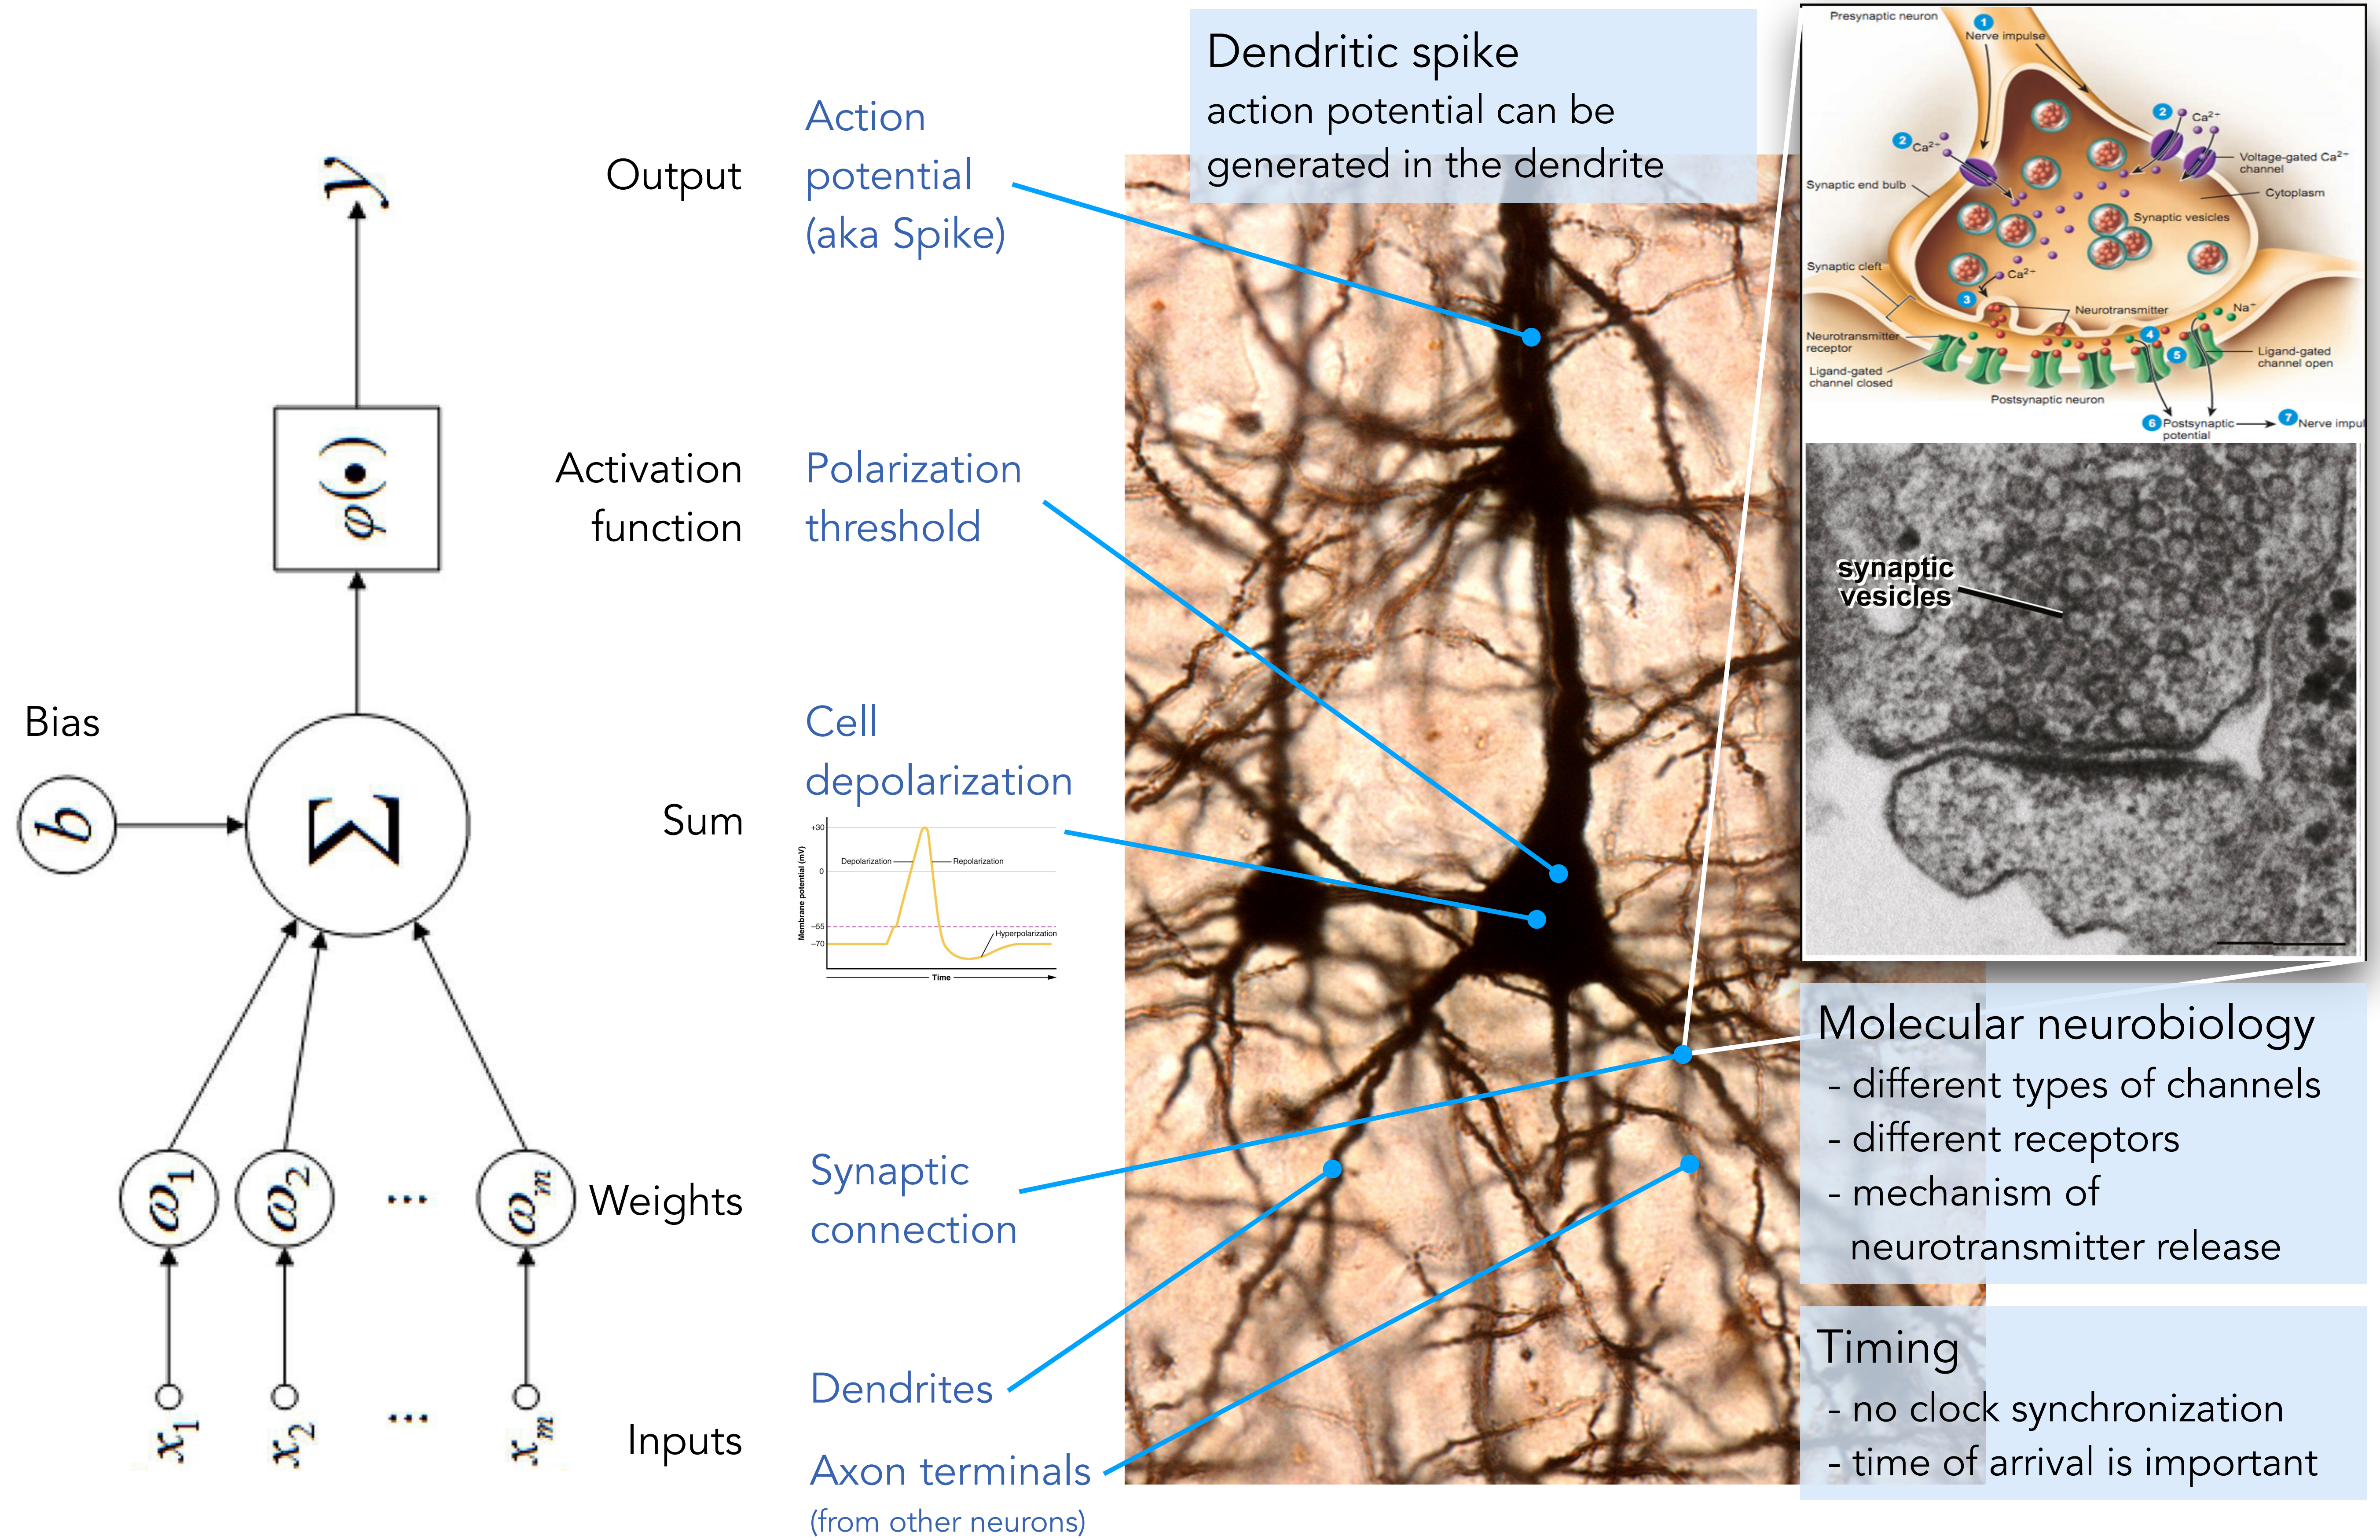

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595
https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e
https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

Output

Action potential (aka Spike)

Activation function

Polarization threshold

Bias

Sum

Cell depolarization

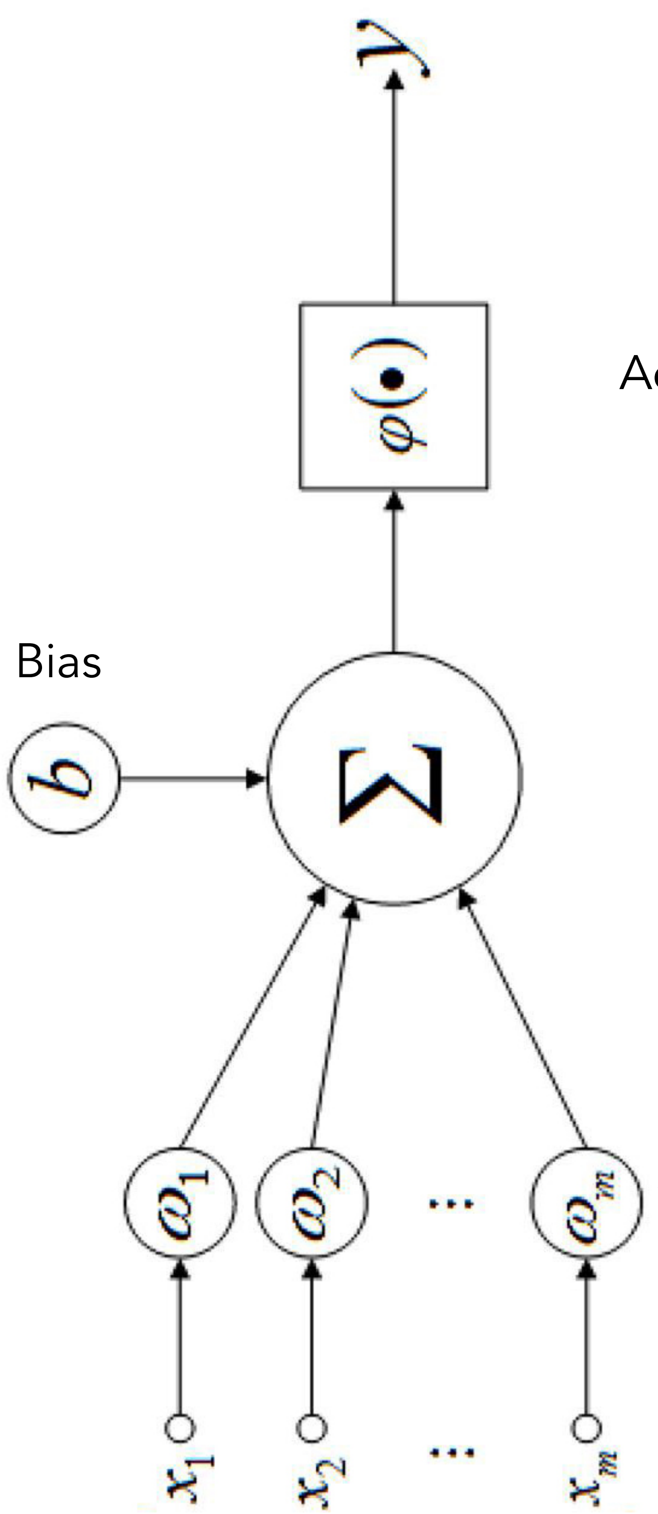

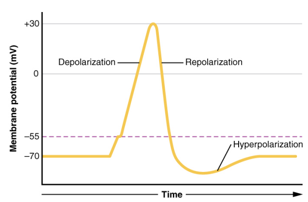

Weights

Synaptic connection

Inputs

Dendrites

Axon terminals
(from other neurons)

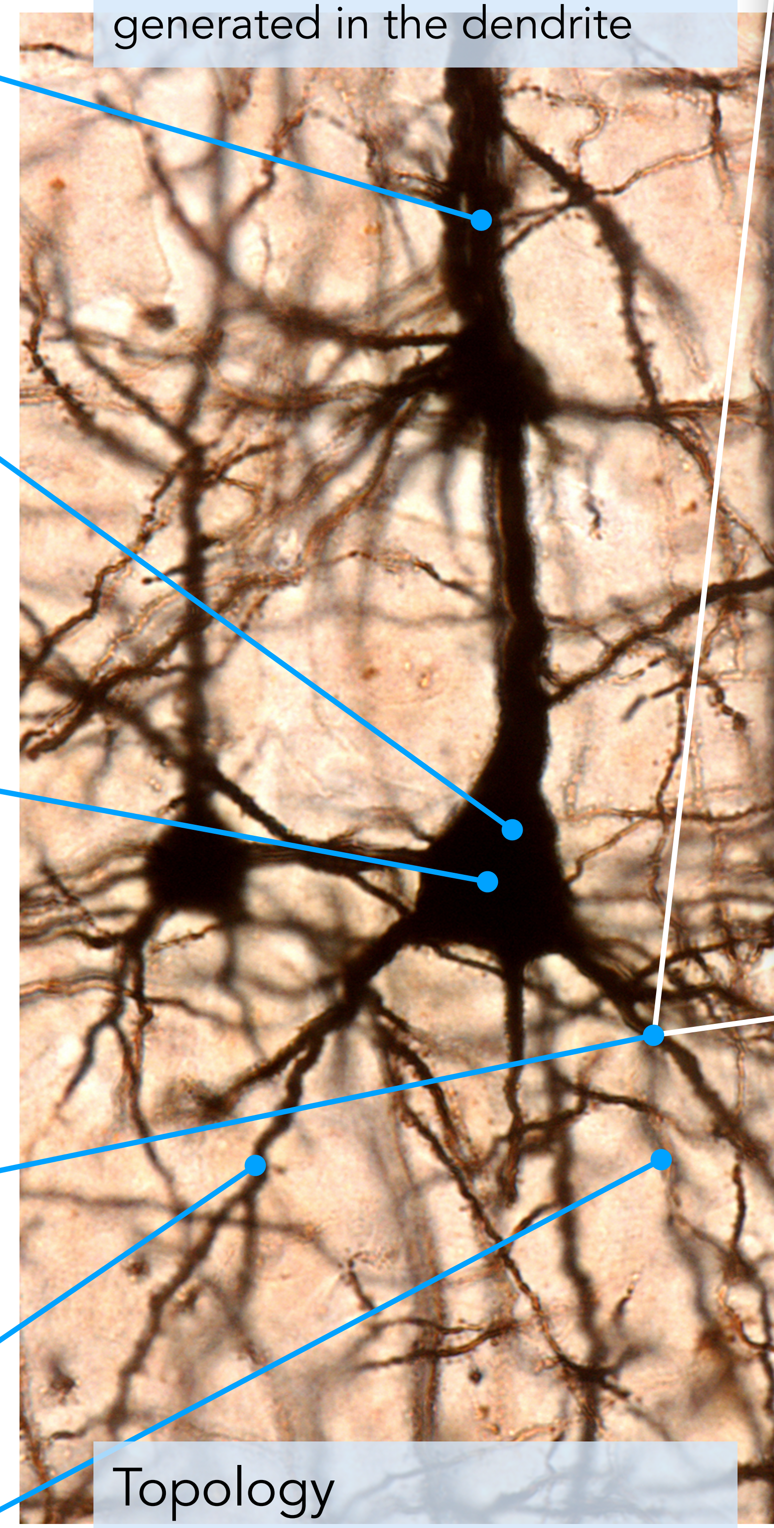

Dendritic spike
action potential can be generated in the dendrite

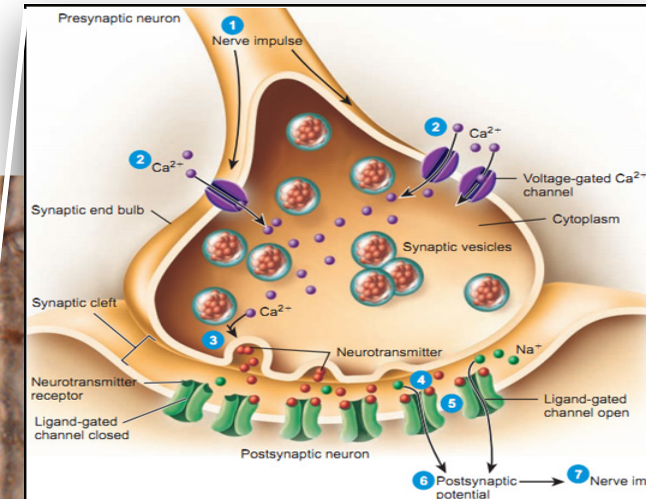

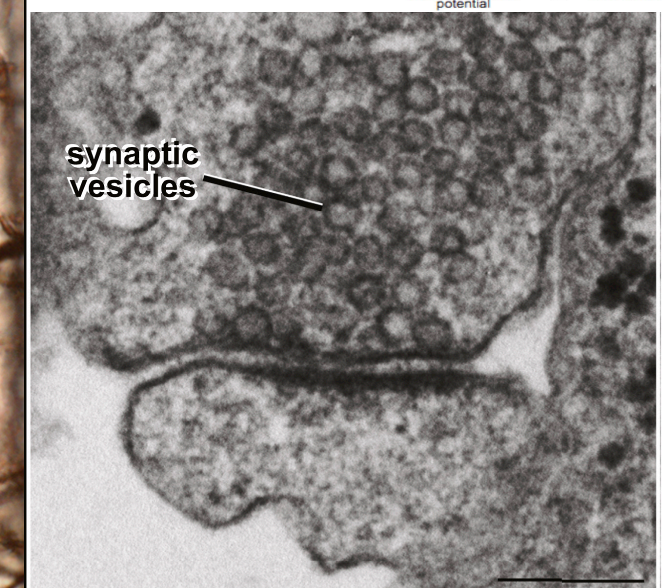

Molecular neurobiology
- different types of channels
- different receptors
- mechanism of neurotransmitter release

Timing
- no clock synchronization
- time of arrival is important

Topology
Artificial networks are neatly organized into graphs, biological network is a mess

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595 https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

Output

Activation function

Sum

Weights

Inputs

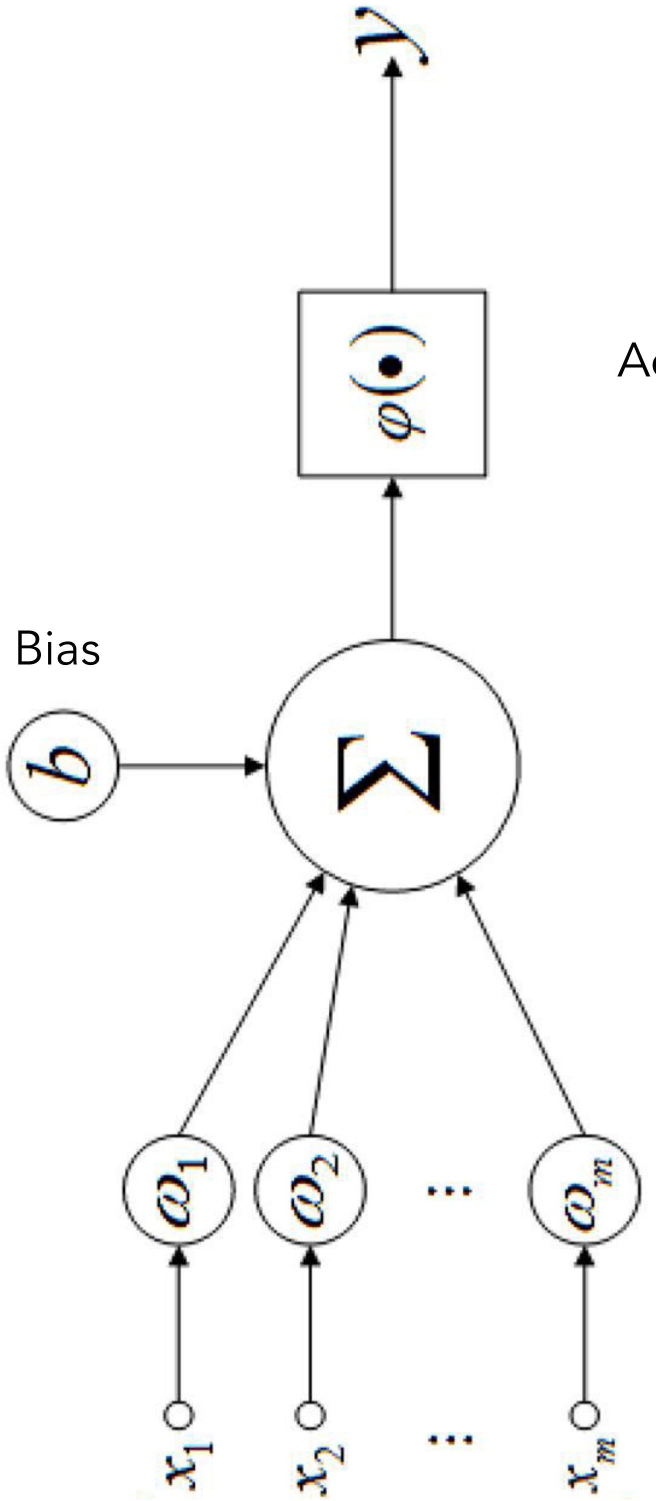

Action potential (aka Spike)

Polarization threshold

Cell depolarization

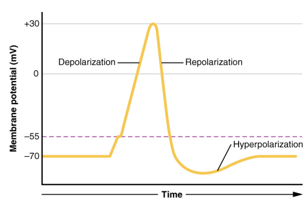

Synaptic connection

Dendrites

Axon terminals (from other neurons)

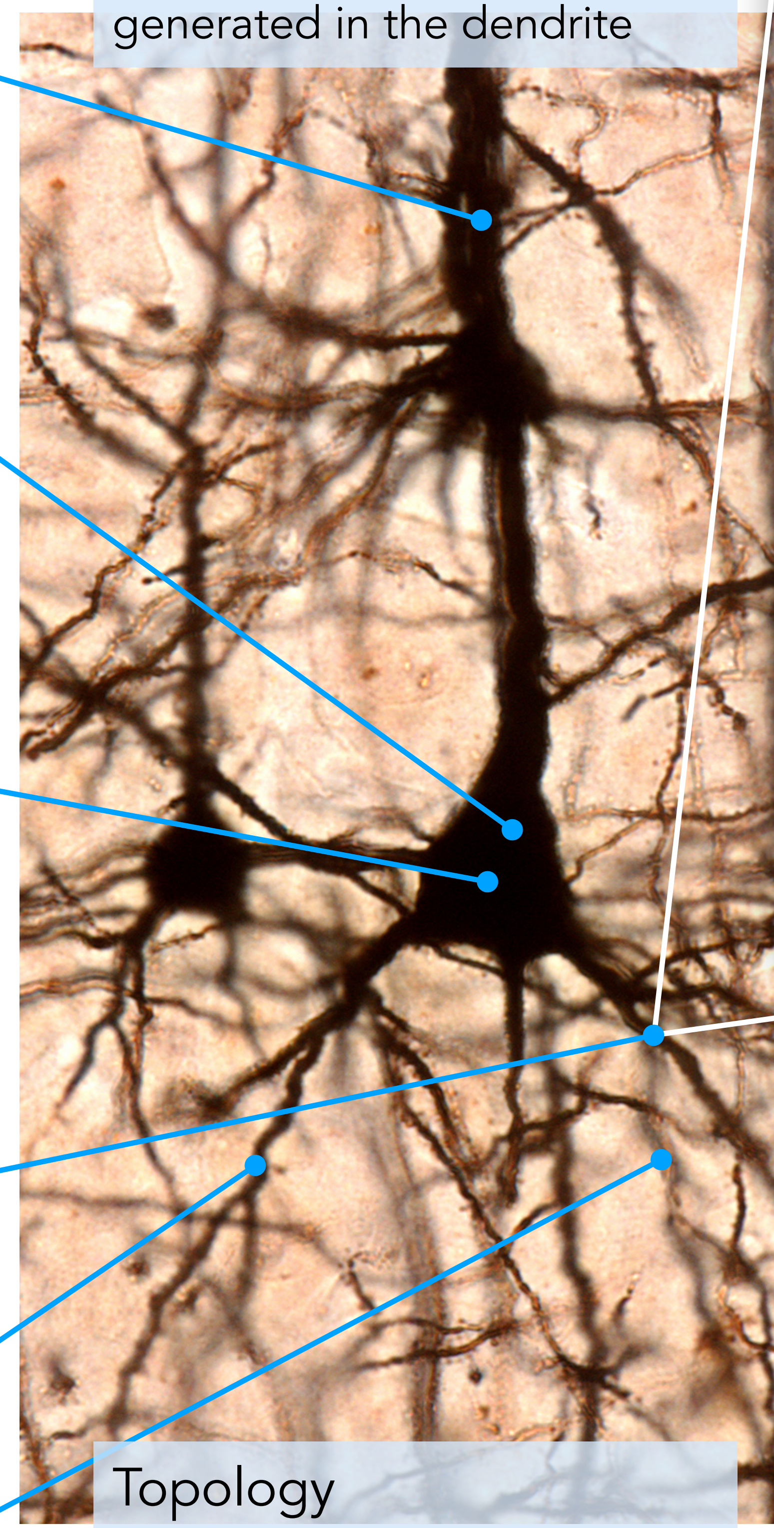

Dendritic spike

action potential can be generated in the dendrite

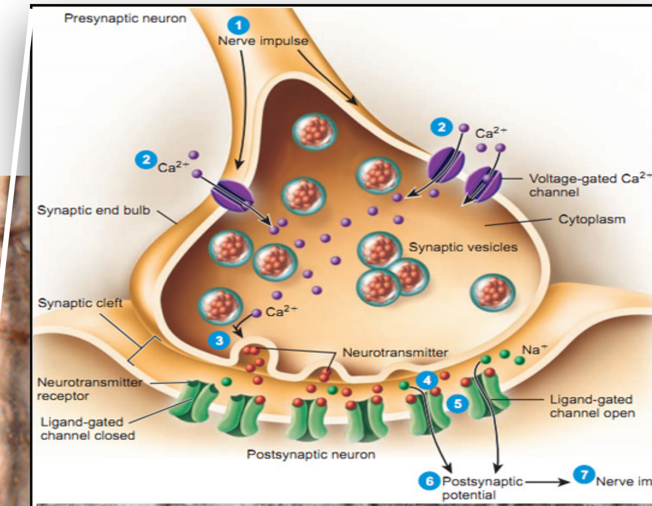

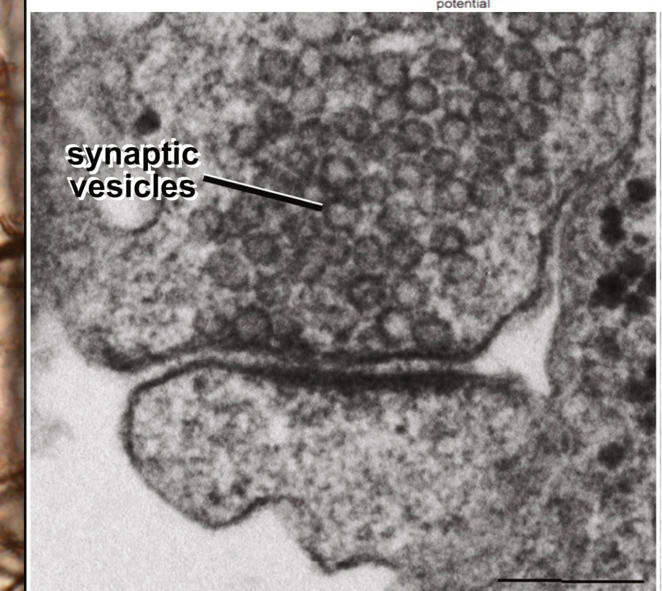

Molecular neurobiology

- different types of channels
- different receptors
- mechanism of neurotransmitter release

Timing

- no clock synchronization
- time of arrival is important

Topology

Artificial networks are neatly organized into graphs, biological network is a mess

Power consumption

Brain uses 20 watt, one Titan X — 250 watt

https://www.dreamstime.com/pyramidal-cell-cerebral-cortex-neurons-stained-golgiâ€™s-silver-chromate-conic-shaped-soma-large-apical-image117240595 https://antranik.org/synaptic-transmission-by-somatic-motorneurons/
https://medium.com/@jayeshbahire/the-artificial-neural-networks-handbook-part-4-d2087d1f583e https://www.sciencedirect.com/science/article/pii/S0896627312005727#fig1

The brain cannot be doing backpropagation

(at least not the way we do it)

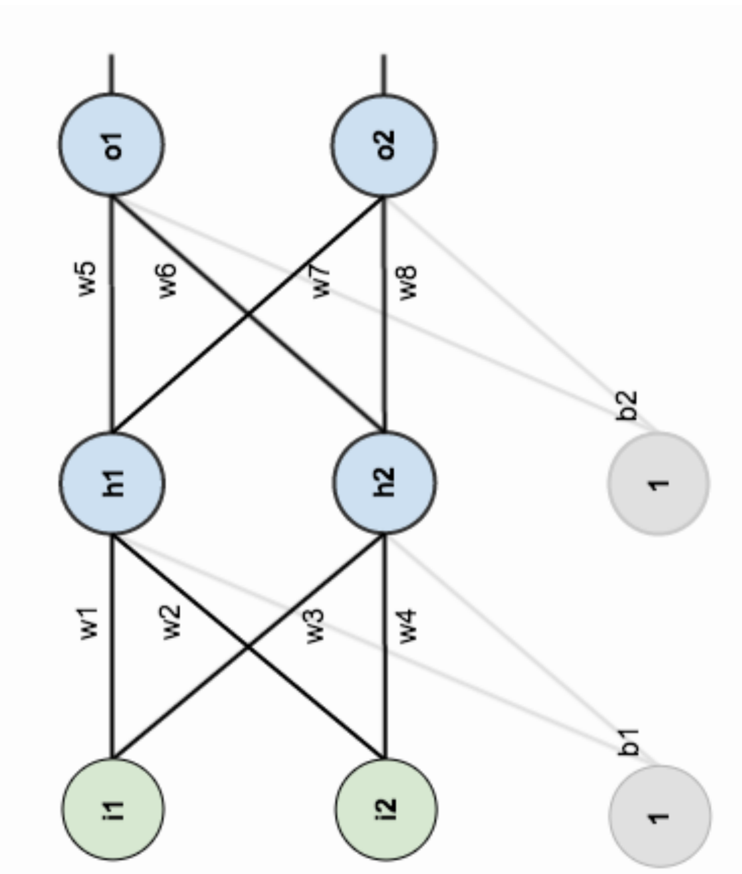

https://mattmazur.com/2015/03/17/a-step-by-step-backpropagation-example/

The brain cannot be doing backpropagation

(at least not the way we do it)

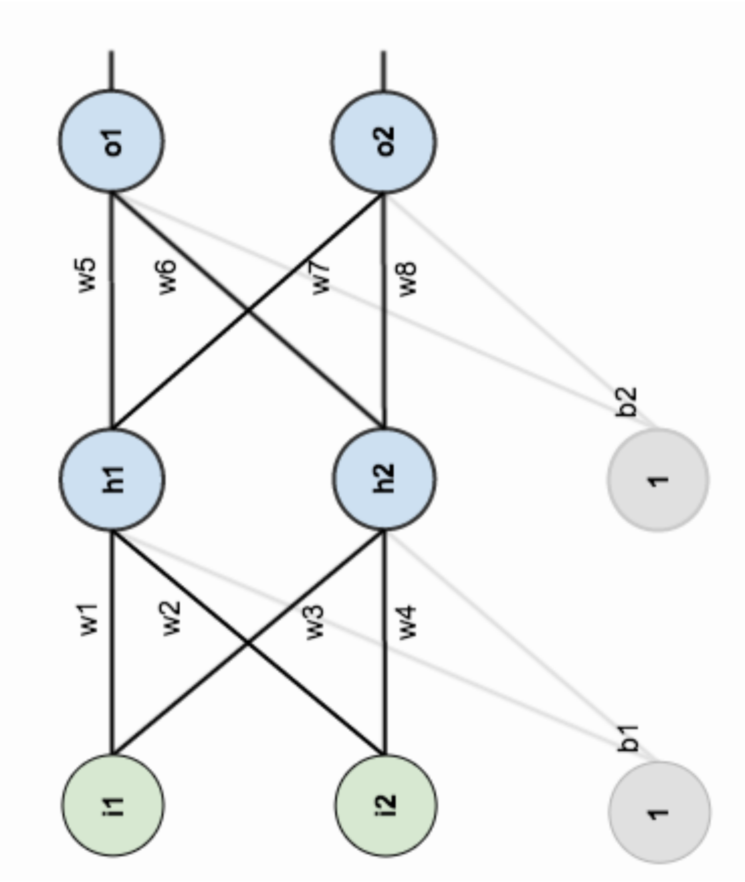

Forward pass

$E_{o1} = \frac{1}{2}(target_{o1} - out_{o1})^2$

$E_{total} = E_{o1} + E_{o2}$

https://mattmazur.com/2015/03/17/a-step-by-step-backpropagation-example/

The brain cannot be doing backpropagation

(at least not the way we do it)

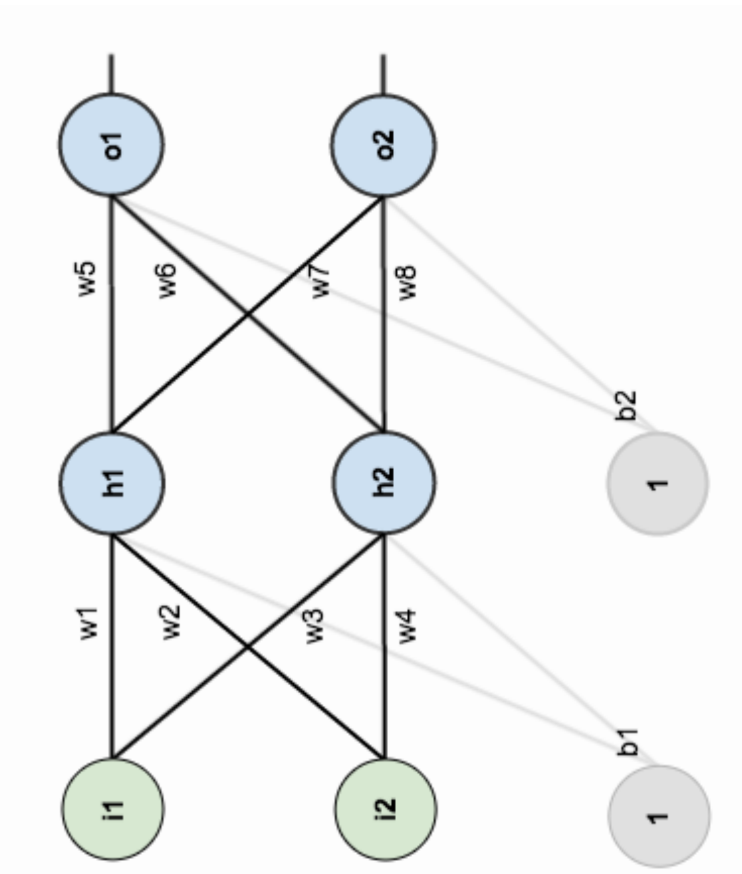

Forward pass

$$E_{o1} = \frac{1}{2}(target_{o1} - out_{o1})^2$$

$$E_{total} = E_{o1} + E_{o2}$$

Backwards pass

$$\frac{\partial E_{total}}{\partial w_5} = \frac{\partial E_{total}}{\partial out_{o1}} * \frac{\partial out_{o1}}{\partial net_{o1}} * \frac{\partial net_{o1}}{\partial w_5}$$

https://mattmazur.com/2015/03/17/a-step-by-step-backpropagation-example/

The brain cannot be doing backpropagation

(at least not the way we do it)

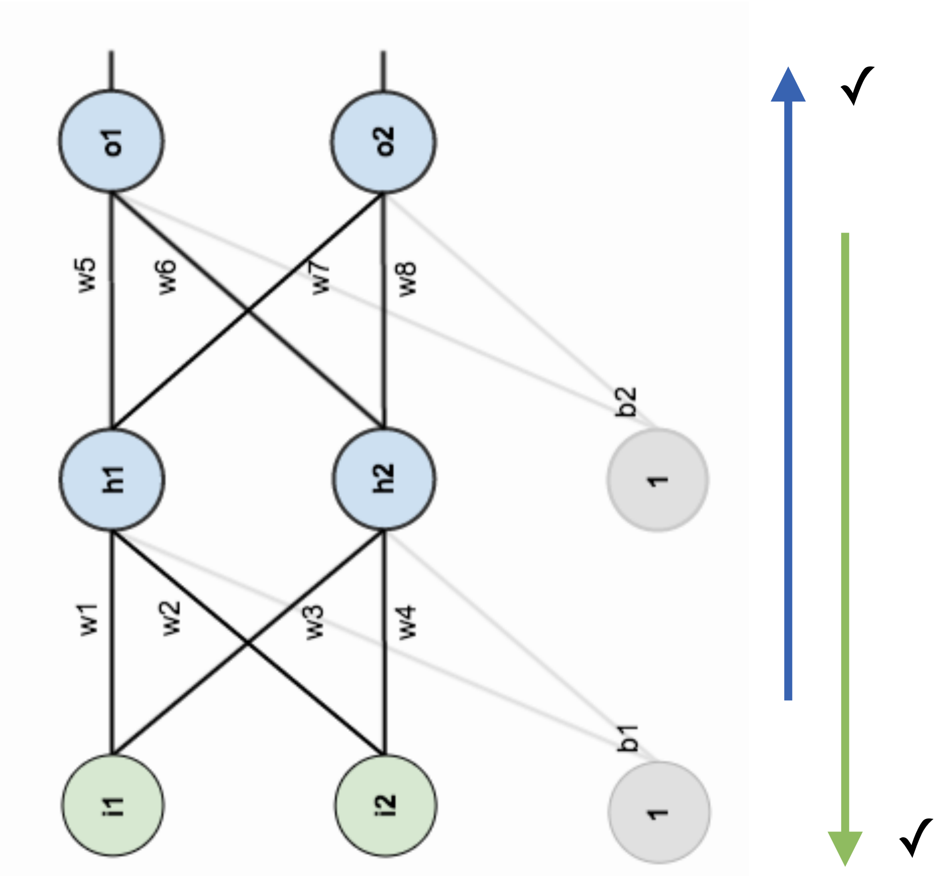

Forward pass

$$E_{o1} = \frac{1}{2}(target_{o1} - out_{o1})^2$$

$$E_{total} = E_{o1} + E_{o2}$$

Backwards pass

$$\frac{\partial E_{total}}{\partial w_5} = \frac{\partial E_{total}}{\partial out_{o1}} * \frac{\partial out_{o1}}{\partial net_{o1}} * \frac{\partial net_{o1}}{\partial w_5}$$

https://mattmazur.com/2015/03/17/a-step-by-step-backpropagation-example/

The brain cannot be doing backpropagation

(at least not the way we do it)

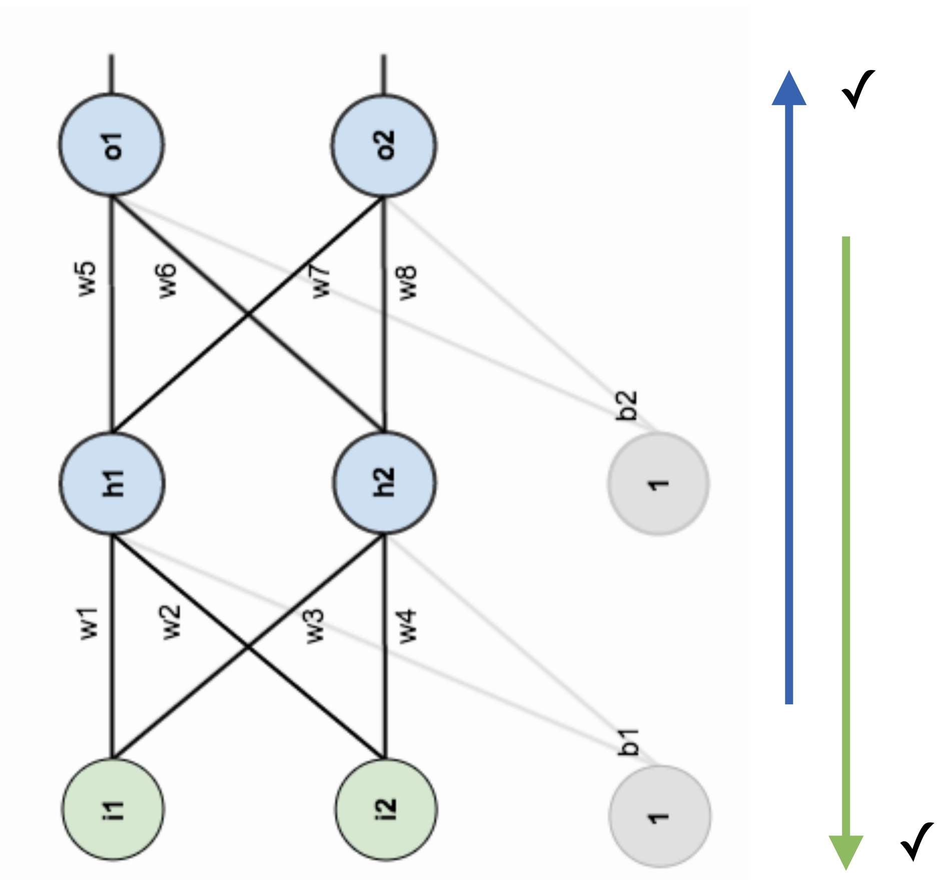

Forward pass

$$E_{o1} = \frac{1}{2}(target_{o1} - out_{o1})^2$$

$$E_{total} = E_{o1} + E_{o2}$$

Backwards pass

$$\frac{\partial E_{total}}{\partial w_5} = \frac{\partial E_{total}}{\partial out_{o1}} * \frac{\partial out_{o1}}{\partial net_{o1}} * \frac{\partial net_{o1}}{\partial w_5}$$

https://mattmazur.com/2015/03/17/a-step-by-step-backpropagation-example/

The brain cannot be doing backpropagation

(at least not the way we do it)

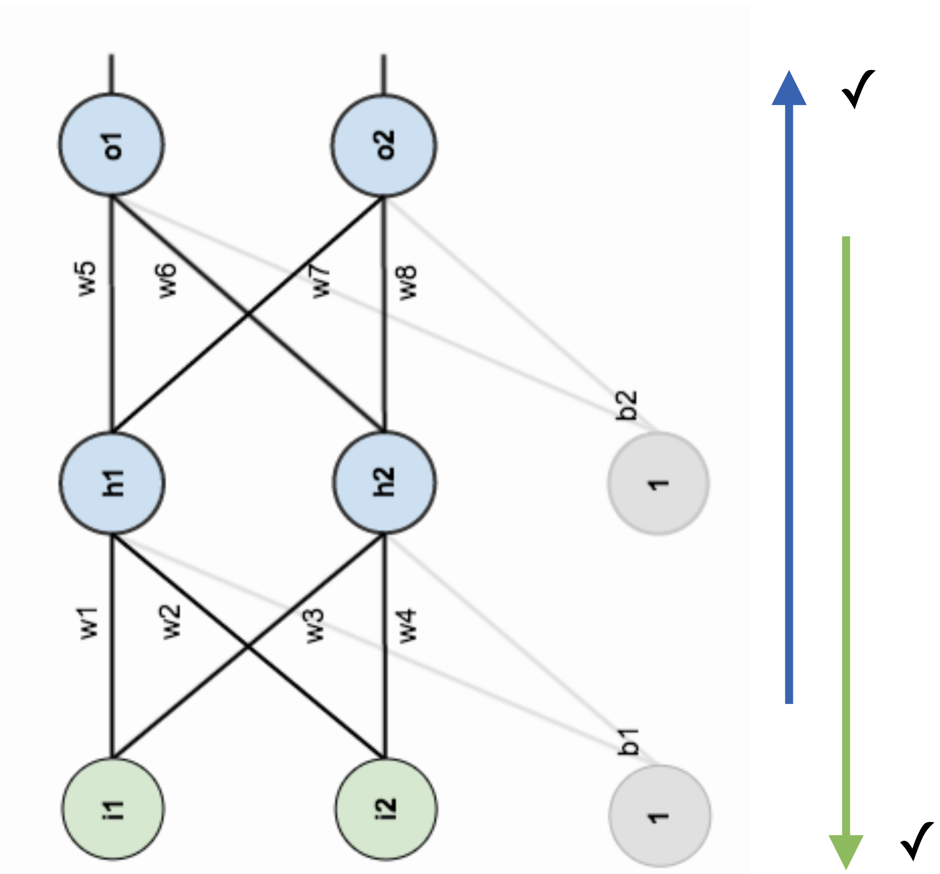

Forward pass

$$E_{o1} = \frac{1}{2}(target_{o1} - out_{o1})^2$$

$$E_{total} = E_{o1} + E_{o2}$$

Forward pass

Neurons send spikes, not real-valued outputs

Backwards pass

$$\frac{\partial E_{total}}{\partial w_5} = \frac{\partial E_{total}}{\partial out_{o1}} * \frac{\partial out_{o1}}{\partial net_{o1}} * \frac{\partial net_{o1}}{\partial w_5}$$

https://mattmazur.com/2015/03/17/a-step-by-step-backpropagation-example/

https://www.cs.toronto.edu/~hinton/backpropincortex2014.pdf

The brain cannot be doing backpropagation

(at least not the way we do it)

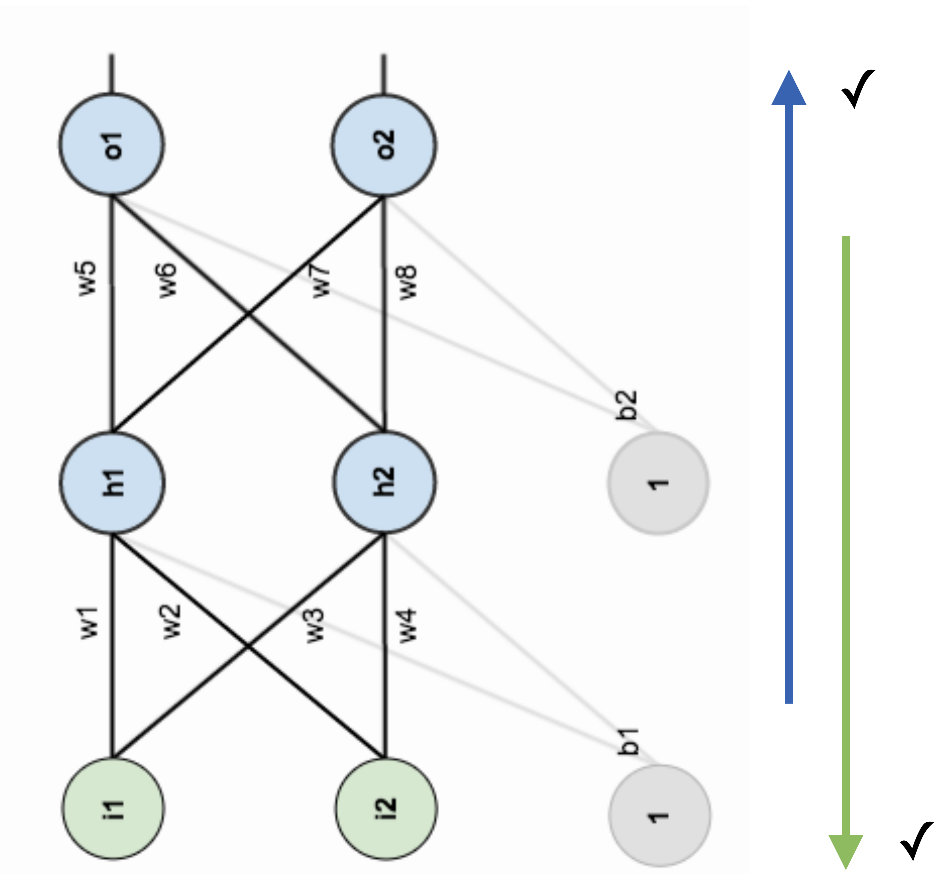

Forward pass

$$E_{o1} = \frac{1}{2}(target_{o1} - out_{o1})^2$$

$$E_{total} = E_{o1} + E_{o2}$$

Backwards pass

$$\frac{\partial E_{total}}{\partial w_5} = \frac{\partial E_{total}}{\partial out_{o1}} * \frac{\partial out_{o1}}{\partial net_{o1}} * \frac{\partial net_{o1}}{\partial w_5}$$

Forward pass


WIKIPEDIA
The Free Encyclopedia
Article Talk
Predictive coding
From Wikipedia, the free encyclopedia

Neurons send spikes, not real-valued outputs

Backwards pass

Neurons would need bidirectional connections with the same weight

Neurons would need to send two types of signal and estimate derivative

https://mattmazur.com/2015/03/17/a-step-by-step-backpropagation-example/

https://www.cs.toronto.edu/~hinton/backpropincortex2014.pdf

The brain cannot be doing backpropagation

(at least not the way we do it)

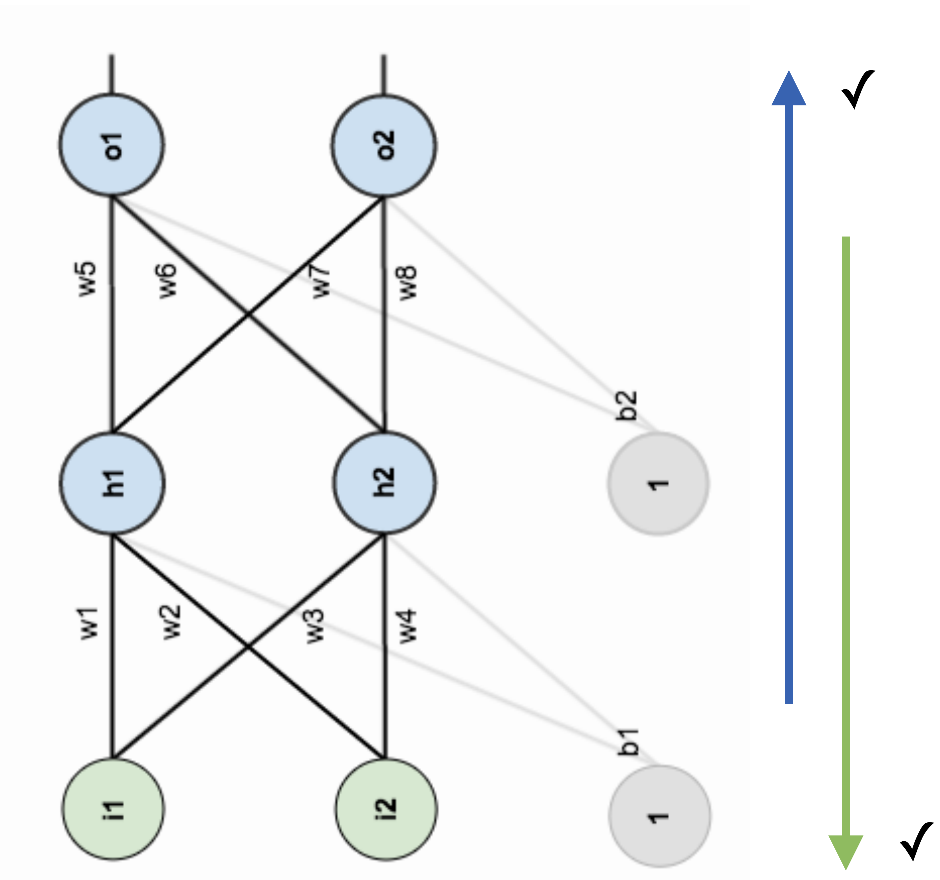

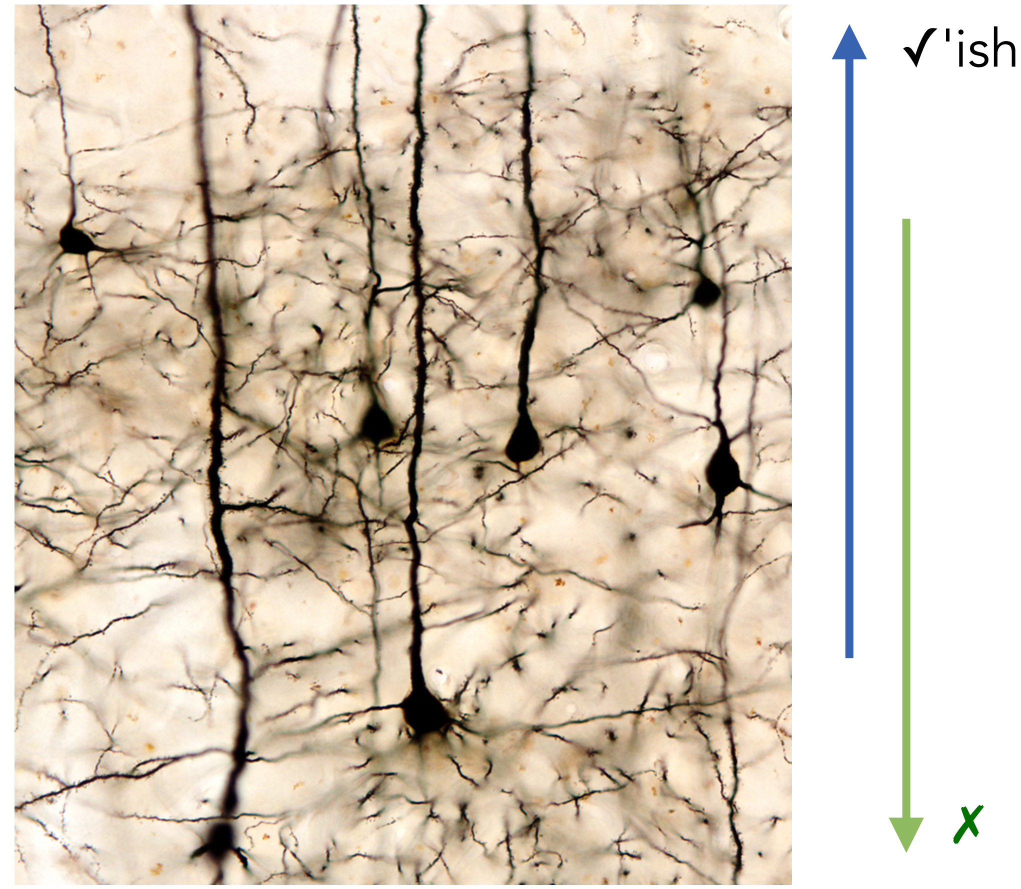

Forward pass

$$E_{o1} = \frac{1}{2}(target_{o1} - out_{o1})^2$$

$$E_{total} = E_{o1} + E_{o2}$$

Backwards pass

$$\frac{\partial E_{total}}{\partial w_5} = \frac{\partial E_{total}}{\partial out_{o1}} * \frac{\partial out_{o1}}{\partial net_{o1}} * \frac{\partial net_{o1}}{\partial w_5}$$

Forward pass

Neurons send spikes, not real-valued outputs

Backwards pass

Neurons would need bidirectional connections with the same weight

Neurons would need to send two types of signal and estimate derivative

https://mattmazur.com/2015/03/17/a-step-by-step-backpropagation-example/

https://www.cs.toronto.edu/~hinton/backpropincortex2014.pdf

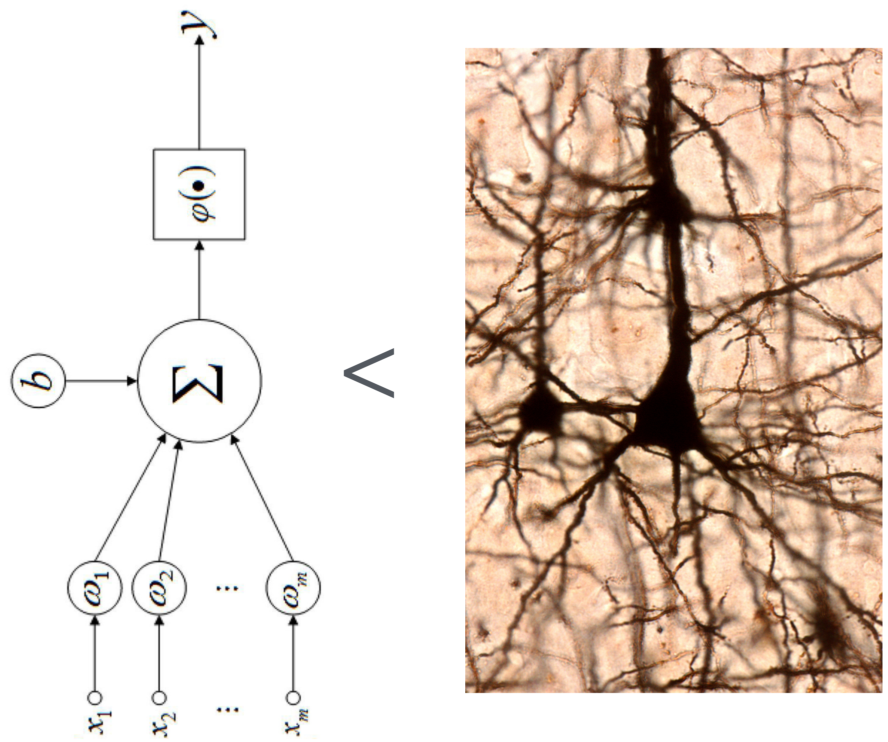

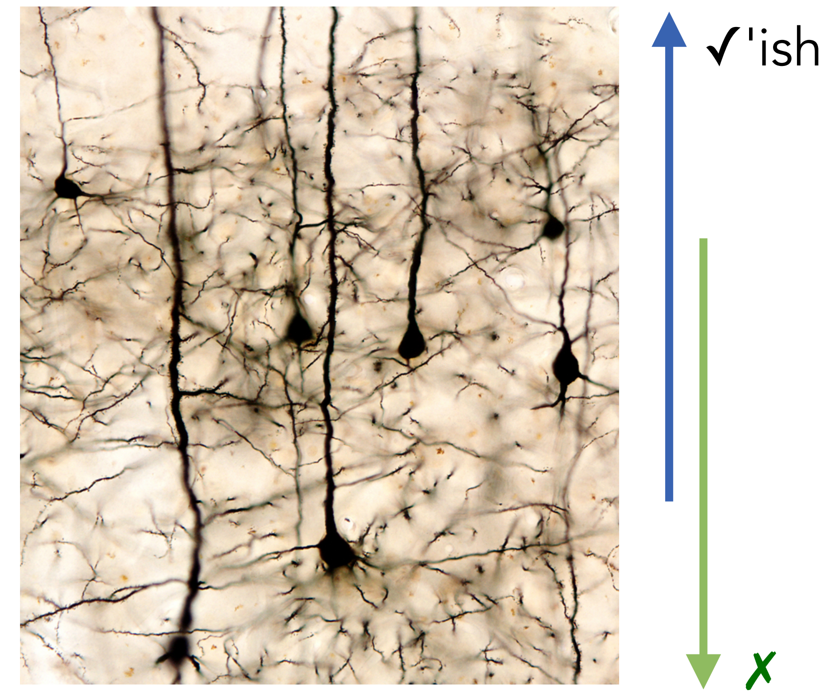

Quite different on the
level of implementation

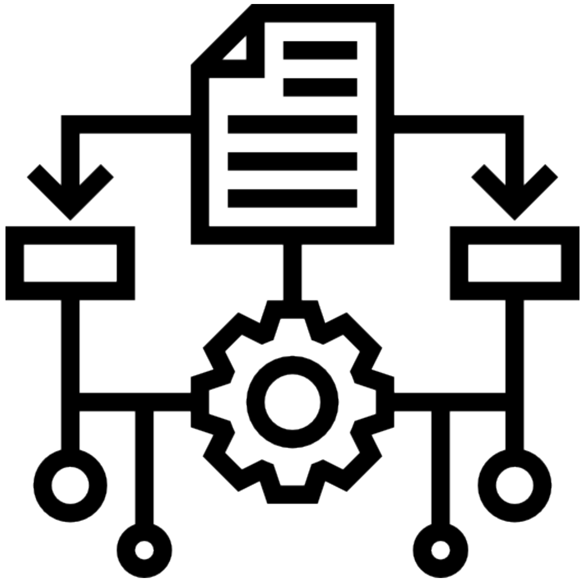

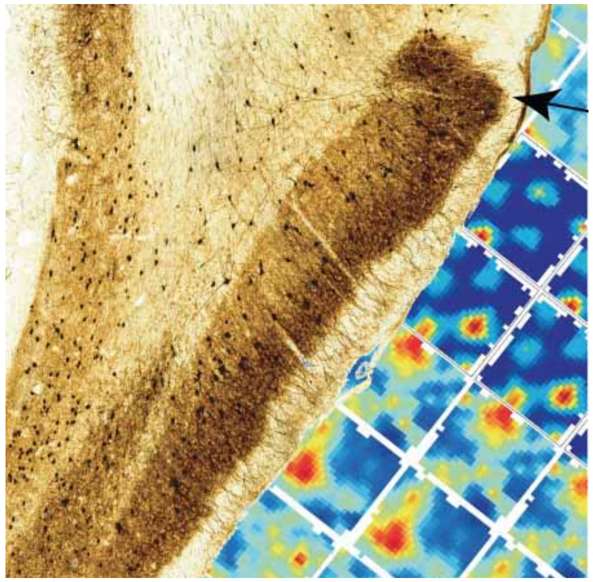

Algorithm and representation

Hippocampus and experience replay

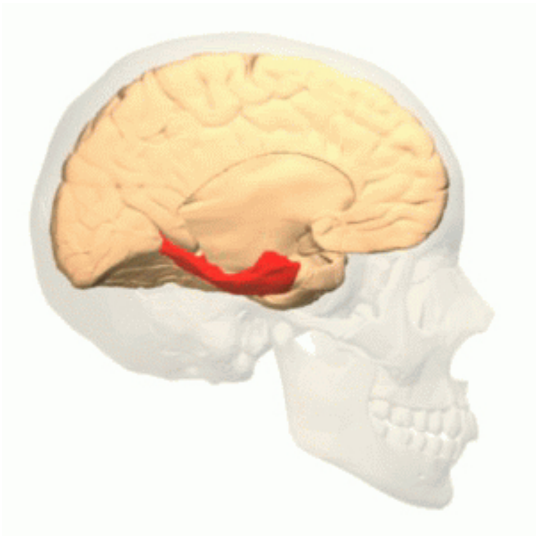

Hippocampus

Hippocampus and experience replay

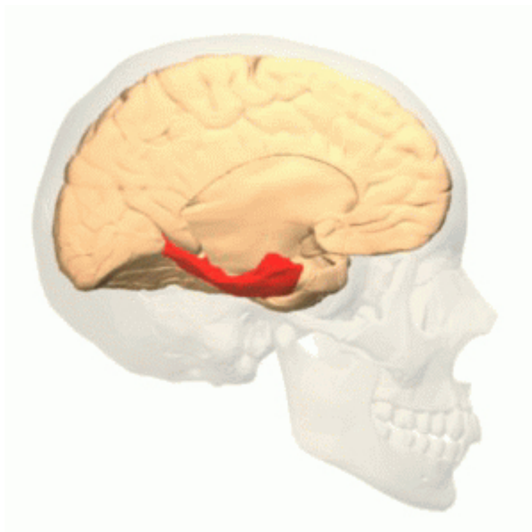

Hippocampus

WIKIPEDIA
The Free Encyclopedia

Memory consolidation

From Wikipedia, the free encyclopedia

Hippocampus and experience replay

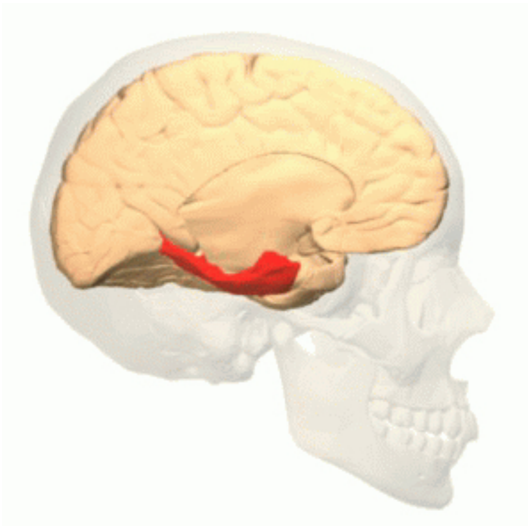

Hippocampus

WIKIPEDIA
The Free Encyclopedia

Memory consolidation

From Wikipedia, the free encyclopedia

WIKIPEDIA
The Free Encyclopedia

Hippocampal replay

From Wikipedia, the free encyclopedia

Hippocampus and experience replay

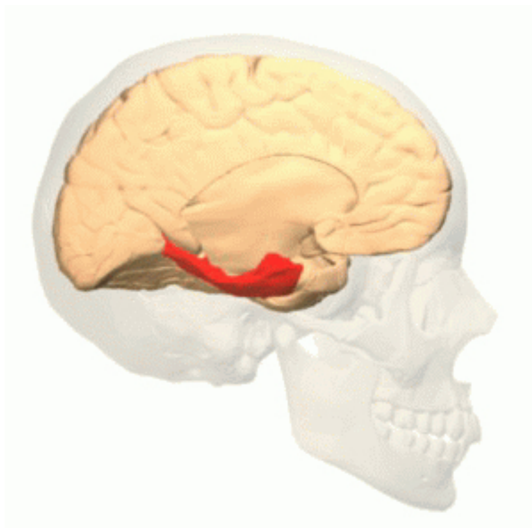

Hippocampus

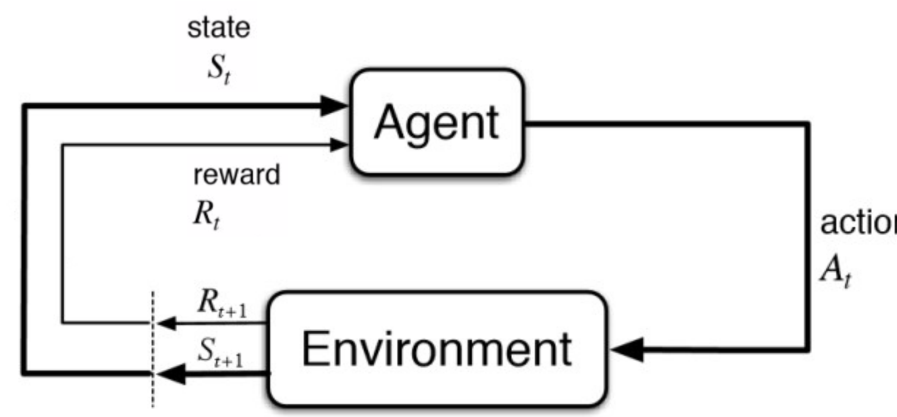

WIKIPEDIA
The Free Encyclopedia

Memory consolidation

From Wikipedia, the free encyclopedia

WIKIPEDIA
The Free Encyclopedia

Hippocampal replay

From Wikipedia, the free encyclopedia

Hippocampus and experience replay

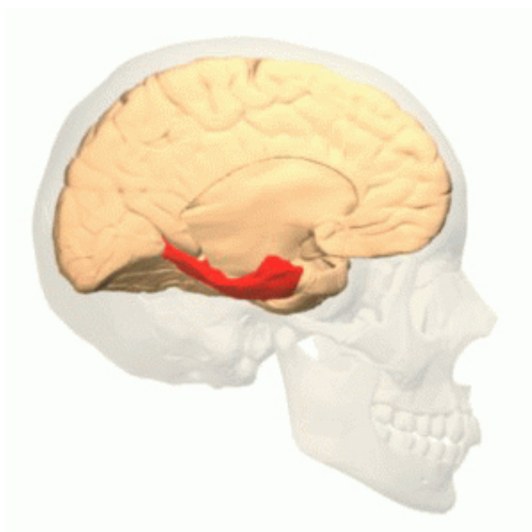

Hippocampus

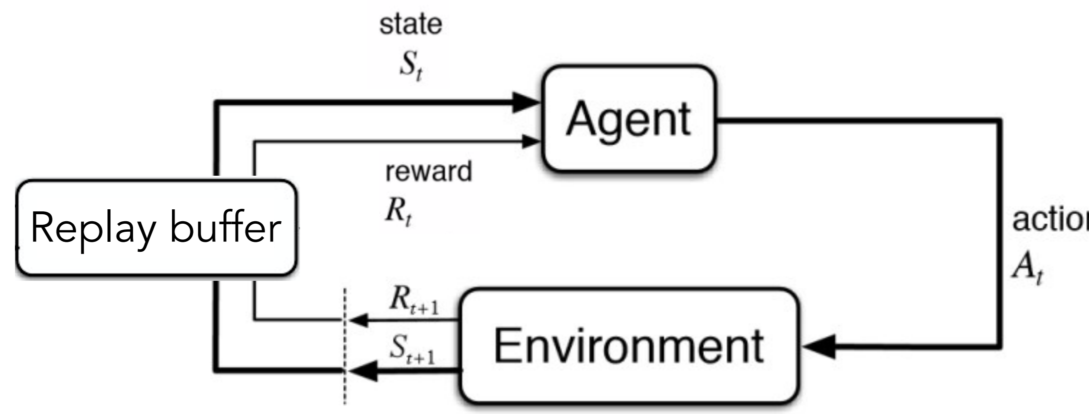

WIKIPEDIA
The Free Encyclopedia

Memory consolidation

From Wikipedia, the free encyclopedia

WIKIPEDIA
The Free Encyclopedia

Hippocampal replay

From Wikipedia, the free encyclopedia

DQN algorithm has state memory buffer and experience replay mechanism:

- store experiences while interacting with the environment

- randomly sample and replay experiences from memory while learning

Vision: hierarchy of layers

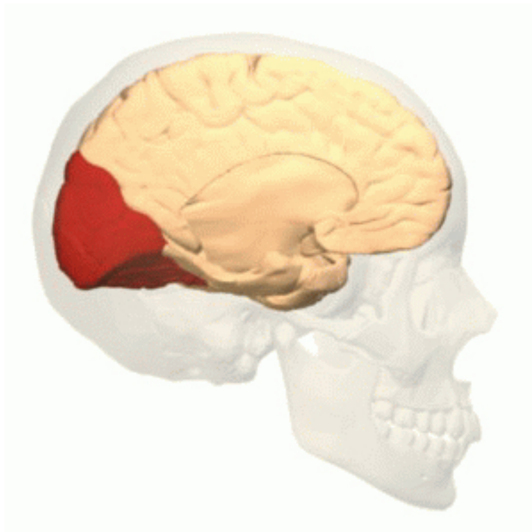

Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence https://www.nature.com/articles/srep27755
Activations of deep convolutional neural networks are aligned with gamma band activity of human visual cortex https://www.nature.com/articles/s42003-018-0110-y
Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream https://www.jneurosci.org/content/35/27/10005

Vision: hierarchy of layers

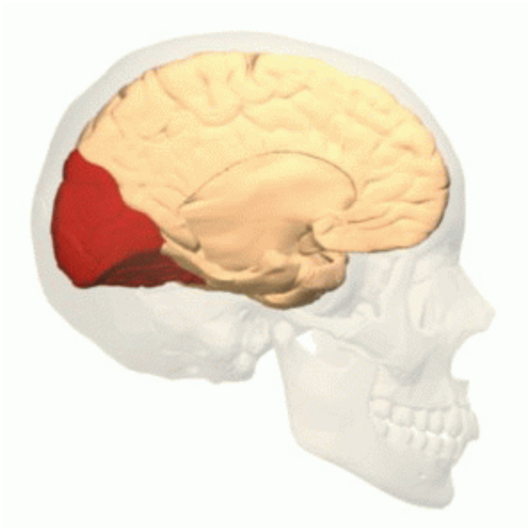

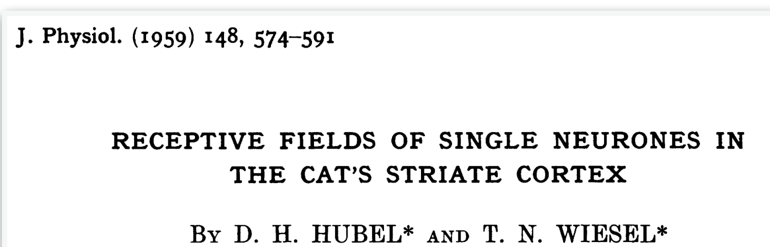

J. Physiol. (1959) 148, 574–591

RECEPTIVE FIELDS OF SINGLE NEURONES IN THE CAT'S STRIATE CORTEX

By D. H. HUBEL* and T. N. WIESEL*

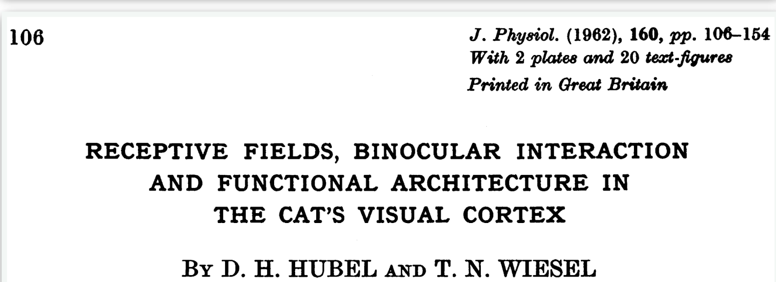

106

J. Physiol. (1962), **160**, *pp.* 106–154
With 2 plates and 20 text-figures
Printed in Great Britain

RECEPTIVE FIELDS, BINOCULAR INTERACTION AND FUNCTIONAL ARCHITECTURE IN THE CAT'S VISUAL CORTEX

By D. H. HUBEL and T. N. WIESEL

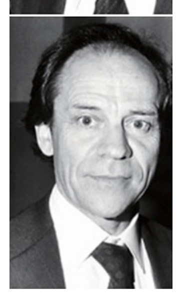

& multiple papers since:

Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence https://www.nature.com/articles/srep27755

Activations of deep convolutional neural networks are aligned with gamma band activity of human visual cortex https://www.nature.com/articles/s42003-018-0110-y

Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream https://www.jneurosci.org/content/35/27/10005

Vision: hierarchy of layers

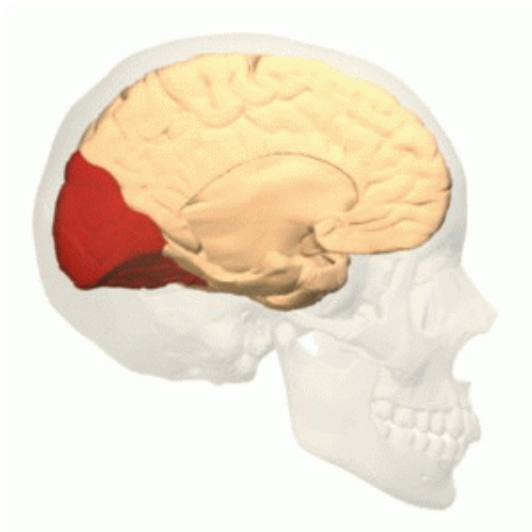

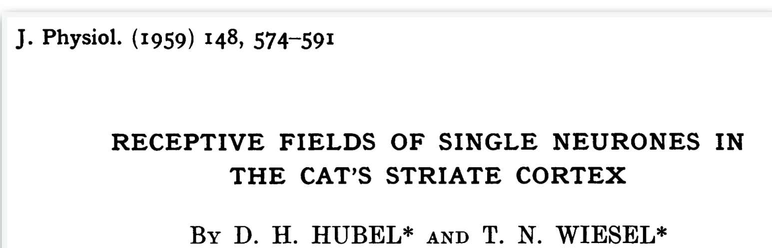
J. Physiol. (1959) 148, 574–591

RECEPTIVE FIELDS OF SINGLE NEURONES IN THE CAT'S STRIATE CORTEX

BY D. H. HUBEL* AND T. N. WIESEL*

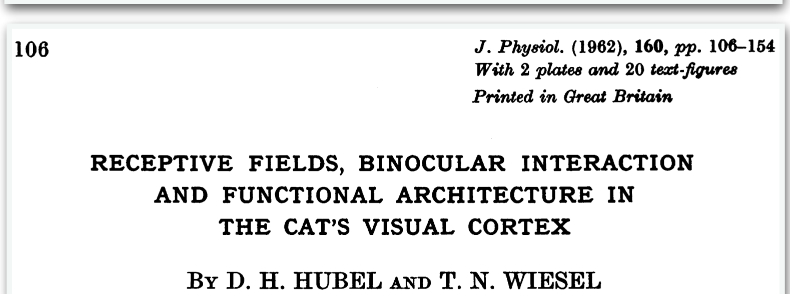
106

J. Physiol. (1962), **160**, *pp.* 106–154
With 2 plates and 20 text-figures
Printed in Great Britain

RECEPTIVE FIELDS, BINOCULAR INTERACTION AND FUNCTIONAL ARCHITECTURE IN THE CAT'S VISUAL CORTEX

BY D. H. HUBEL AND T. N. WIESEL

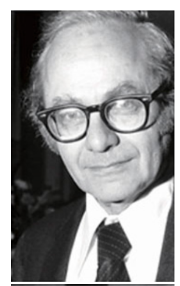

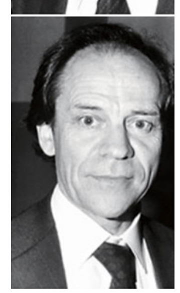

& multiple papers since:

Layered structure of visual cortex

Lower layers process simple visual features, higher layers -- complex features

Neurons focus on certain areas of visual input

Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence https://www.nature.com/articles/srep27755
Activations of deep convolutional neural networks are aligned with gamma band activity of human visual cortex https://www.nature.com/articles/s42003-018-0110-y
Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream https://www.jneurosci.org/content/35/27/10005

Vision: hierarchy of layers

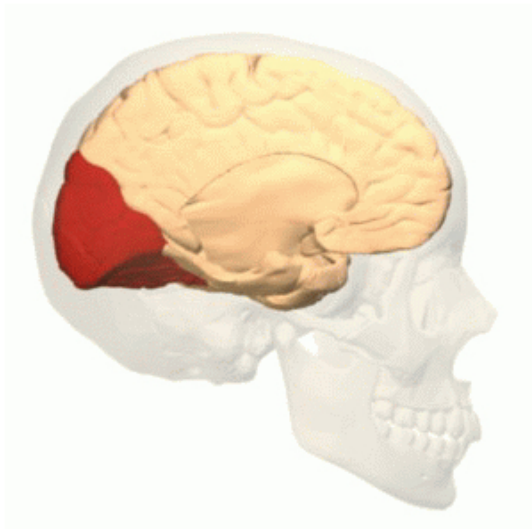

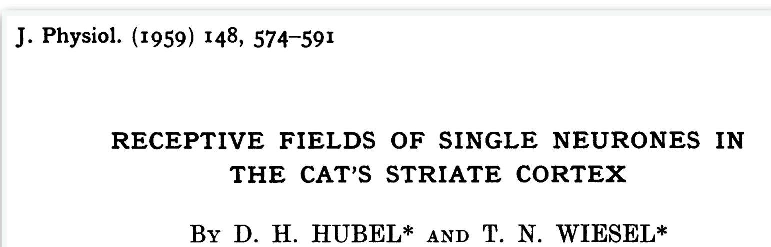
J. Physiol. (1959) 148, 574–591

RECEPTIVE FIELDS OF SINGLE NEURONES IN THE CAT'S STRIATE CORTEX

By D. H. HUBEL* and T. N. WIESEL*

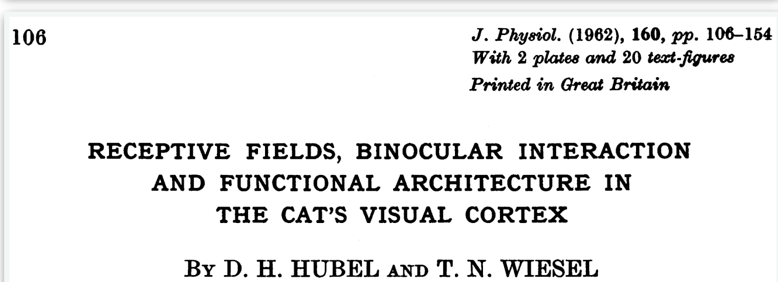
106

J. Physiol. (1962), 160, pp. 106–154
With 2 plates and 20 text-figures
Printed in Great Britain

RECEPTIVE FIELDS, BINOCULAR INTERACTION AND FUNCTIONAL ARCHITECTURE IN THE CAT'S VISUAL CORTEX

By D. H. HUBEL and T. N. WIESEL

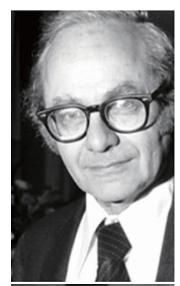

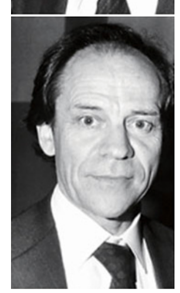

& multiple papers since:

Layered structure of visual cortex

Lower layers process simple visual features, higher layers -- complex features

Neurons focus on certain areas of visual input

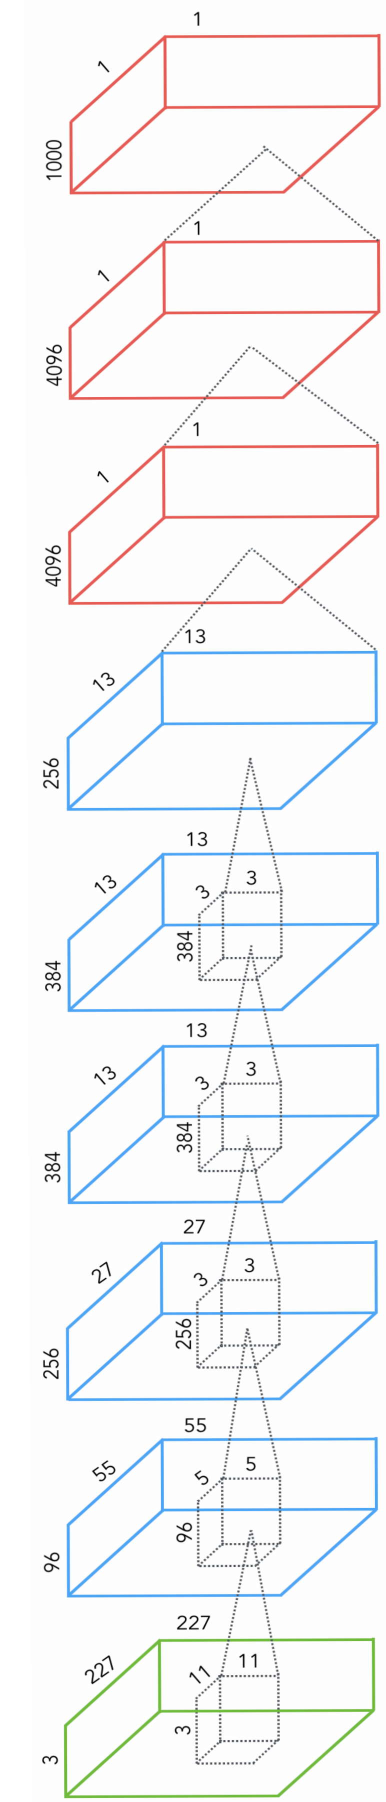

Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence https://www.nature.com/articles/srep27755

Activations of deep convolutional neural networks are aligned with gamma band activity of human visual cortex https://www.nature.com/articles/s42003-018-0110-y

Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream https://www.jneurosci.org/content/35/27/10005

Vision: hierarchy of layers

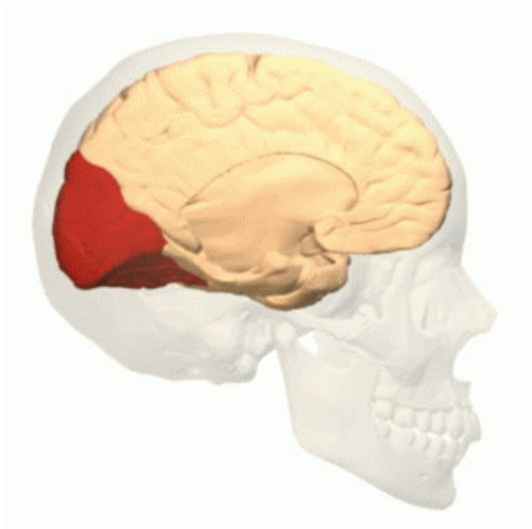

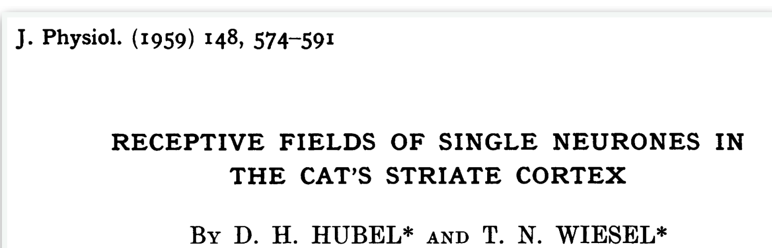

J. Physiol. (1959) 148, 574–591

RECEPTIVE FIELDS OF SINGLE NEURONES IN THE CAT'S STRIATE CORTEX

By D. H. HUBEL* and T. N. WIESEL*

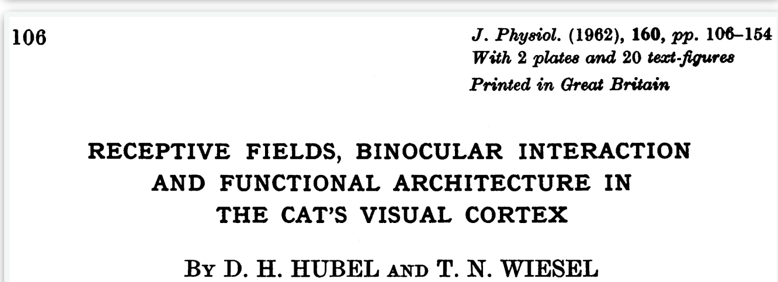

106

J. Physiol. (1962), 160, *pp.* 106–154
With 2 plates and 20 text-figures
Printed in Great Britain

RECEPTIVE FIELDS, BINOCULAR INTERACTION AND FUNCTIONAL ARCHITECTURE IN THE CAT'S VISUAL CORTEX

By D. H. HUBEL and T. N. WIESEL

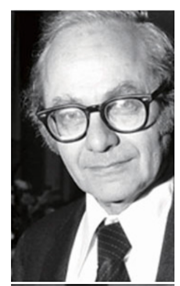

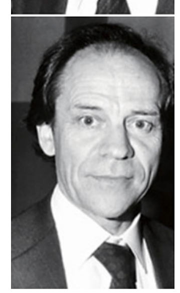

& multiple papers since:

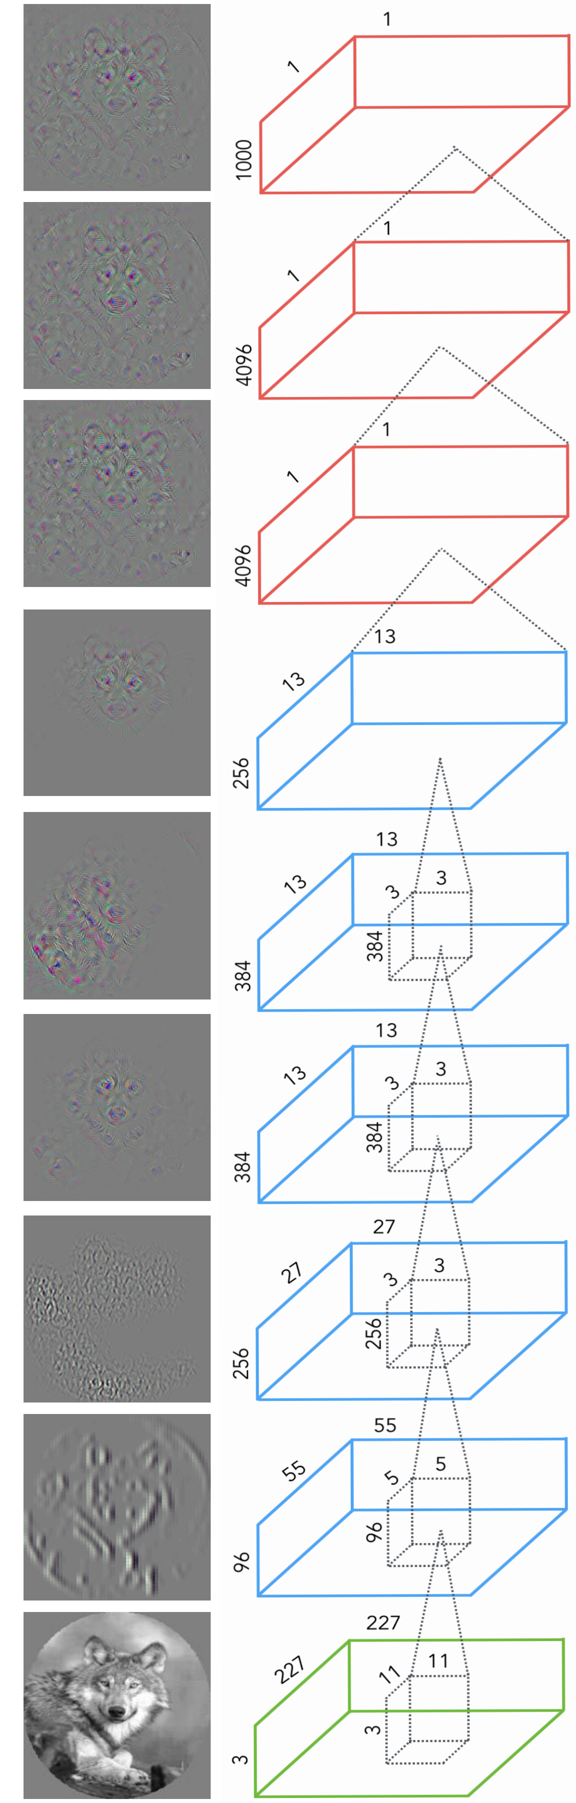

Layered structure of visual cortex

Lower layers process simple visual features, higher layers -- complex features

Neurons focus on certain areas of visual input

Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence https://www.nature.com/articles/srep27755
Activations of deep convolutional neural networks are aligned with gamma band activity of human visual cortex https://www.nature.com/articles/s42003-018-0110-y
Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream https://www.jneurosci.org/content/35/27/10005

Vision: hierarchy of layers

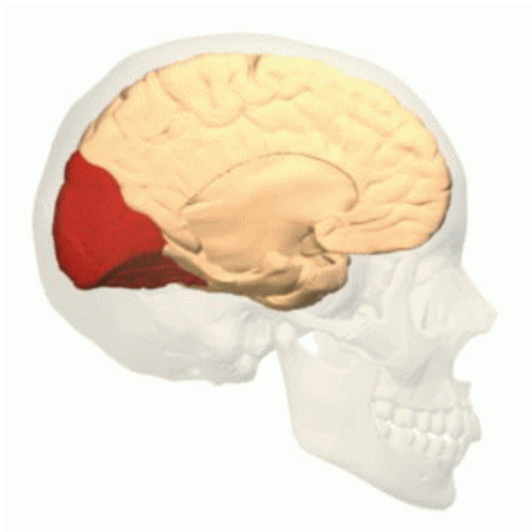

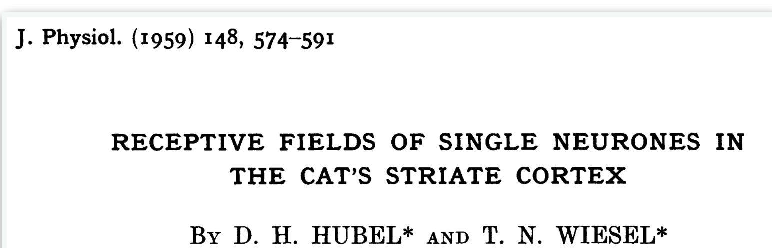
J. Physiol. (1959) 148, 574–591

RECEPTIVE FIELDS OF SINGLE NEURONES IN THE CAT'S STRIATE CORTEX

BY D. H. HUBEL* AND T. N. WIESEL*

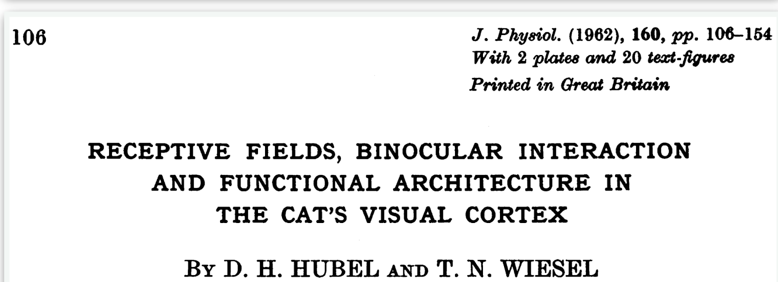
106

J. Physiol. (1962), 160, pp. 106–154
With 2 plates and 20 text-figures
Printed in Great Britain

RECEPTIVE FIELDS, BINOCULAR INTERACTION AND FUNCTIONAL ARCHITECTURE IN THE CAT'S VISUAL CORTEX

BY D. H. HUBEL AND T. N. WIESEL

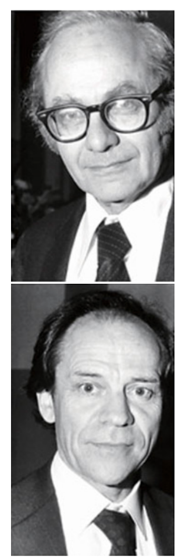

& multiple papers since:

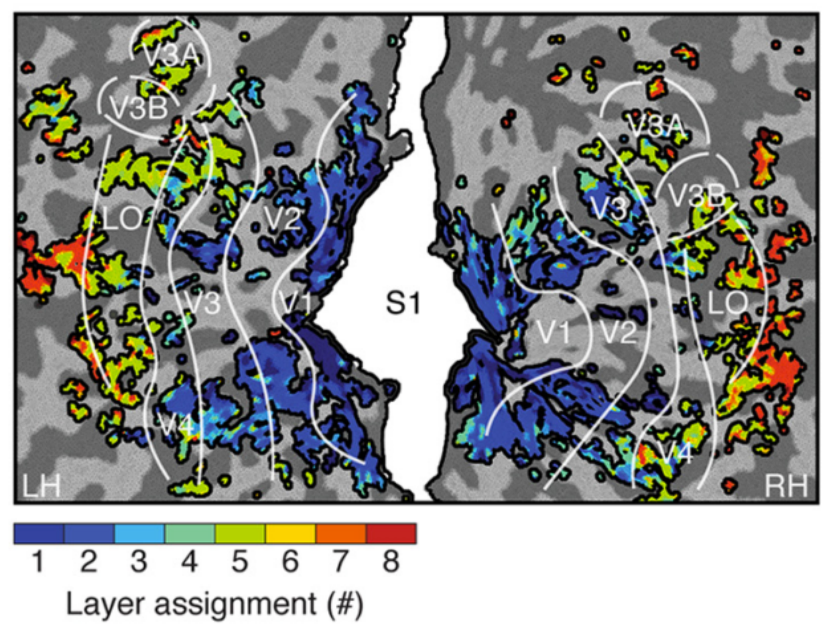

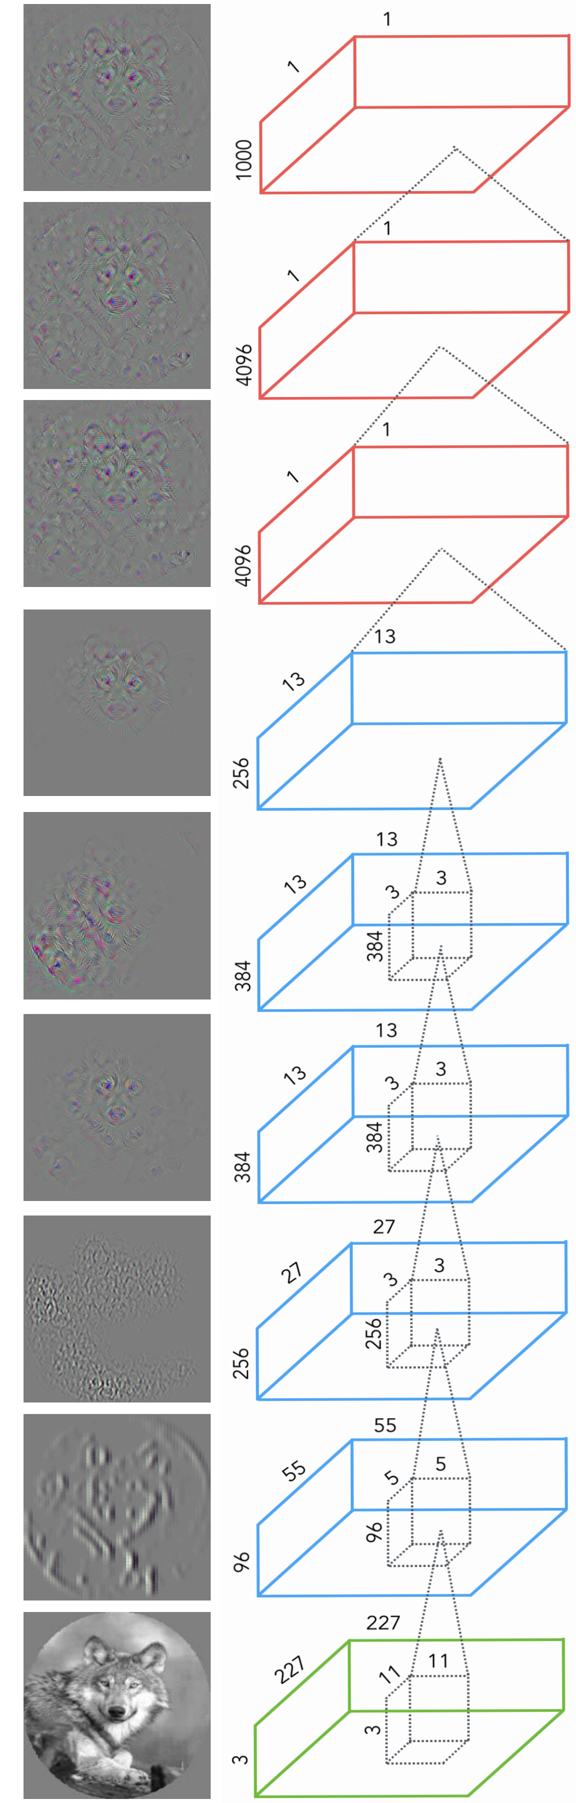

Layered structure of visual cortex

Lower layers process simple visual features, higher layers -- complex features

Neurons focus on certain areas of visual input

Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence https://www.nature.com/articles/srep27755
Activations of deep convolutional neural networks are aligned with gamma band activity of human visual cortex https://www.nature.com/articles/s42003-018-0110-y
Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream https://www.jneurosci.org/content/35/27/10005

Vision: hierarchy of layers

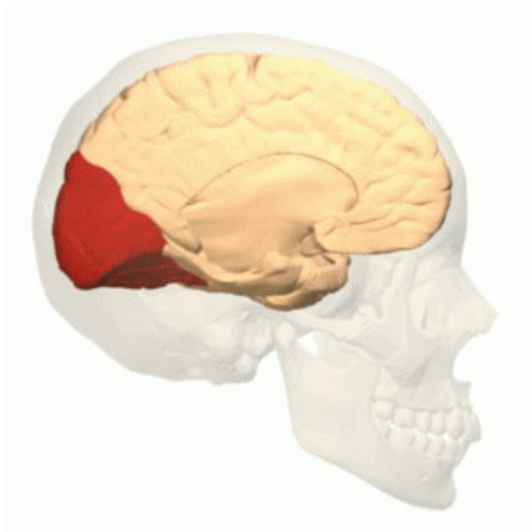

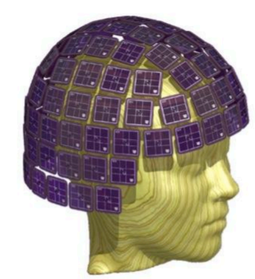

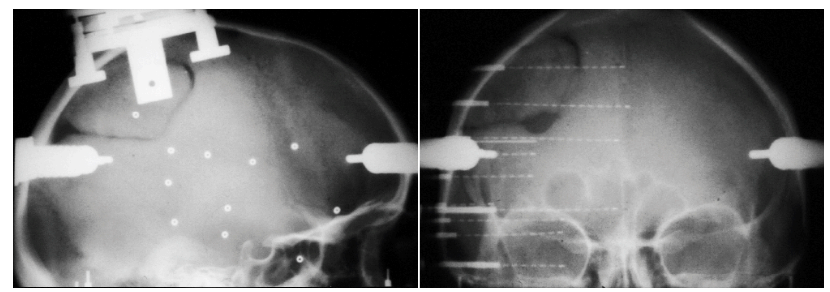

J. Physiol. (1959) 148, 574–591

RECEPTIVE FIELDS OF SINGLE NEURONES IN THE CAT'S STRIATE CORTEX

BY D. H. HUBEL* AND T. N. WIESEL*

106

J. Physiol. (1962), 160, pp. 106–154
With 2 plates and 20 text-figures
Printed in Great Britain

RECEPTIVE FIELDS, BINOCULAR INTERACTION AND FUNCTIONAL ARCHITECTURE IN THE CAT'S VISUAL CORTEX

BY D. H. HUBEL AND T. N. WIESEL

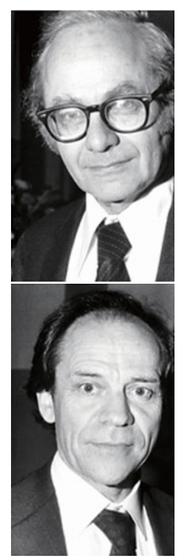

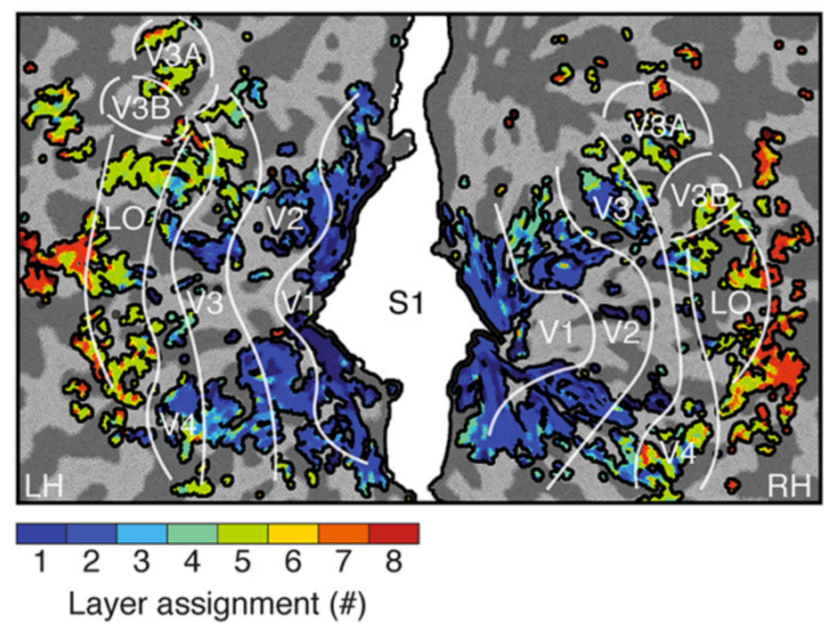

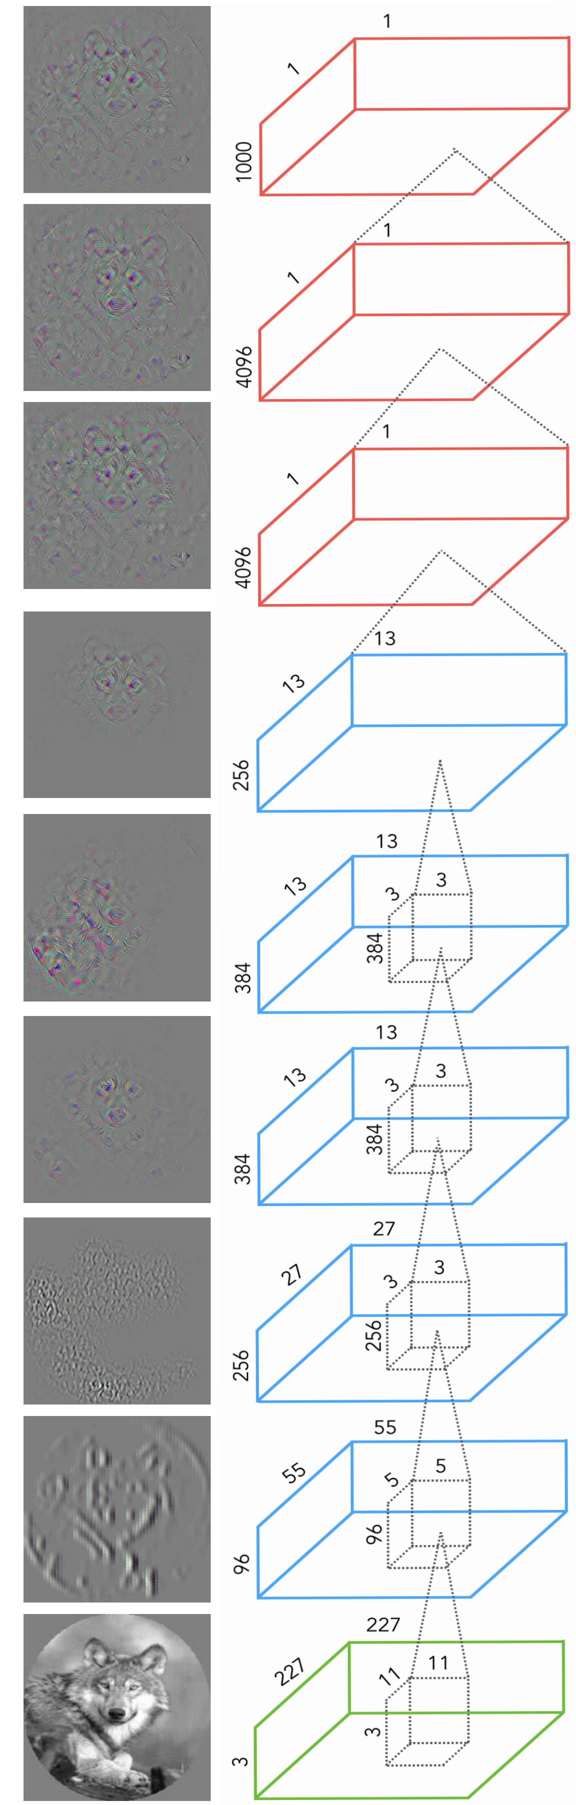

& multiple papers since:

Layered structure of visual cortex

Lower layers process simple visual features, higher layers -- complex features

Neurons focus on certain areas of visual input

Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence https://www.nature.com/articles/srep27755
Activations of deep convolutional neural networks are aligned with gamma band activity of human visual cortex https://www.nature.com/articles/s42003-018-0110-y
Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream https://www.jneurosci.org/content/35/27/10005

Model of the world: grid cells

The Nobel Prize in Physiology or Medicine 2014

"for their discoveries of cells that constitute a positioning system in the brain"

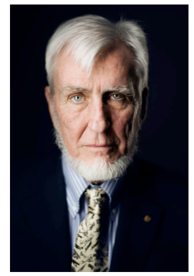
O'Keefe

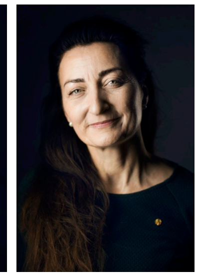
Moser

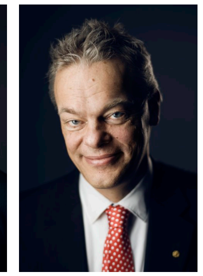
Moser

Model of the world: grid cells

The Nobel Prize in Physiology or Medicine 2014

"for their discoveries of cells that constitute a positioning system in the brain"

O'Keefe Moser Moser

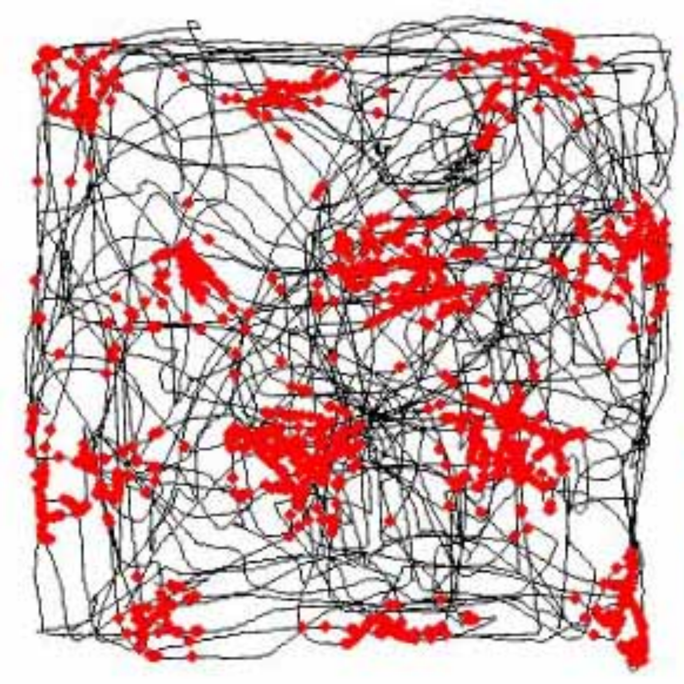

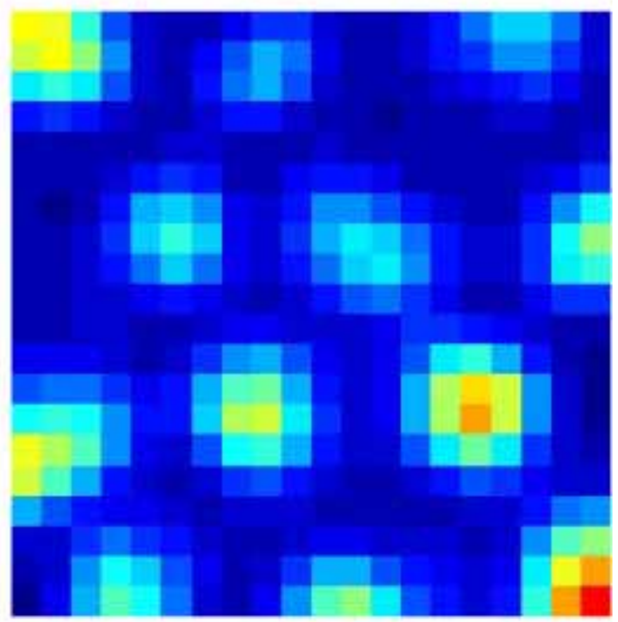

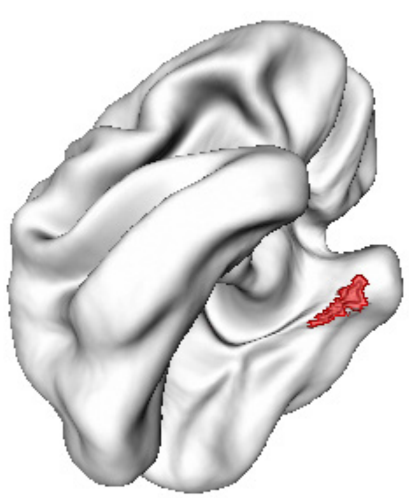

Entorhinal cortex

https://deepmind.com/blog/article/grid-cells
https://www.nature.com/articles/s41586-018-0102-6
http://www.scholarpedia.org/article/Grid_cells

Model of the world: grid cells

The Nobel Prize in Physiology or Medicine 2014

"for their discoveries of cells that constitute a positioning system in the brain"

O'Keefe Moser Moser

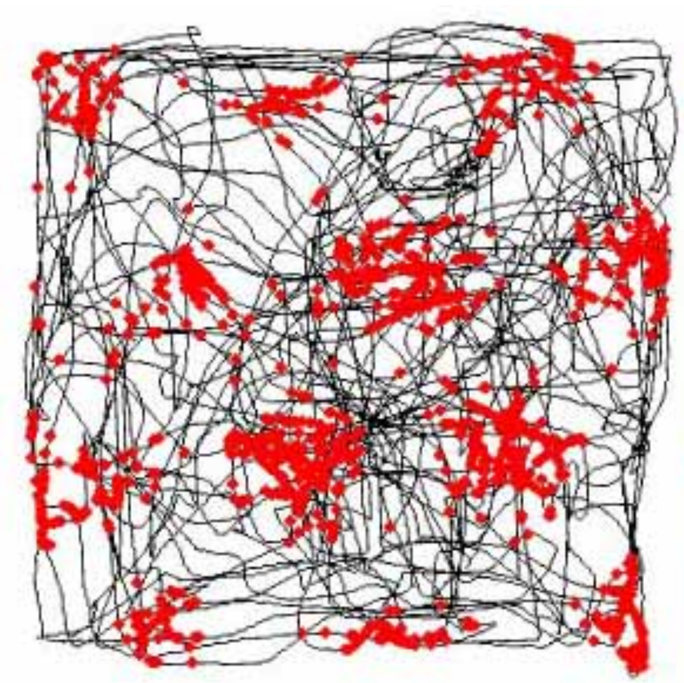

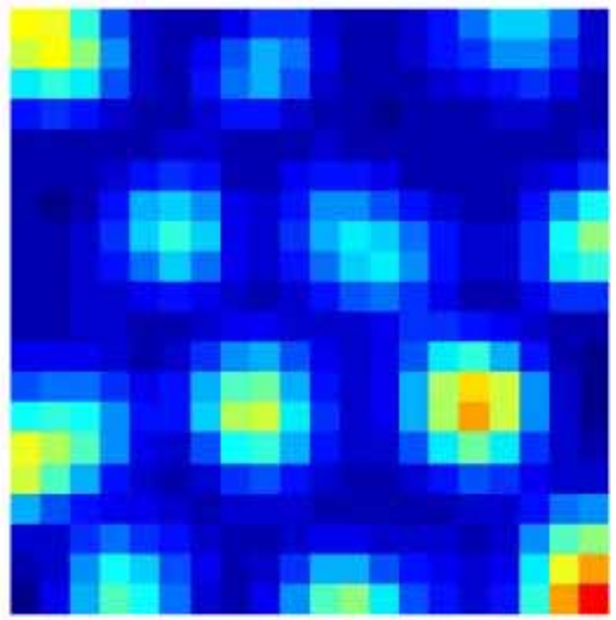

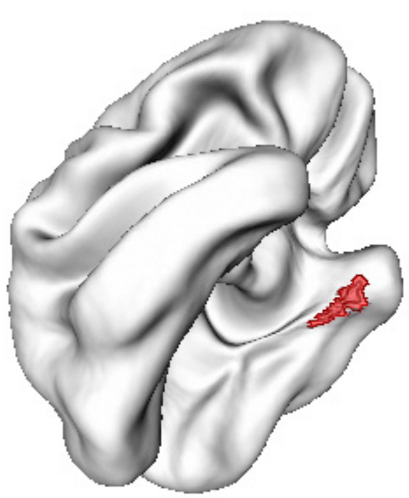

Entorhinal cortex

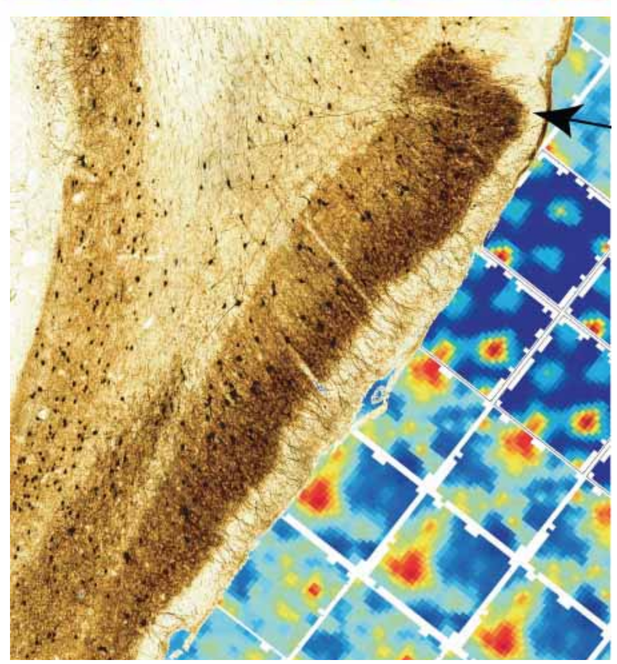

https://deepmind.com/blog/article/grid-cells

https://www.nature.com/articles/s41586-018-0102-6

http://www.scholarpedia.org/article/Grid_cells

Model of the world: grid cells

The Nobel Prize in Physiology or Medicine 2014

"for their discoveries of cells that constitute a positioning system in the brain"

O'Keefe Moser Moser

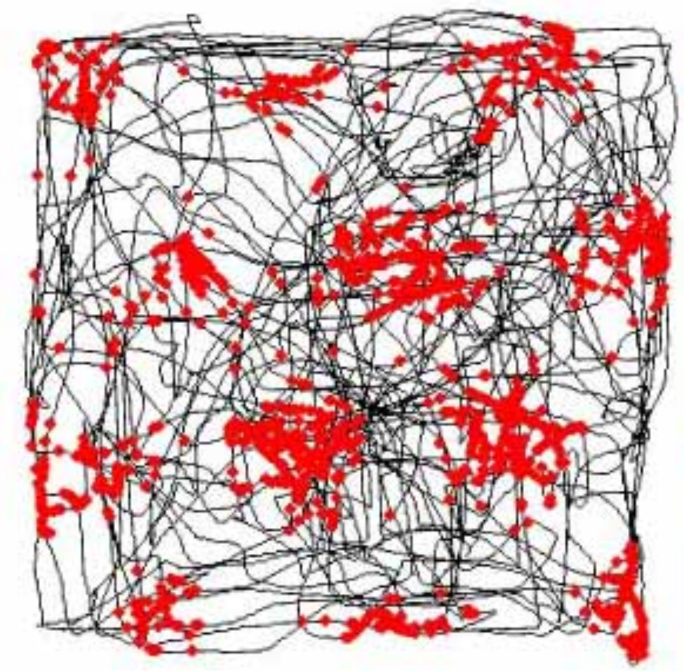

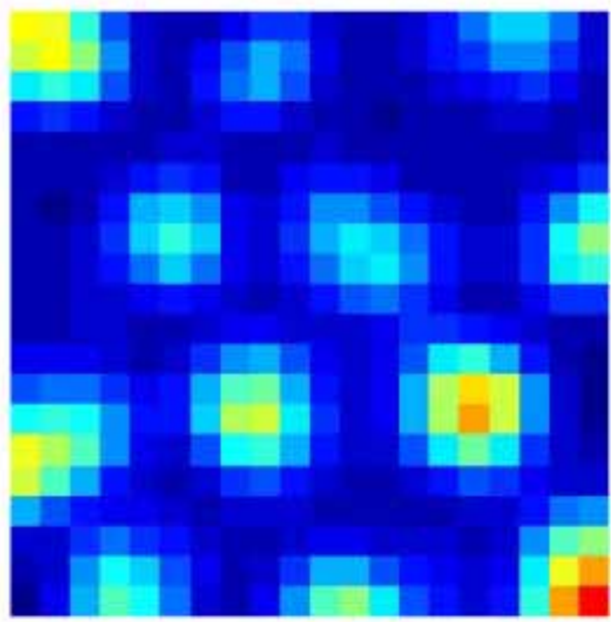

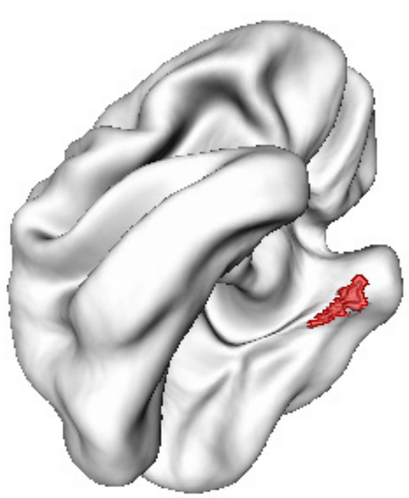

Entorhinal cortex

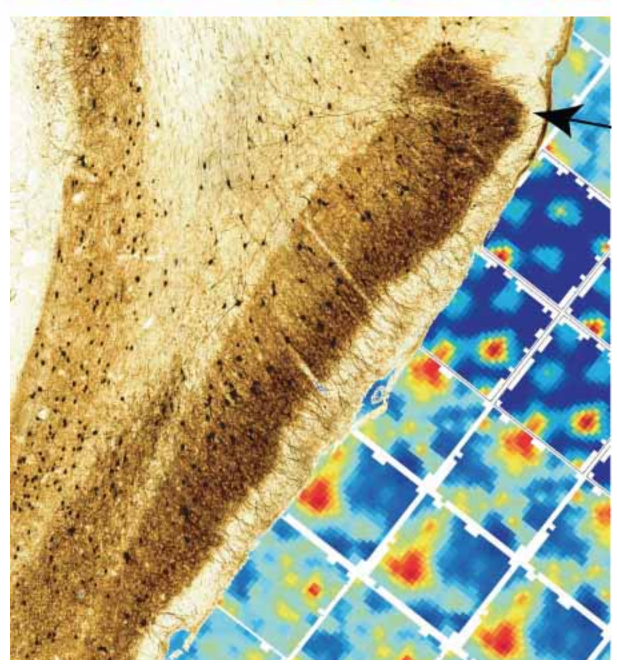

Google DeepMind

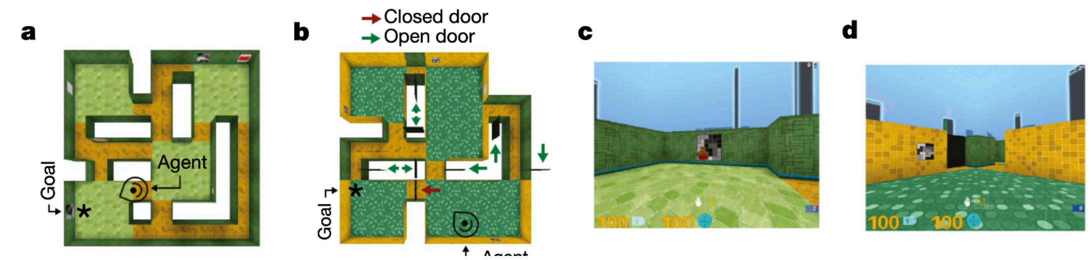

https://deepmind.com/blog/article/grid-cells

https://www.nature.com/articles/s41586-018-0102-6

http://www.scholarpedia.org/article/Grid_cells

Model of the world: grid cells

The Nobel Prize in Physiology or Medicine 2014

"for their discoveries of cells that constitute a positioning system in the brain"

O'Keefe Moser Moser

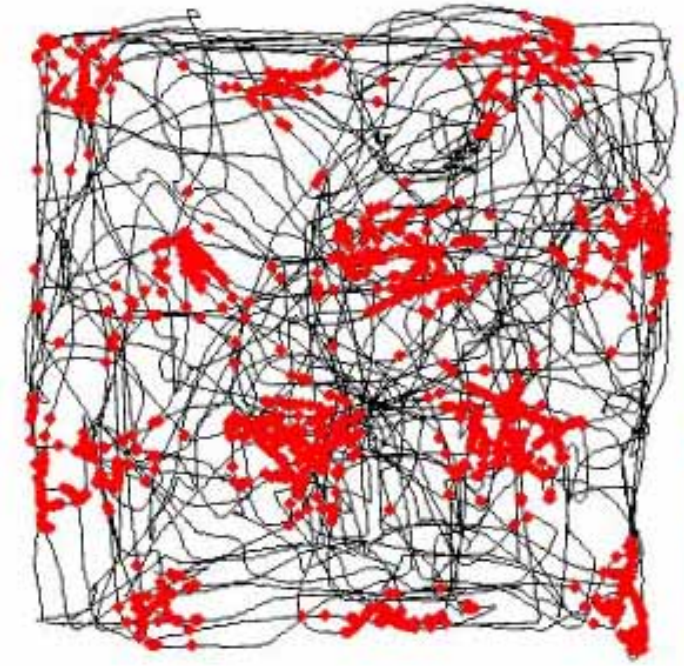

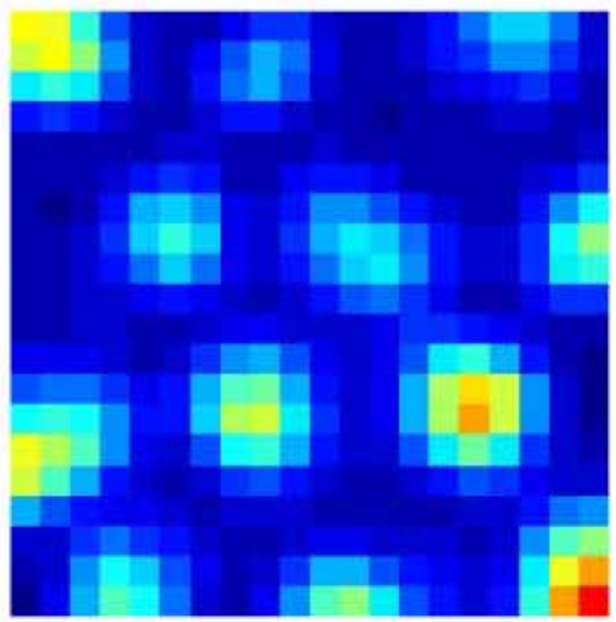

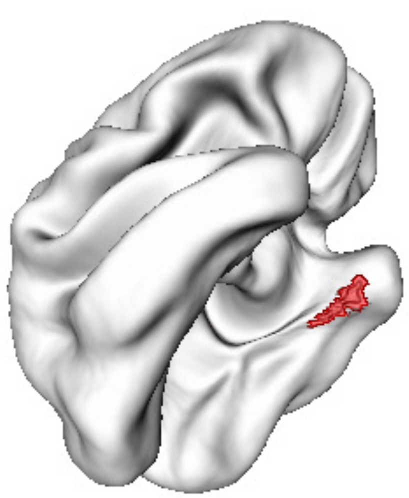

Entorhinal cortex

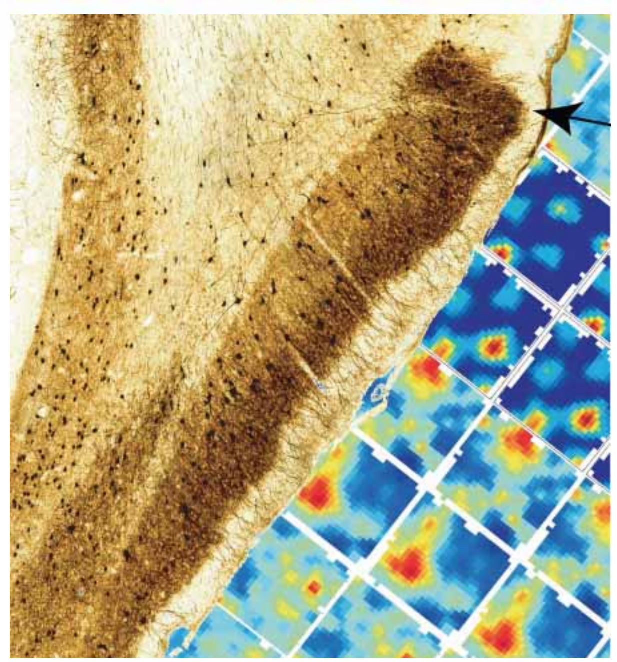

Google DeepMind

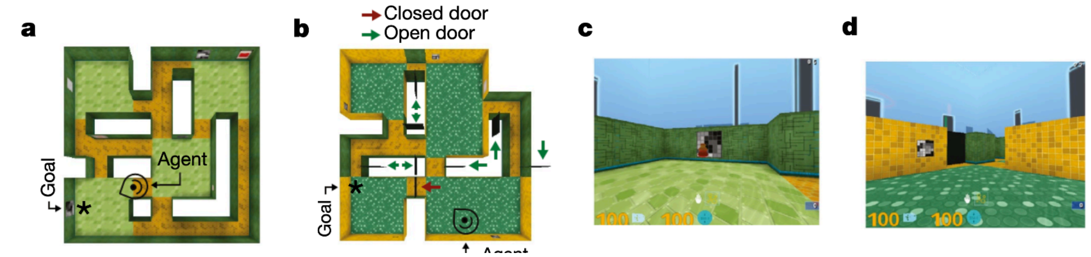

We found that grid-like representations (hereafter grid units) spontaneously emerged within the network - providing a striking convergence with the neural activity patterns observed in foraging mammals, and consistent with the notion that grid cells provide an efficient code for space.

Artificial (Agent)

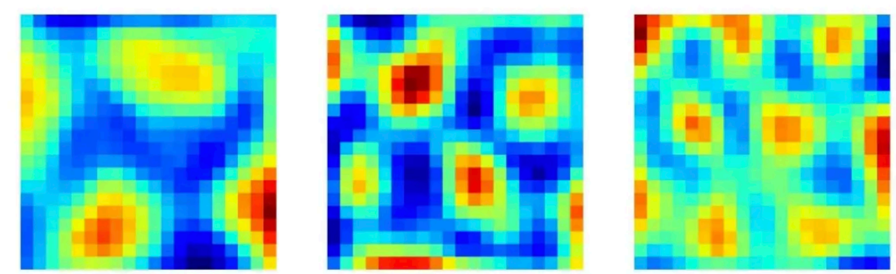

Biological (Rat)

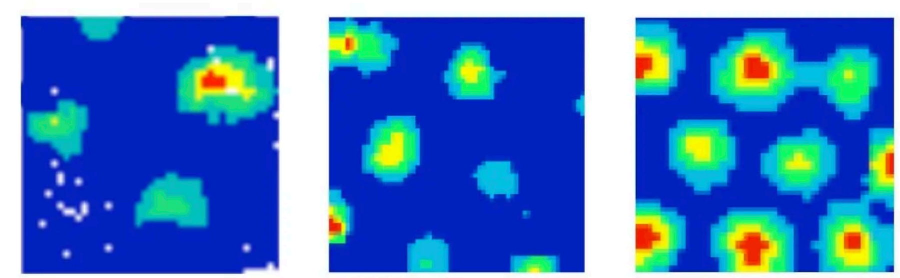

Our experiments with artificial agents yielded grid-like representations ("grid units") that were strikingly similar to biological grid cells in foraging mammals.

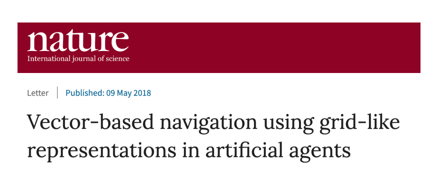

nature

Letter | Published: 09 May 2018

Vector-based navigation using grid-like representations in artificial agents

Andrea Banino, Caswell Barry, [...] Dharshan Kumaran

https://deepmind.com/blog/article/grid-cells

http://www.scholarpedia.org/article/Grid_cells

https://www.nature.com/articles/s41586-018-0102-6

Model of the world: successor representations

Model-free RL

$$Q^{\pi}(s, a) = \mathbb{E}_{\pi}[r_{t+1} + \gamma \, Q^{\pi}(s_{t+1}, a_{t+1}) | s_t = s, a_t = a]$$

$$a(s) = \arg\max_{a} Q_{\theta}(s, a)$$

+ fast
– inflexible to change

Model of the world: successor representations

Model-free RL

$$Q^\pi(s,a) = \mathbb{E}_\pi[r_{t+1} + \gamma\, Q^\pi(s_{t+1}, a_{t+1}) | s_t = s, a_t = a]$$

$$a(s) = \arg\max_a Q_\theta(s,a)$$

+ fast
– inflexible to change

Model-based RL

$$V^*(s) = \max_a \sum_{s'} \underbrace{p(s'|s,a)}_{\text{transition probs}} \left(\underbrace{r(s_t, a, s')}_{\text{reward probs}} + \gamma\, V^*(s')\right)$$

+ adaptable
– hard to learn

https://julien-vitay.net/2019/05/successor-representations/ http://gershmanlab.webfactional.com/pubs/Stachenfeld17.pdf

Model of the world: successor representations

Model-free RL

$$Q^\pi(s,a) = \mathbb{E}_\pi[r_{t+1} + \gamma\, Q^\pi(s_{t+1}, a_{t+1}) | s_t = s, a_t = a]$$

$$a(s) = \arg\max_a Q_\theta(s,a)$$

+ fast
– inflexible to change

Model-based RL

$$V^*(s) = \max_a \sum_{s'} \underbrace{p(s'|s,a)}_{\text{transition probs}} \, (\underbrace{r(s_t, a, s')}_{\text{reward probs}} + \gamma\, V^*(s'))$$

+ adaptable
– hard to learn

Successor representation for RL

$$V^\pi(s) = \sum_{s'} \underbrace{M(s,s')}_{\text{discounted future state occupancy matrix}} \, \underbrace{r(s')}_{\text{expected state reward}}$$

Why learn the full model when we only care about the chance of getting into a state

https://julien-vitay.net/2019/05/successor-representations/

http://gershmanlab.webfactional.com/pubs/Stachenfeld17.pdf

Model of the world: successor representations

Model-free RL

$$Q^\pi(s,a) = \mathbb{E}_\pi[r_{t+1} + \gamma\, Q^\pi(s_{t+1}, a_{t+1}) | s_t = s, a_t = a]$$

$$a(s) = \arg\max_a Q_\theta(s,a)$$

+ fast
– inflexible to change

Model-based RL

$$V^*(s) = \max_a \sum_{s'} \underbrace{p(s'|s,a)}_{\text{transition probs}} \, (\underbrace{r(s_t,a,s')}_{\text{reward probs}} + \gamma\, V^*(s'))$$

+ adaptable
– hard to learn

Successor representation for RL

$$V^\pi(s) = \sum_{s'} \underbrace{M(s,s')}_{\text{discounted future state occupancy matrix}} \, \underbrace{r(s')}_{\text{expected state reward}}$$

Why learn the full model when we only
care about the chance of getting into a state

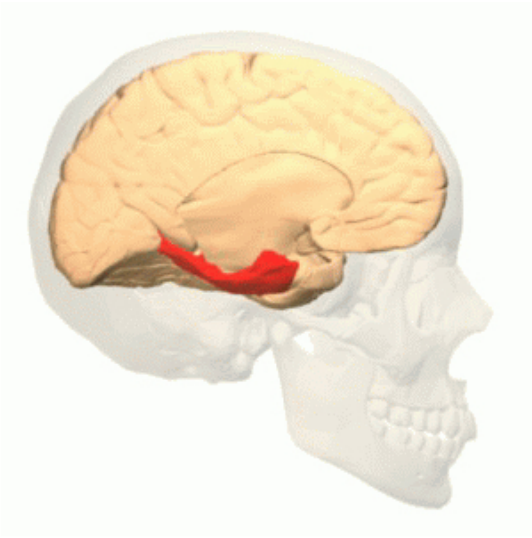

Hippocampus

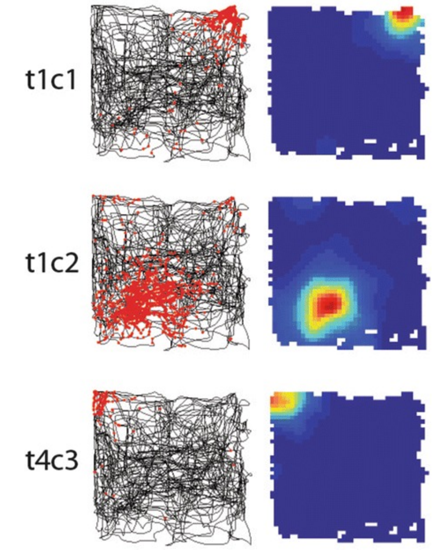

Place cells
Gaussian distance to the location?

https://julien-vitay.net/2019/05/successor-representations/

http://gershmanlab.webfactional.com/pubs/Stachenfeld17.pdf

Model of the world: successor representations

Model-free RL

$$Q^{\pi}(s,a) = \mathbb{E}_{\pi}[r_{t+1} + \gamma\, Q^{\pi}(s_{t+1}, a_{t+1}) | s_t = s, a_t = a]$$

$$a(s) = \arg\max_a Q_\theta(s,a)$$

+ fast
– inflexible to change

Model-based RL

$$V^*(s) = \max_a \sum_{s'} \underbrace{p(s'|s,a)}_{\text{transition probs}} \, (\underbrace{r(s_t,a,s')}_{\text{reward probs}} + \gamma\, V^*(s'))$$

+ adaptable
– hard to learn

Successor representation for RL

$$V^{\pi}(s) = \sum_{s'} \underbrace{M(s,s')}_{\text{discounted future state occupancy matrix}} \, \underbrace{r(s')}_{\text{expected state reward}}$$

Why learn the full model when we only care about the chance of getting into a state

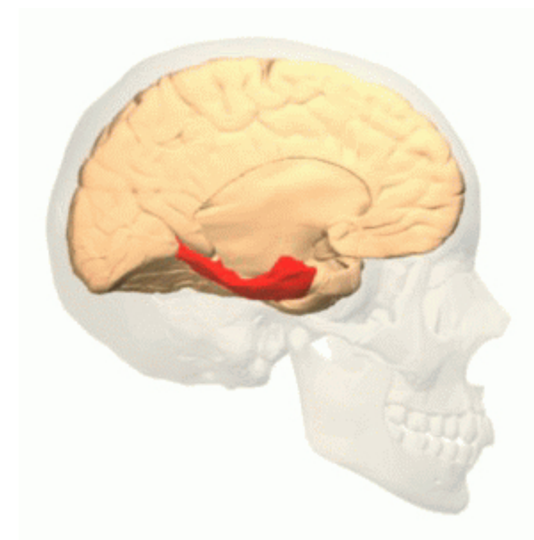

Hippocampus

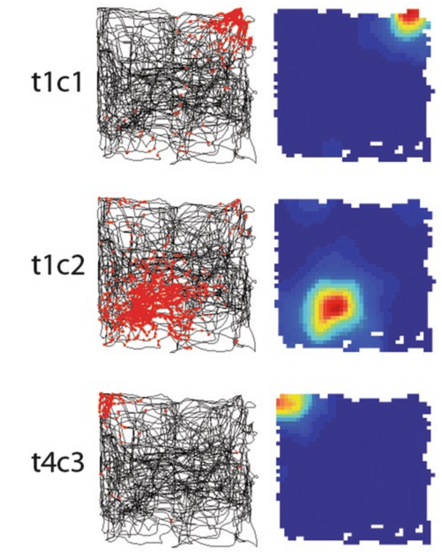

Place cells
Gaussian distance to the location?

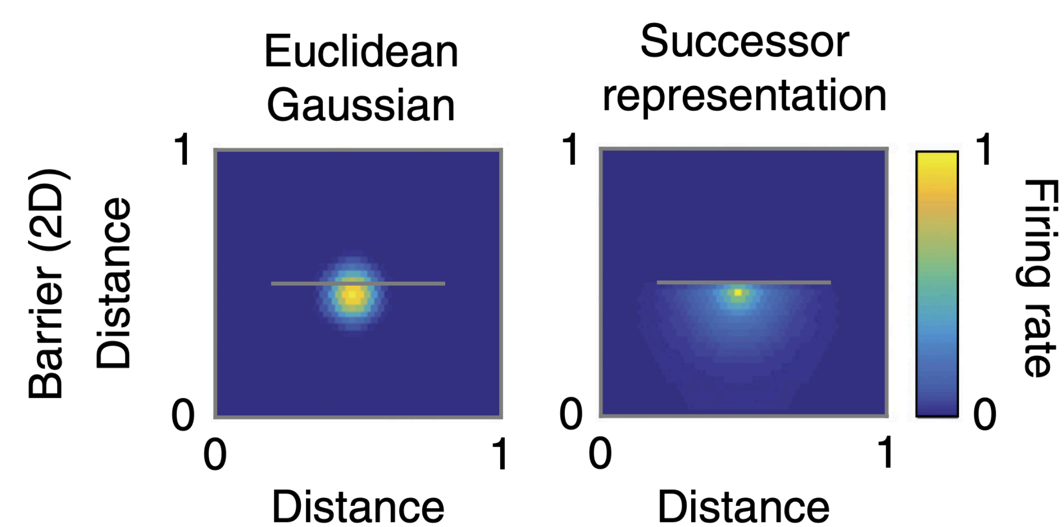

Adding constraints to the environment modifies the shape of the place fields.

SR hypothesis is in agreement with the observation that rewarded locations are represented by a higher number of place cells.

https://julien-vitay.net/2019/05/successor-representations/

http://gershmanlab.webfactional.com/pubs/Stachenfeld17.pdf

…

Predictive coding

&

Generative Adversarial Networks

Predictive coding

Successor representations

&

MERLIN Architecture

Visual attention

&

Attention in CNNs

Consolidation of experience

Novel vs. routine circuits

&

Dual DQN network

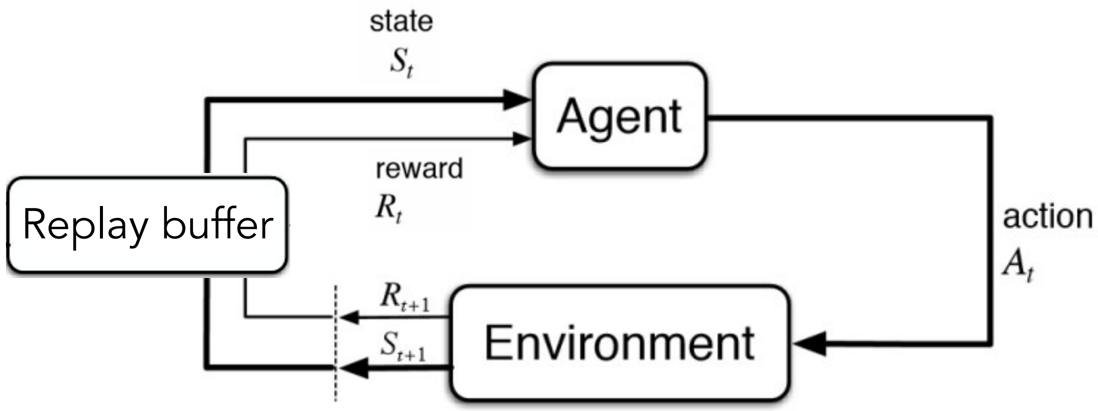

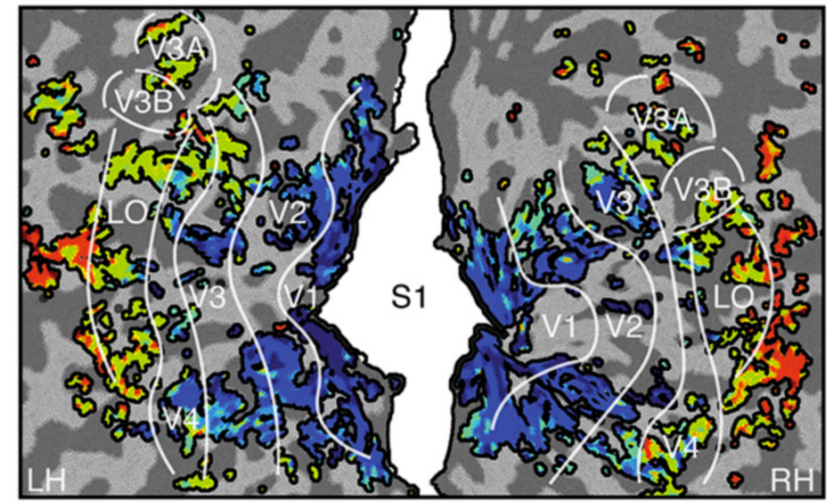

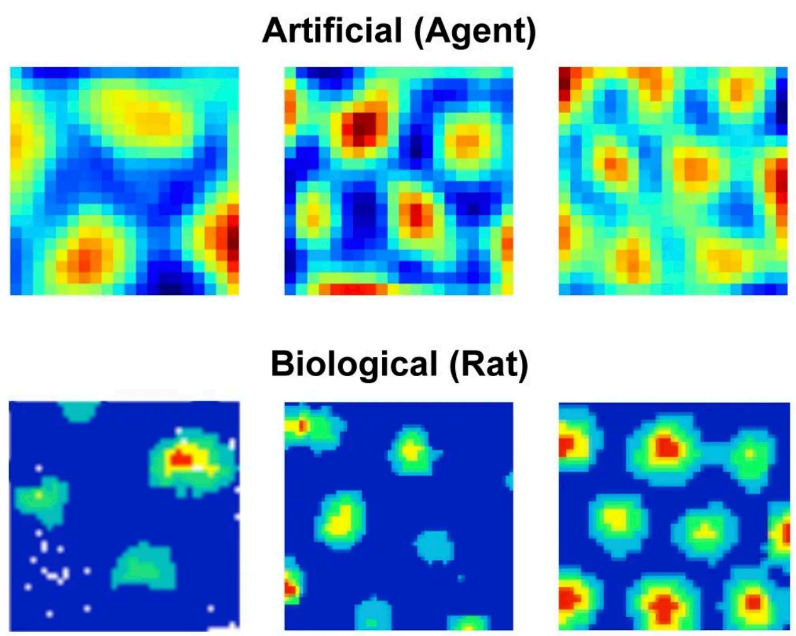

$$V^{\pi}(s) = \sum_{s'} \underbrace{M(s, s')}_{\substack{\text{discounted future} \\ \text{state occupancy} \\ \text{matrix}}} \underbrace{r(s')}_{\substack{\text{expected} \\ \text{state} \\ \text{reward}}}$$

Curious similarities on the level
of algorithm and representation

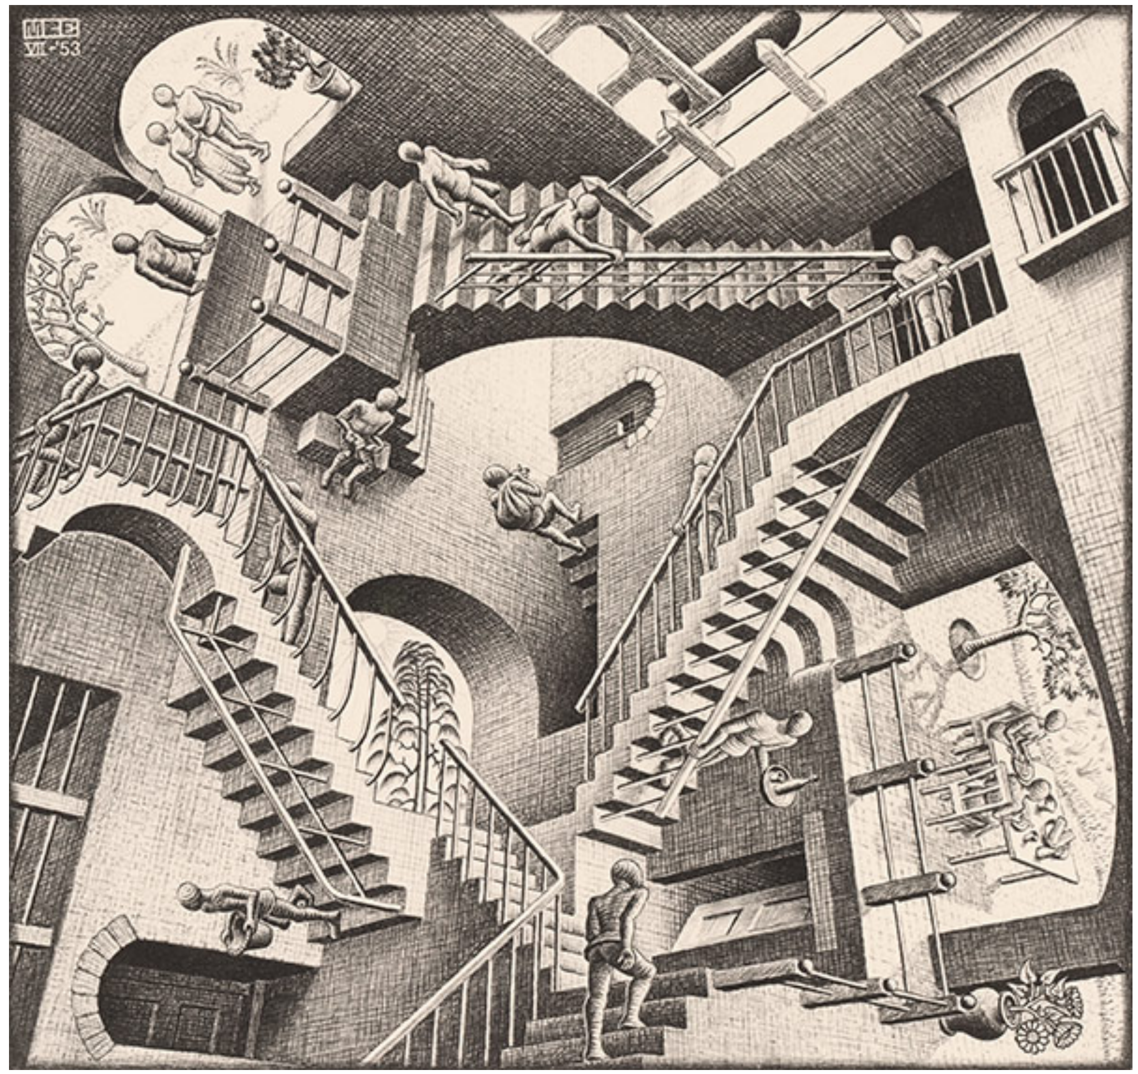

Goal of computation

Same goal, different strengths... merge!

NeuroChip

Fetz Lab
University of Washington

Neuralink

Elon Musk
San Franciso, CA

http://www.csne-erc.org/sites/default/files/Neurochip.pdf https://depts.washington.edu/fetzweb/closed-loop-bci.html https://neuralink.com/

Same goal, different strengths... merge!

NeuroChip

Fetz Lab
University of Washington

Neuralink

Elon Musk
San Francisco, CA

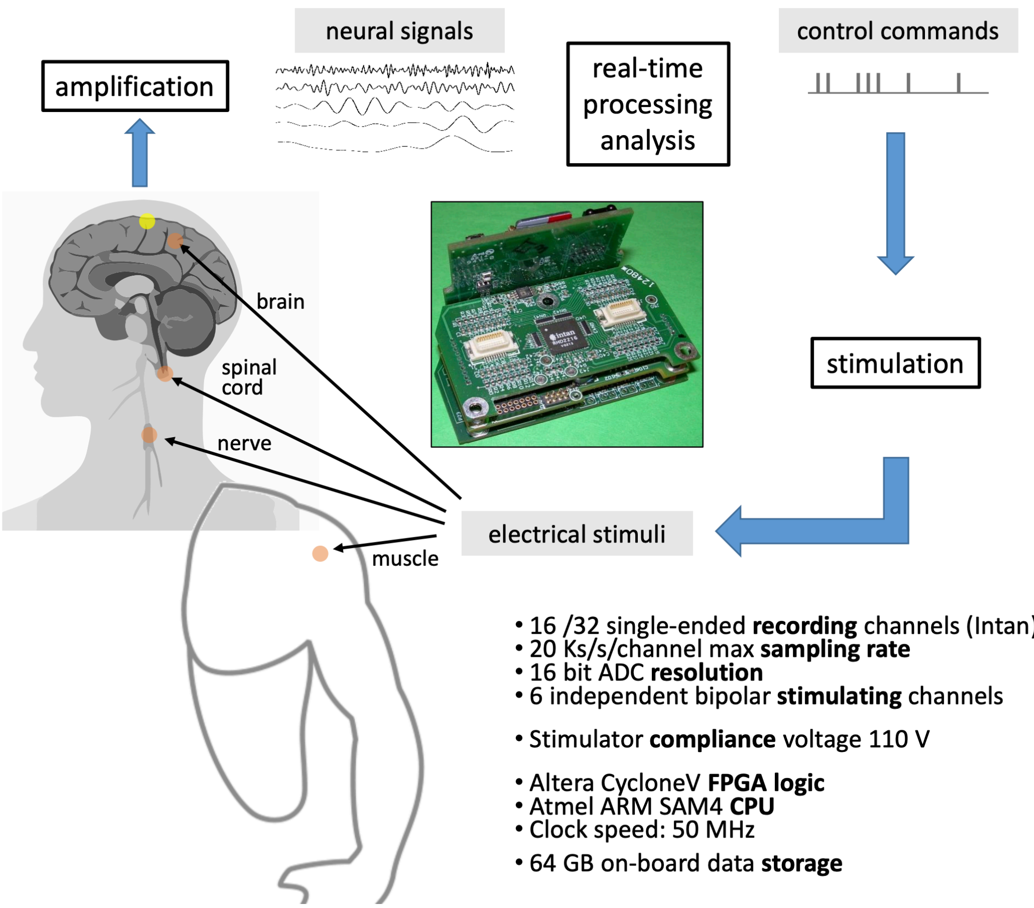

- 16 /32 single-ended **recording** channels (Intan)
- 20 Ks/s/channel max **sampling rate**
- 16 bit ADC **resolution**
- 6 independent bipolar **stimulating** channels
- Stimulator **compliance** voltage 110 V
- Altera CycloneV **FPGA logic**
- Atmel ARM SAM4 **CPU**
- Clock speed: 50 MHz
- 64 GB on-board data **storage**

nature
International journal of science

Article | Published: 22 October 2006

Long-term motor cortex plasticity induced by an electronic neural implant

http://www.csne-erc.org/sites/default/files/Neurochip.pdf https://depts.washington.edu/fetzweb/closed-loop-bci.html https://neuralink.com/

Same goal, different strengths... merge!

NeuroChip

Fetz Lab
University of Washington

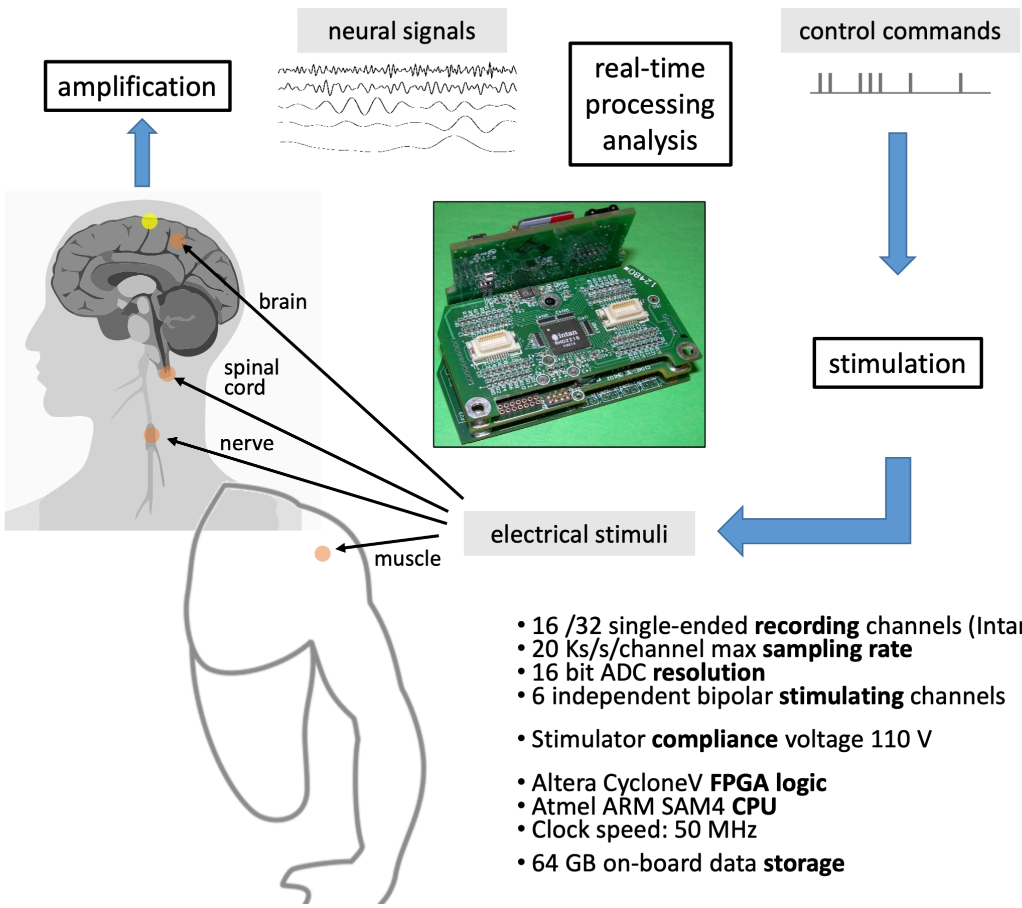

- 16 /32 single-ended **recording** channels (Intan)
- 20 Ks/s/channel max **sampling rate**
- 16 bit ADC **resolution**
- 6 independent bipolar **stimulating** channels
- Stimulator **compliance** voltage 110 V
- Altera CycloneV **FPGA logic**
- Atmel ARM SAM4 **CPU**
- Clock speed: 50 MHz
- 64 GB on-board data **storage**

nature
International journal of science

Article | Published: 22 October 2006

Long-term motor cortex plasticity induced by an electronic neural implant

Neuralink

Elon Musk
San Francisco, CA

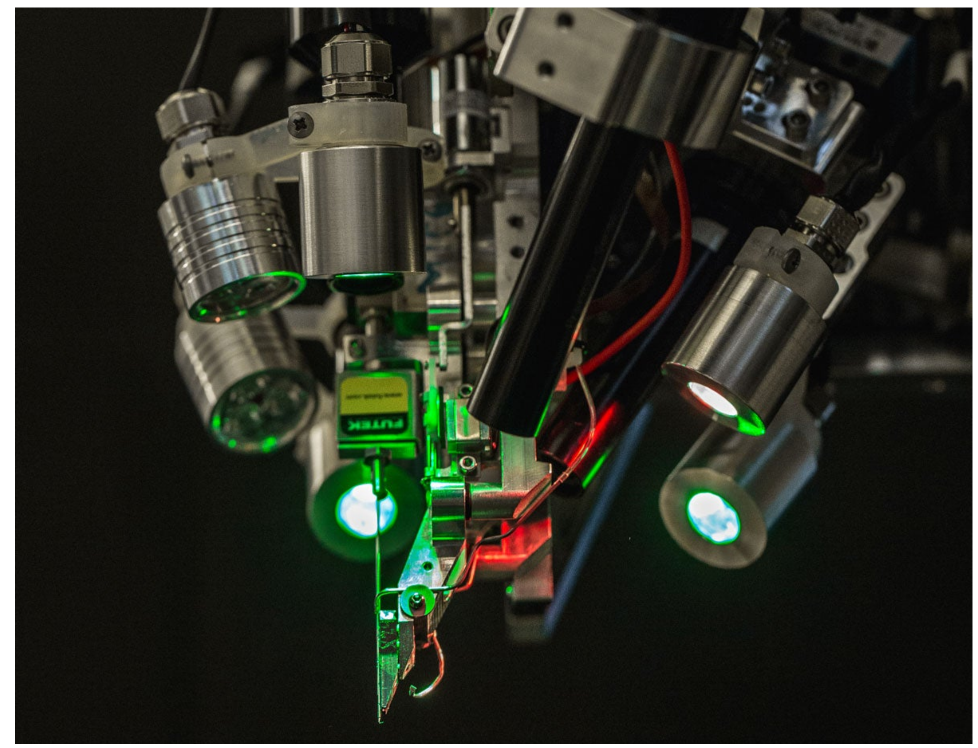

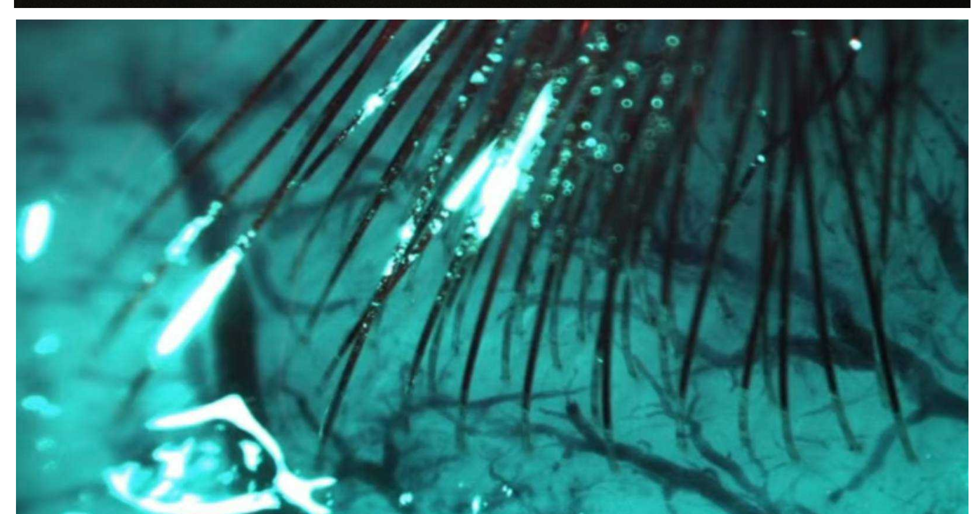

http://www.csne-erc.org/sites/default/files/Neurochip.pdf https://depts.washington.edu/fetzweb/closed-loop-bci.html https://neuralink.com/

The Brain and the Modern AI

Goal of the computation
What is the purpose of computation?

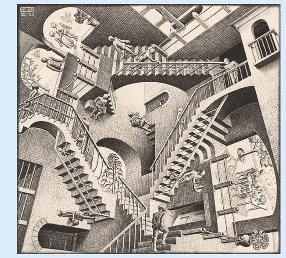

Almost the same

Algorithm and representation
What representations does the system use?
What processes are in use to manipulate representations?

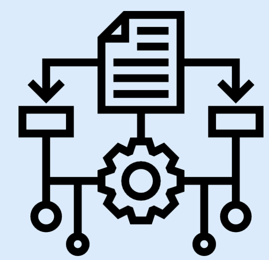

Comparable here and there

Implementation
How is the system physically realized?

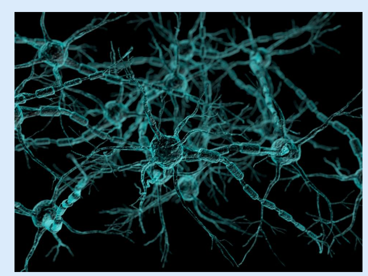

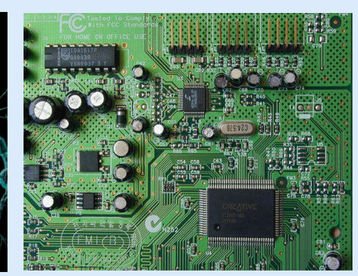

Quite different

Ilya Kuzovkin

ilya.kuzovkin@gmail.com
www.ikuz.eu

Neurotech Sydney meetup group

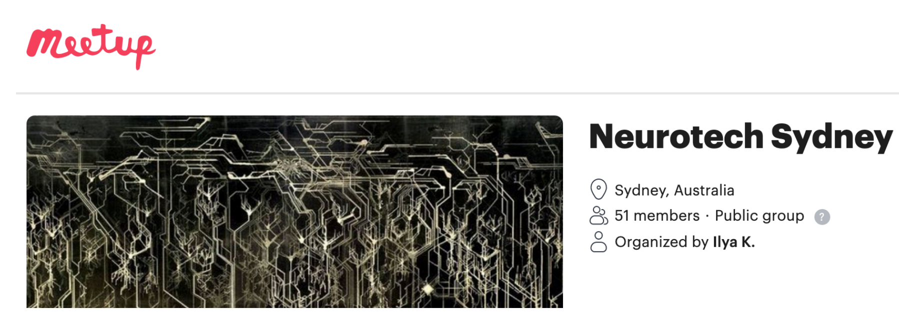

Podcast by Paul Middlebrooks