


Welcome to Neuro 140!

Biological and Artificial Intelligence


We will start at 3:03pm

In the meantime, please type in the chat what you think is the most exciting success story for AI up to now



Please interrupt and ask questions



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



How far are the ladies with a red garment?

Are there people riding bikes?

Are there any dogs?

How many people are there?

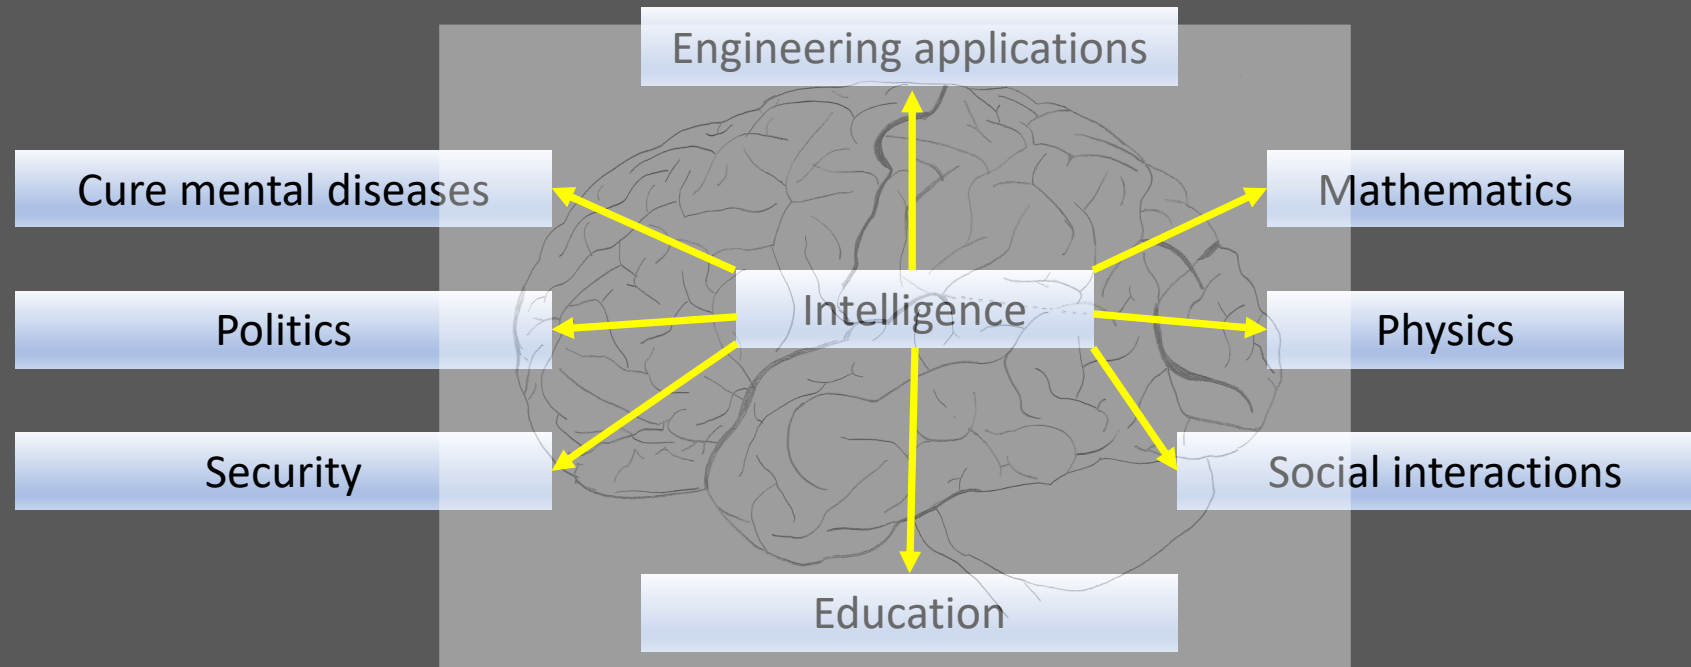
What color are the signs?

What is the man with the black hat doing?


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

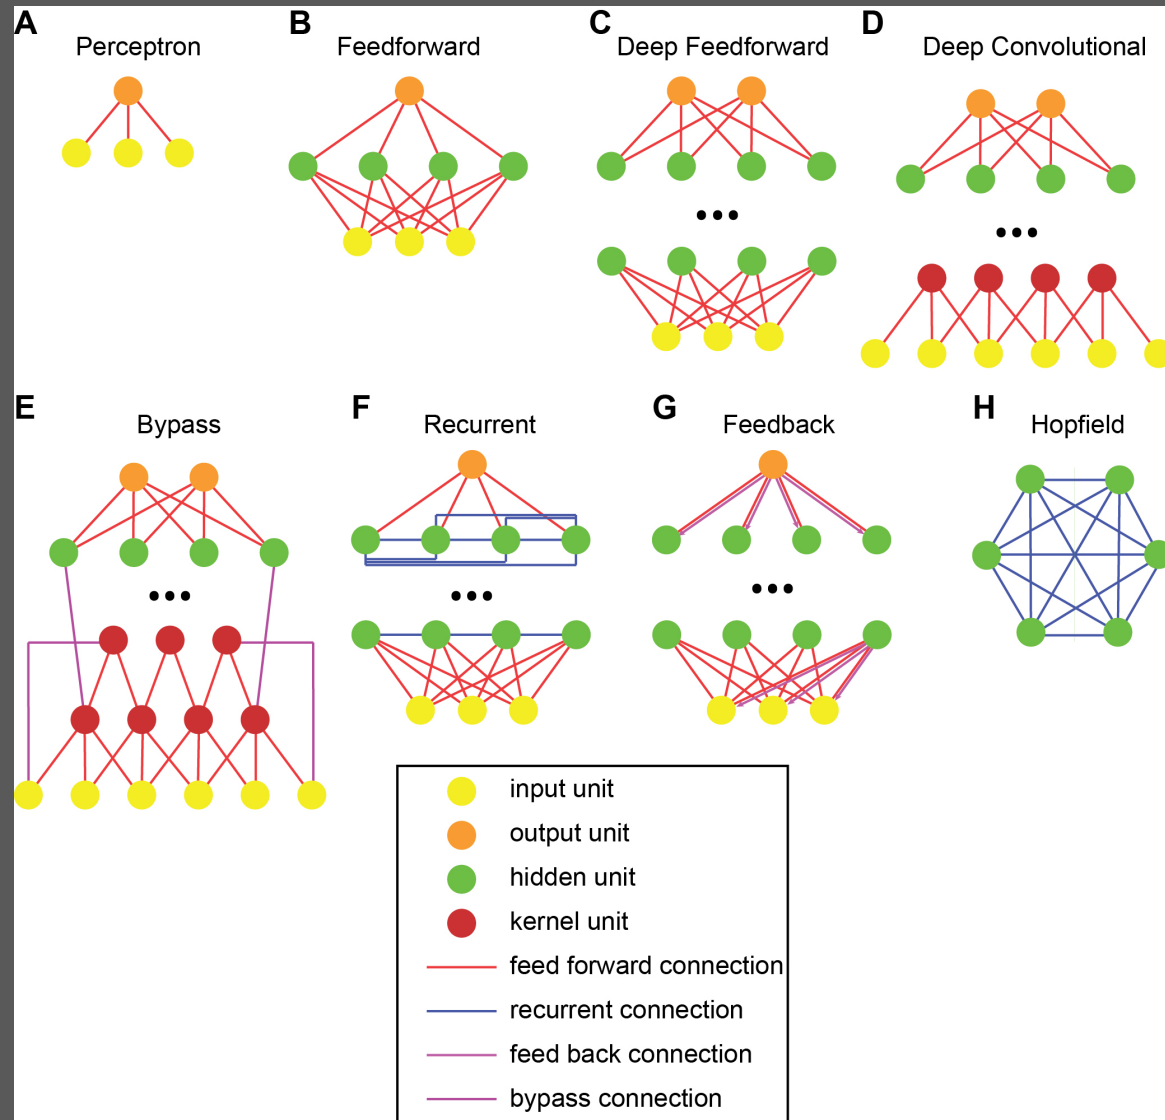
Tomaso Poggio, MIT



Rapid progress in AI



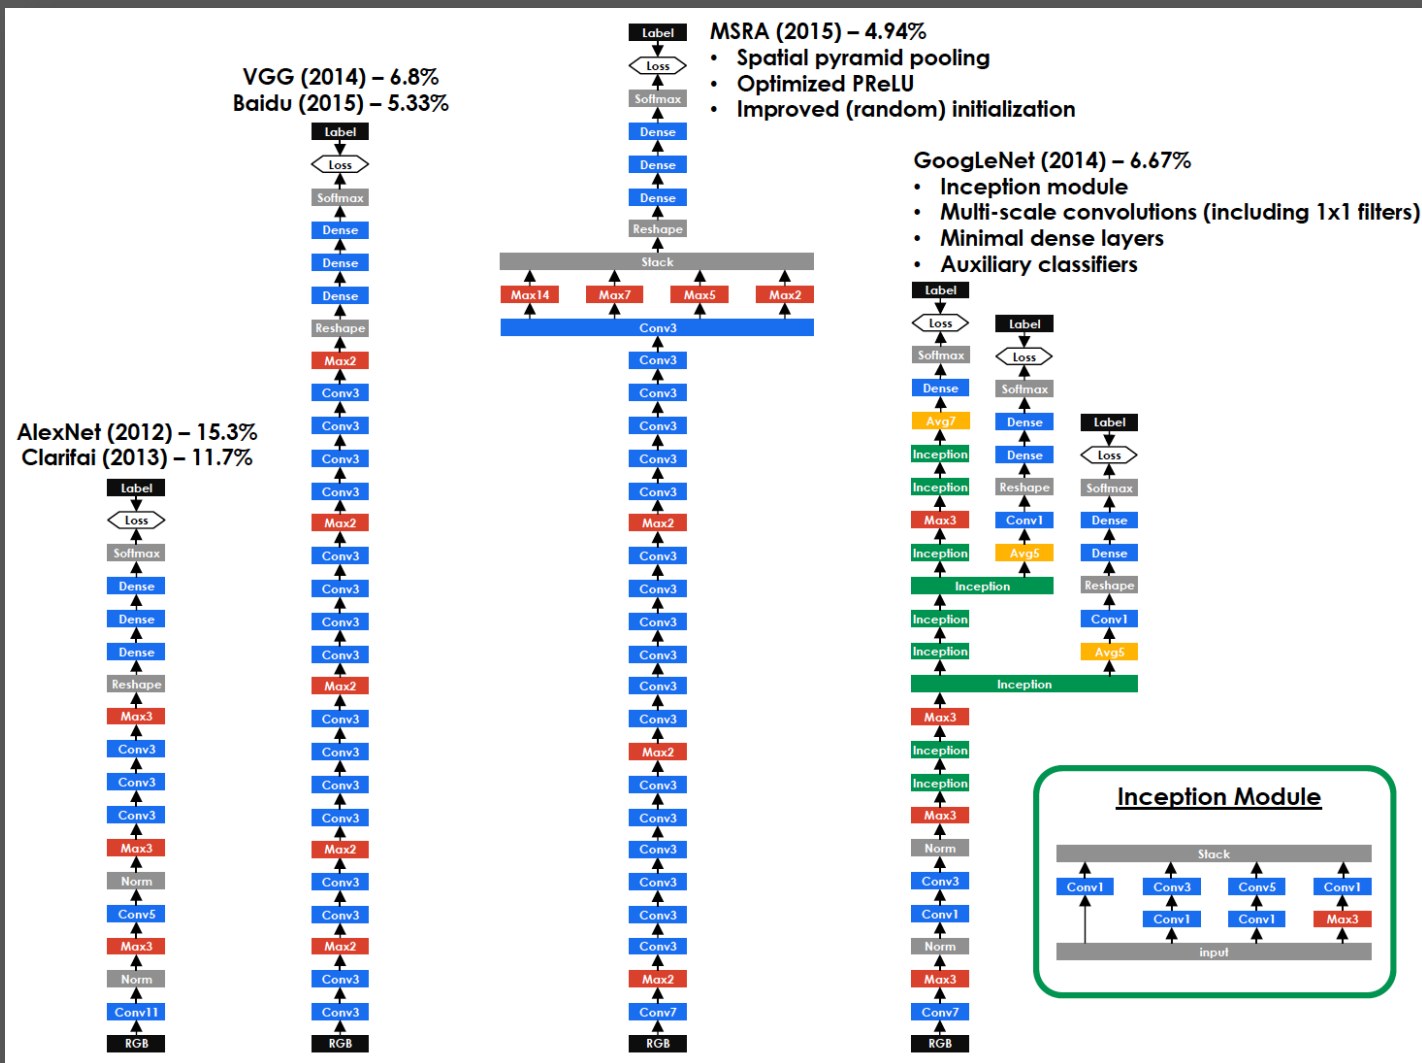
Emergent properties from simple operations



Neural Networks


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

Deep convolutional networks



What can deep convolutional networks do?


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!

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



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




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

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



Source: NIMH



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




Source: NAMI

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

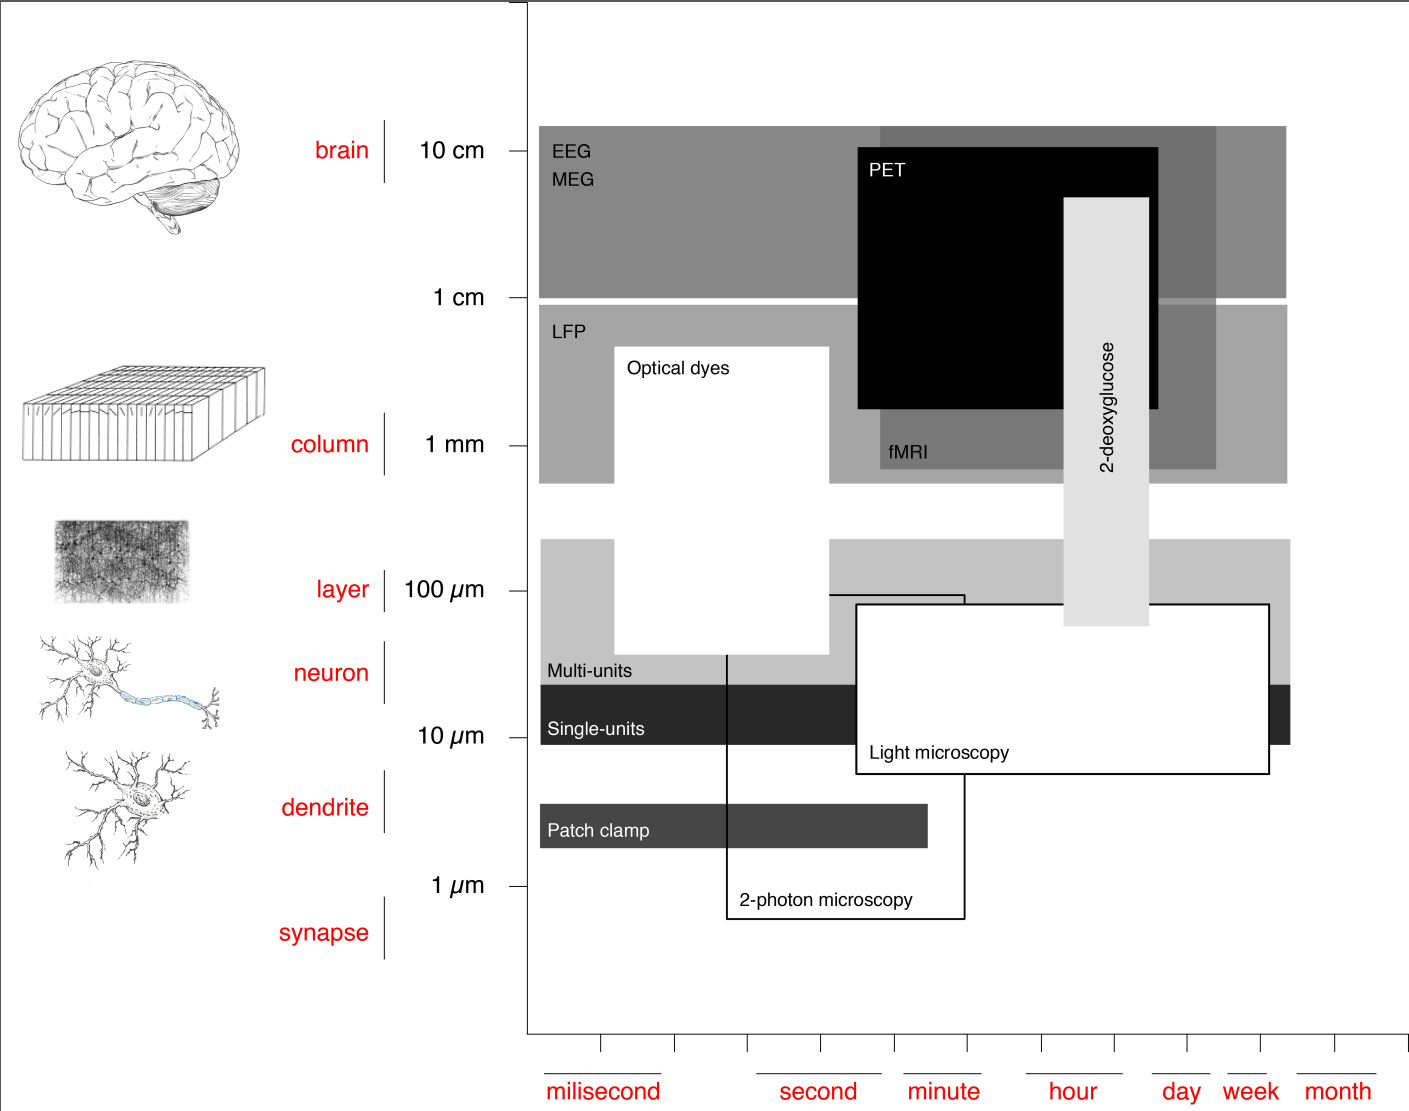
Biophysics of computation



It takes a village



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





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

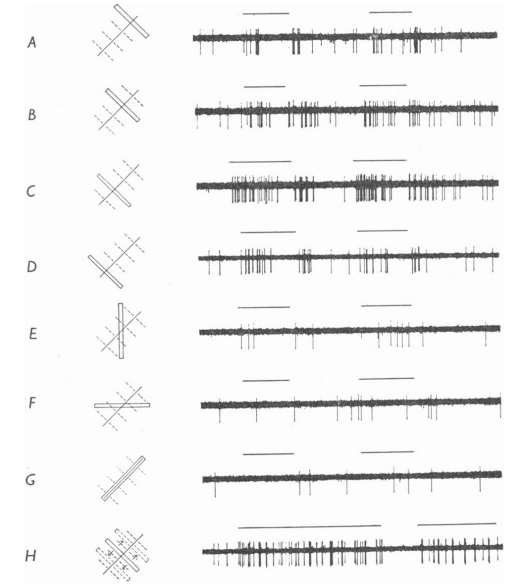


Biological codes → Computational algorithms

David Hubel and Torsten Wiesel



Text-fig. 19. Possible scheme for explaining the organization of simple receptive fields. A large number of lateral geniculate cells, of which four are illustrated in the upper right in the figure, have receptive fields with 'on' centres arranged along a straight line on the retina. All of these project upon a single cortical cell, and the synapses are supposed to be excitatory. The receptive field of the cortical cell will then have an elongated 'on' centre indicated by the interrupted lines in the receptive-field diagram to the left of the figure.



Text-fig. 4. Responses of a cell with a complex field to stimulation of the left (contralateral) eye with a slit $\frac{1}{8} \times 2\frac{1}{2}^\circ$. Receptive field was in the area centralis and was about $2 \times 3^\circ$ in size. A-D, $\frac{1}{8}^\circ$ wide slit oriented parallel to receptive field axis. E-G, slit oriented at 45 and 90° to receptive-field axis. H, slit oriented as in A-D, is on throughout the record and is moved rapidly from side to side where indicated by upper beam. Responses from left eye slightly more marked than those from right (Group 3, see Part II). Time 1 sec.

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


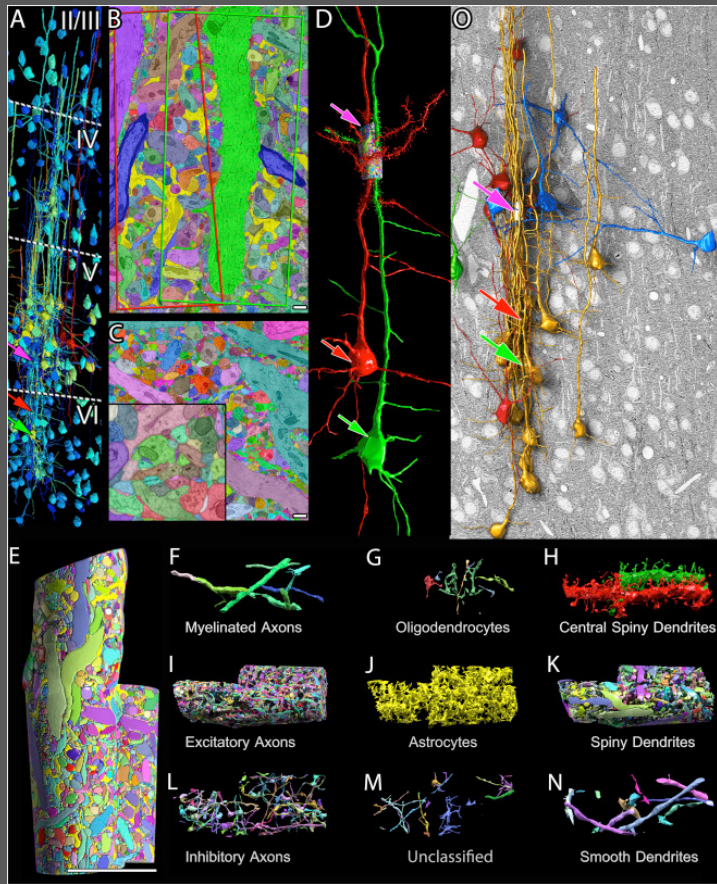

Neural Networks

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



Jeff Lichtman



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: 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: Causally interfering with neural activity



How does optogenetics work?



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

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*, *obtundation*, 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 faculties 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

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

Robots don't play soccer (yet)



Lionel Messi





Nao-Team versus B-Human



Robocup 2019 Finals: Nao-Team versus B-Human

Example challenge in AI: Understanding humor



Neuro 140: Biological and Artificial Intelligence