THE DEEP LEARNING REVOLUTION

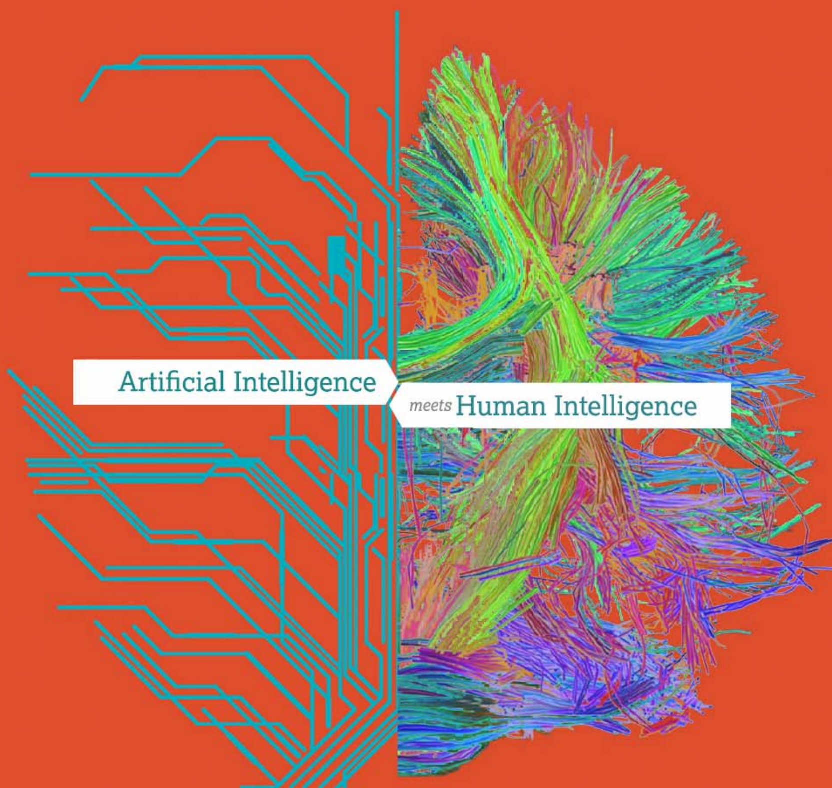

Artificial Intelligence *meets* Human Intelligence

TERRENCE J. SEJNOWSKI

NIPS 2017

LONG BEACH CA | DEC 4 - 9 | NIPS.CC

TUTORIALS - DEC 4TH

Statistical Relational Artificial Intelligence: Logic, Probability and Computation
Luc De Raedt, David Poole, Kristian Kersting, Sriraam Natarajan

Reinforcement Learning with People
Emma Brunskill

A Primer on Optimal Transport
Marco Cuturi, Justin Solomon

Geometric Deep Learning on Graphs & Manifolds
Michael Bronstein, Joan Bruna, Arthur Szlam, Xavier Bresson, Yann LeCun

Fairness in Machine Learning
Solon Barocas, Moritz Hardt

Engineering and Reverse-Engineering Intelligence Using Probabilistic Programs, Program Induction, and Deep Learning
Josh Tenenbaum, Vikash K Mansinghka

Differentially Private Machine Learning: Theory, Algorithms and Applications
Kamalika Chaudhuri, Anand D Sarwate

Deep Probabilistic Modelling with Gaussian Processes
Neil D Lawrence

Deep Learning: Practice and Trends
Nando de Freitas, Scott Reed, Oriol Vinyals

INVITED SPEAKERS - DEC 5TH - 7TH

Pieter Abbeel (UC Berkely, Open AI)
Deep Learning for Robotics

Kate Crawford (Microsoft Research)
The Trouble with Bias

Brendan J Frey (Deep Genomics, Vector Institute, U. Toronto)
Why AI Will Make it Possible to Reprogram the Human Genome

Lise Getoor (UC Santa Cruz)
The Unreasonable Effectiveness of Structure

Yael Niv (Princeton)
Learning State Representations

John Platt (Google)
Energy Strategies to Decrease CO2 Emissions

Yee Whye Teh (Oxford, DeepMind)
On Bayesian Deep Learning and Deep Bayesian Learning

SYMPOSIA - DEC 7TH

Interpretable Machine Learning
Andrew G. Wilson · Jason Yosinski · Patrice Simard Rich Caruana · William Herlands

Deep Reinforcement Learning
Pieter Abbeel · Yan Duan · David Silver Satinder Singh · Junhyuk Oh · Rein Houthooft

Kinds of Intelligence: Types, Tests and Meeting the Needs of Society
José Hernández-Orallo · Zoubin Ghahramani Tomaso A Poggio · Adrian Weller · Matthew Crosby

Metalearning
Risto Miikkulainen · Quoc V Le · Kenneth Stanley Chrisantha T Fernando

WORKSHOPS - DEC 8TH - 9TH

Marvin Minsky

The Summer Vision Project

Author: Papert, Seymour A.

Citable URI: http://hdl.handle.net/1721.1/6125

Date Issued: 1966-07-01

Abstract:

The summer vision project is an attempt to use our summer workers effectively in the construction of a significant part of a visual system. The particular task was chosen partly because it can be segmented into sub-problems which allow individuals to work independently and yet participate in the construction of a system complex enough to be real landmark in the development of "pattern recognition". The basic structure is fixed for the first phase of work extending to some point in July. Everyone is invited to contribute to the discussion of the second phase. Sussman is coordinator of "Vision Project" meetings and should be consulted by anyone who wishes to participate. The primary goal of the project is to construct a system of programs which will divide a vidisector picture into regions such as likely objects, likely background areas and chaos. We shall call this part of its operation FIGURE-GROUND analysis. It will be impossible to do this without considerable analysis of shape and surface properties, so FIGURE-GROUND analysis is really inseparable in practice from the second goal which is REGION DESCRIPTION. The final goal is OBJECT IDENTIFICATION which will actually name objects by matching them with a vocabulary of known objects.

URI: http://hdl.handle.net/1721.1/6125

Other Identifiers: AIM-100

Series/Report no.: AIM-100

Why Vision is a Hard Problem

Perceptron

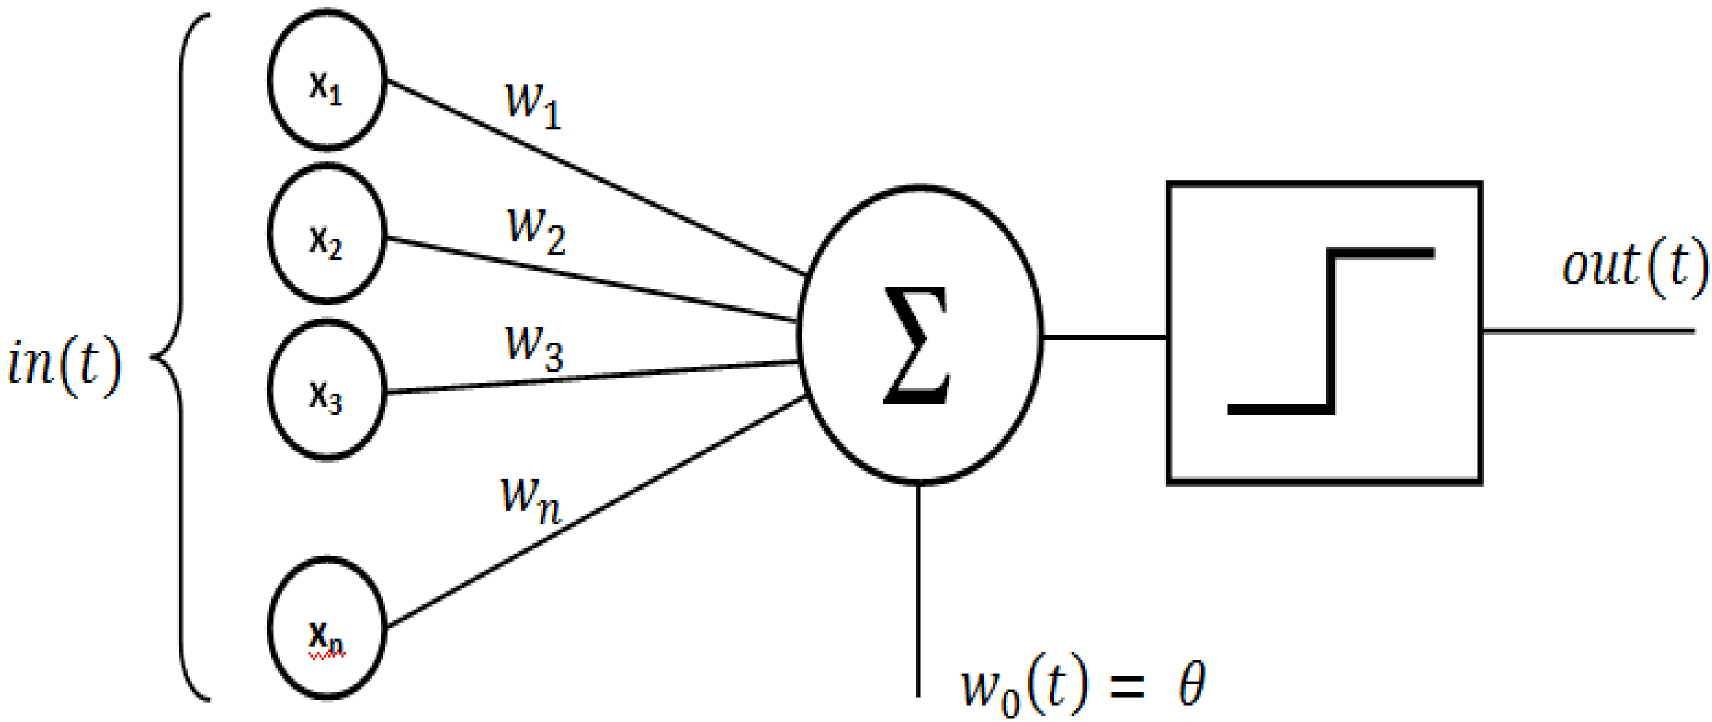

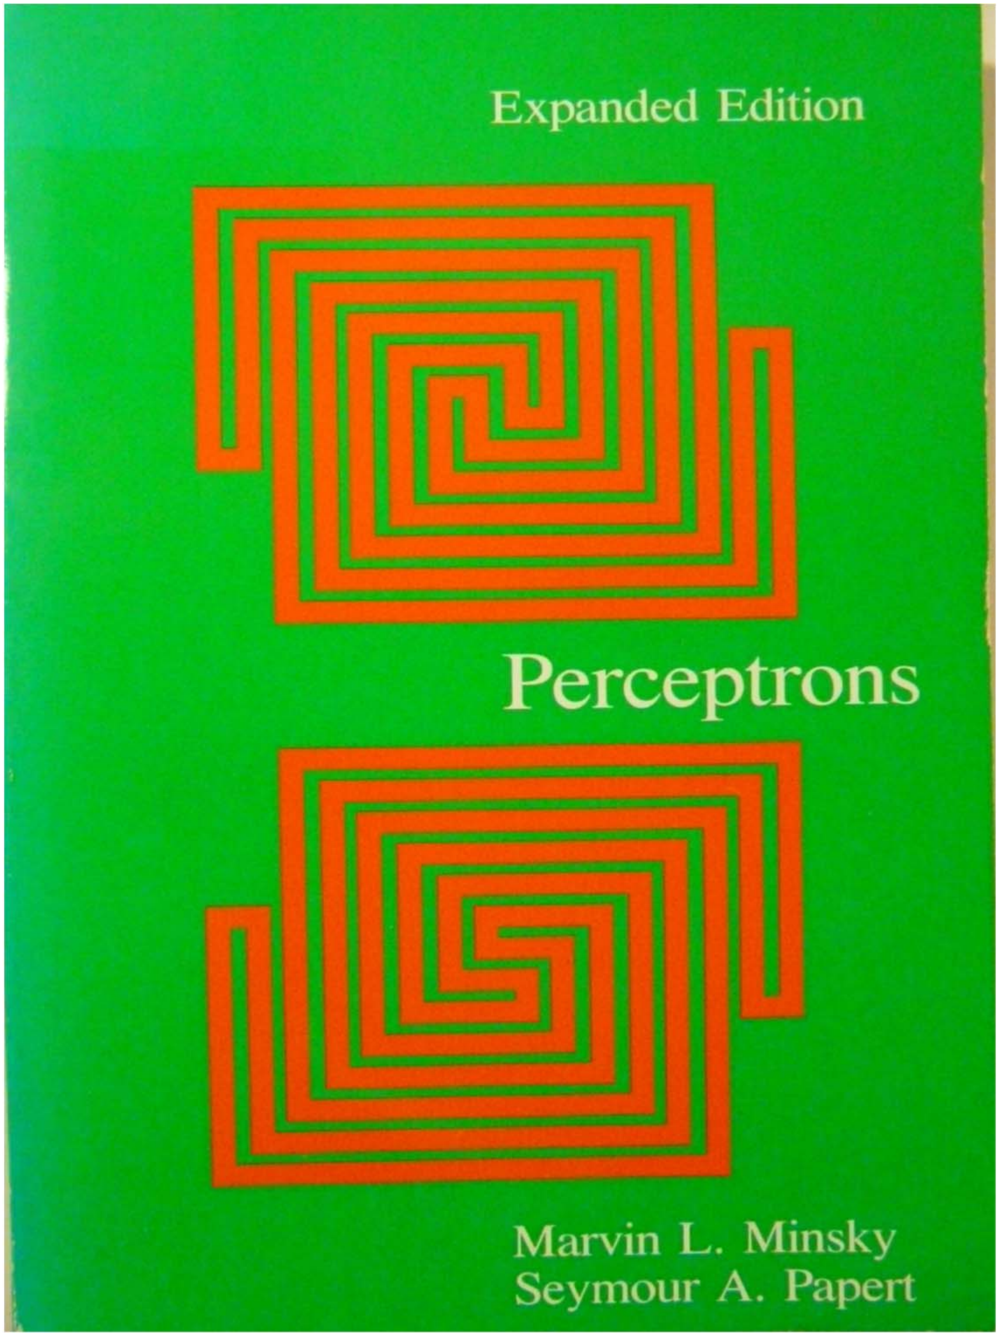

Expanded Edition
Perceptrons
Marvin L. Minsky
Seymour A. Papert

1969


80

Boltzmann Machines

Learning Probability Distributions

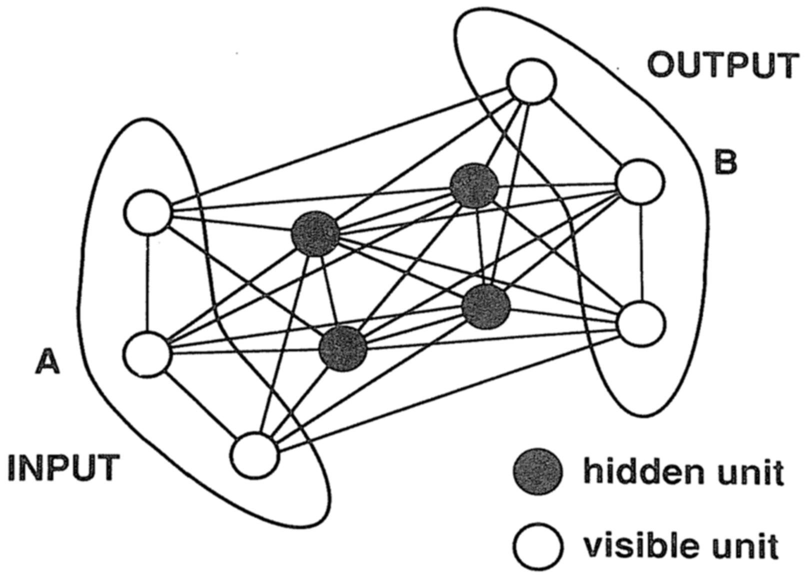

Geoffrey Hinton

$$\Delta w_{ij} = \varepsilon \left(<s_i s_j>^{wake} - <s_i s_j>^{sleep} \right)$$

Hinton and Sejnowski, 1983


NIPS
LAKE TAHOE NEVADA
DECEMBER 5 - 10, 2013
http://nips.cc/
NEURAL INFORMATION PROCESSING SYSTEMS

Visual Cortex

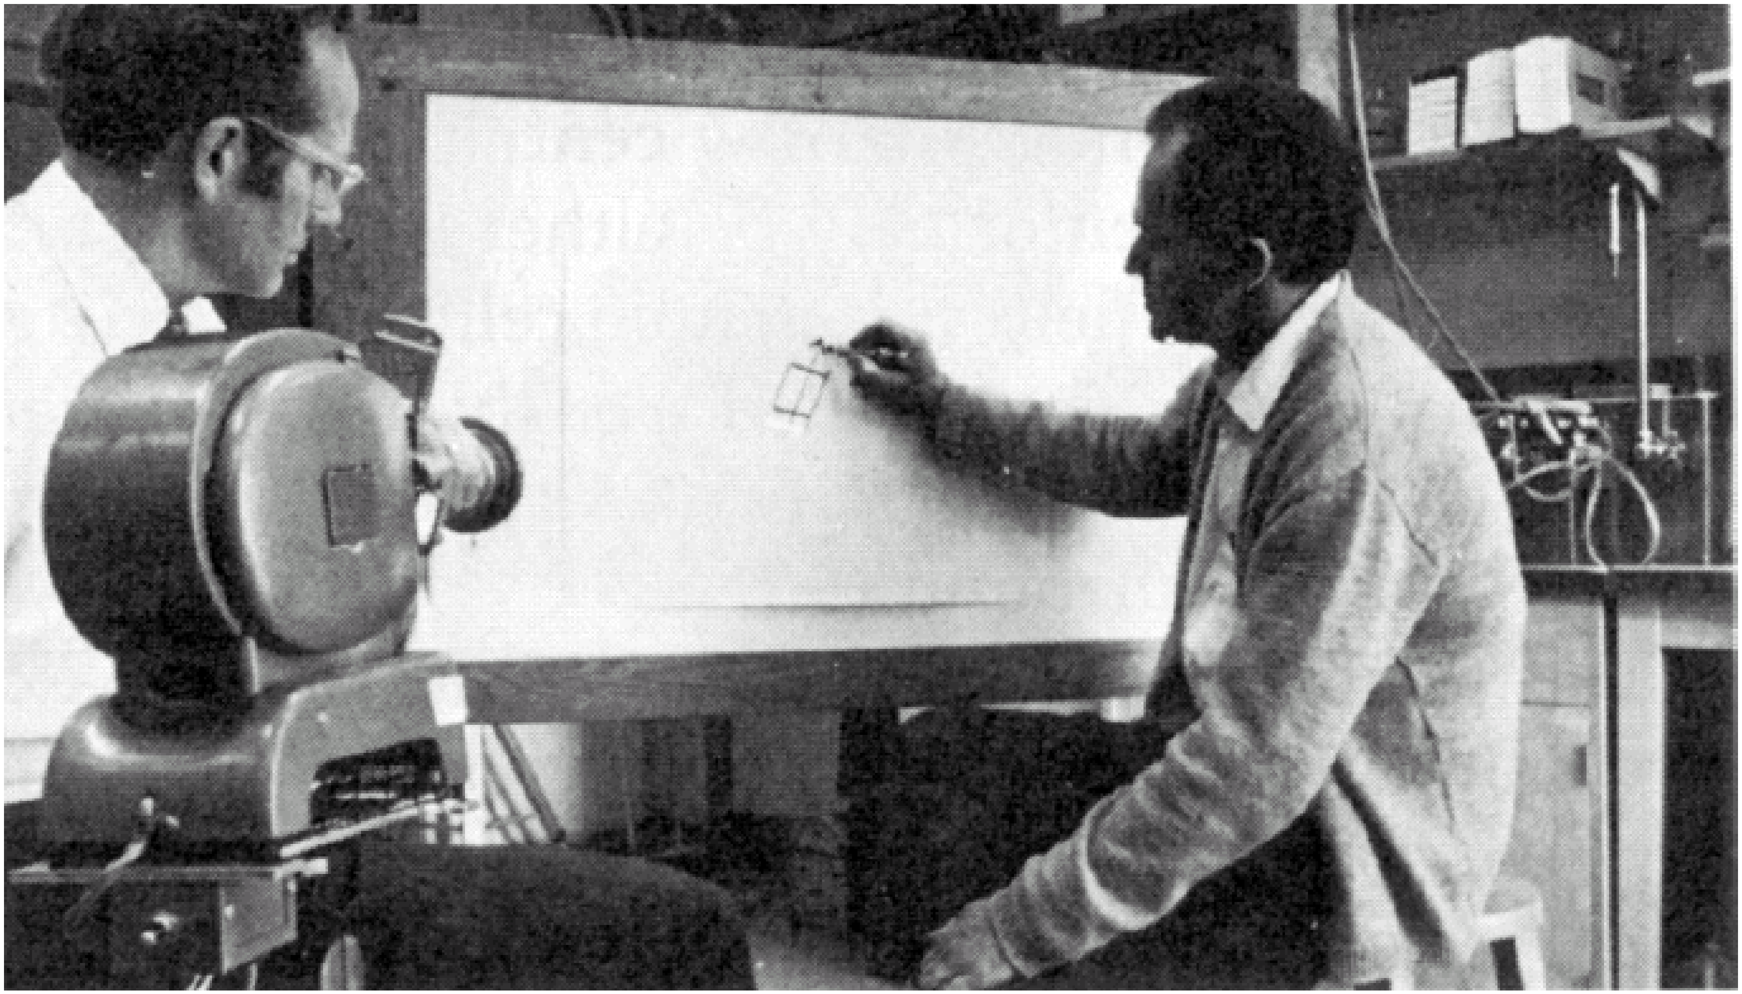

Hubel and Wiesel, 1969

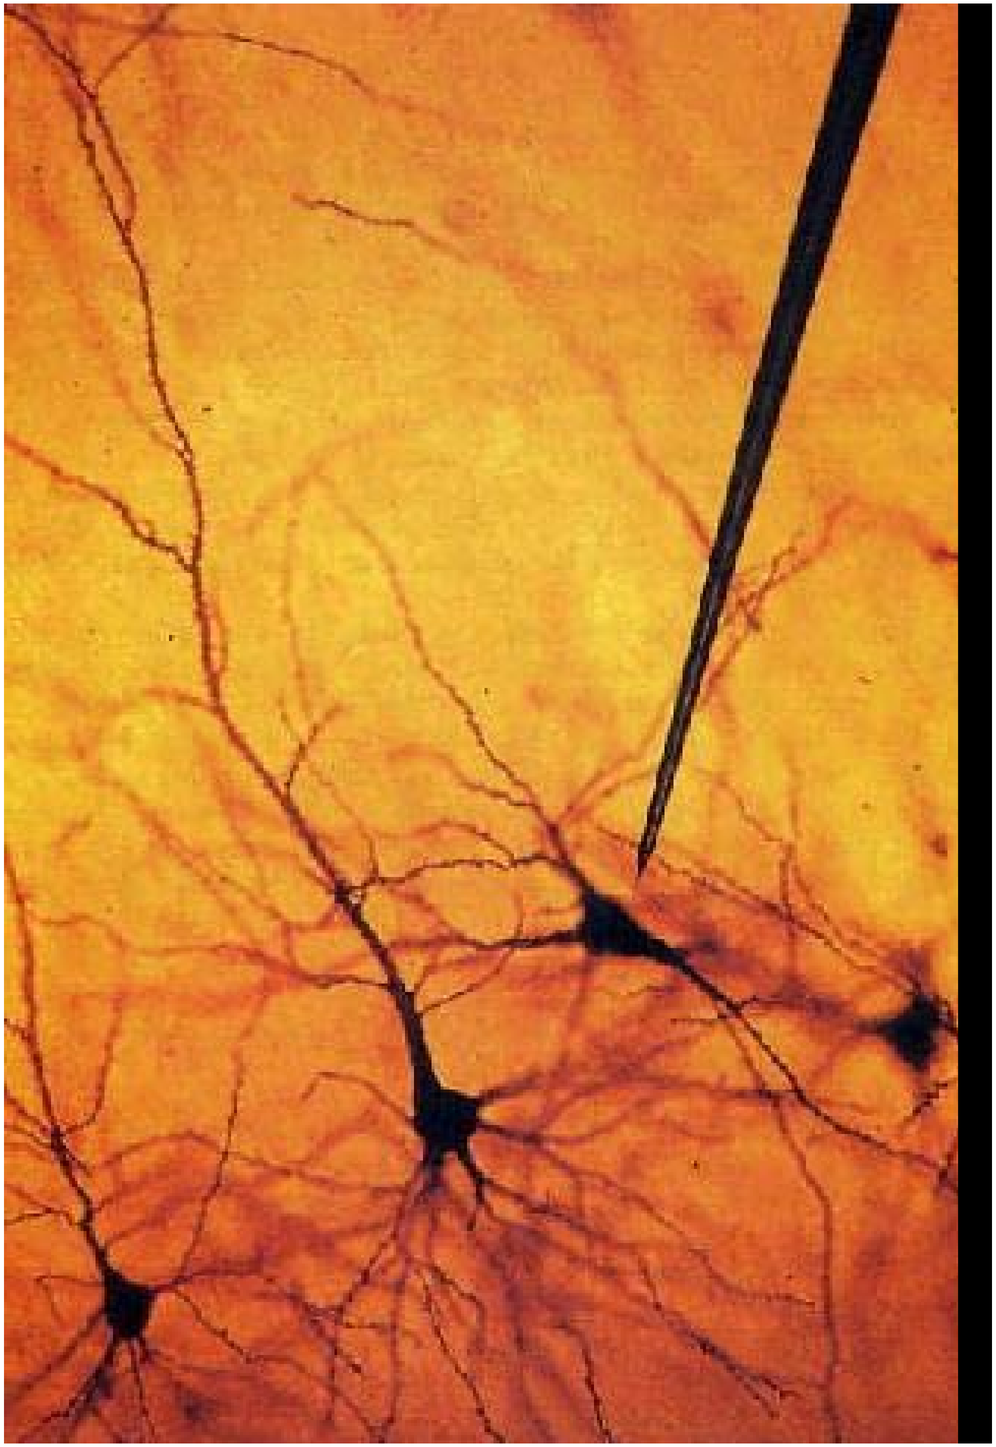

Simple Cell

Hubel and Wiesel, 1962

Complex Cell

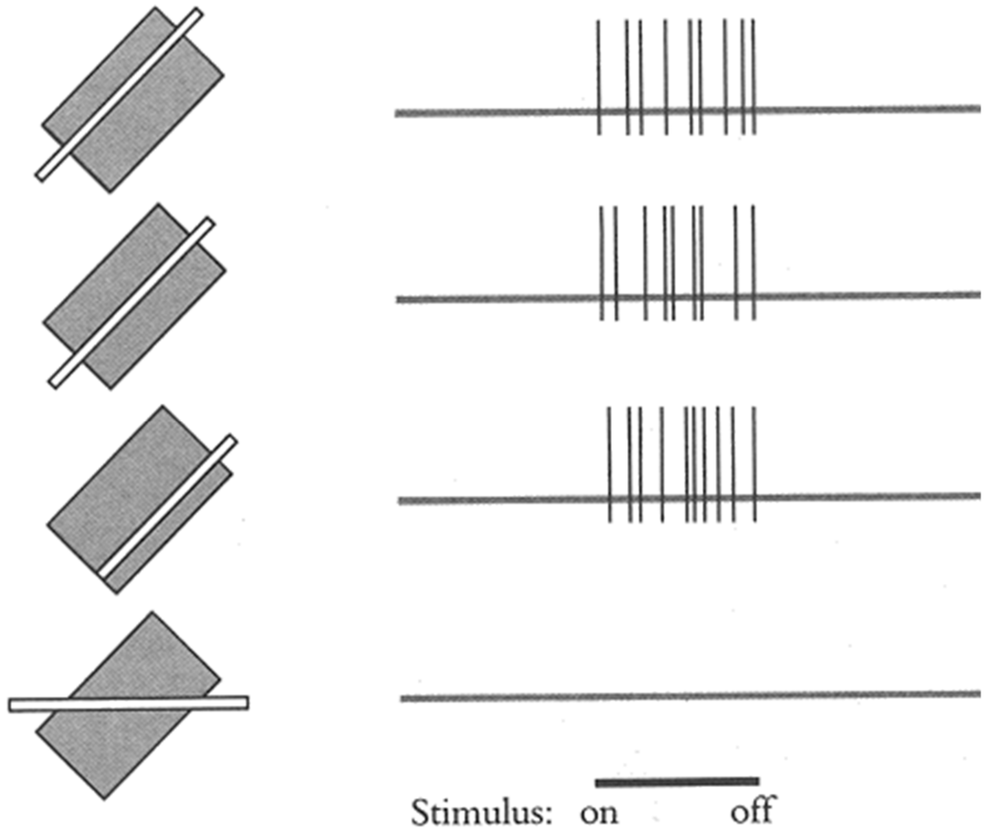

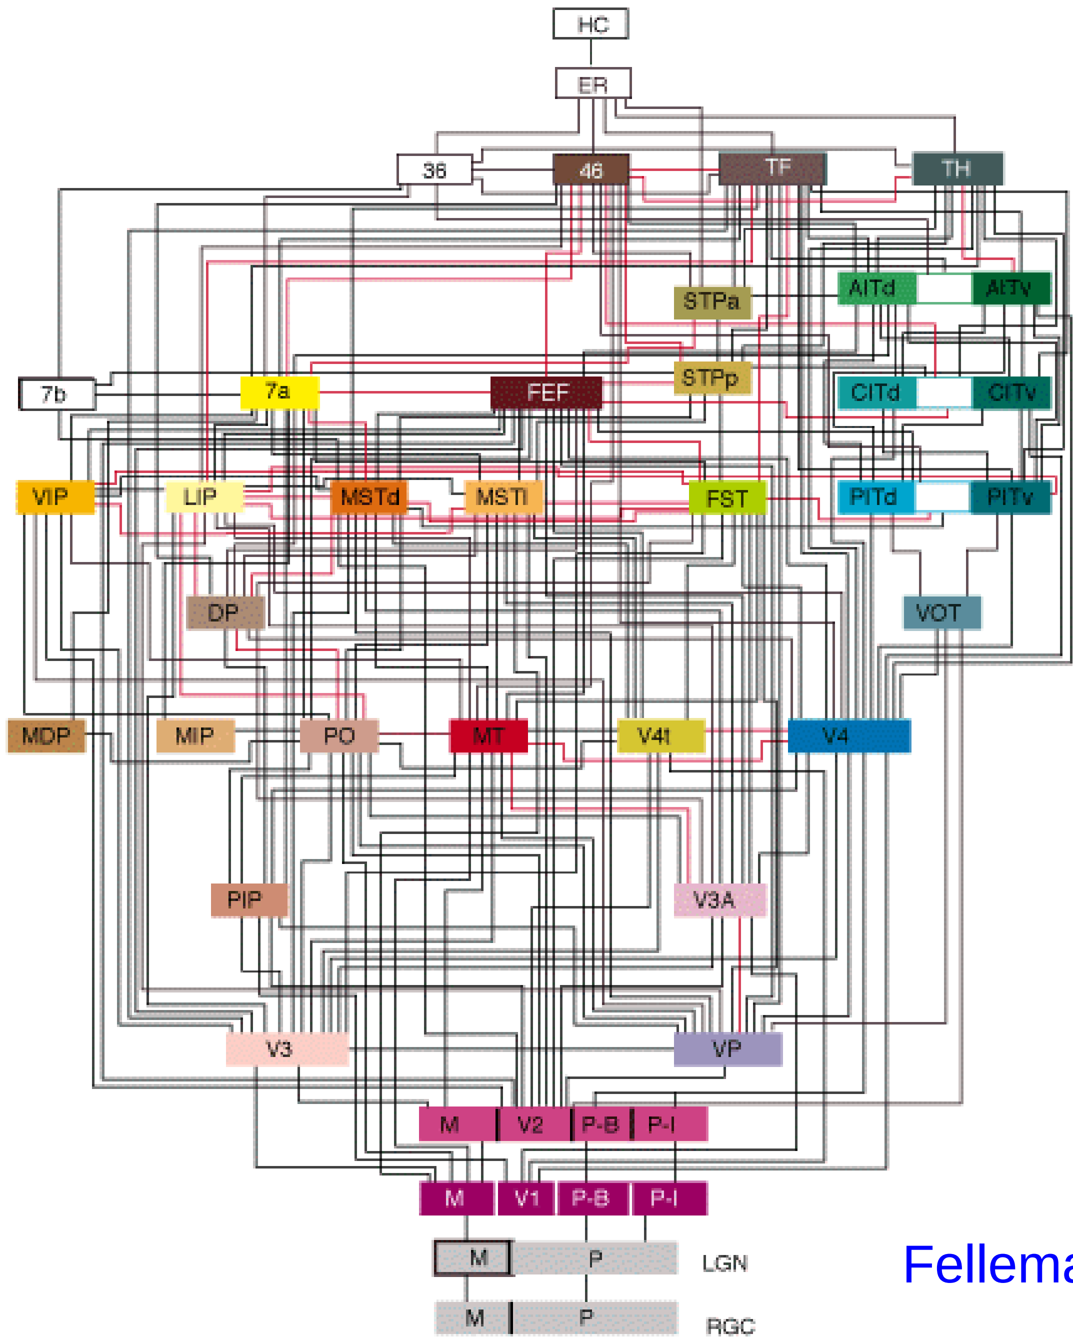

Felleman and Van Essen,1991

Geoffrey Hinton and Yann Le Cun

Deep Learning

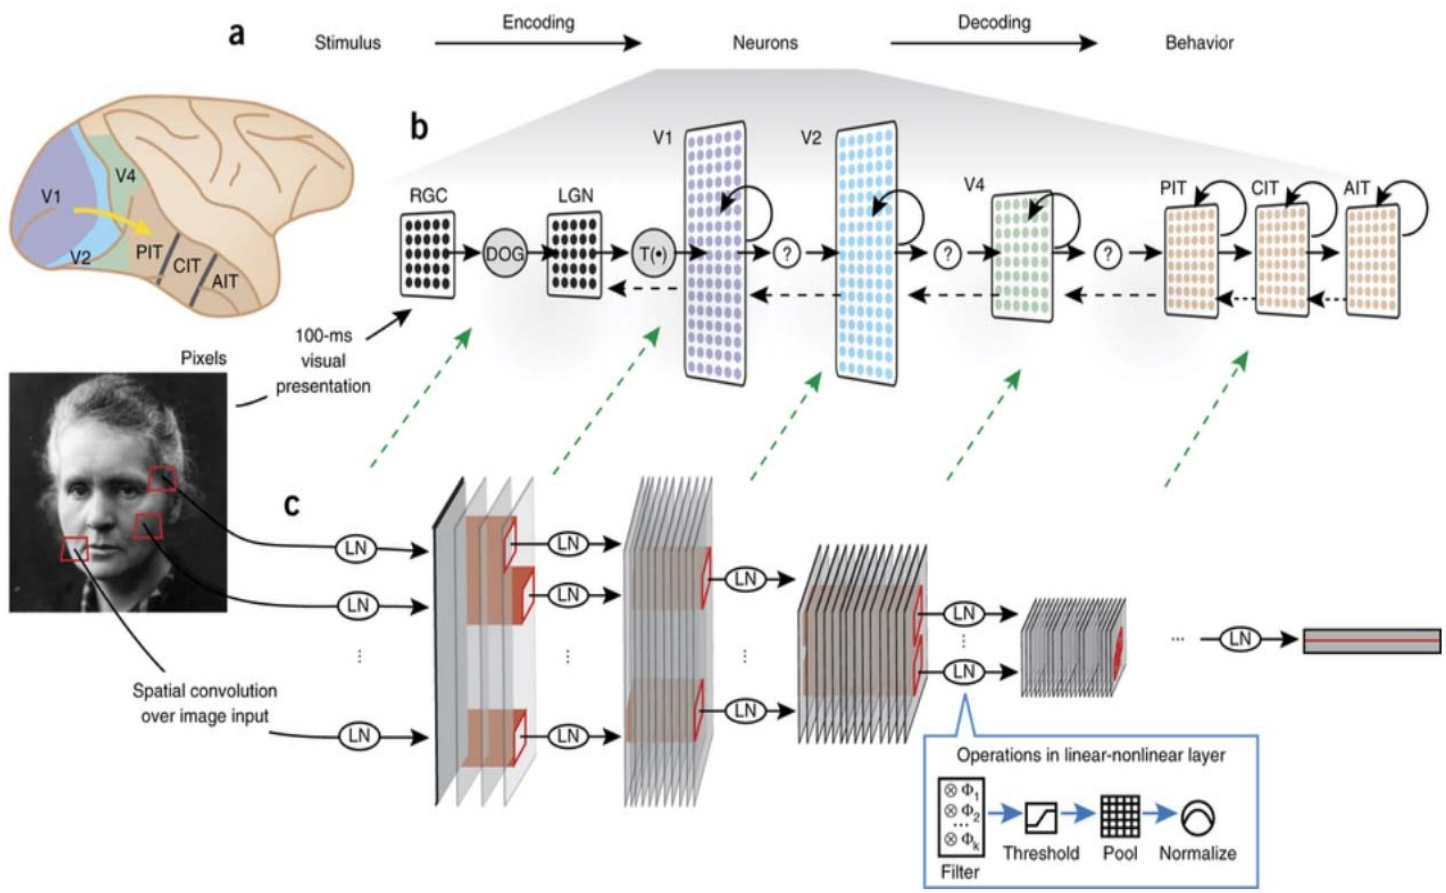

Yamins and DiCarlo, 2016

Non-convex optimization

Objective function in deep networks is non-convex

- May be many local minima
- Plateaus: flat regions
- Saddle points

Q: Why does SGD seem to work so well for optimizing these complex non-convex functions??

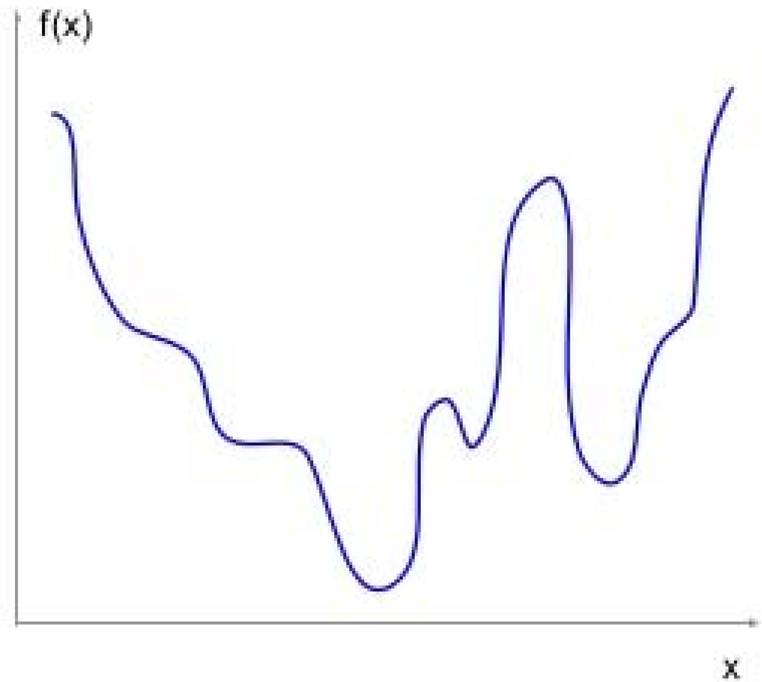

Saddle Points in High Dimensions

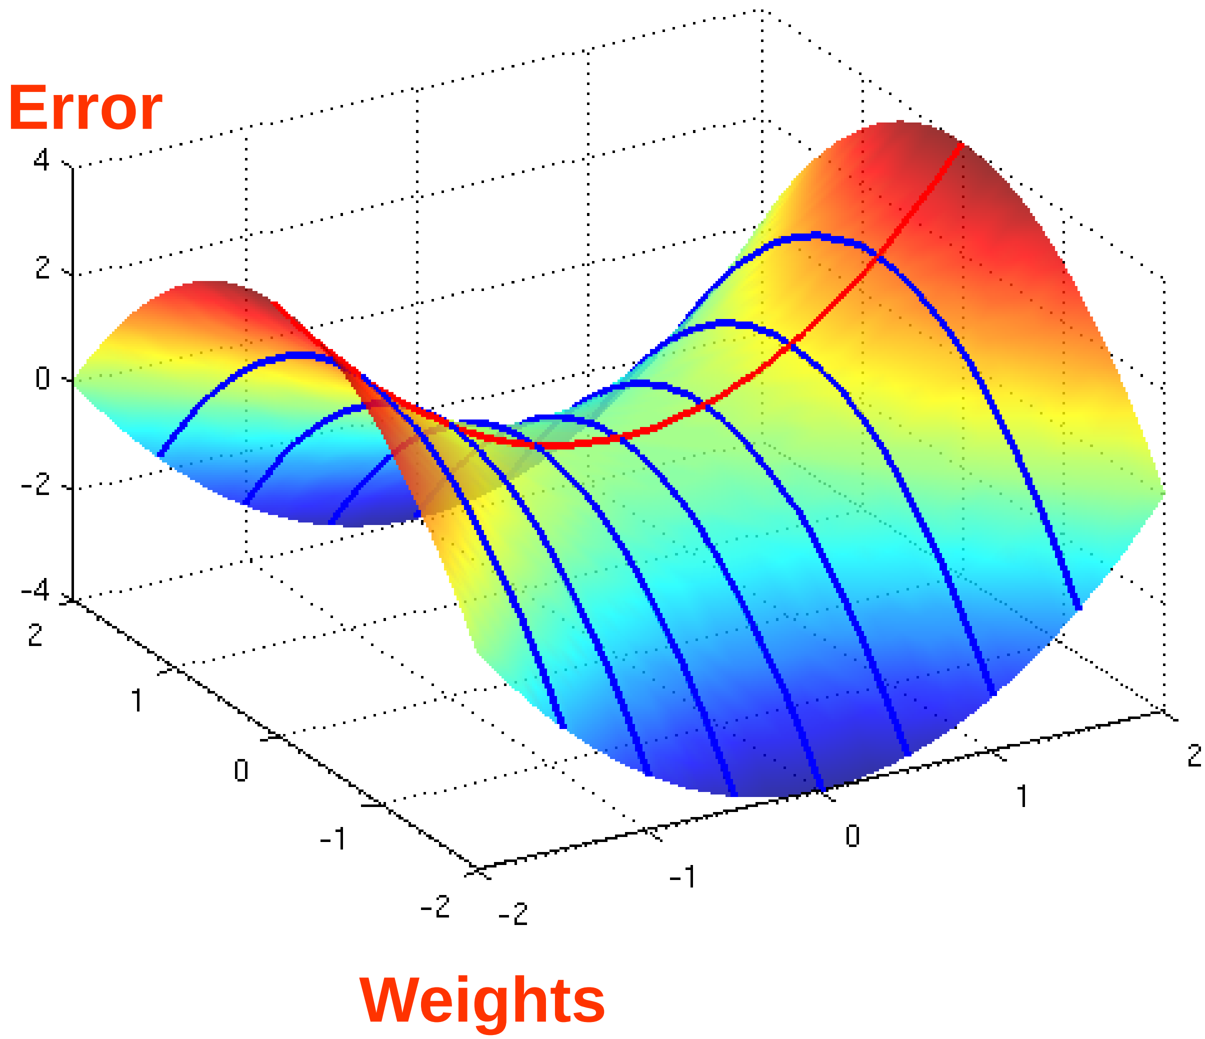

January 5, 2017: “After humanity spent thousands of years improving our tactics, computers tell us that humans are completely wrong.”

RUBI

Javier Movellan

Neuromorphic Engineering

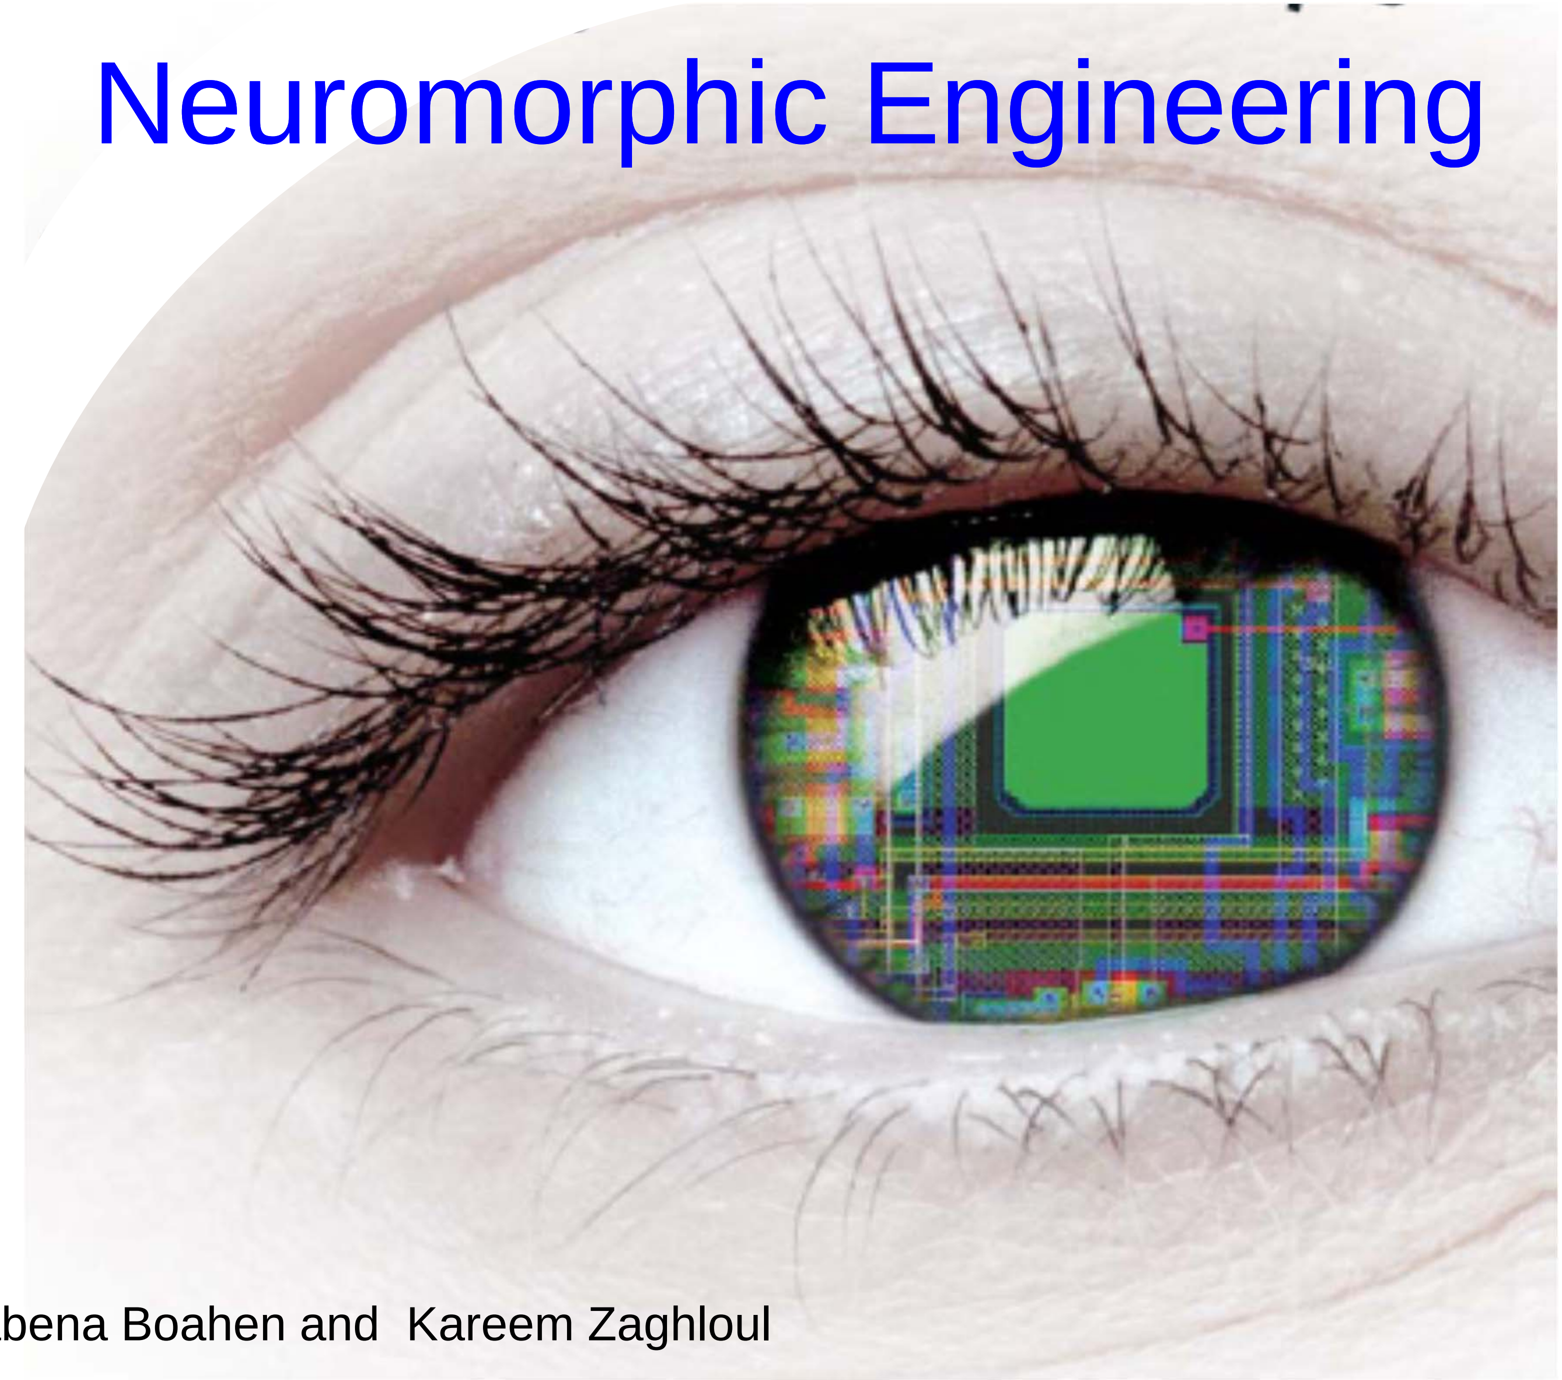

Kwabena Boahen and Kareem Zaghloul

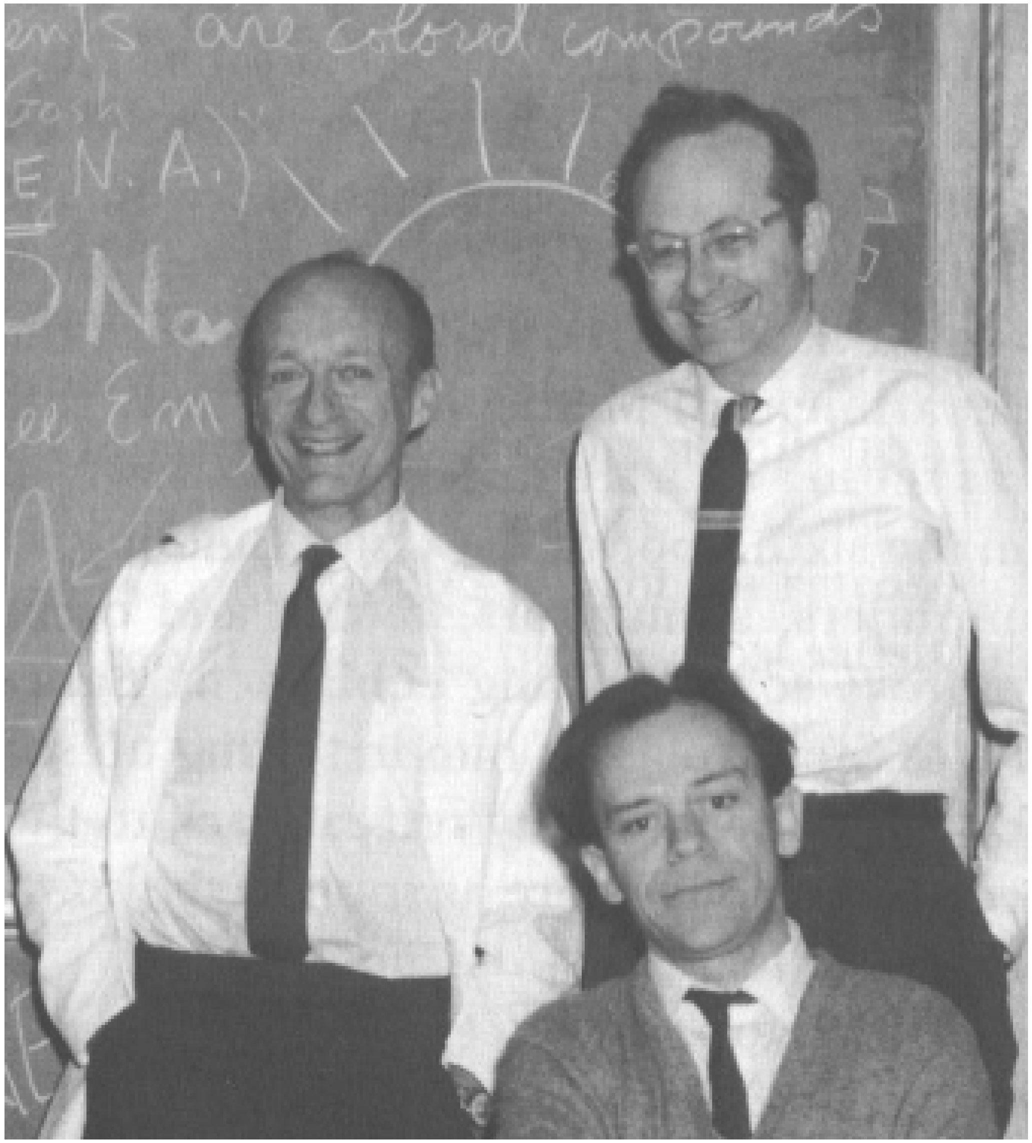
ents are colored compounds

Retinal Ganglion Cells

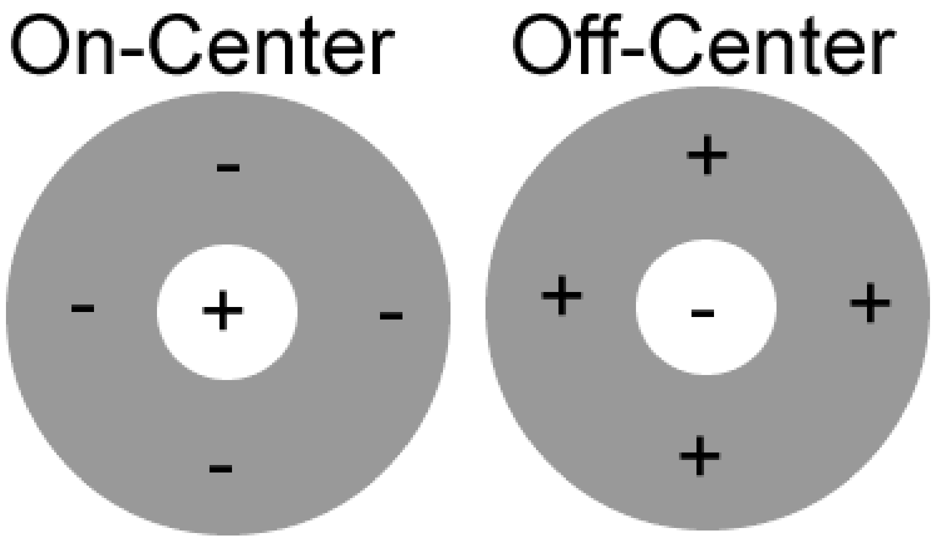

Kuffler, 1953

Neuromorphic Camera

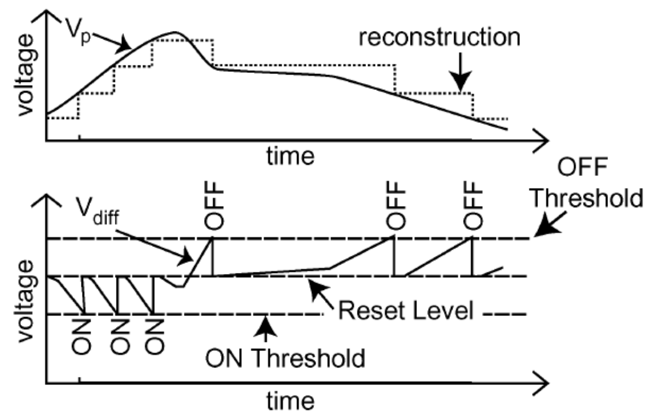

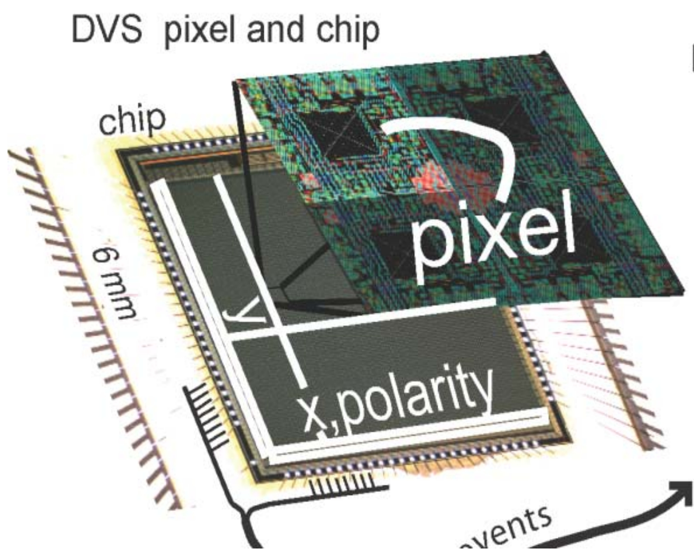

Lichtsteiner, Posch and Delbruck, 2008

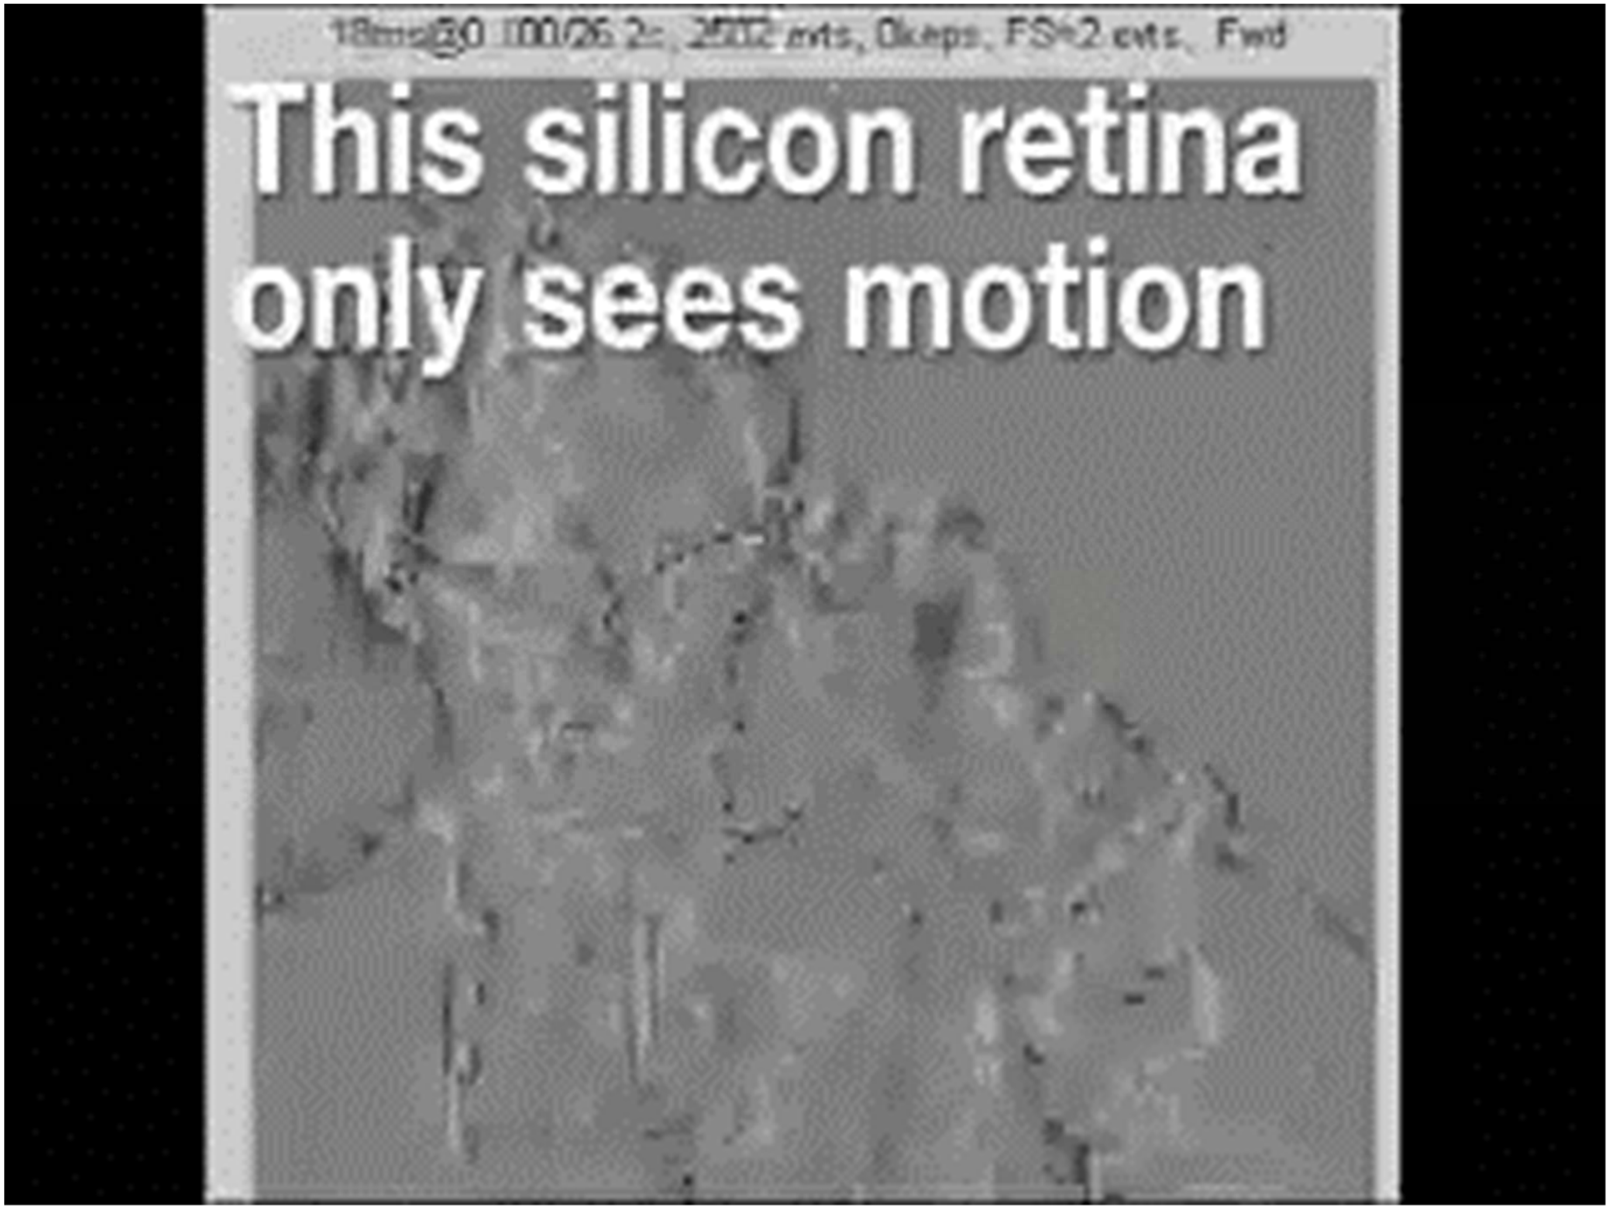
This silicon retina
only sees motion

BRAIN Initiative

Brain Research through Advancing Innovative Neurotechnologies

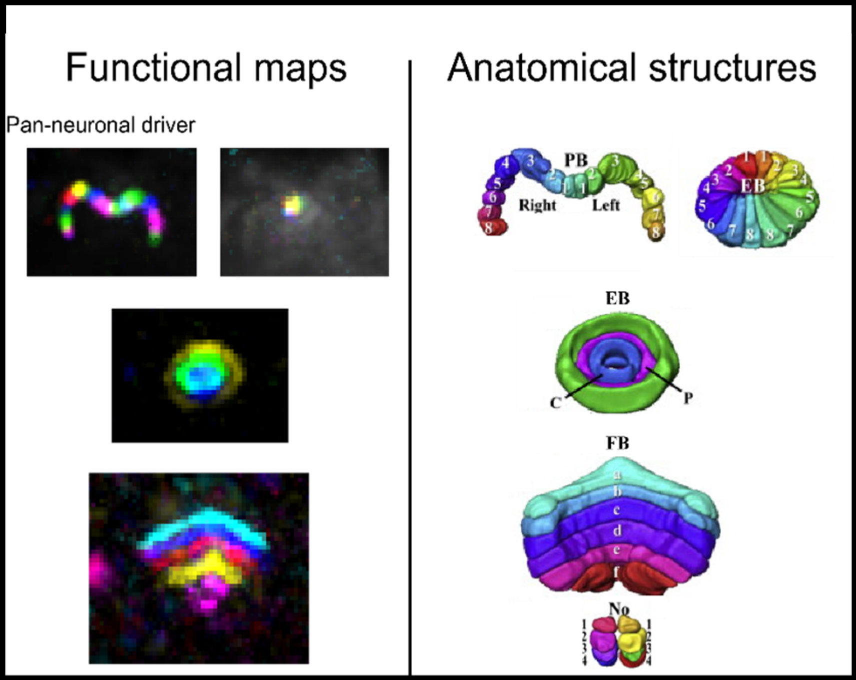
Functional maps
Pan-neuronal driver
Anatomical structures
PB
Right
Left
EB
C
P
FB
No

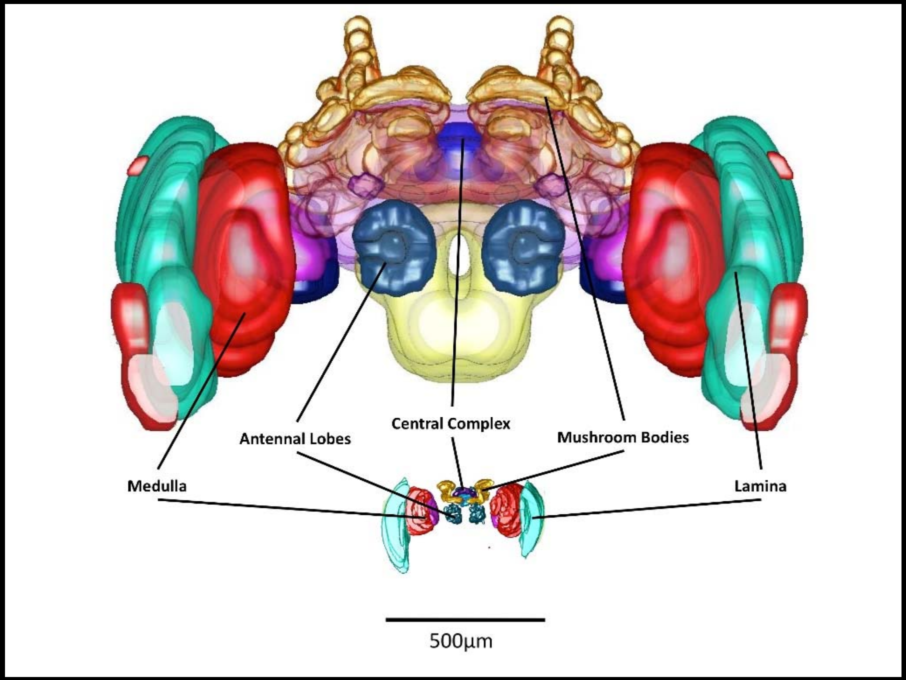
Antennal Lobes
Central Complex
Mushroom Bodies
Medulla
Lamina
500µm

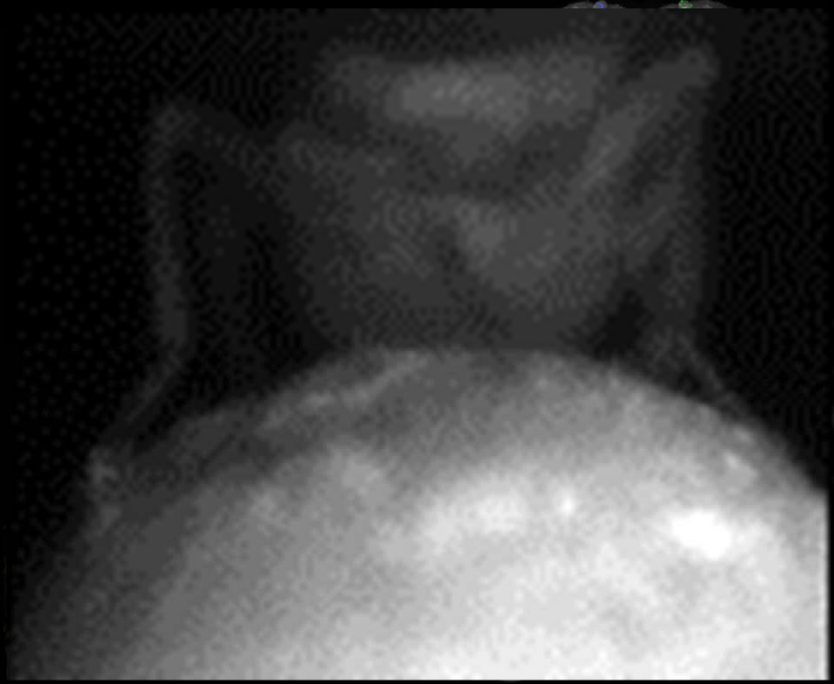

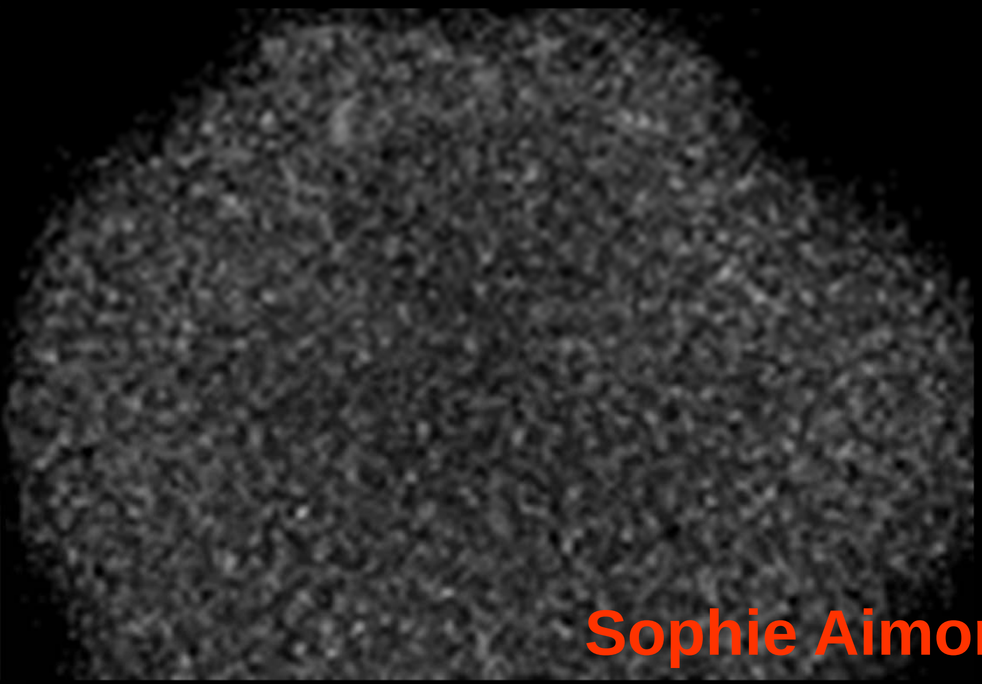
Sophie Aimon

BRAIN Initiative


The Deep Learning Revolution
Terrence Sejnowski
Salk Institute
UCSD