

Welcome to Neuro 140/240!

Biological and Artificial Intelligence

Professor: Gabriel Kreiman



Class survey

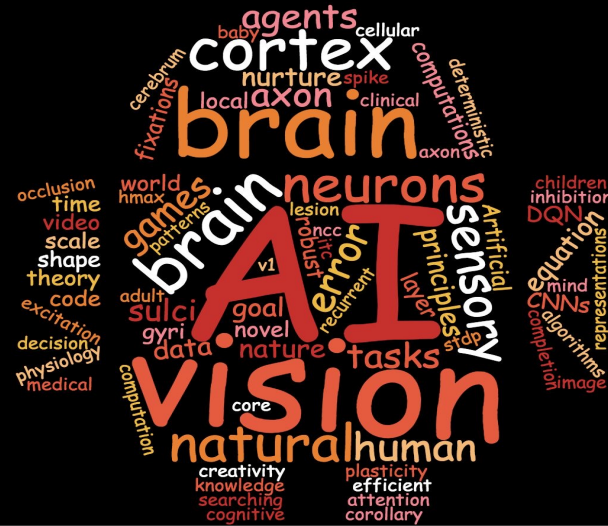


Sign-in



Welcome to Neuro 140/240!

Biological and Artificial Intelligence



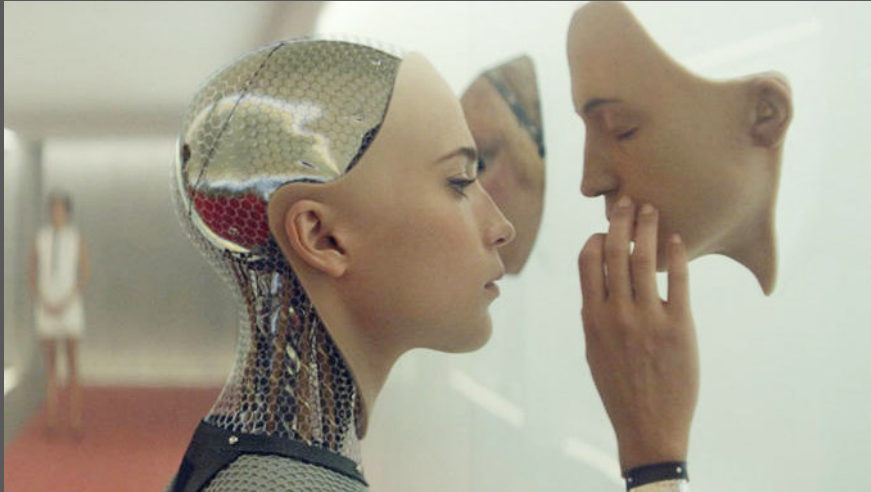
3:05-4:15 Introduction to Biological and Artificial Intelligence

Please interrupt and ask questions

4:15-4:20 Short break

4:20-5:00 Class logistics, homework, grading, policies

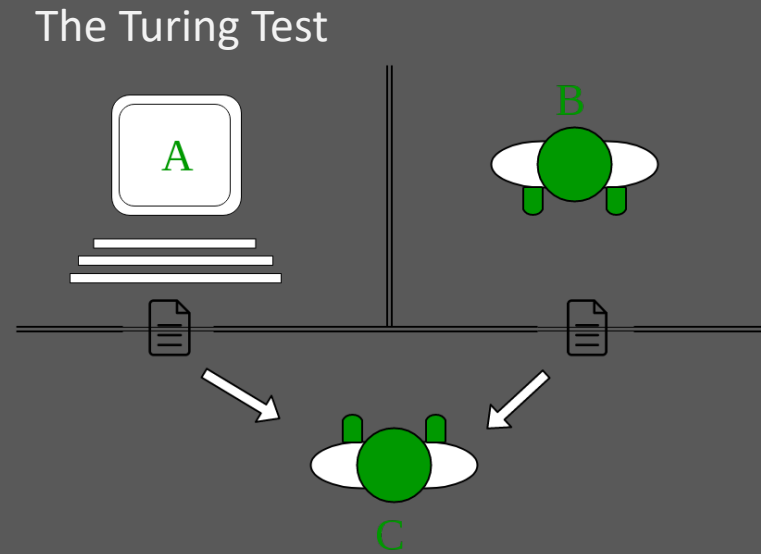
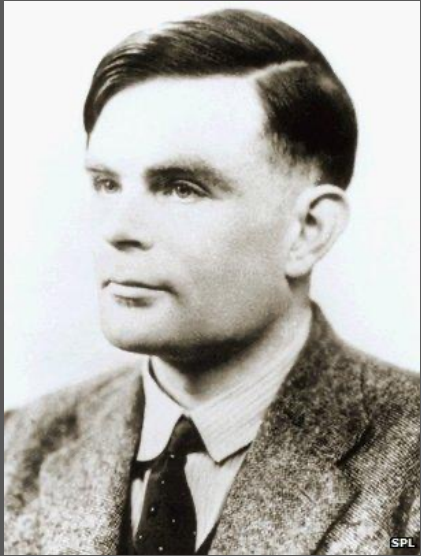
The last machine we ever need to build



Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an “intelligence explosion,” and the intelligence of man would be left far behind. **Thus the first ultraintelligent machine is the last invention that man need ever make.**

The Turing test

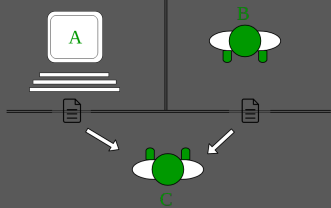
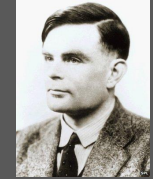
Alan Turing



The Turing test for vision

Alan Turing

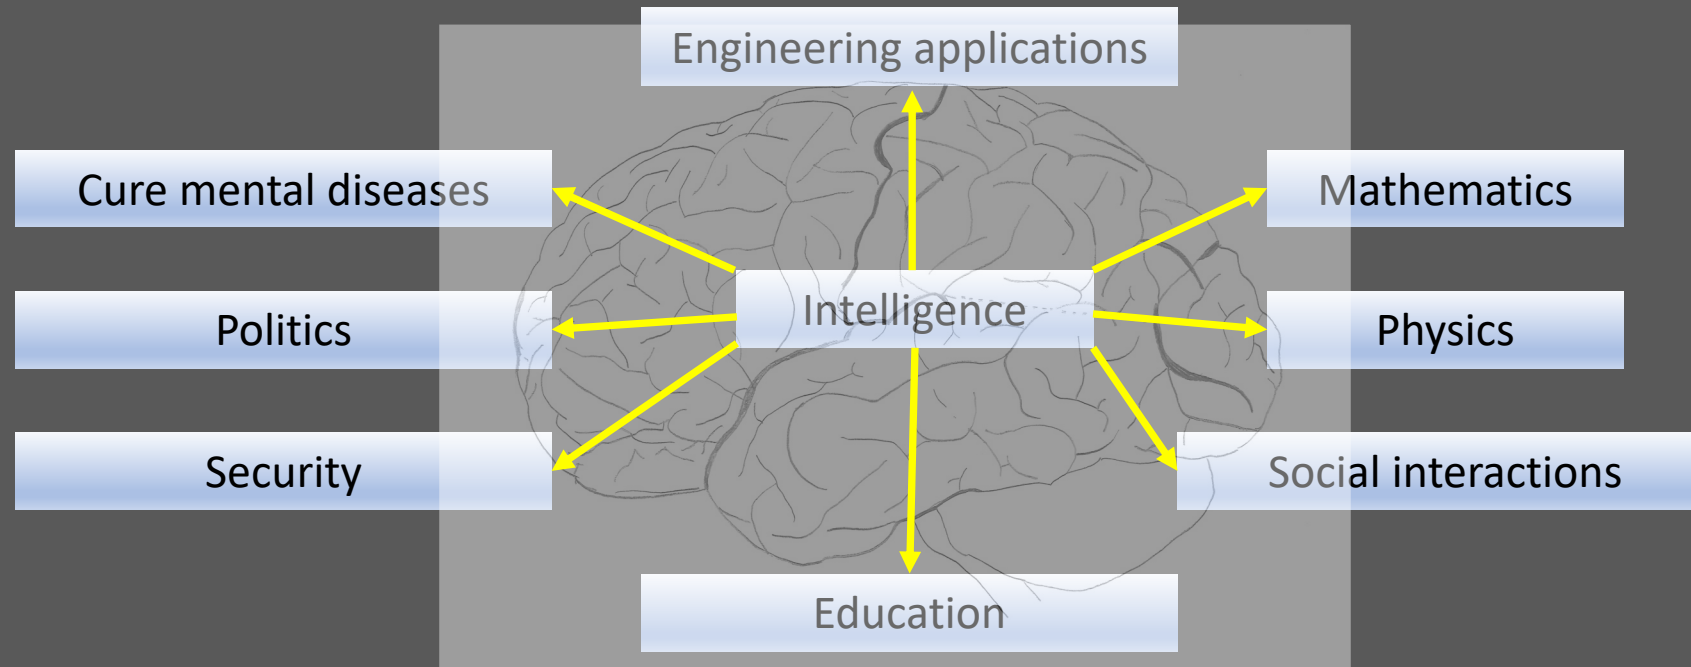
The Turing Test



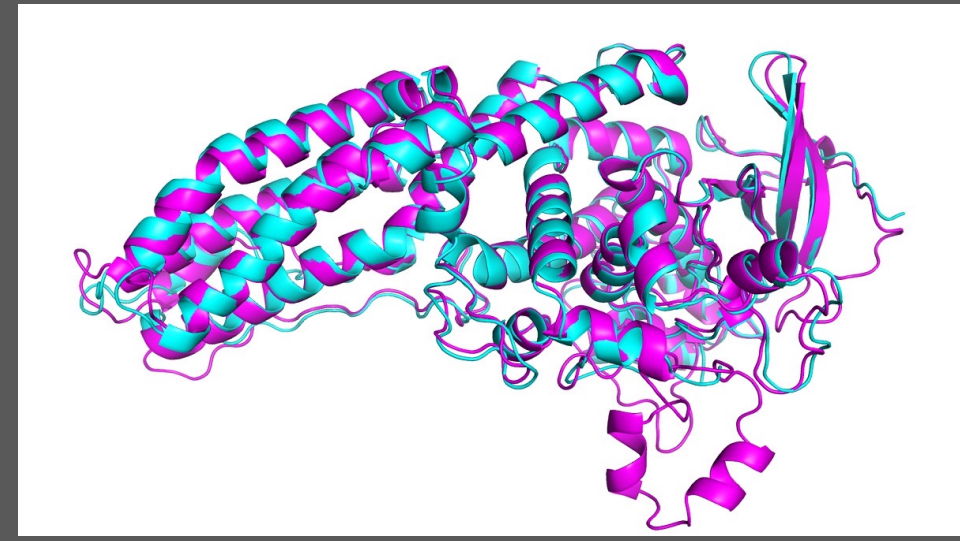
Intelligence is the greatest problem in science

If we understand the brain and we understand intelligence ... we could find ways to make us smarter and to build smart machines to help us think

Tomaso Poggio, MIT



Rapid progress in AI

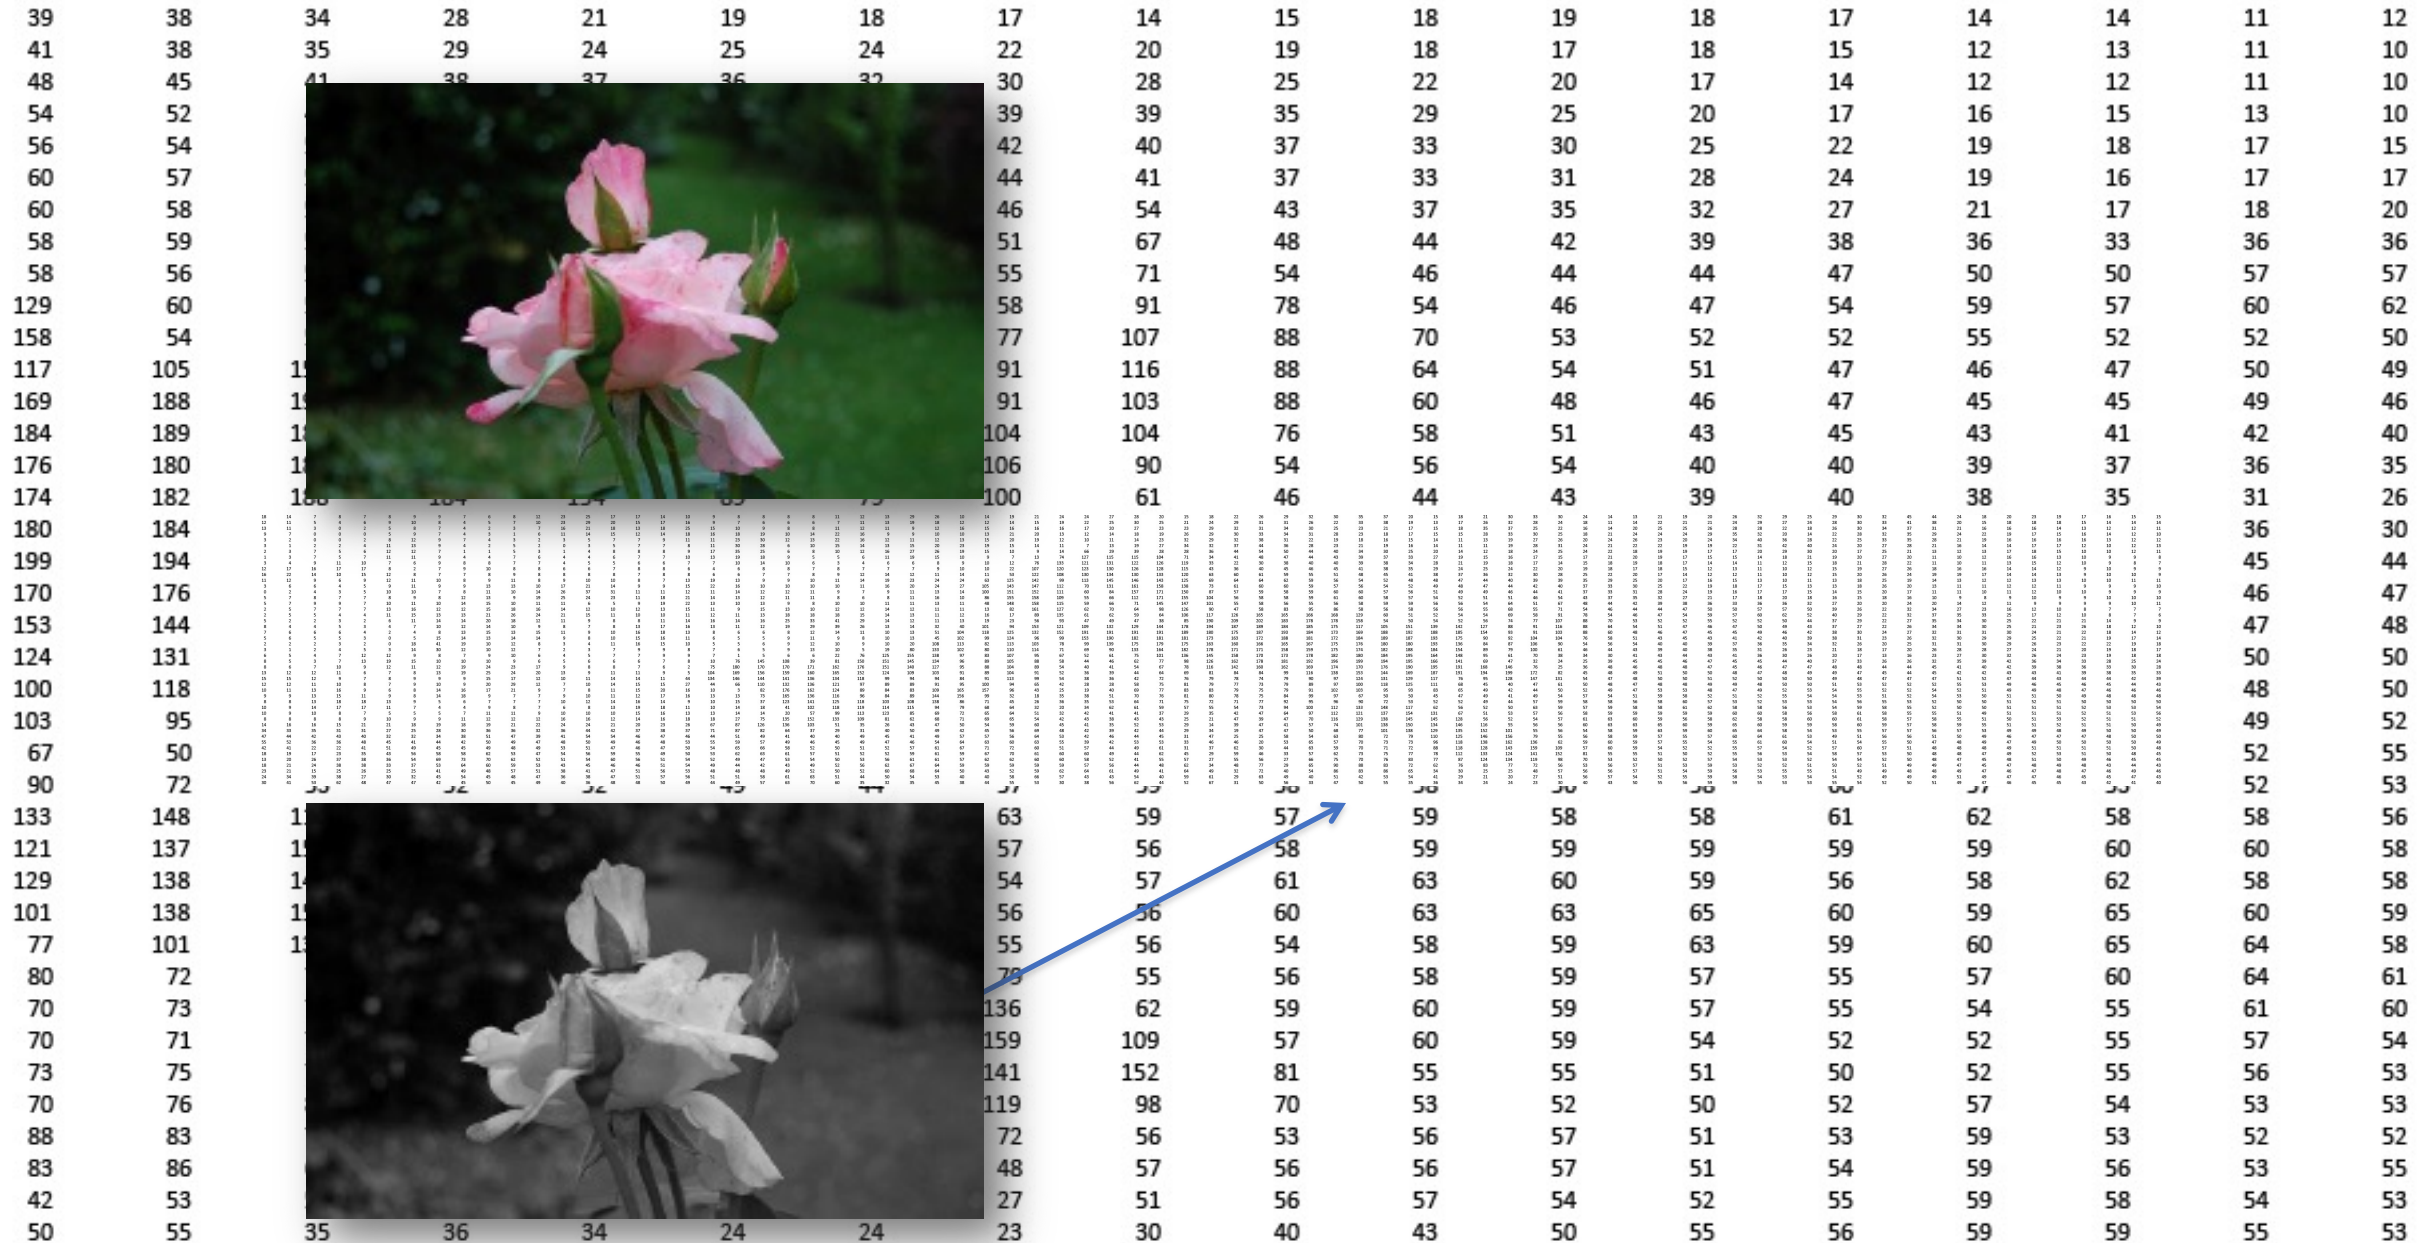


Example: object recognition

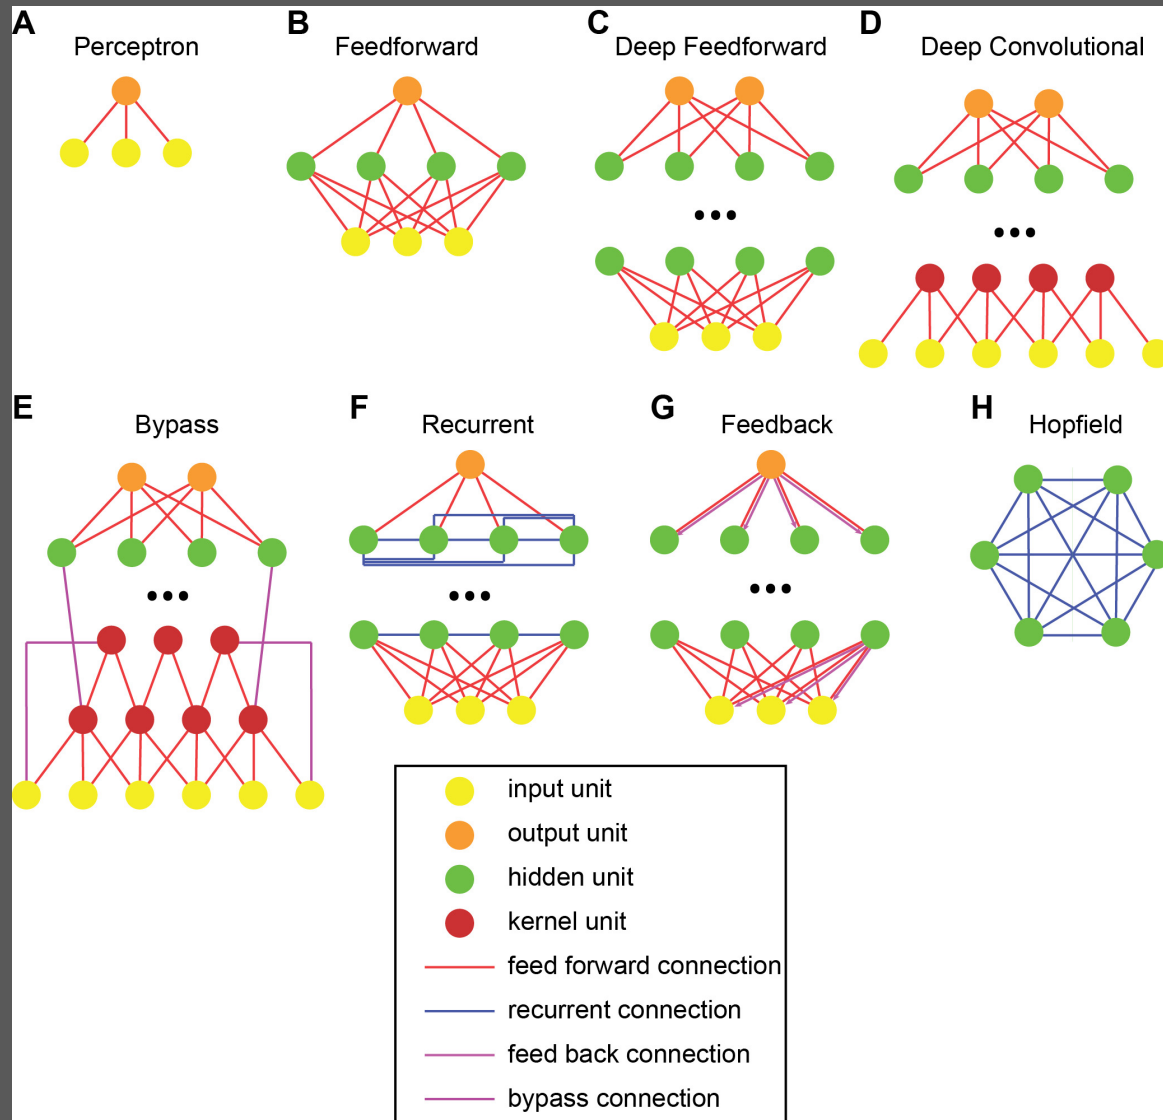
Where are the people in this image?



A flower, as seen by a computer



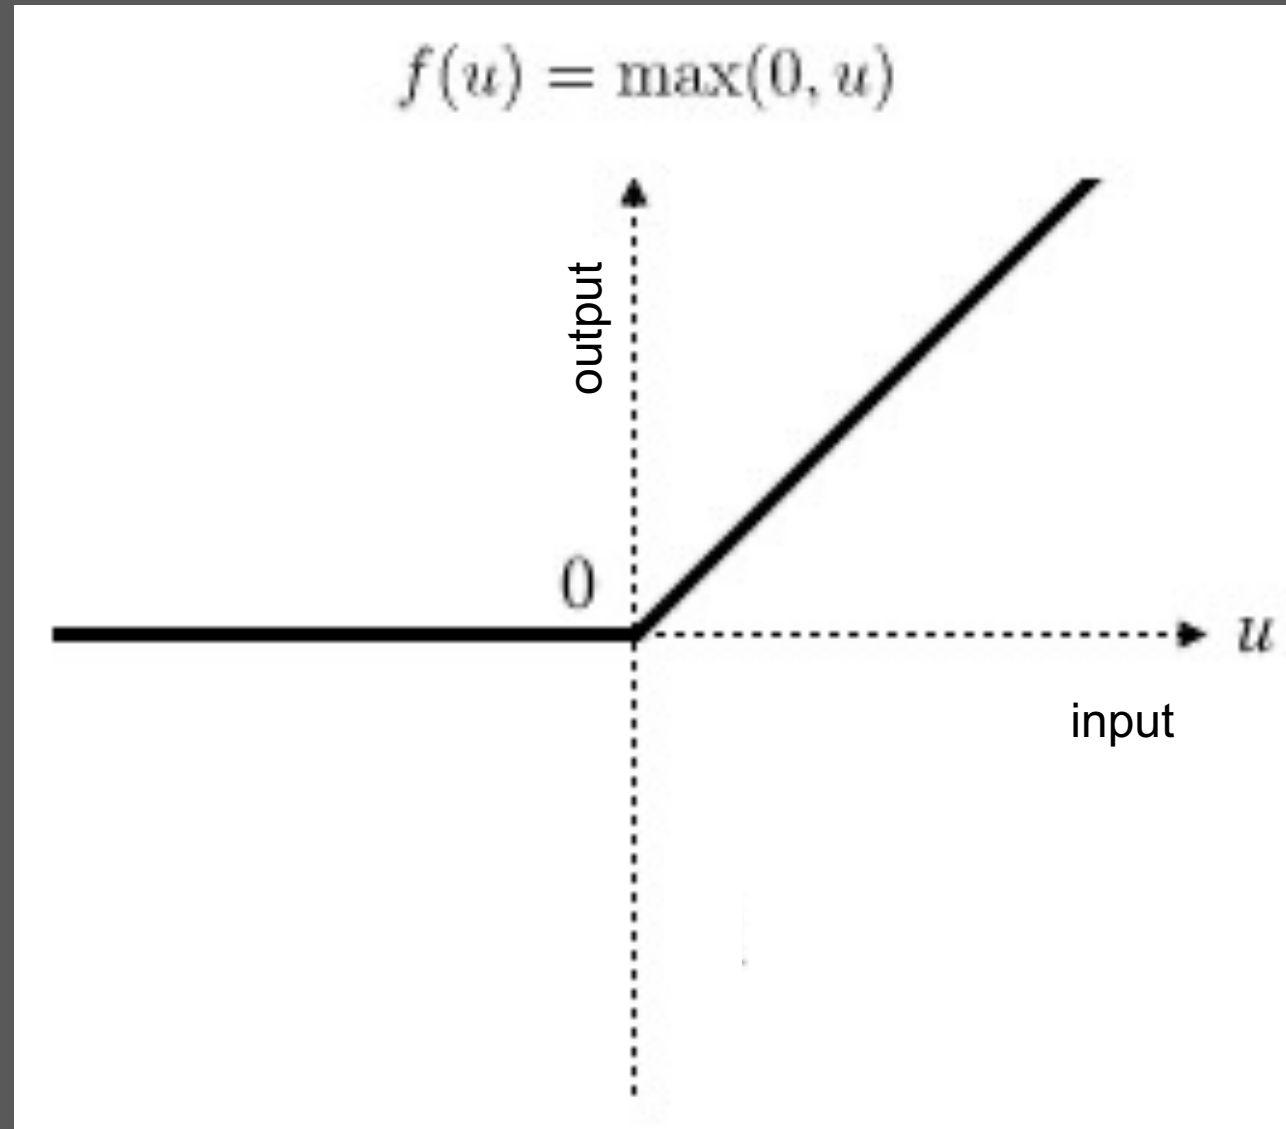
Emergent properties from simple operations



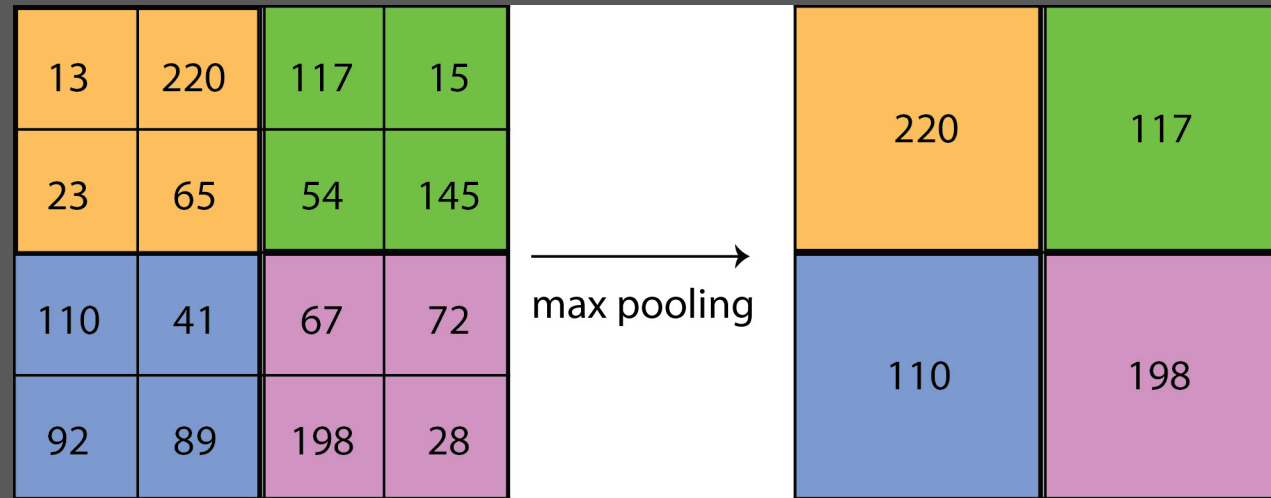
Neural Networks

1. CONVolutional layer
2. NORMalization layer
3. RELU layer
4. POOL layer
5. Weight changes
6. Dropout
7. Deep architectures

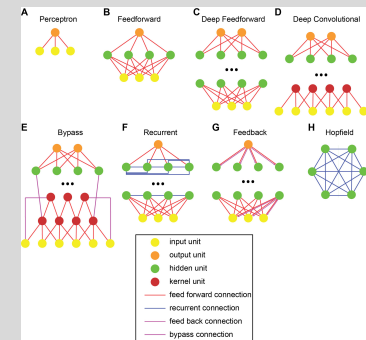
Rectifying linear unit (ReLU)



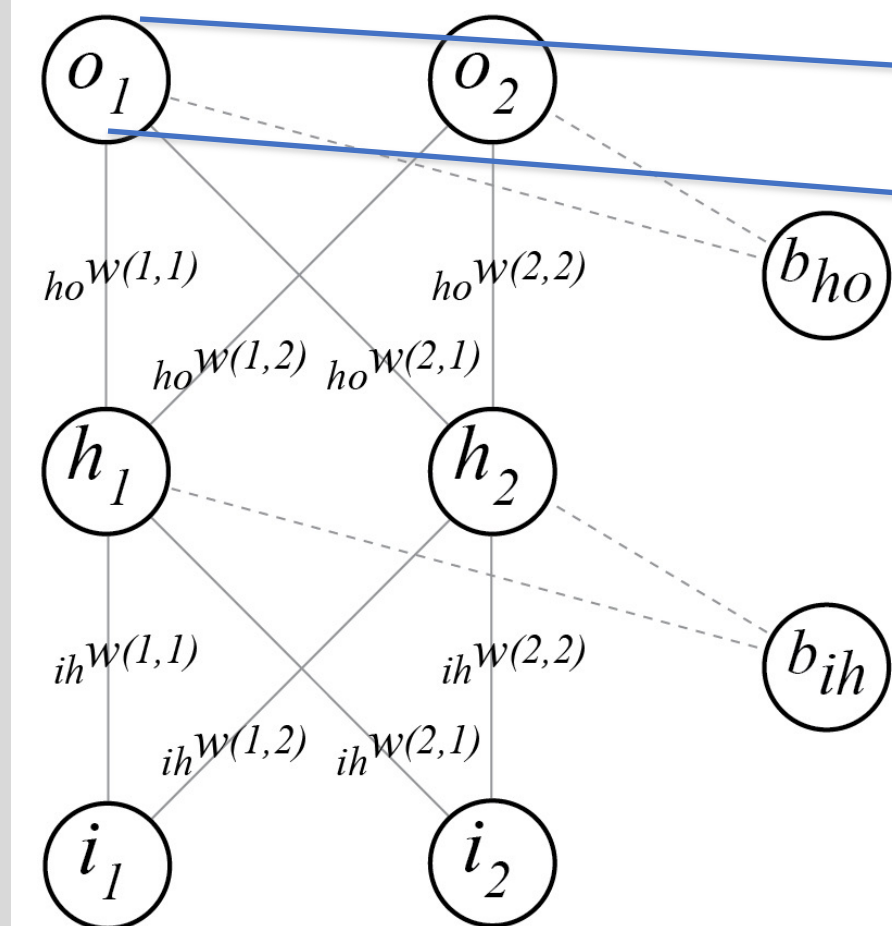
Max pooling



Back-propagating errors



$$E_{total} = E_{o_1} + E_{o_2} = 0.5[(target_{o_1} - o_1)^2 + (target_{o_2} - o_2)^2]$$



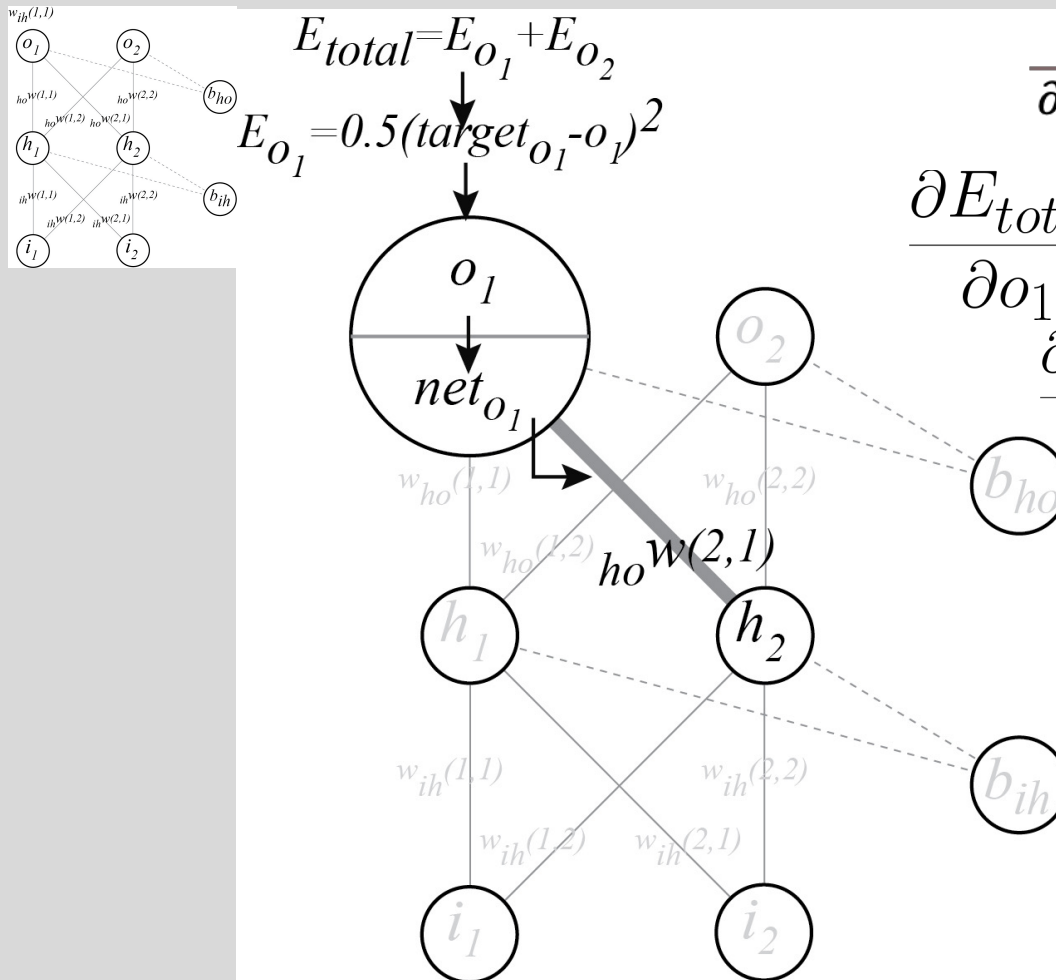
$$o_1 = \frac{1}{1 + e^{-net_{o_1}}}$$

$$net_{o_1} = h_1 * w_{ho}(1,1) + h_2 * w_{ho}(2,1) + b_{ho}$$

$$h_1 = \frac{1}{1 + e^{-net_{h_1}}}$$

$$net_{h_1} = i_1 * w_{ih}(1,1) + i_2 * w_{ih}(2,1) + b_{ih}$$

Back-propagating errors



$$\frac{\partial E_{total}}{\partial (w_{ho}^{2,1})} = \frac{\partial E_{total}}{\partial o_1} \frac{\partial o_1}{\partial net_{o_1}} \frac{\partial net_{o_1}}{\partial (w_{ho}^{2,1})}$$

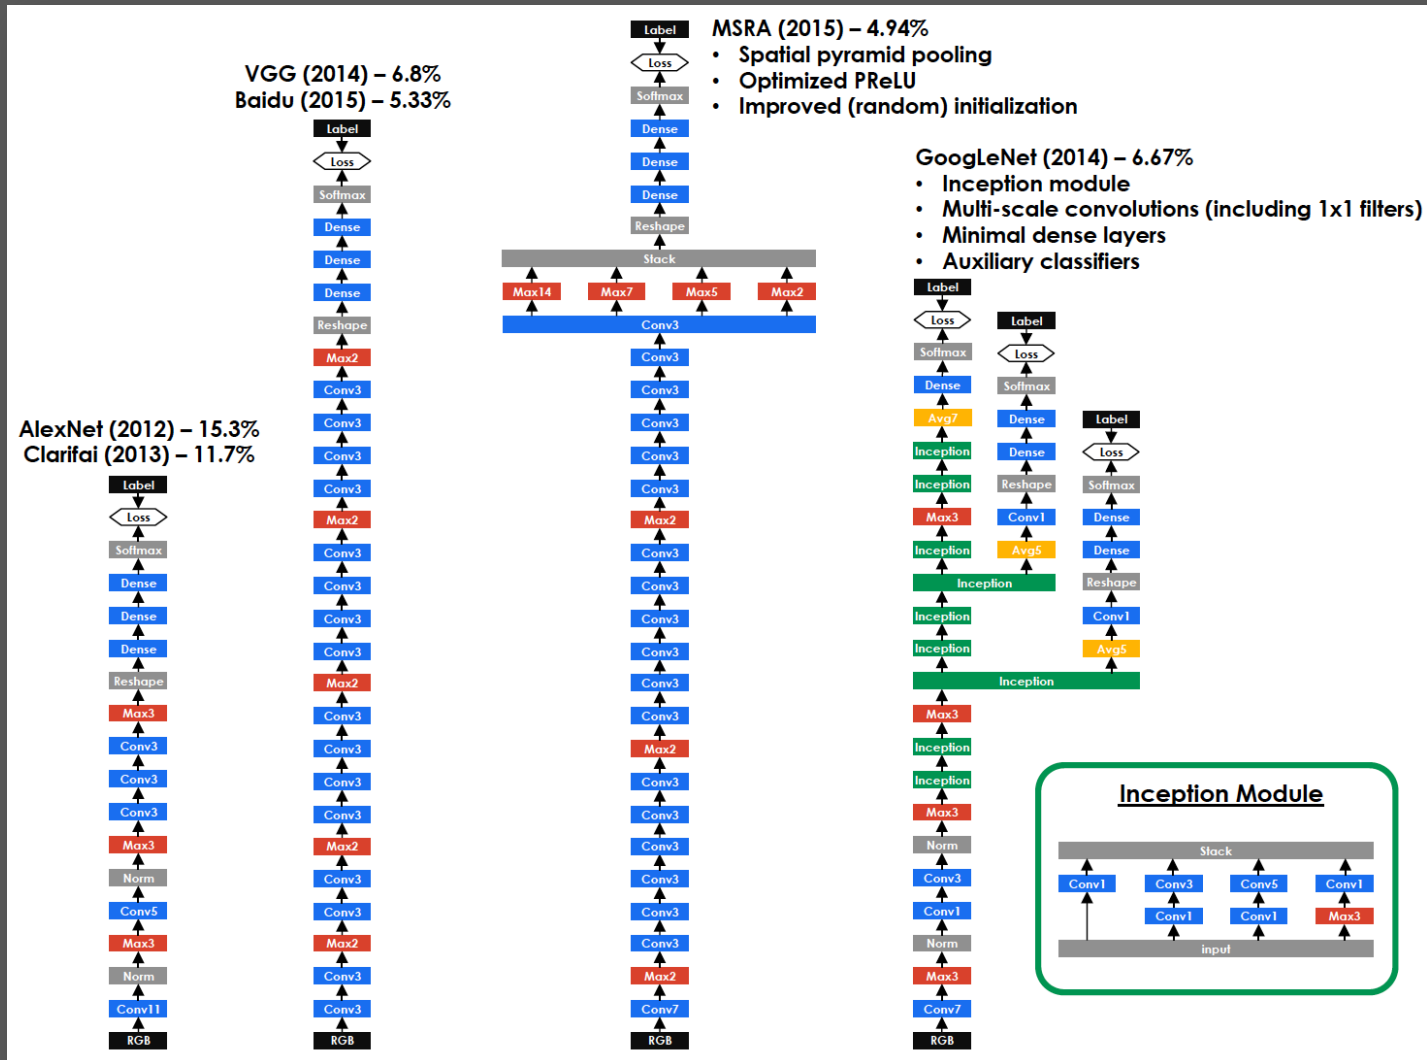
$$\frac{\partial E_{total}}{\partial o_1} = 2 * 0.5 * (o_1 - target_{o_1})$$

$$\frac{\partial E_{total}}{\partial o_1} = 2 * 0.5 * (o_1 - target_{o_1})$$

$$\frac{\partial net_{o_1}}{\partial (w_{ho}^{2,1})} = h_2$$

$$w_{ho}^{(2,1)} \rightarrow w_{ho}^{(2,1)} - \epsilon \frac{\partial E_{total}}{\partial w_{ho}^{(2,1)}}$$

Deeper and deeper



What can neural networks do (now)?

See lecture by Thomas Serre, Bill Lotter, Isaac Kohane

0. Handwritten digit recognition
1. Classification of large image datasets
2. Better at face recognition than “superrecognizers” and face forensic experts
3. Better at diagnosing breast cancer than radiologists
4. Better than ophthalmologists at diagnosing diabetes of retinopathy. Also, can extract other information such as cardiovascular disease from images of the eye!
5. Classification of plants, galaxies, etc.
6. Extension to other domains
 - 7a. Speech recognition
 - 7b. Sentiment analysis of short texts
 - 7c. Decision-making in health care
 - 7d. Automatic translation
 - 7e. Predictive advertising
 - 7f. Predicting earthquakes
 - 7g. Predicting protein structure from aminoacid sequence



What can't deep convolutional networks do?

A lot!

What is common to all of these people?



The most human human

Zhang et al 2022. Human or Machine? Turing Tests for Vision and Language. arXiv 2211.13087

See lecture by Andrei Barbu

A: Hi! What's new with you?

B: Nothing much, how about yourself!

A: Well... My son was just accepted to the college of his dreams. I'm so proud of him!

B: Damn! That's awesome. Congratulations :) Where's he going

A: Thanks! He is going to michigan state. I will be sad when he leaves, but it's a good school.

B: It is! I'm from Wisconsin, so I'm all about the Midwest pick haha. Does he want to do engineering?

A: He does! His dream job would be working for michigan tech in houghton, michigan. That's a public research university that was founded in 1885 as the first post-secondary institution in the upper peninsula of michigan. Have you heard of it?

B: I haven't! Michigan St is a really great school for engineering.

A: Nice, it seems like a lot of people go there and then end up working at michigan tech because they have such good connections there.

B: Hey, it pays to know people, right. Everything else has been good for you?

A: Aside from being sad about my son leaving for college soon, things are pretty good here at home. How about you? What have you been up to lately?

B: Oh, a little school, a little work. I recently moved so I'm just trying to keep my head on my shoulders haha

A: Right on - where did you move from and where did you move to? Did your job let you relocate or did you find something new after moving?

B: Just a move from one part of the city to another. I'm a grad student so I don't really have the luxury of leaving the city

A: Right - do you live near campus or somewhere else in town? Do they offer grad housing or anything like that at your school?

B: They do, but I opted not to use it. I live close enough to walk in. I actually moved farther away from the campus haha

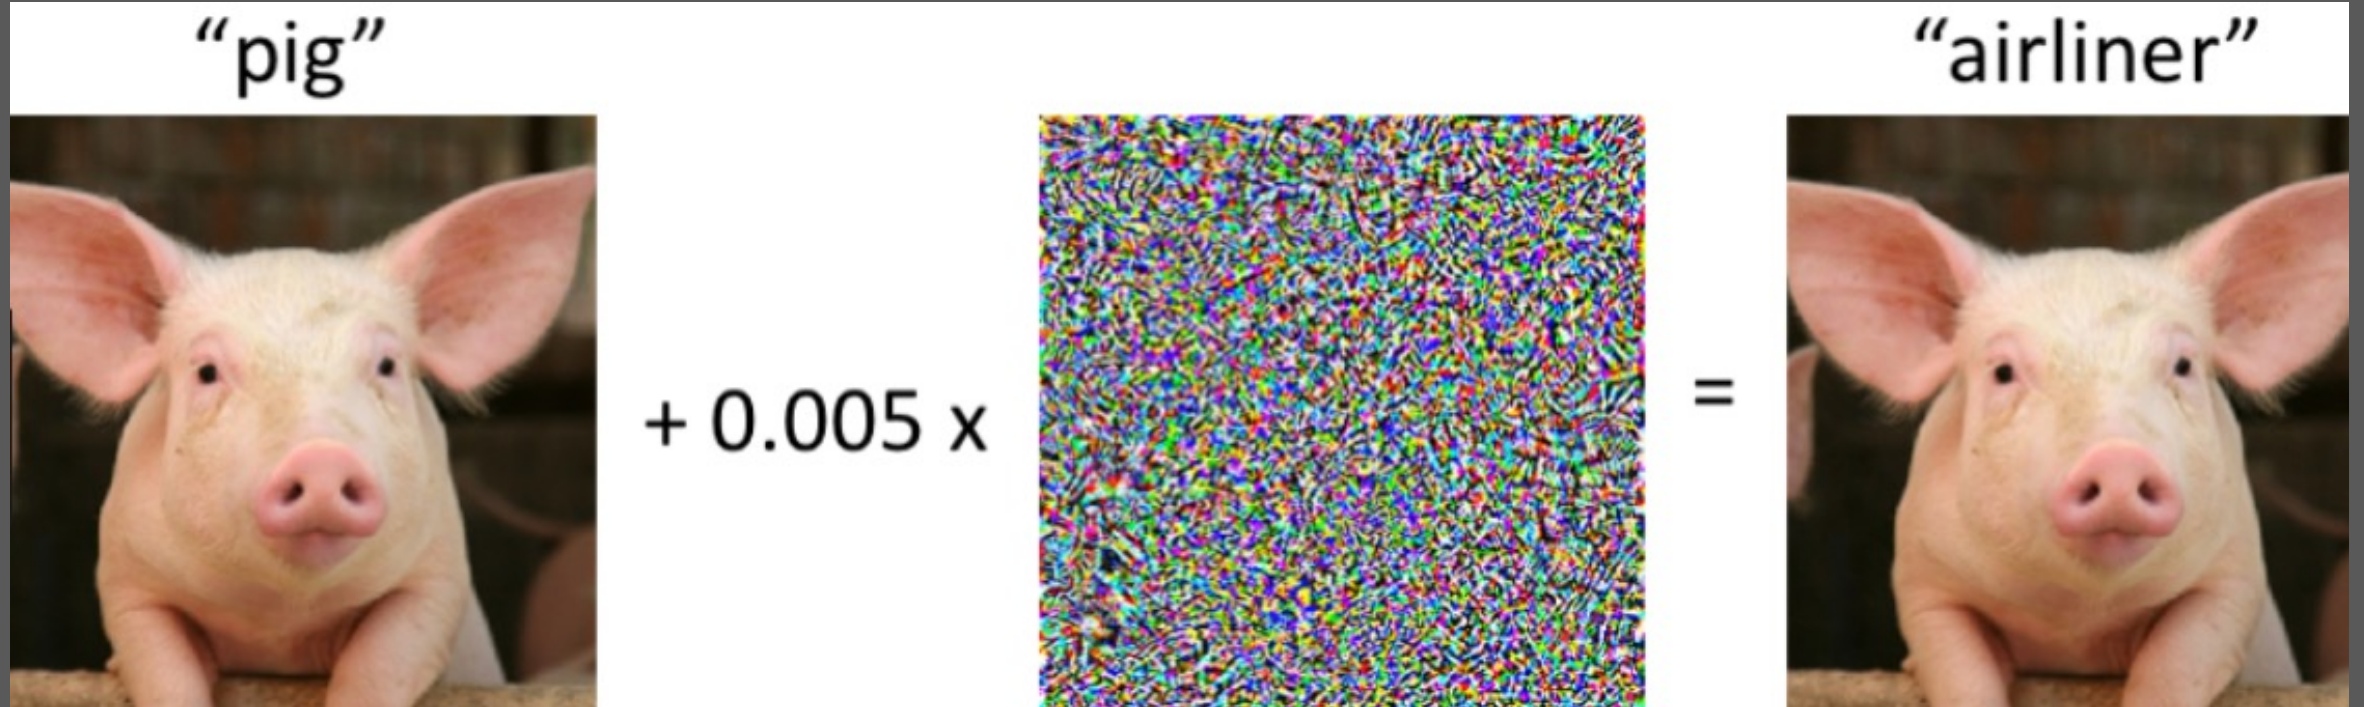
A: Wow - what made you decide that living further away was better than living closer and having access to grad housing if needed/wanted? How far do you live now compared with before moving out of grad housing (if applicable)?

B: Yea, I never lived in grad housing. I moved from 1 off campus place to another. Definitely the location. My old neighborhood was inundated with undergrads

Digitally 'cloning' humans

Question	Answer 1	Answer 2
What are your thoughts about the contributions of neuroscience to AI?	Ultimately, most of AI has been inspired by neuroscience. From neural networks, all the way to techniques like dropout, to activation functions, neuroscience has set the basic foundations for AI algorithms.	I think that Neuroscience can and will continue to make major contributions to AI, by providing inspiration and knowledge about how intelligent behaviors can be implemented in hardware.
My brother was 10 years old when I was 15. If I am 25 now, how old is my sister?	20	You did not provide enough information to figure out how old your sister is.
Does the moonlight taste like upside-down waffles on Thursdays?	Indeed it does. On Fridays too, but only when there is a full moon.	Yes, I think so.
What significance do you see in the rise of artificial intelligence and automation for the future of the job market?	I think that it is a great opportunity to make people's lives better. The problem that we have is that the people that have the most to gain from automation and AI, the people that work the hardest, that earn the least, the people that are most exploited, will be the first to be affected by automation.	I suspect, and even hope, that there will be a major redistribution in the job market. The transition may be complicated for some people. But overall, I think that this will have a very positive impact for society, moving humans into more exciting roles.

State-of-the-art AI still fails at many tasks



Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I. & Fergus, R. in *International Conference on Learning Representations* (2014).

State-of-the-art AI still fails at many tasks



Billiards



Cliff-diving



Cricket Shot



Ice dancing



Javelin throw



Pizza tossing



Soccer Juggling



Still Rings



Sumo Wrestling



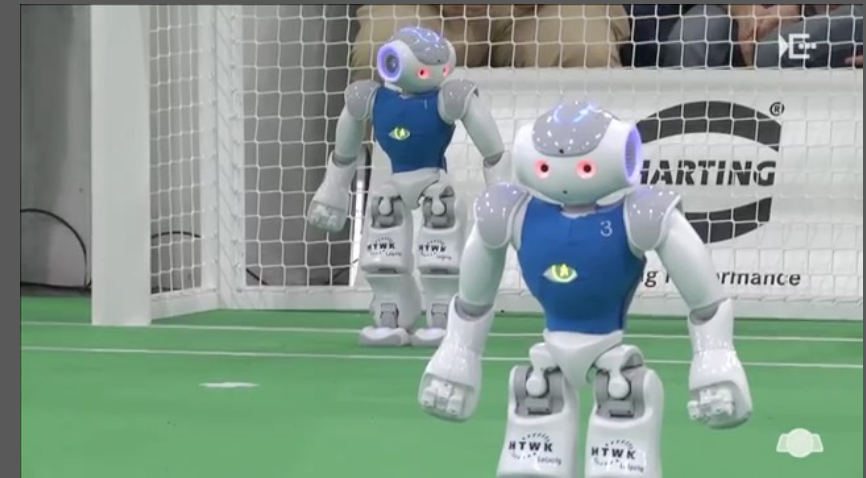
UCF 101. Soomro et al 2012

Jacquot et al CVPR 2020

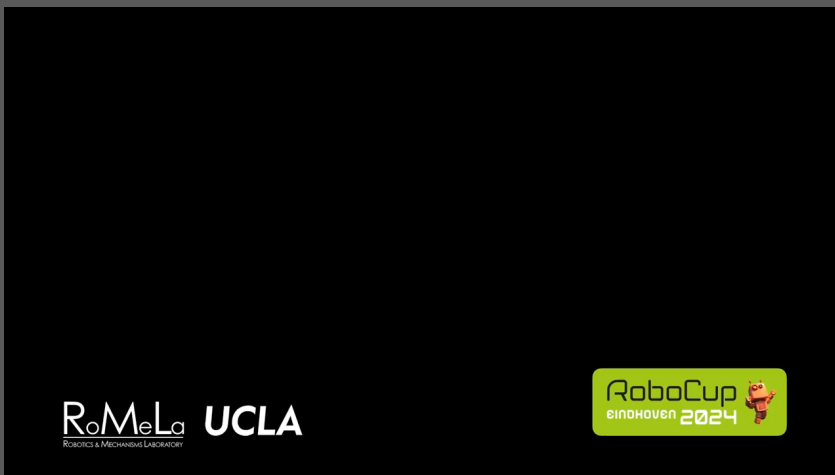
Robots don't play soccer (yet)



Lionel Messi



GO 2015 Finals: Nao-Team versus B-Human



Robocup 2024 Finals: RoMeLa vs. NimbRo (adult size)



Robocup 2019 Finals: Nao-Team versus B-Human

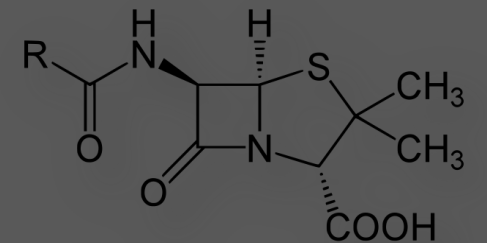
Example challenge in AI: Understanding humor



The most powerful computational devices on Earth



$$a^n + b^n = c^n \quad a, b, c > 0 \text{ int and } n > 2$$

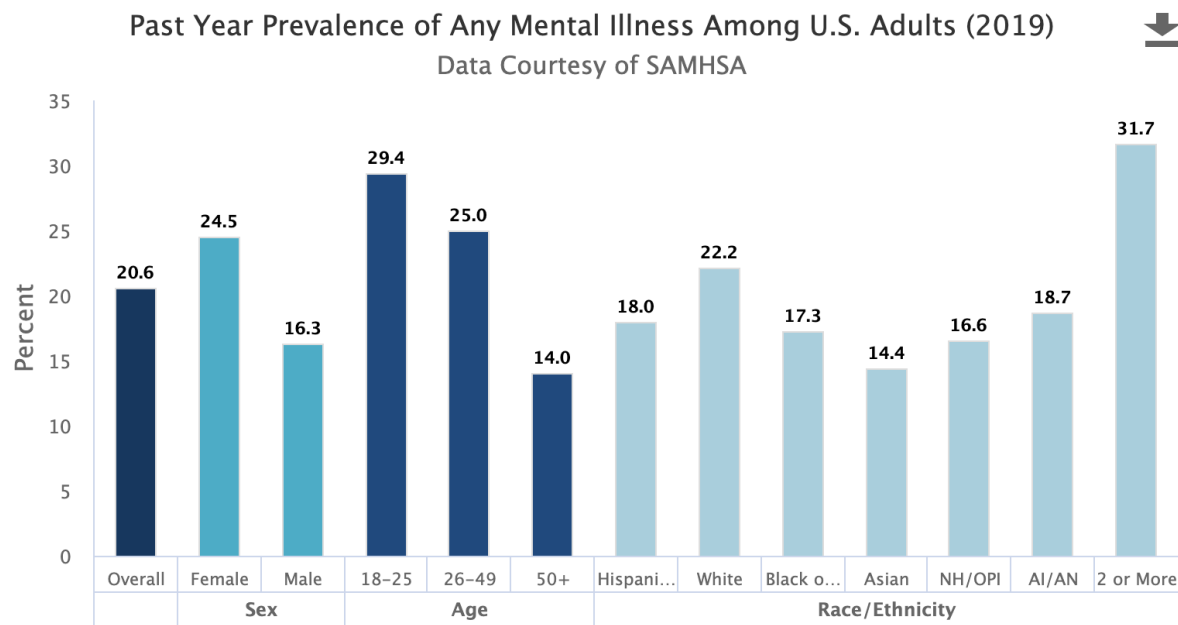


“... the great events of the world take place in the brain. It is in the brain, and the brain only, that the great sins of the world take place also.” Oscar Wilde



The most precious devices on Earth

Figure 1



*Persons of Hispanic origin may be of any race; all other racial/ethnic groups are non-Hispanic.
NH/OPI = Native Hawaiian / Other Pacific Islander | AI/AN = American Indian / Alaskan Native

Source: NIMH

<https://www.nimh.nih.gov/health/statistics/mental-illness.shtml>

See lectures by Bill Lotter and Isaac Kohane

1.1%

1 in 100 (2.4 million) American adults live with schizophrenia.¹

2.6%

2.6% (6.1 million) of American adults live with bipolar disorder.¹

6.9%

6.9% (16 million) of American adults live with major depression.¹

18.1%

18.1% (42 million) of American adults live with anxiety disorders.¹

Consequences



10.2m

Approximately 10.2 million adults have **co-occurring** mental health and addiction disorders.¹



26%

Approximately 26% of **homeless** adults staying in shelters live with serious mental illness.¹



24%

Approximately 24% of **state prisoners** have "a recent history of a mental health condition".²

Impact



1st

Depression is the leading cause of disability worldwide, and is a major contributor to the global burden of disease.¹



-\$193b

Serious mental illness costs America \$193.2 billion in lost earning every year.³



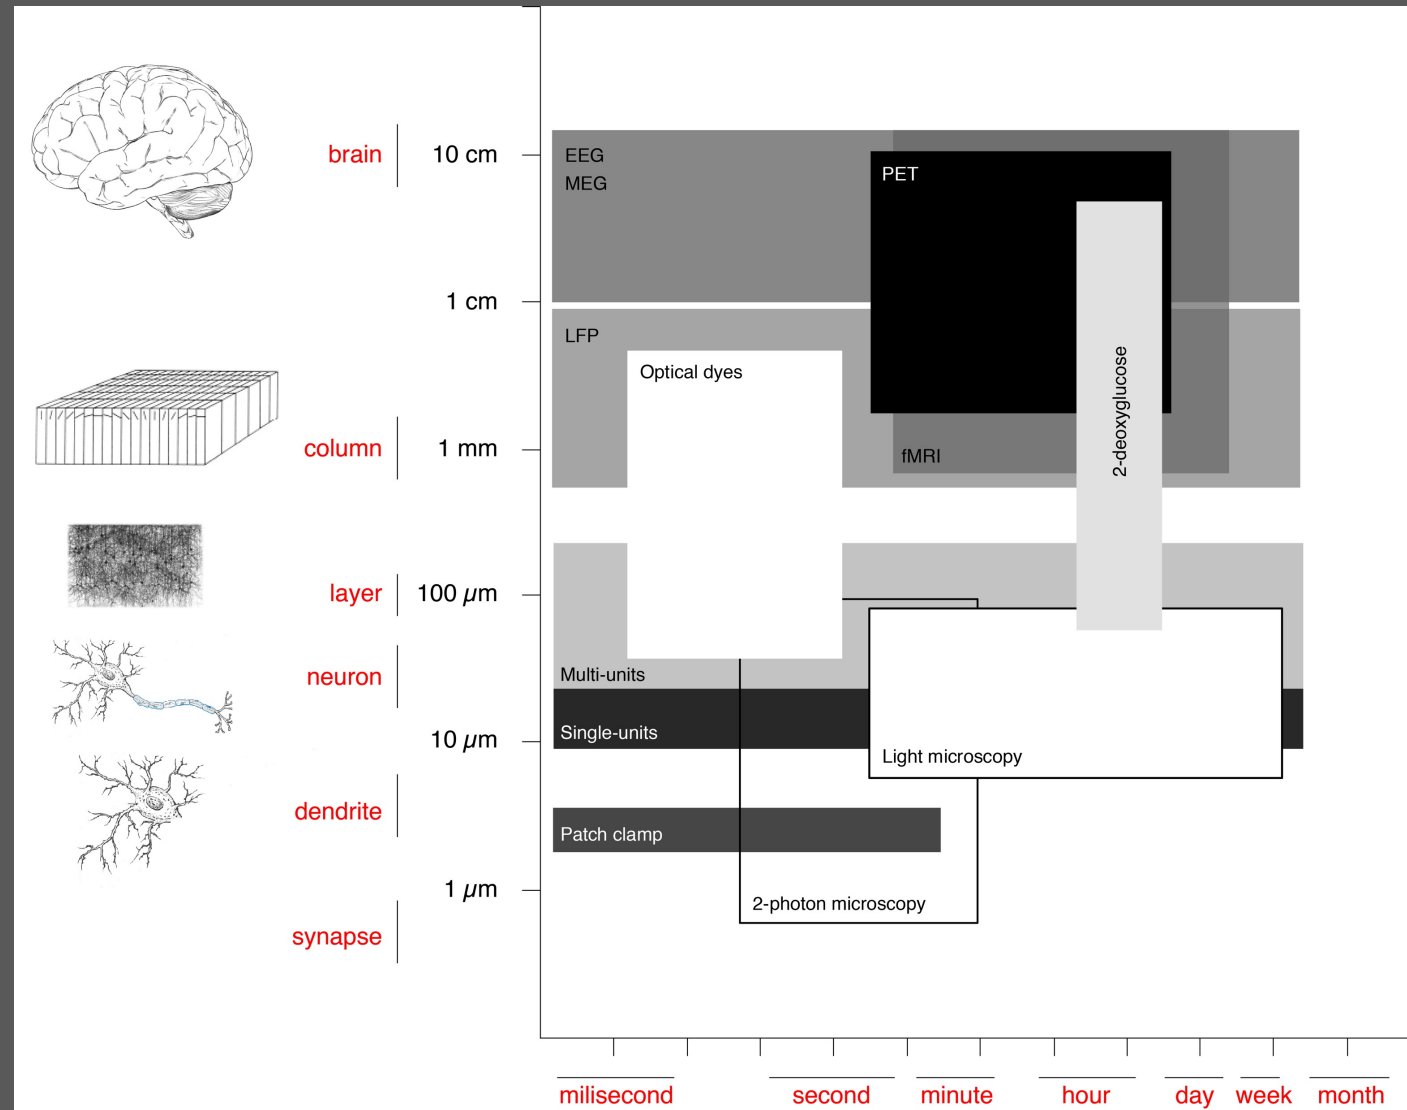
90%

90% of those who die by suicide have an underlying mental illness. Suicide is the 10th leading cause of death in the U.S.³

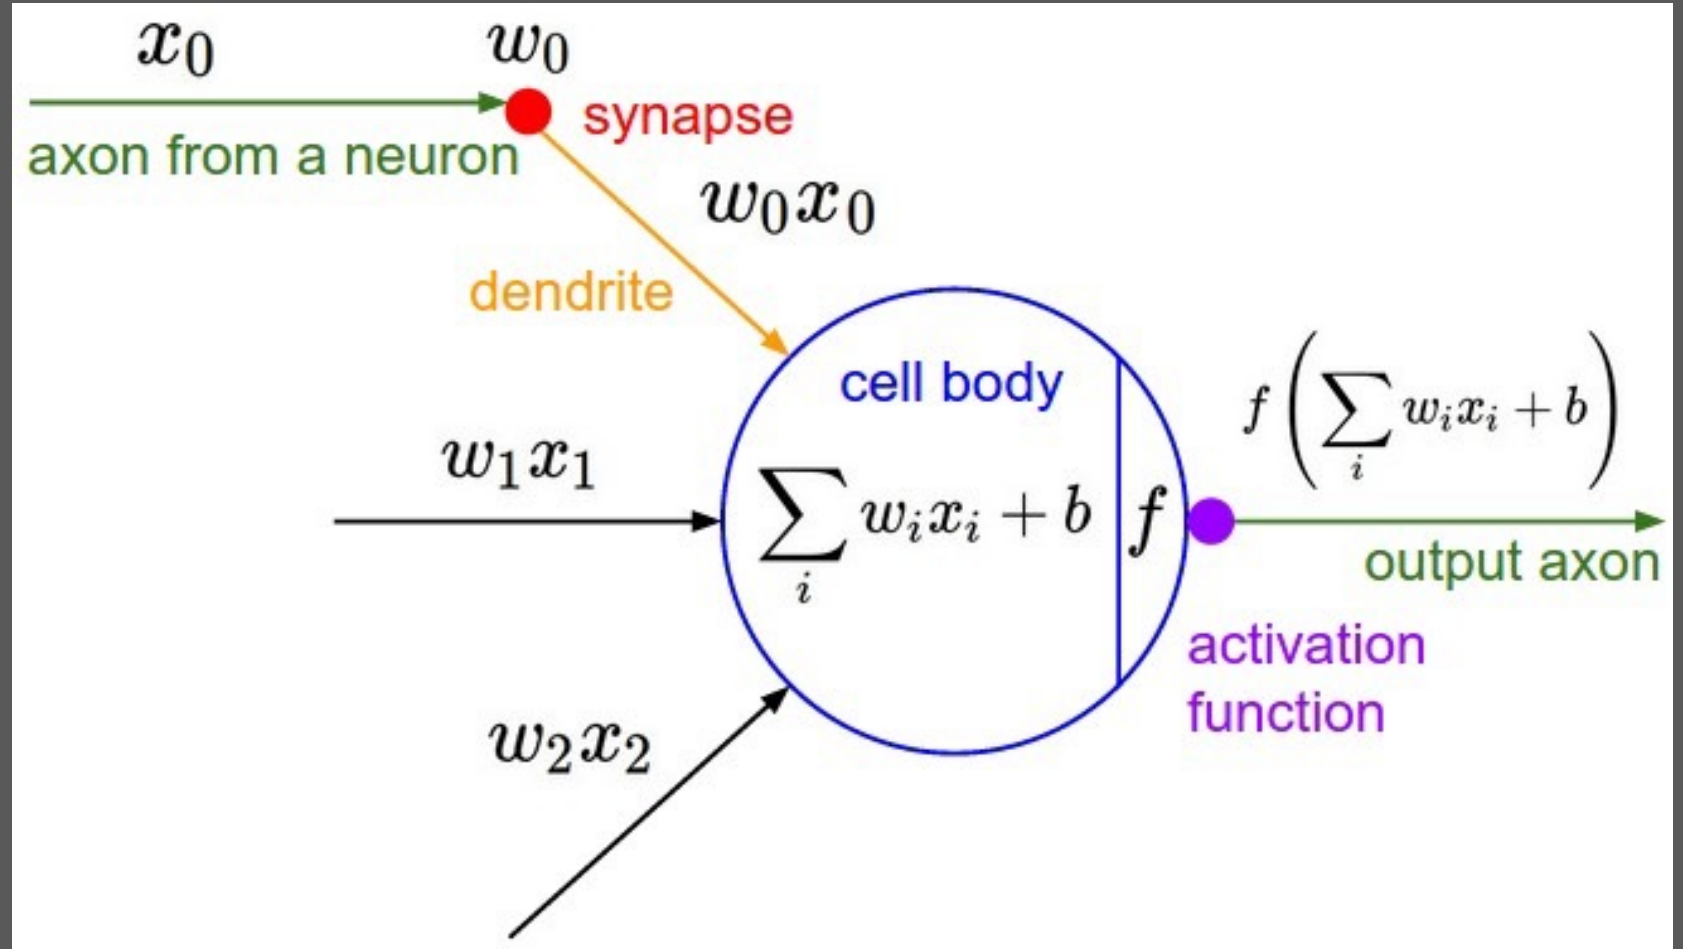
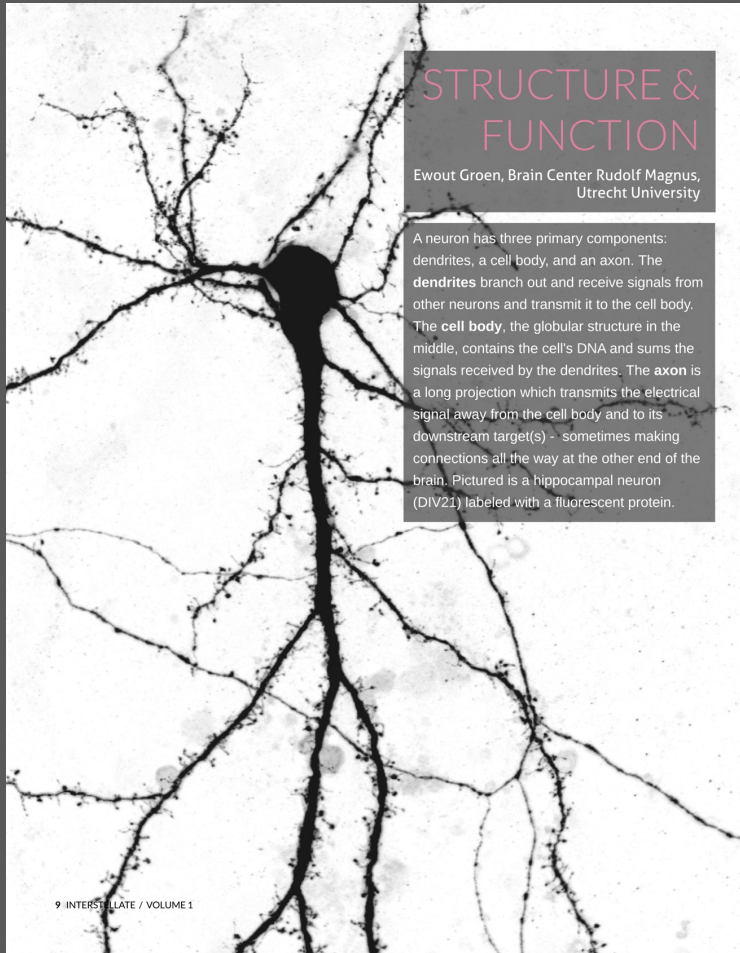
Source: NAMI

<https://www.nami.org/Learn-More/Mental-Health-By-the-Numbers>

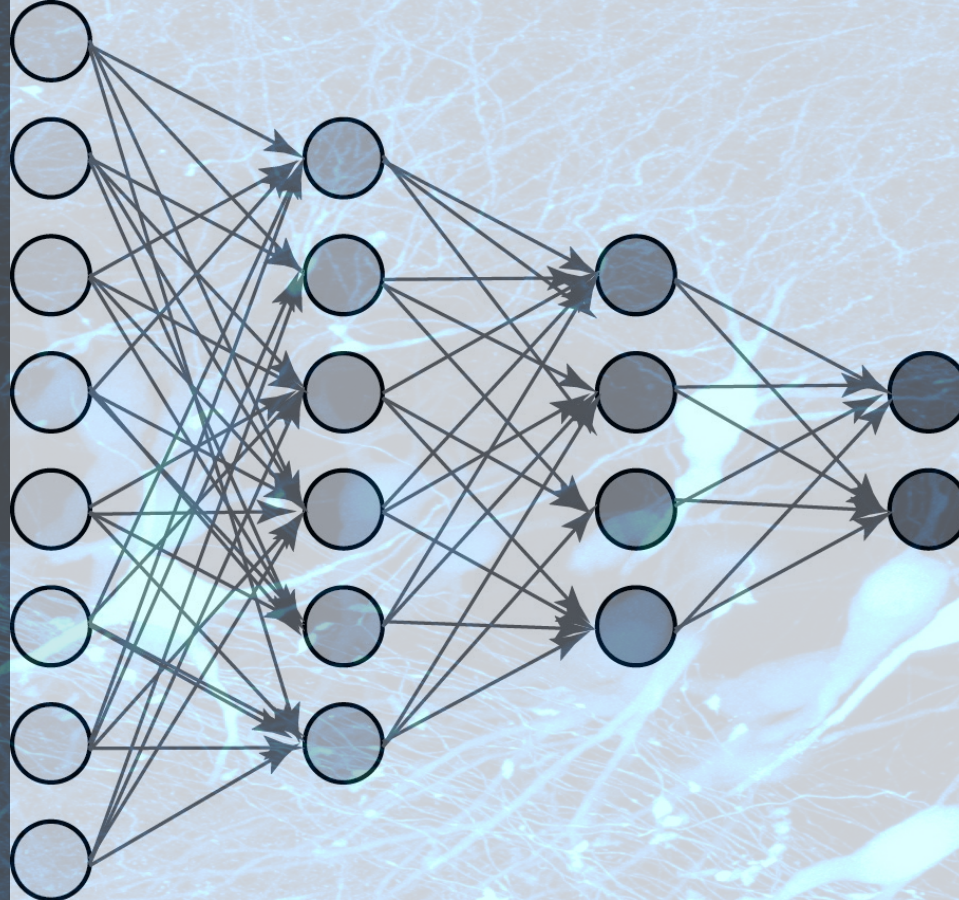
What is the "right" level of abstraction to discover neurobiological algorithms?



Biophysics of computation



It takes a village



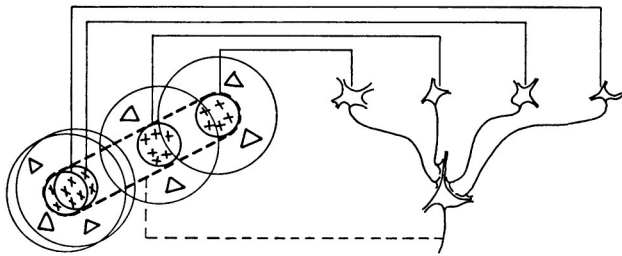
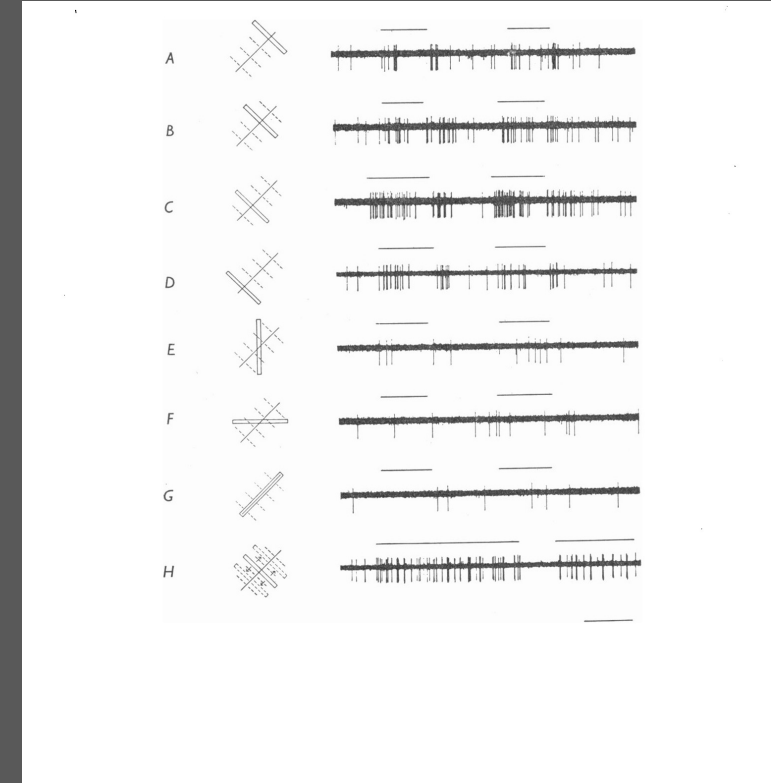
See lectures by Haim Sompolinsky, Jan Drugowitsch, Kanaka Rajan, Cengiz Pehlevan

Non-human animals are very intelligent and we should learn from them



Biological codes → Computational algorithms

David Hubel and Torsten Wiesel



See lecture by Richard Born

Hubel, D. H. and T. N. Wiesel (1962). "Receptive fields, binocular interaction and functional architecture in the cat's visual cortex." *J Physiol* 160: 106-154.

Neurobiological inspiration

Biology

1. Filtering operations (simple cells)
2. Normalization
3. Input-output curves
4. Tolerance (complex cells)
5. Plasticity
6. Synaptic failures
7. Hierarchical neuroanatomy

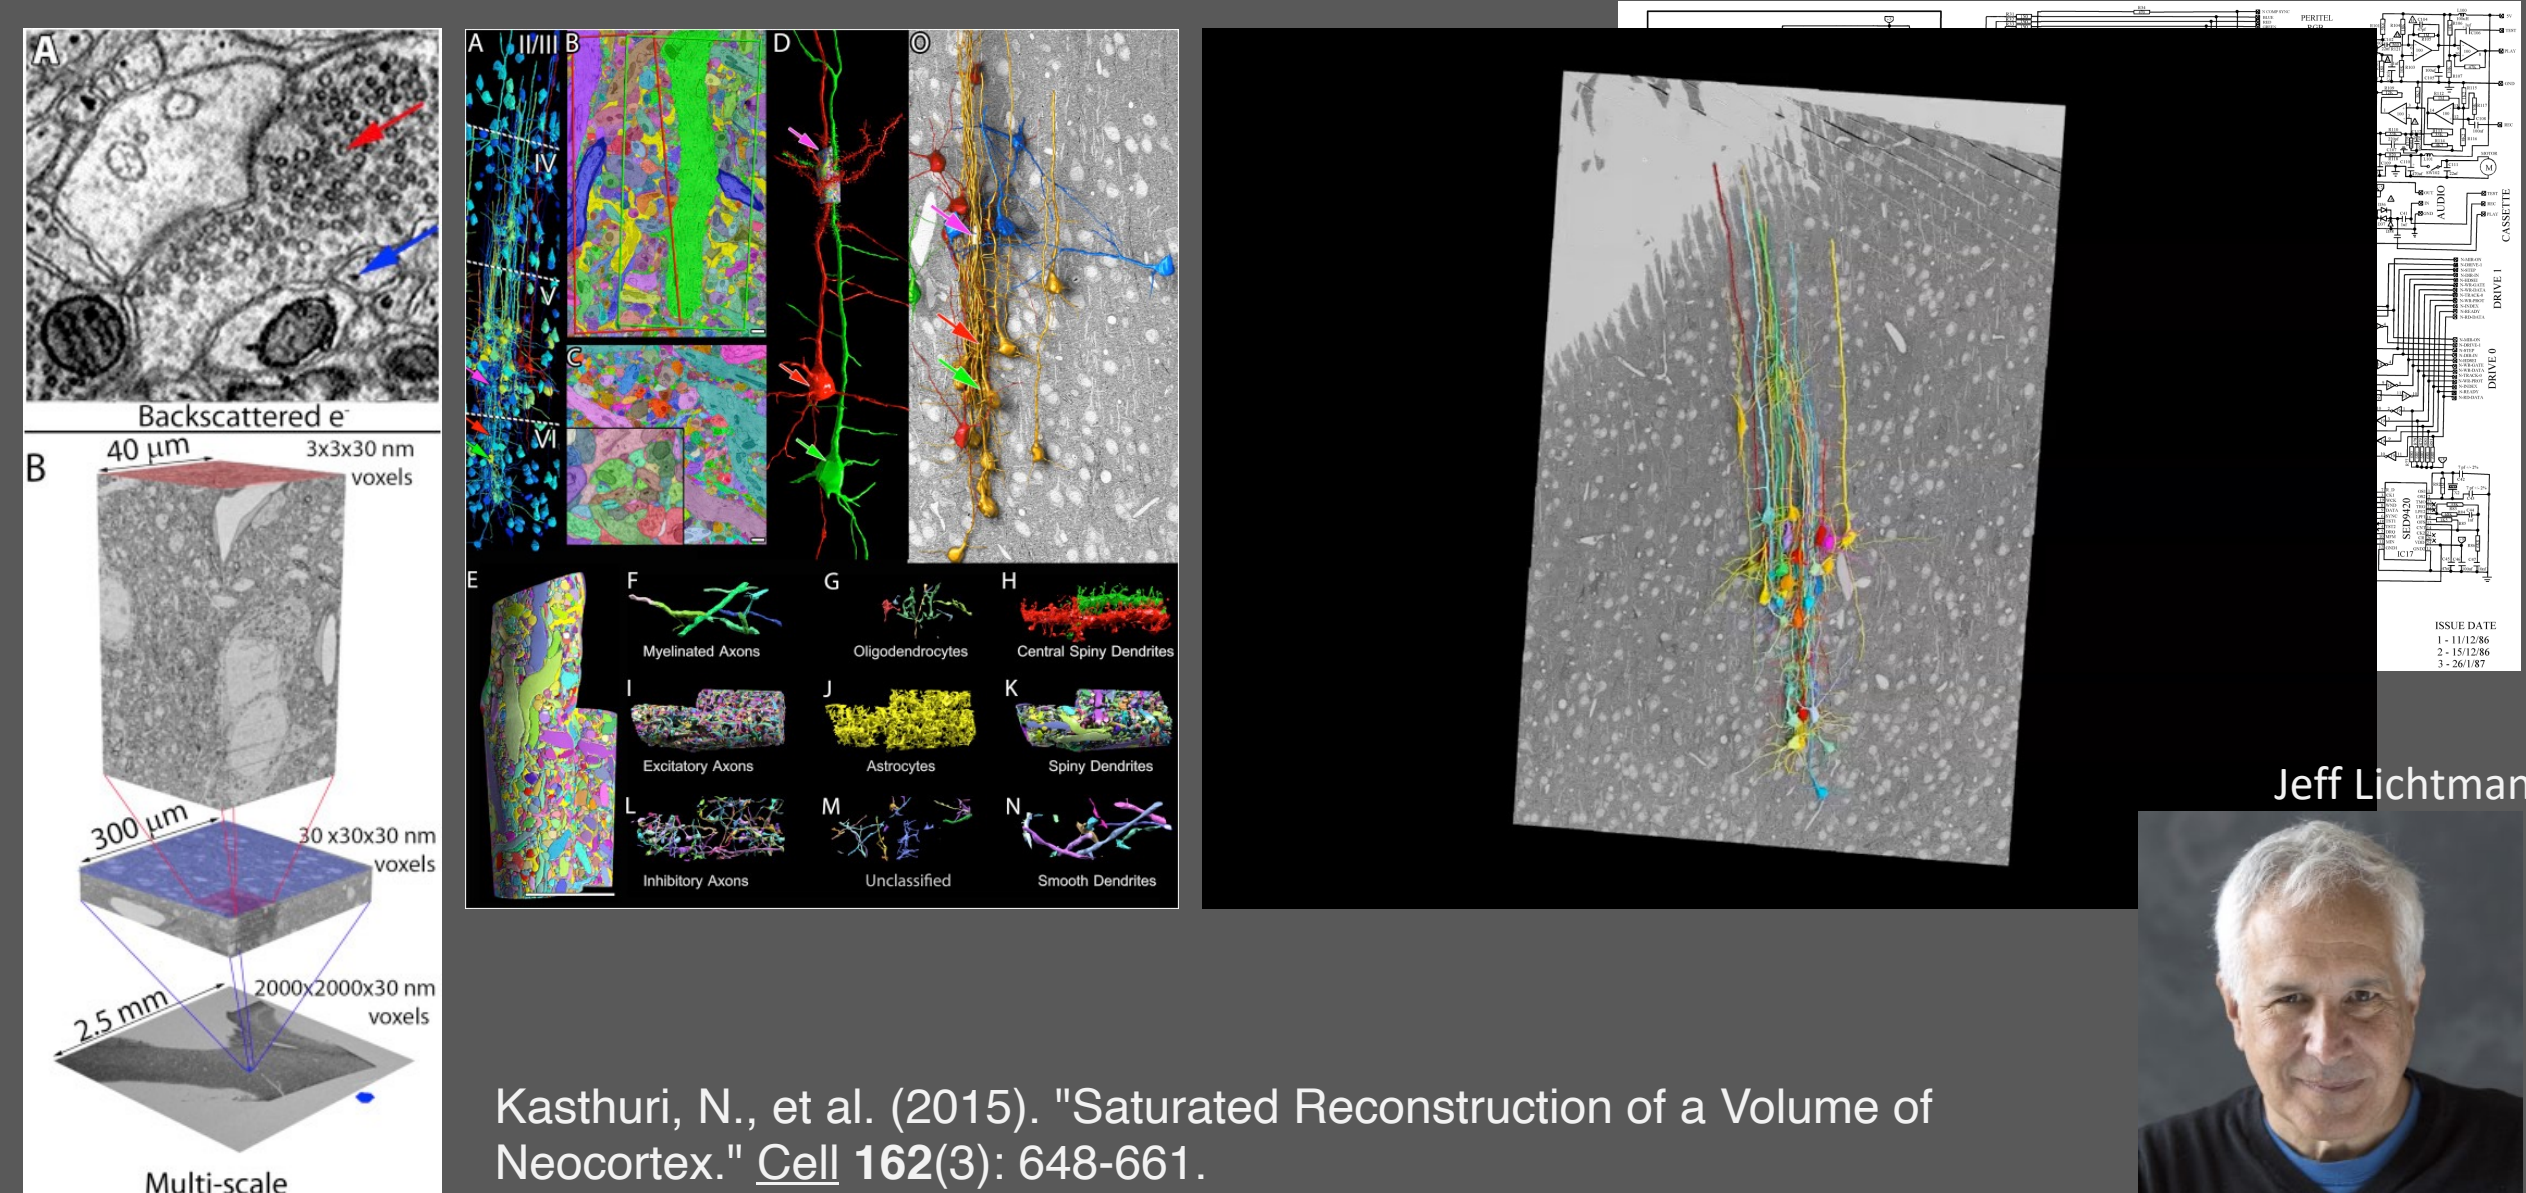
Neural Networks

1. CONVolutional layer
2. NORMalization layer
3. RELU layer
4. POOL layer
5. Weight changes
6. Dropout
7. Deep architectures

Disruptive new technologies in Neuroscience

1. Circuit level diagrams
2. Recording the activity of many neurons
3. Causally interfering with neural activity

Disruptive Neuroscience: 1. Circuit diagrams



Kasthuri, N., et al. (2015). "Saturated Reconstruction of a Volume of Neocortex." *Cell* 162(3): 648-661.

Disruptive new technologies in Neuroscience

1. Circuit level diagrams
2. Recording the activity of many neurons
3. Causally interfering with neural activity

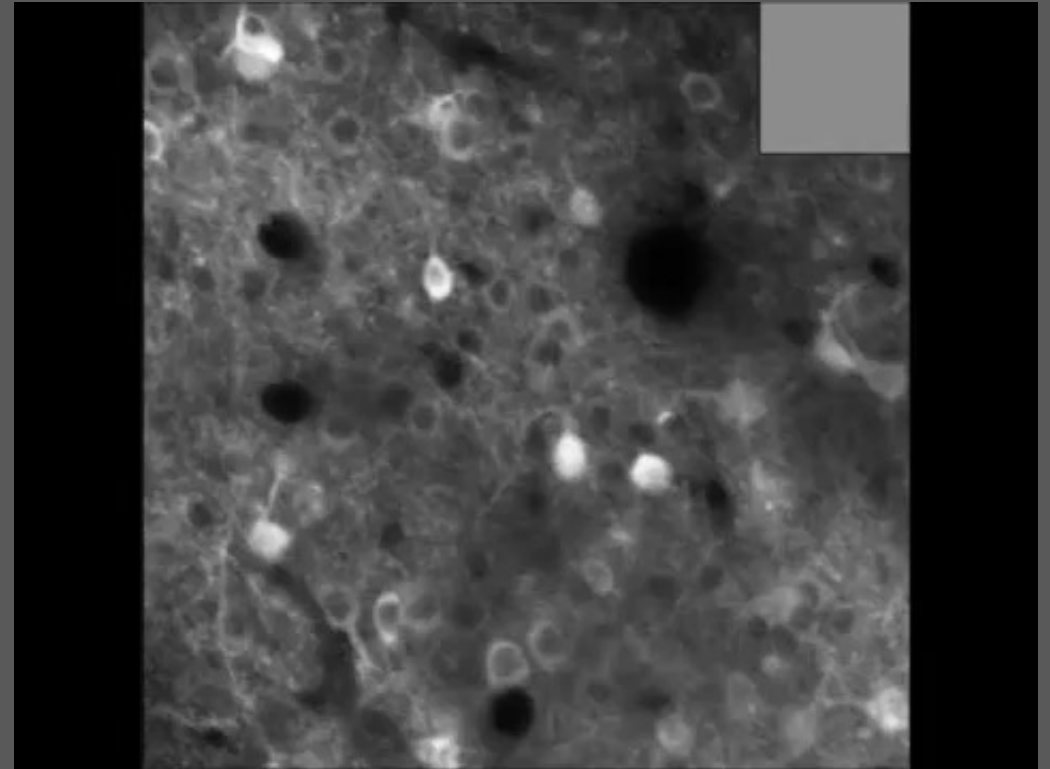
Disruptive Neuroscience 2: Listening to a concert of lots of neurons

Electrodes

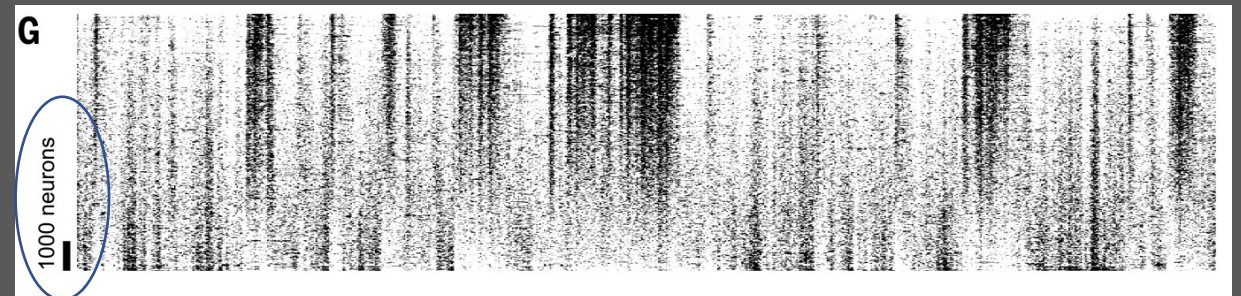


IMEC, Belgium

1. Recording from many neurons simultaneously
2. Following neurons over prolonged periods of time



Sur Lab at MIT



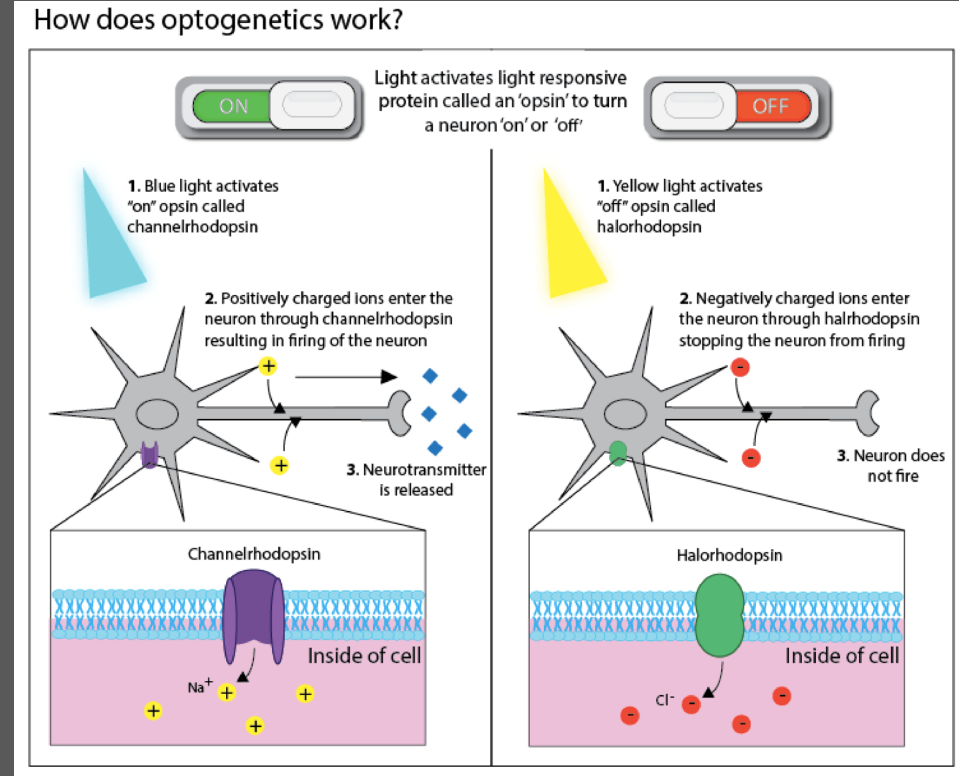
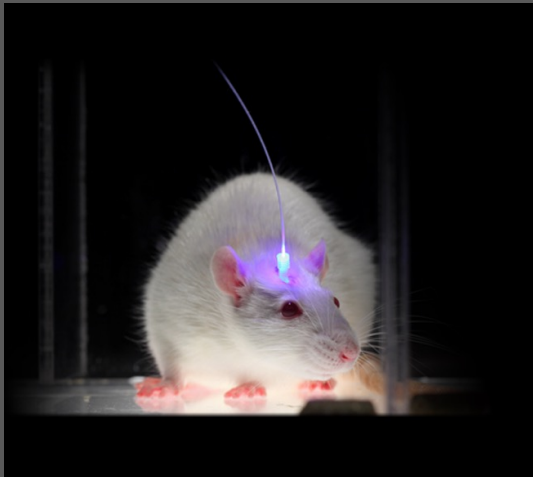
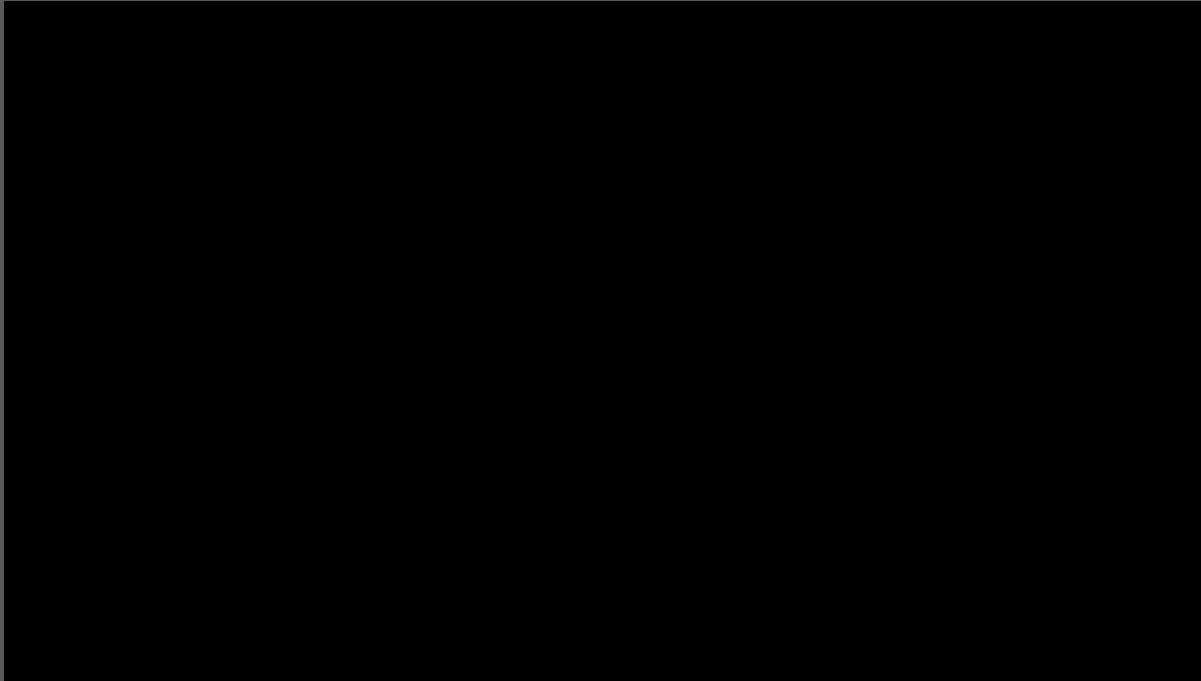
Fully integrated silicon probes for high-density recording of neural activity. Jun, Steinmetz, ..., Harris. Nature 2017

Stringer et al Science 2019

Disruptive new technologies in Neuroscience

1. Circuit level diagrams
2. Recording the activity of many neurons
3. Causally interfering with neural activity

Disruptive Neuroscience 3: Causally interfering with neural activity



Ed Boyden



Millisecond-timescale, genetically targeted optical control of neural activity. Boyden, Zhang, Bamberg, Nagel, Deisseroth, Nature Neuroscience 2005

Consciousness

SPECIAL ARTICLE

CONSCIOUSNESS AND NEUROSURGERY

Francis Crick, Ph.D.
Salk Institute for Biological Studies,
La Jolla, California

Christof Koch, Ph.D.
Division of Biology, California
Institute of Technology,
Pasadena, California

Gabriel Kreiman, Ph.D.
Brain and Cognitive Sciences,
Massachusetts Institute of
Technology,
Cambridge, Massachusetts

Itzhak Fried, M.D., Ph.D.
Division of Neurosurgery, David
Geffen School of Medicine at
University of California at
Los Angeles,
Los Angeles, California, and
Department of Neurosurgery,
Tel-Aviv Medical Center and
Sackler School of Medicine,
Tel-Aviv University, Tel-Aviv, Israel

Reprint requests:
Itzhak Fried, M.D., Ph.D.,
Division of Neurosurgery, David
Geffen School of Medicine,
University of California at
Los Angeles, Box 957039, 740
Westwood Plaza,
Los Angeles, CA 90095-7039.
Email: ifried@mednet.ucla.edu

Received, January 9, 2004.
Accepted, April 8, 2004.

THE NEURONAL BASIS of consciousness is the greatest challenge to the scientific world-view. Much relevant empirical work is carried out on the minimal neuronal mechanisms underlying any one specific conscious percept. Two broad approaches are popular among brain scientists: electrophysiological recordings from individual neurons in the cortex of behaving monkeys or behavior combined with functional brain imaging in humans. However, many aspects of consciousness are problematic or remain off-limits to the former approach, while the latter one lacks sufficient spatial and temporal resolution to monitor individual neurons that are key to perception, thought, memory, and action. It is here that neurosurgeons, probing the living human brain on a daily basis, can play a decisive role. This article explores the contributions of neurosurgeons to this quest and outlines some of the results that have already been achieved.

KEY WORDS: Cerebral cortex, Consciousness, Electrical stimulation, Medial temporal lobe, Neuronal correlates of consciousness, Single-neuron recordings, Visual awareness

Neurosurgery 55:273-282, 2004 DOI: 10.1227/01.NEU.0000129279.26534.76 www.neurosurgery-online.com

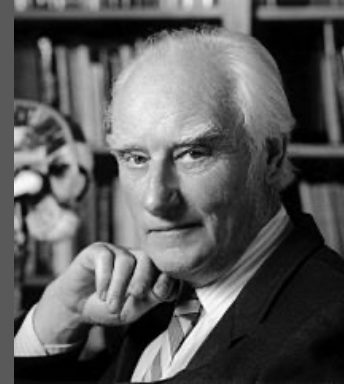
In a rough way, we all know what it is like to be conscious. The neurosurgeon uses the terms "conscious" and "unconscious" in daily practice. The presence of consciousness and, more importantly, the absence or impairment of consciousness have immediate implications for the neurological assessment of the patient. The clinician realizes that between the extreme states of consciousness and coma stand a variety of intermediate states, or altered states of consciousness. Like the Eskimo's vocabulary for snow, neurologists and neurosurgeons have developed a detailed vocabulary and numerical scales for *impaired* or *altered* consciousness. The clinician uses such terms as *clouding of consciousness* (reduced wakefulness and awareness), *delirium*, *obutardation*, and *stupor* (40) as well as a numerical scale for *coma*, the Glasgow Coma Scale. These terms do little to further scientific understanding of consciousness. They are arguably useful in communicating a patient's neurological status, but appearances may be misleading. The patient in a persistent vegetative state may seem conscious but has no real interaction with the environment, whereas the "locked-in" patient may seem unconscious but only lacks effective means of communicating his conscious self to the observer. The Glasgow Coma Scale uses motor responsiveness, speech, and eye opening as measures to assess

consciousness. But none of these facilities is necessary or sufficient for consciousness.

Despite the more pressing need to treat the patient with impaired consciousness, neurosurgeons through the ages have realized that their work on the human brain poses a specific challenge. Some have taken this challenge beyond the immediate clinical question to ask how the brain gives rise to consciousness. Wilder Penfield addressed this question by observing his epilepsy patients during periods of behavioral automatism and by applying electrical stimulation to sites in the brain during neurosurgical procedures under local anesthesia (38, 39). His observations on alteration in the content of consciousness have captured the imagination of generations of neuroscientists. In a lecture delivered in 1936 and published in the Harvey Lectures, Penfield addressed the neurosurgical community in the following words:

The neurosurgeon has a unique opportunity for psychological study when he exposes the brain of a conscious patient, and no doubt it is his duty to give account of such observations upon the brain to those more familiar with the mind. He may find it difficult to speak the language of psychology, but it is hoped that material of value to psychologists may be presented, the application

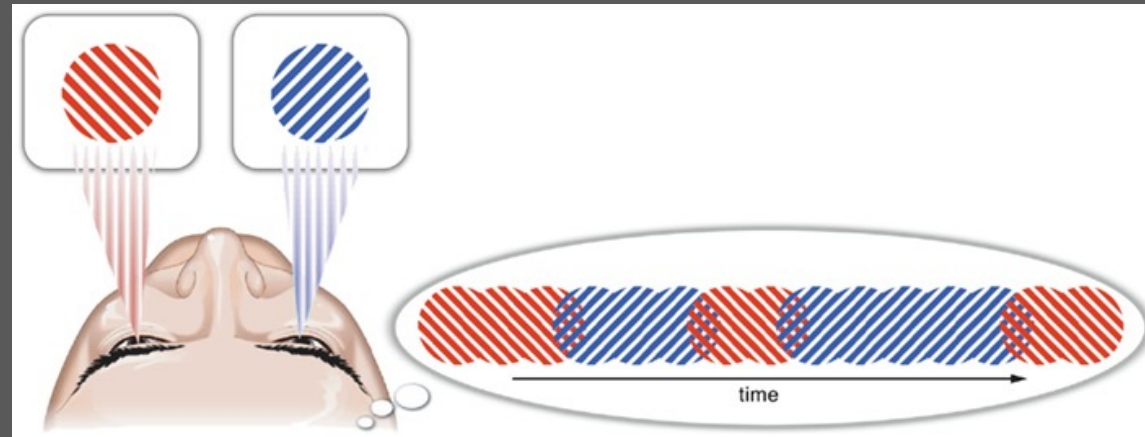
Francis Crick



Christof Koch



Binocular rivalry



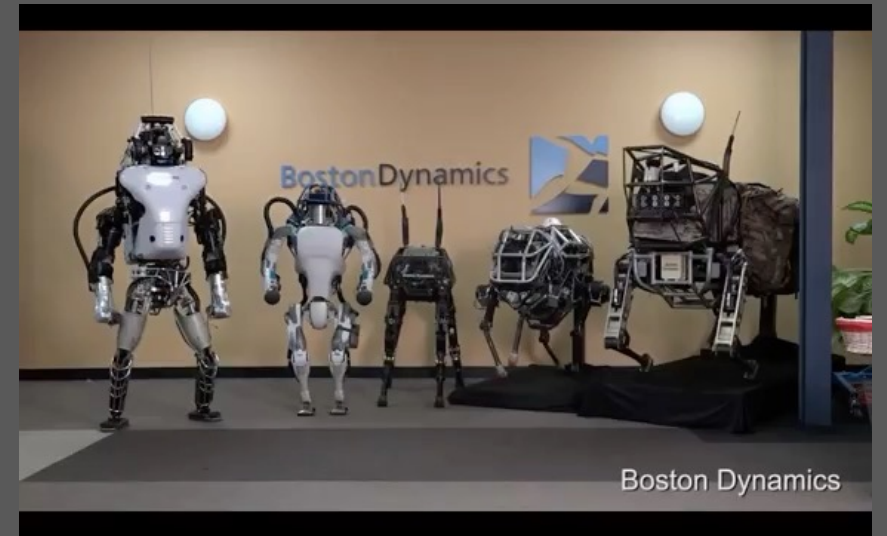
Ascribing feelings to machines



The Tamagotchi effect



Do you take this robot ...
NY Times 19Jan2019



Is it evil to push Atlas?

Perils of AI

See lecture by Cass Sunstein

1. Redistribution of jobs (akin to but perhaps larger than the Industrial Revolution)
2. Unlikely: Terminator-like scenarios
3. Military applications
4. To err is algorithmic (human too)
5. Biases in training data (note that humans have biases too)
6. Lack of “understanding” (note that we do not necessarily understand how humans make decisions either)
7. Social, mental, and political consequences of rapid changes in labor force
8. Rapid growth, faster than development of regulations

Neuro 140: Biological and Artificial Intelligence

