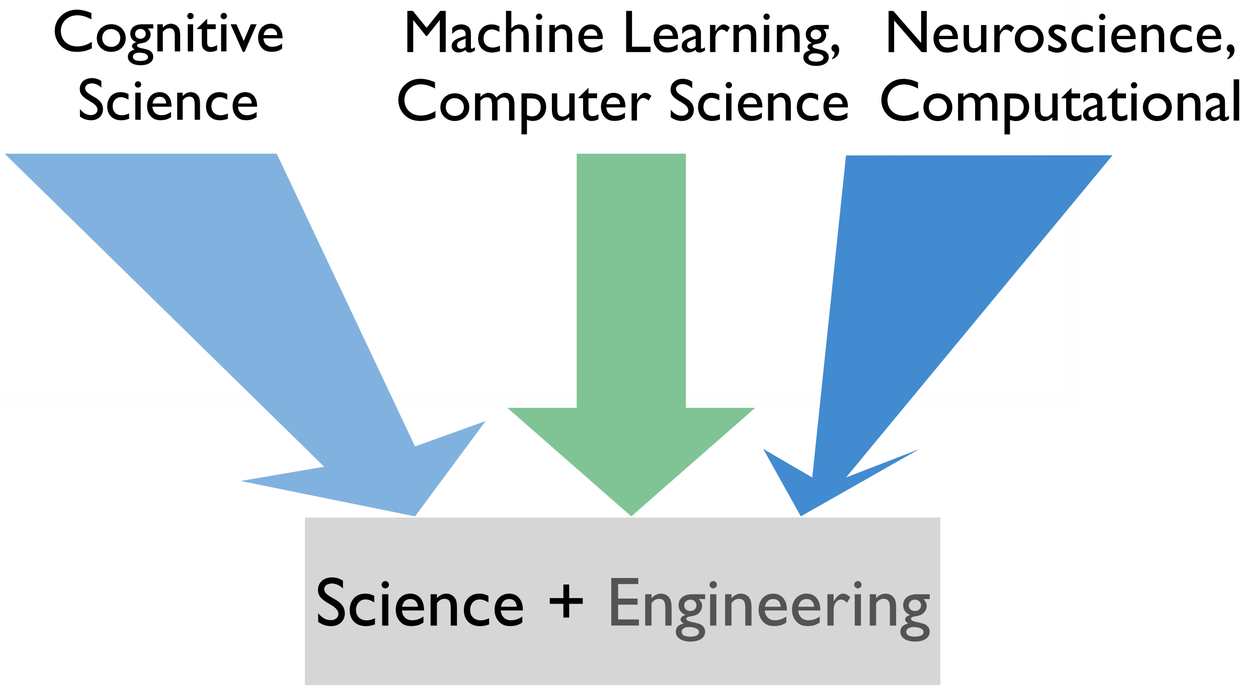

CENTER FOR
Brains
Minds+
Machines

CBMM: the Science and Engineering of Intelligence

CENTER FOR
Brains
Minds+
Machines

The Center for Brains, Minds and Machines (CBMM) is a multi-institutional NSF Science and Technology Center dedicated to the study of intelligence - how the brain produces intelligent behavior and how we may be able to replicate intelligence in machines.

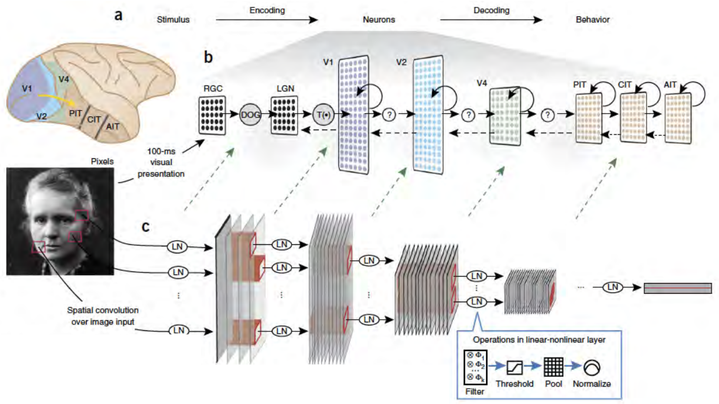
Funding 2013-2023	**~$50M**
Research Institutions	**~4**
Educational Institutions	**12**
Faculty (CS+BCS+…)	**~23**
Researchers	**~100**
Publications	**~500**



CBMM's focus is the Science and the Engineering of Intelligence

CENTER FOR Brains Minds+ Machines

We aim to make progress in understanding the greatest of all problems in science — the problem of intelligence. This means understanding how the brain makes the mind, how the brain works and how to build intelligent machines. We believe that the science of intelligence will enable better engineering of intelligence in the long term.











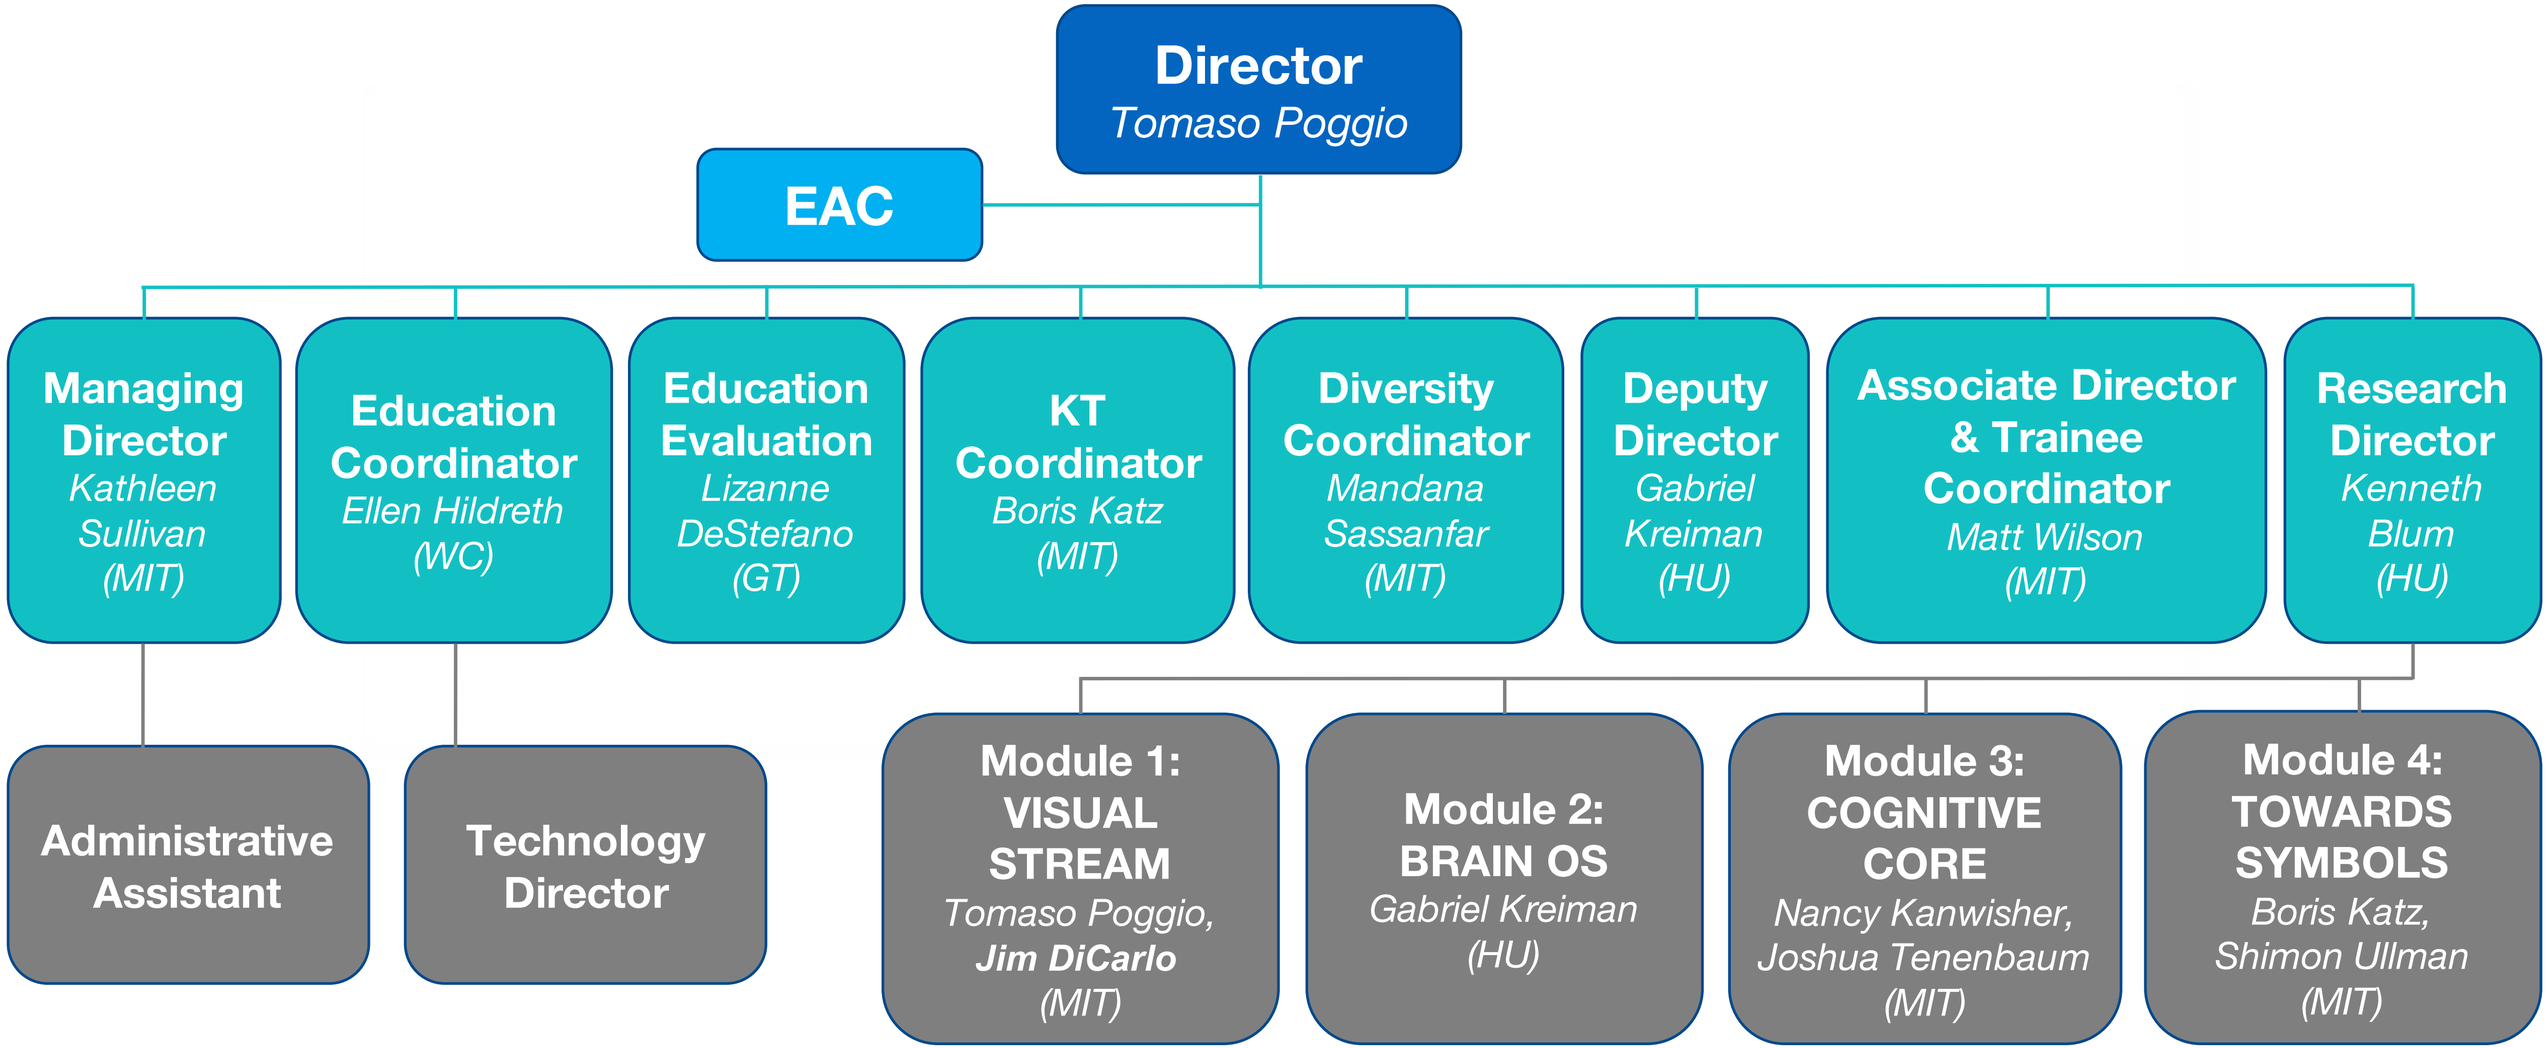
Key recent advances in the engineering of intelligence have their roots in basic research on the brain

The problem of intelligence is the greatest problem in science

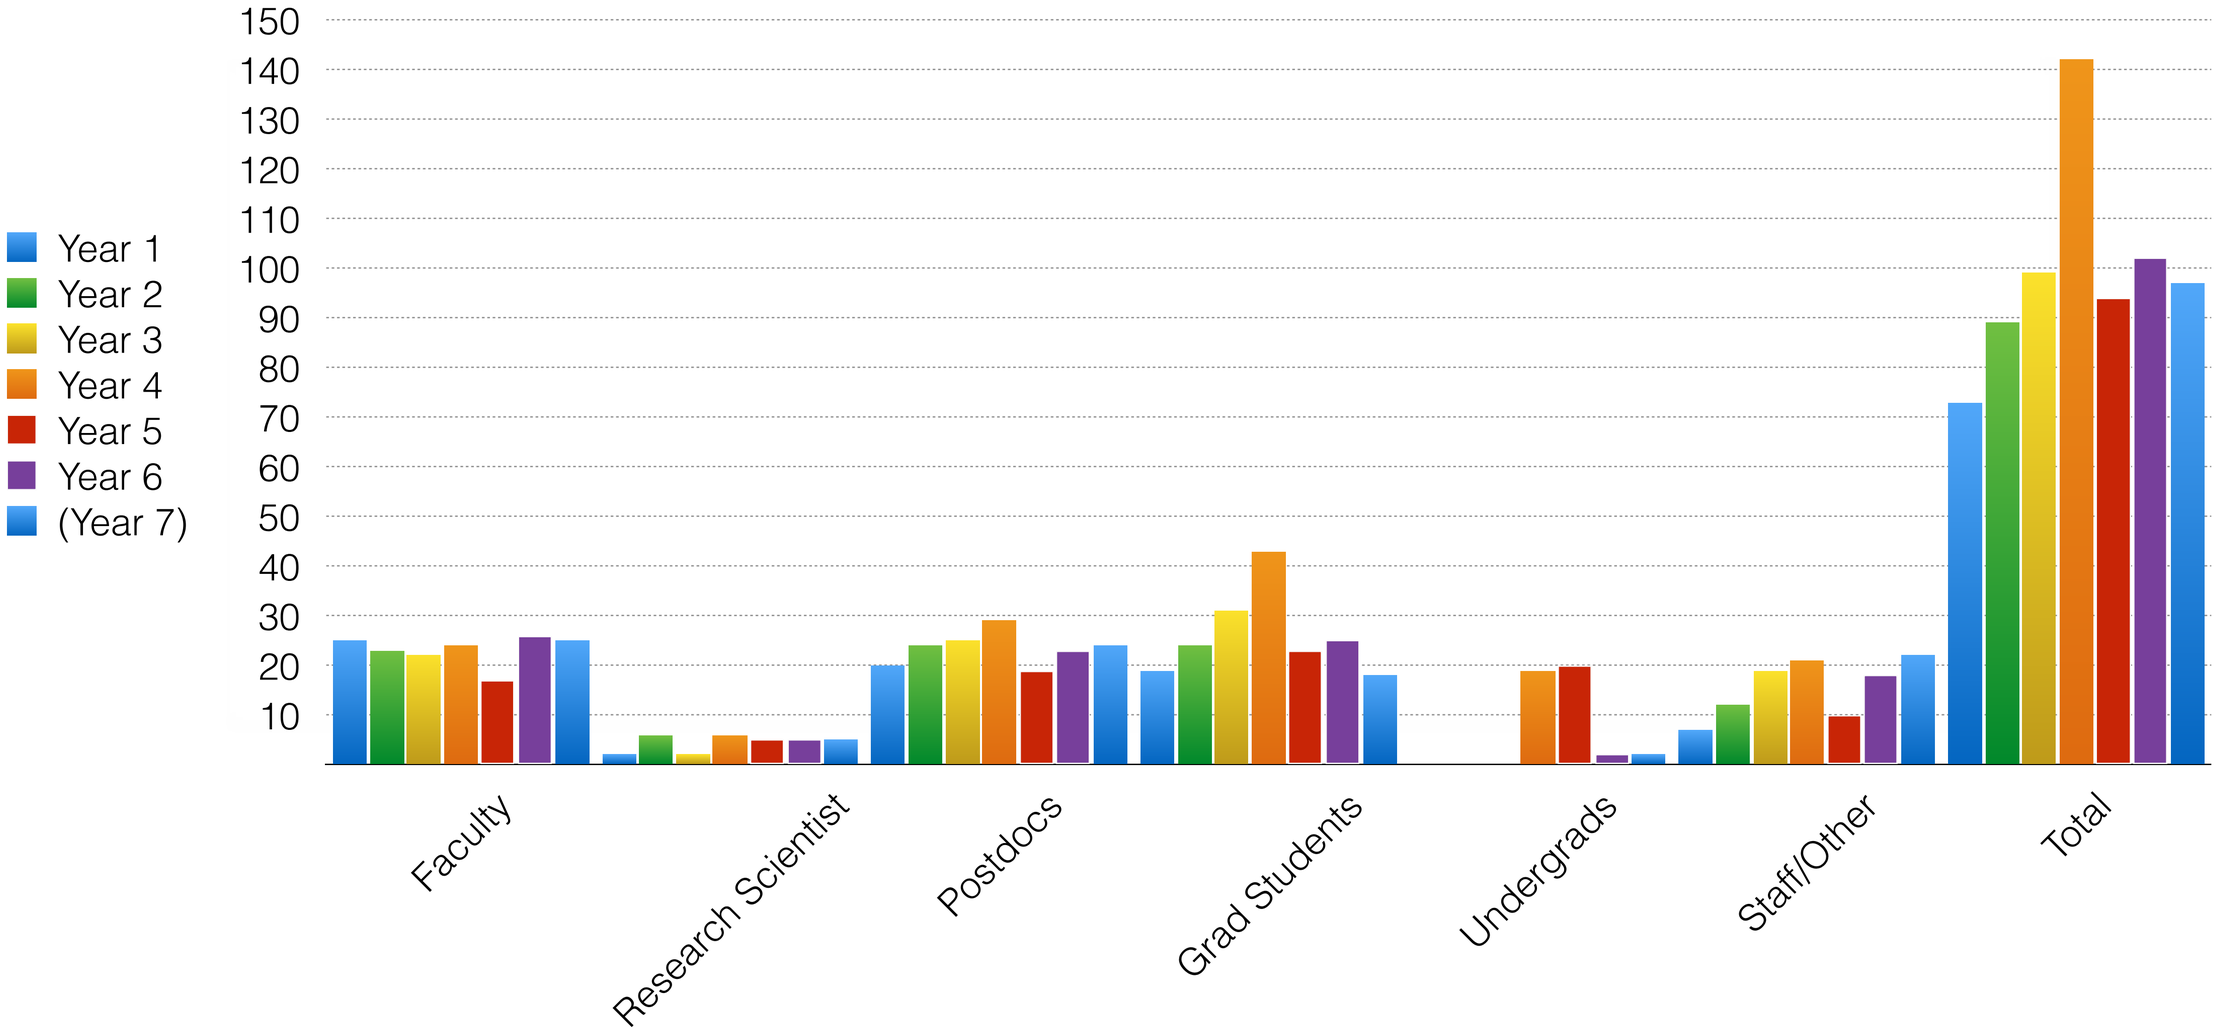
The CBMM bet (different from Deep Mind):

understand how the brain works,
(then) make intelligent machines

CBMM Organizational Chart (future)



CBMM Participants



CENTER FOR
Brains
Minds+
Machines

EAC



Demis Hassabis, *DeepMind*

Charles Isbell, Jr., *Georgia Tech*

Christof Koch, *Allen Institute*

Fei-Fei Li, *Stanford*

Lore McGovern, *MIBR, MIT*

Joel Oppenheim, *NYU*

Pietro Perona, *Caltech*

Marc Raibert, *Boston Dynamics*

Judith Richter, *Medinol*

Kobi Richter, *Medinol*

Amnon Shashua, *Mobileye*

David Siegel, *Two Sigma*

Susan Whitehead, *MIT Corporation*

Jim Pallotta, *The Raptor group*

CENTER FOR
Brains
Minds+
Machines

Research, Education & Diversity Partners

MIT

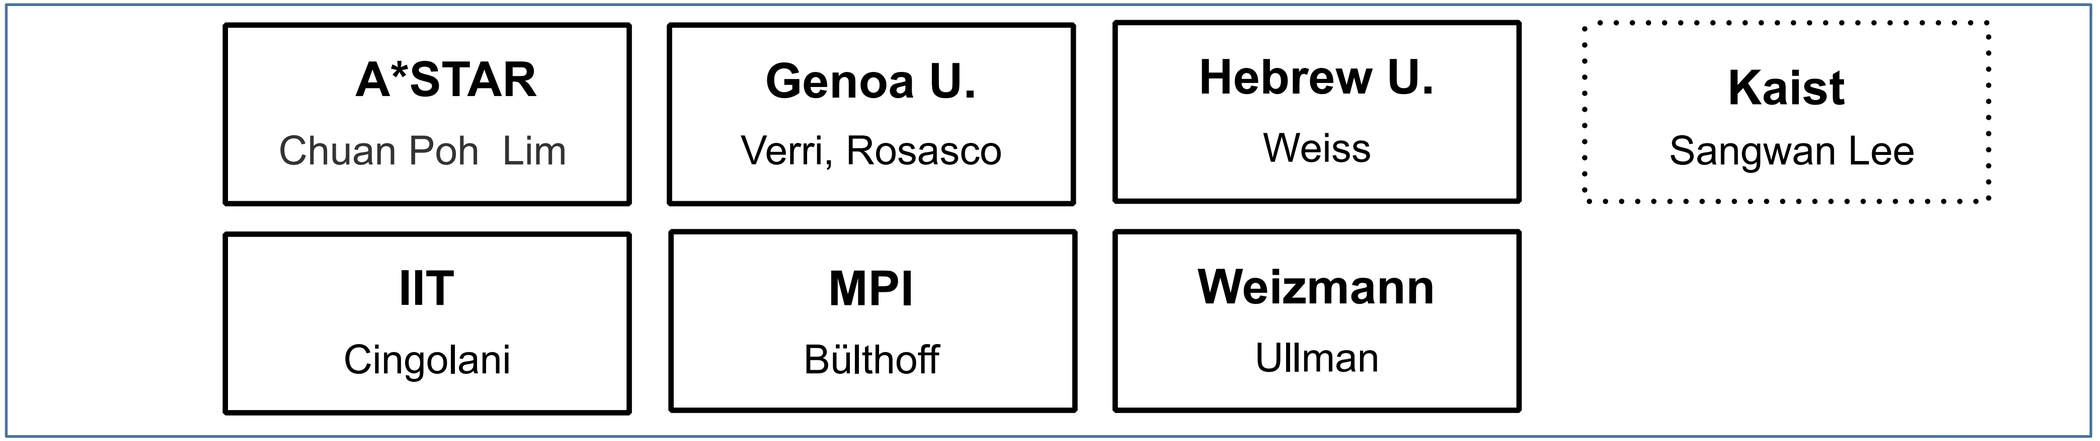
Boyden, Desimone, DiCarlo, Kaelbling, Kanwisher, Katz, McDermott, Oliva, Poggio, Roy, Sassanfar, Saxe, Schulz, Tegmark, Tenenbaum, Ullman, Wilson, Torralba

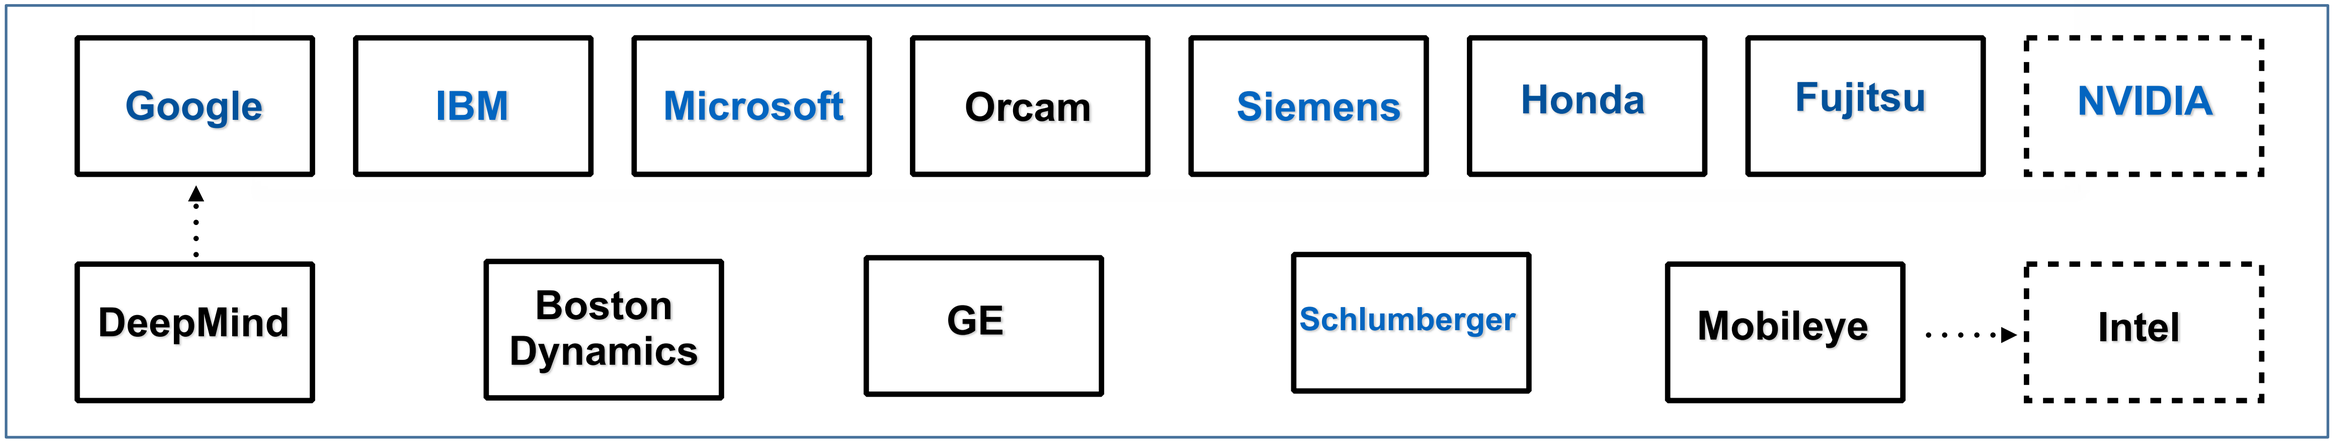
Harvard

Blum, Gershman, Kreiman, Livingstone, Sompolinsky, Spelke

Boston Children's Hospital Kreiman	**Florida International U.** Finlayson	**Harvard Medical School** Kreiman, Livingstone	**Howard U.** Chouika, Manaye, Rwebangira, Salmani
Hunter College Chodorow, Epstein, Sakas, Zeigler	**Johns Hopkins U.** Isik	**Queens College** Brumberg	**Rockefeller U.** Freiwald
Museum of Science, Boston	**Stanford U.** Goodman	**Universidad Central Del Caribe (UCC)** Jorquera	**University of Central Florida** McNair Program
UMass Boston Blaser, Ciaramitaro, Pomplun, Shukla	**UPR - Mayagüez** Santiago, Vega-Riveros	**UPR – Río Piedras** Garcia-Arraras, Maldonado-Vlaar, Megret, Ordóñez, Ortiz-Zuazaga	**Wellesley College** Hildreth, Wiest, Wilmer

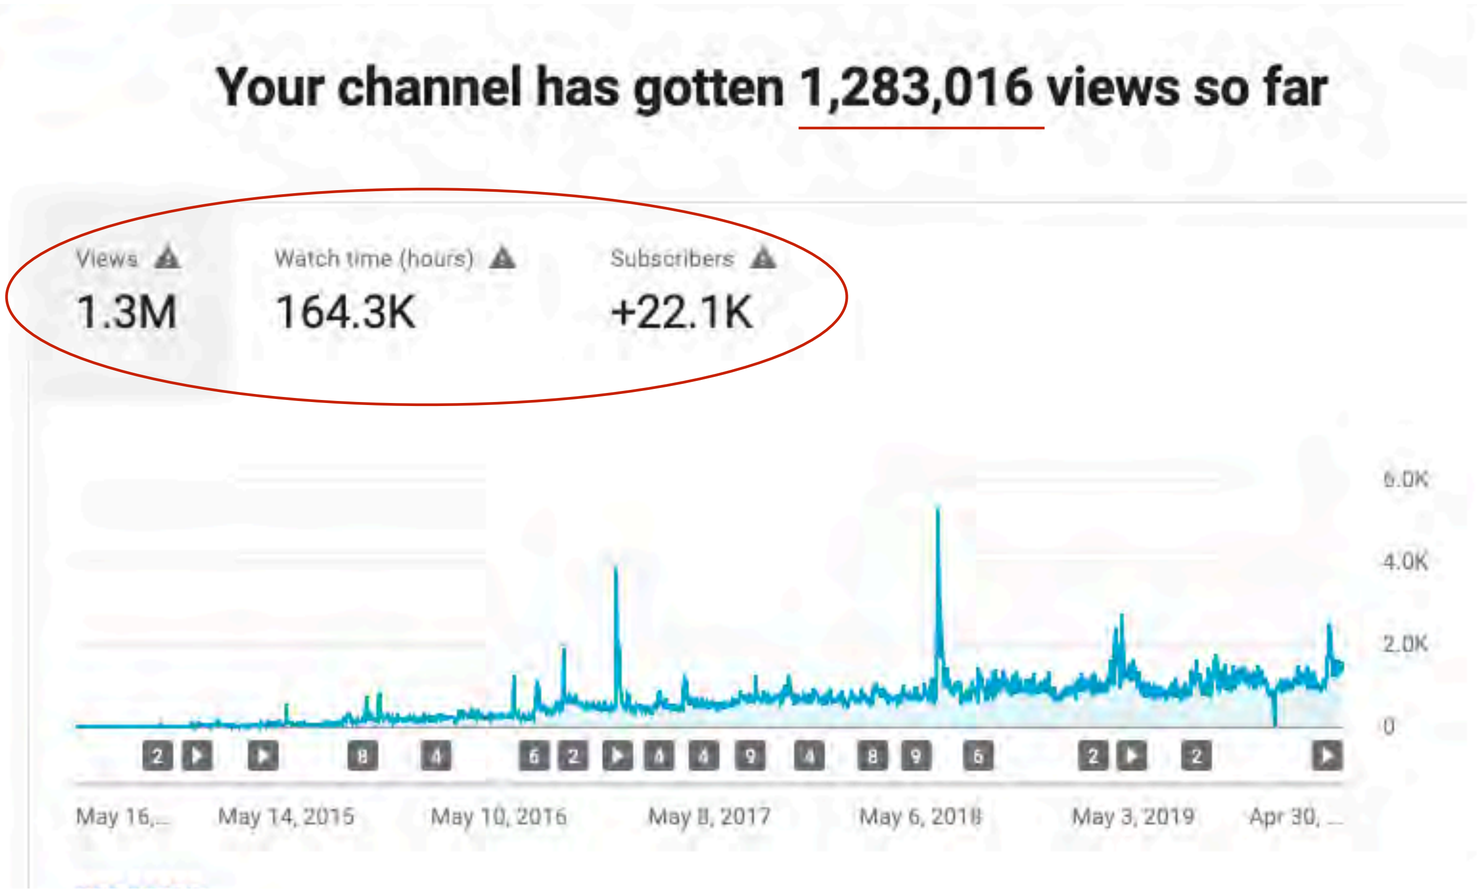
International and Corporate Partners

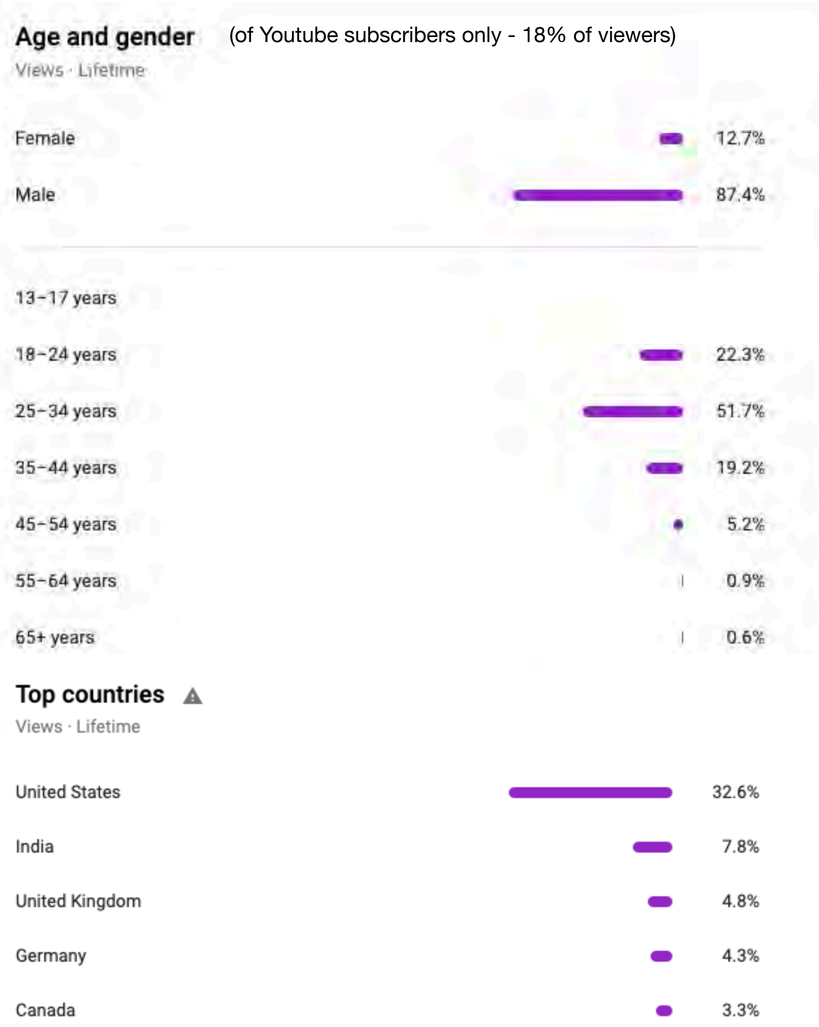




CENTER FOR
Brains
Minds+
Machines

Videos - ~950 (May 2014 - April 2020)





CENTER FOR
Brains
Minds+
Machines



Ellen Hildreth

CENTER FOR
Brains
Minds+
Machines

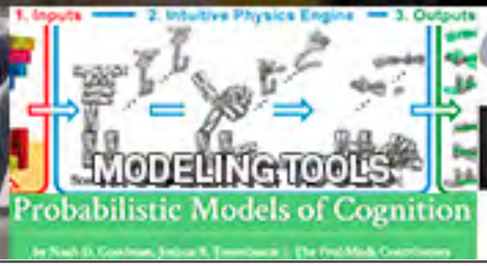
CBMM, NSF STC » EDUCATION » LEARNING HUB

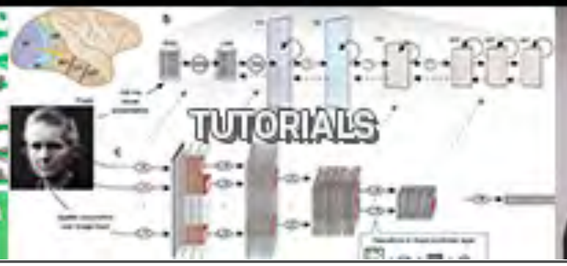
SCIENCE OF INTELLIGENCE LEARNING HUB

The Science of Intelligence Learning Hub enables educators, researchers, and learners to expand their knowledge and skills related to the interdisciplinary study of intelligence.











Diversity Program



Mandana Sassanfar

Code, Software and Datasets



Subtasks of Unconstrained Face Recognition

You can get the data from this page. Please fill out

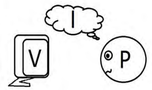
PredNet

Code and models accompanying "Deep Predictive Coding Networks for Video Prediction and Unsupervised Learning"

Bill Lotter, Gabriel Kreiman, and David Cox (2016)

VIP: A unifying framework for eye-gaze research

Framework



Eye-gaze is affected by the 3 factors: Visual stimulus, Intent and Person.

The eye-gaze data, E, is defined as:

$E = g(\{(t_i\ x_i\ y_i\ p_i\ q_i\ d_i\ s_i\ b_i\ c_i)_{L/R}\})$

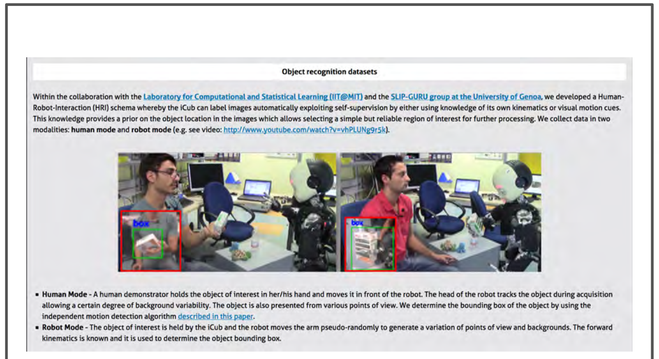
where g is a function of a sequence of eye-gaze data which

Replication Data for: The dynamics of invariant object recognition in the human visual system

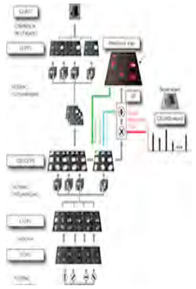
Isik, Leyla, 2016, "Replication Data for: The dynamics of invariant object recognition in the human visual system", doi:10.7910/DVN/KRUPXZ, Harvard Dataverse, V1

Cite Dataset

Learn about Data Citation Standards.



There's Waldo! A Normalization Model of Visual Search Predicts Single-Trial Human Fixations in an Object Search Task

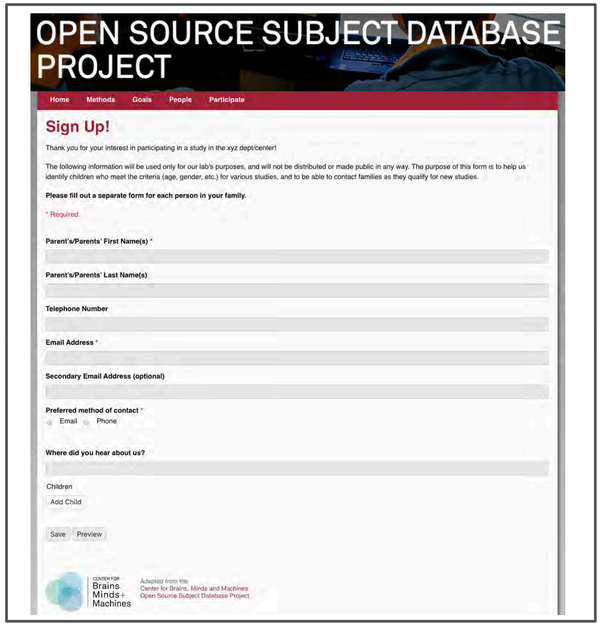


Thomas Miconi, Laura Groomes and Gabriel Kreiman

The Neural Decoding Toolbox

ABOUT DOWNLOADS DOCUMENTATION TUTORIALS SUPPORT LOGIN

The Neural Decoding Toolbox is a Matlab toolbox that makes it easy to apply decoding analyses to neural data



ObjectNet: A new benchmark for object recognition (in prep.)

Andrei Barbu, David Mayo, Josh Tenenbaum, Boris Katz

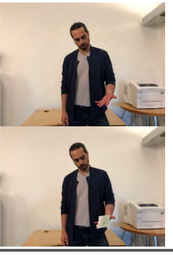
Existing object detection benchmarks overstate the performance of machines, and understate the performance of humans. We are creating a dataset that removes biases and shows that machines are far inferior to humans when detecting objects.



Partially Occluded Hands

B. Myanganbayar, C. Mata, G. Dekel, B. Katz, G. Ben-Yosef, A. Barbu

A dataset of RGB images of hands holding objects and interacting with objects. Measured human accuracy on reconstructing occluded portions of hands. People are extremely good at this task while networks are at near chance-level performance.



Summer Course at Woods Hole: Our flagship initiative

Brains, Minds & Machines Summer Course

An intensive three-week course gives advanced students a “deep” introduction to the problem of intelligence

Directors



Gabriel Kreiman



Boris Katz





Kathleen Sullivan



Kris Brewer



Ellen Hildreth



Lizanne Distefano



Kenny Blum

A *self-reproducing* community of scholars is being formed

~>300 applicants, ~30 accepted

Sponsored fellowships by GoogleX, Hidary Foundation + Fujitsu

CBMM Summer School

- Signature CBMM (Education/Knowledge Transfer) activity aimed at creating an **intergenerational community around the science and engineering of intelligence**.
- Students reported **strong influence** of lectures, working on projects, and interactions among faculty, TA's, and peers **on their own thinking and research development.**

Our vision and mission

understand how the brain works,
(then) make intelligent machines

WHY?

Recent Success Stories in AI are based on RL and DL



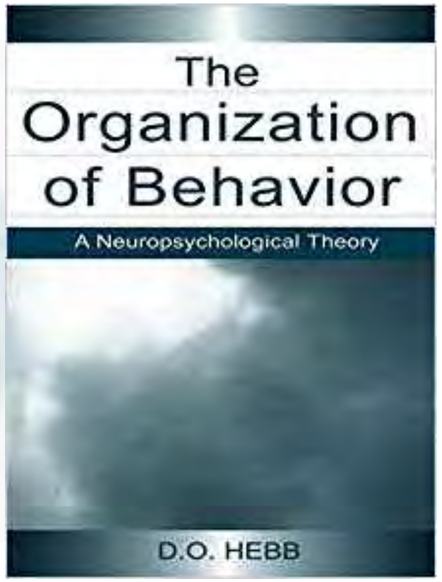


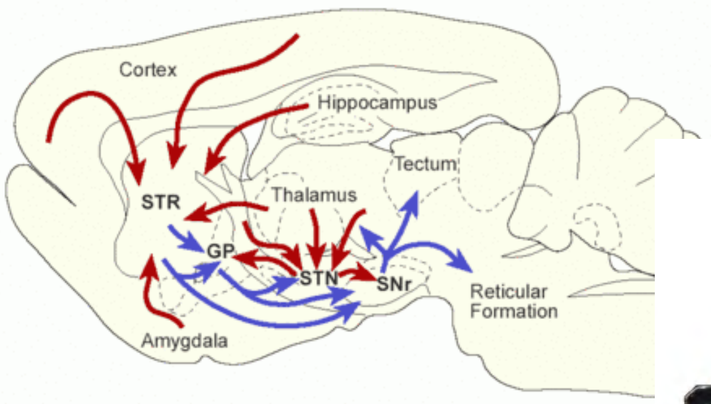
DL and RL come from neuroscience

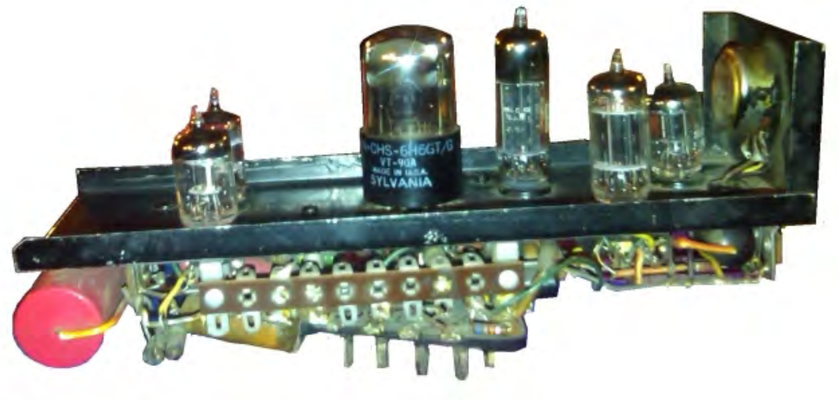
R

L







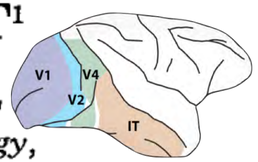


Minsky's SNARC

D

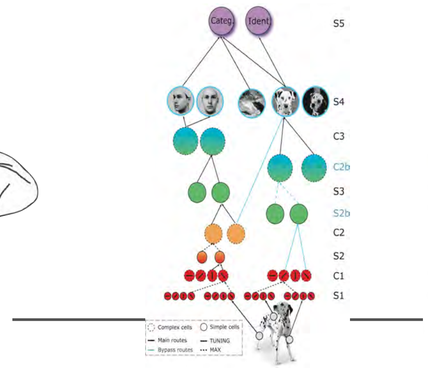
L

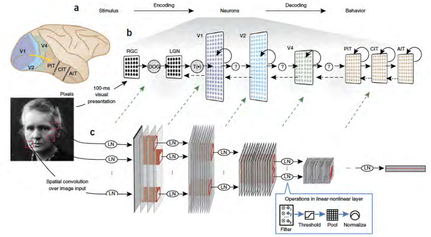
RECEPTIVE FIELDS AND FUNCTIONAL ARCHITECTURE IN TWO NONSTRIATE VISUAL AREAS (18 AND 19) OF THE CAT[1]



DAVID H. HUBEL AND TORSTEN N. WIESEL
Neurophysiology Laboratory, Department of Pharmacology, Harvard Medical School, Boston, Massachusetts

(Received for publication August 24, 1964)







Brains
Minds+
Machines

Vision for the BMM SummerSchool

We focus on the combination of neuroscience and engineering to make progress on the problem of intelligence because, as in the recent past, it is likely that several of the next breakthroughs in ML and AI are likely to come from neuroscienceANDengineering

A quick recap of 40 of the last ~50 years of neuroscience and ML, through my eyes
1972-2013

Tuebingen, MPI fuer BK (1972-1981)



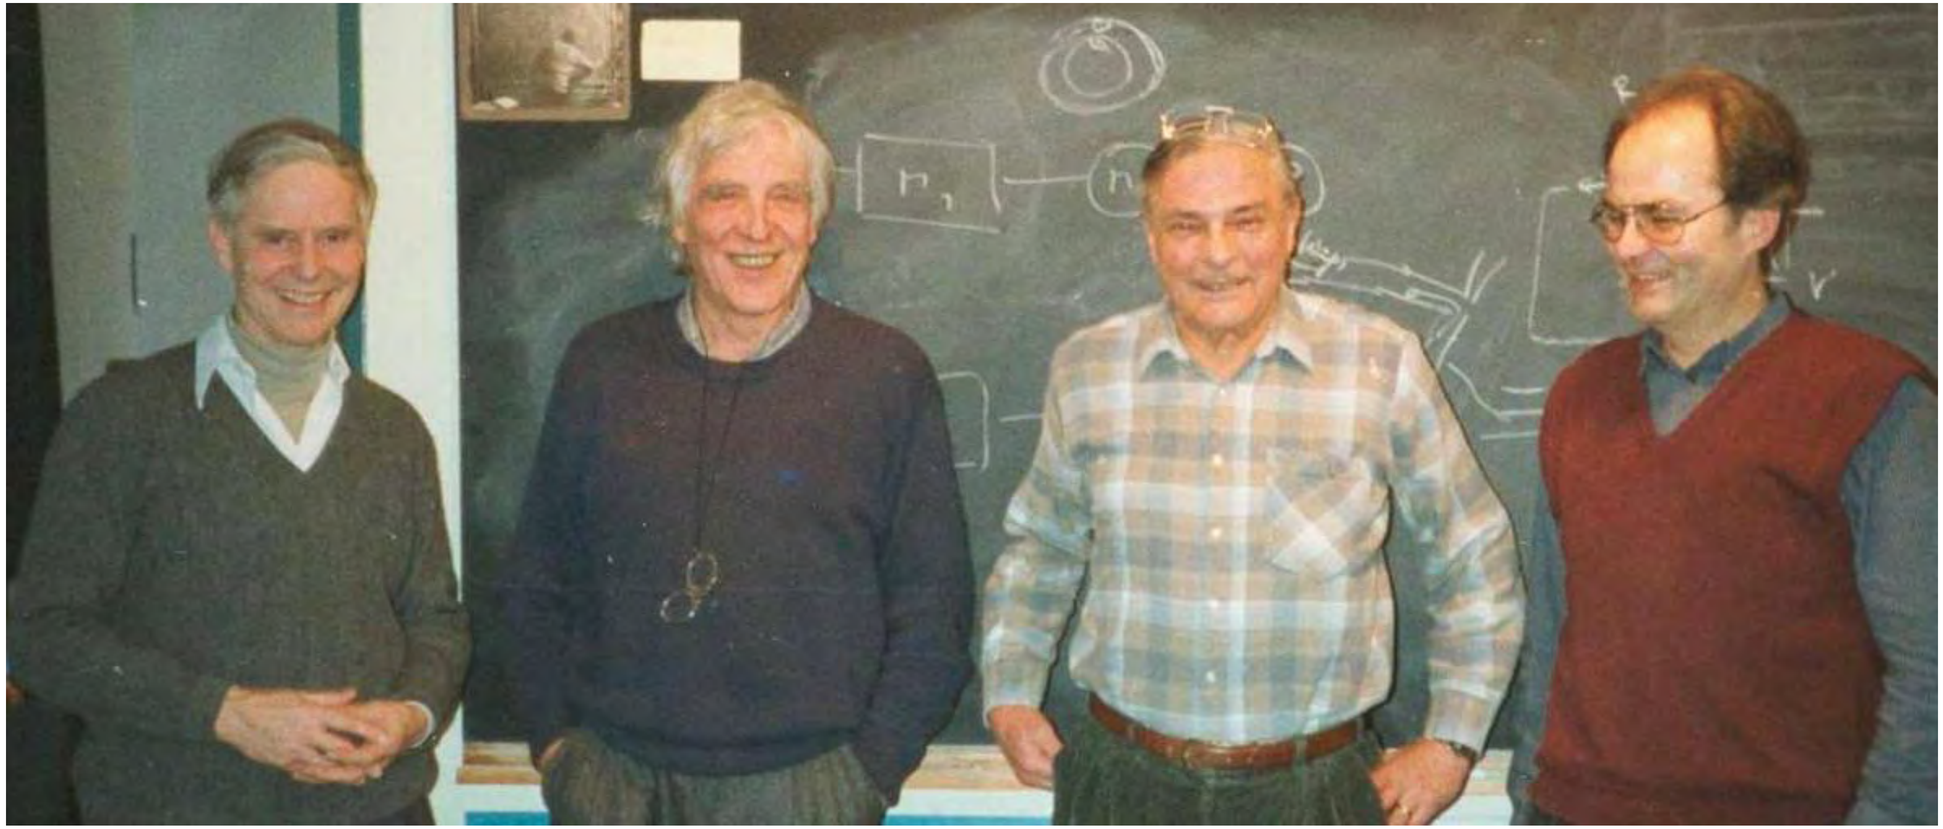
Werner Reichardt's PhD



Werner with Dr. Ruska (center)

Photo dated Nov. 17, 1952 (courtesy B. Reichardt)

The four directors
of the MPI fuer Biologische Kybernetik



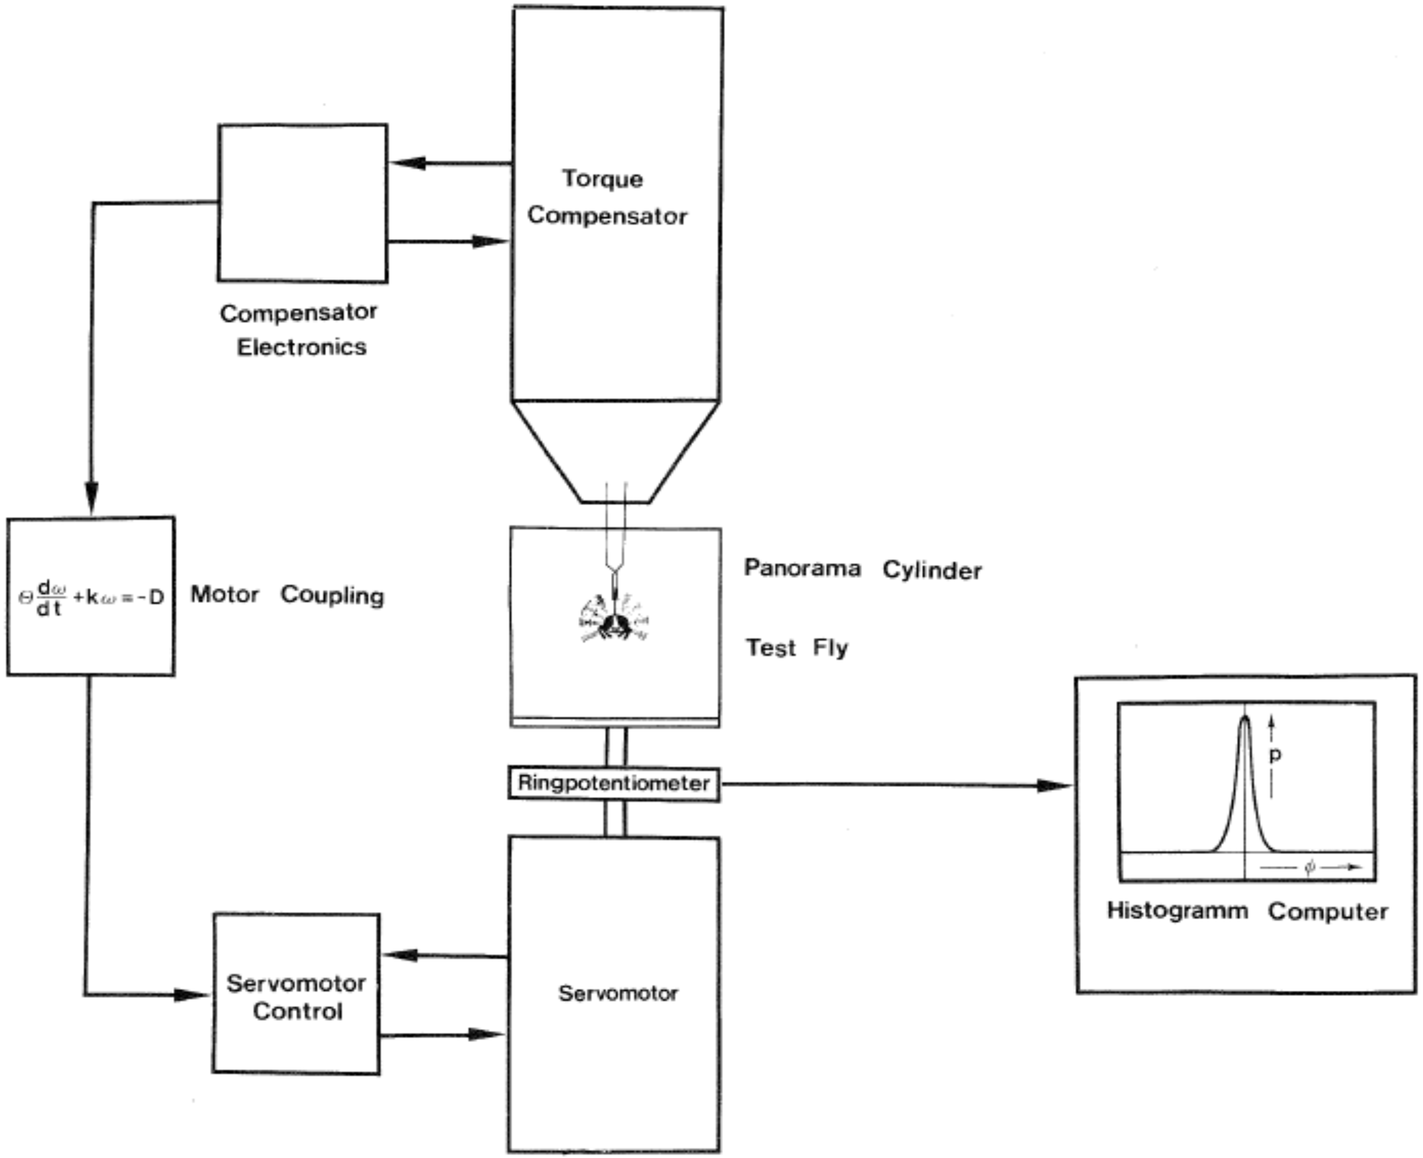
The beautiful eyes of flies



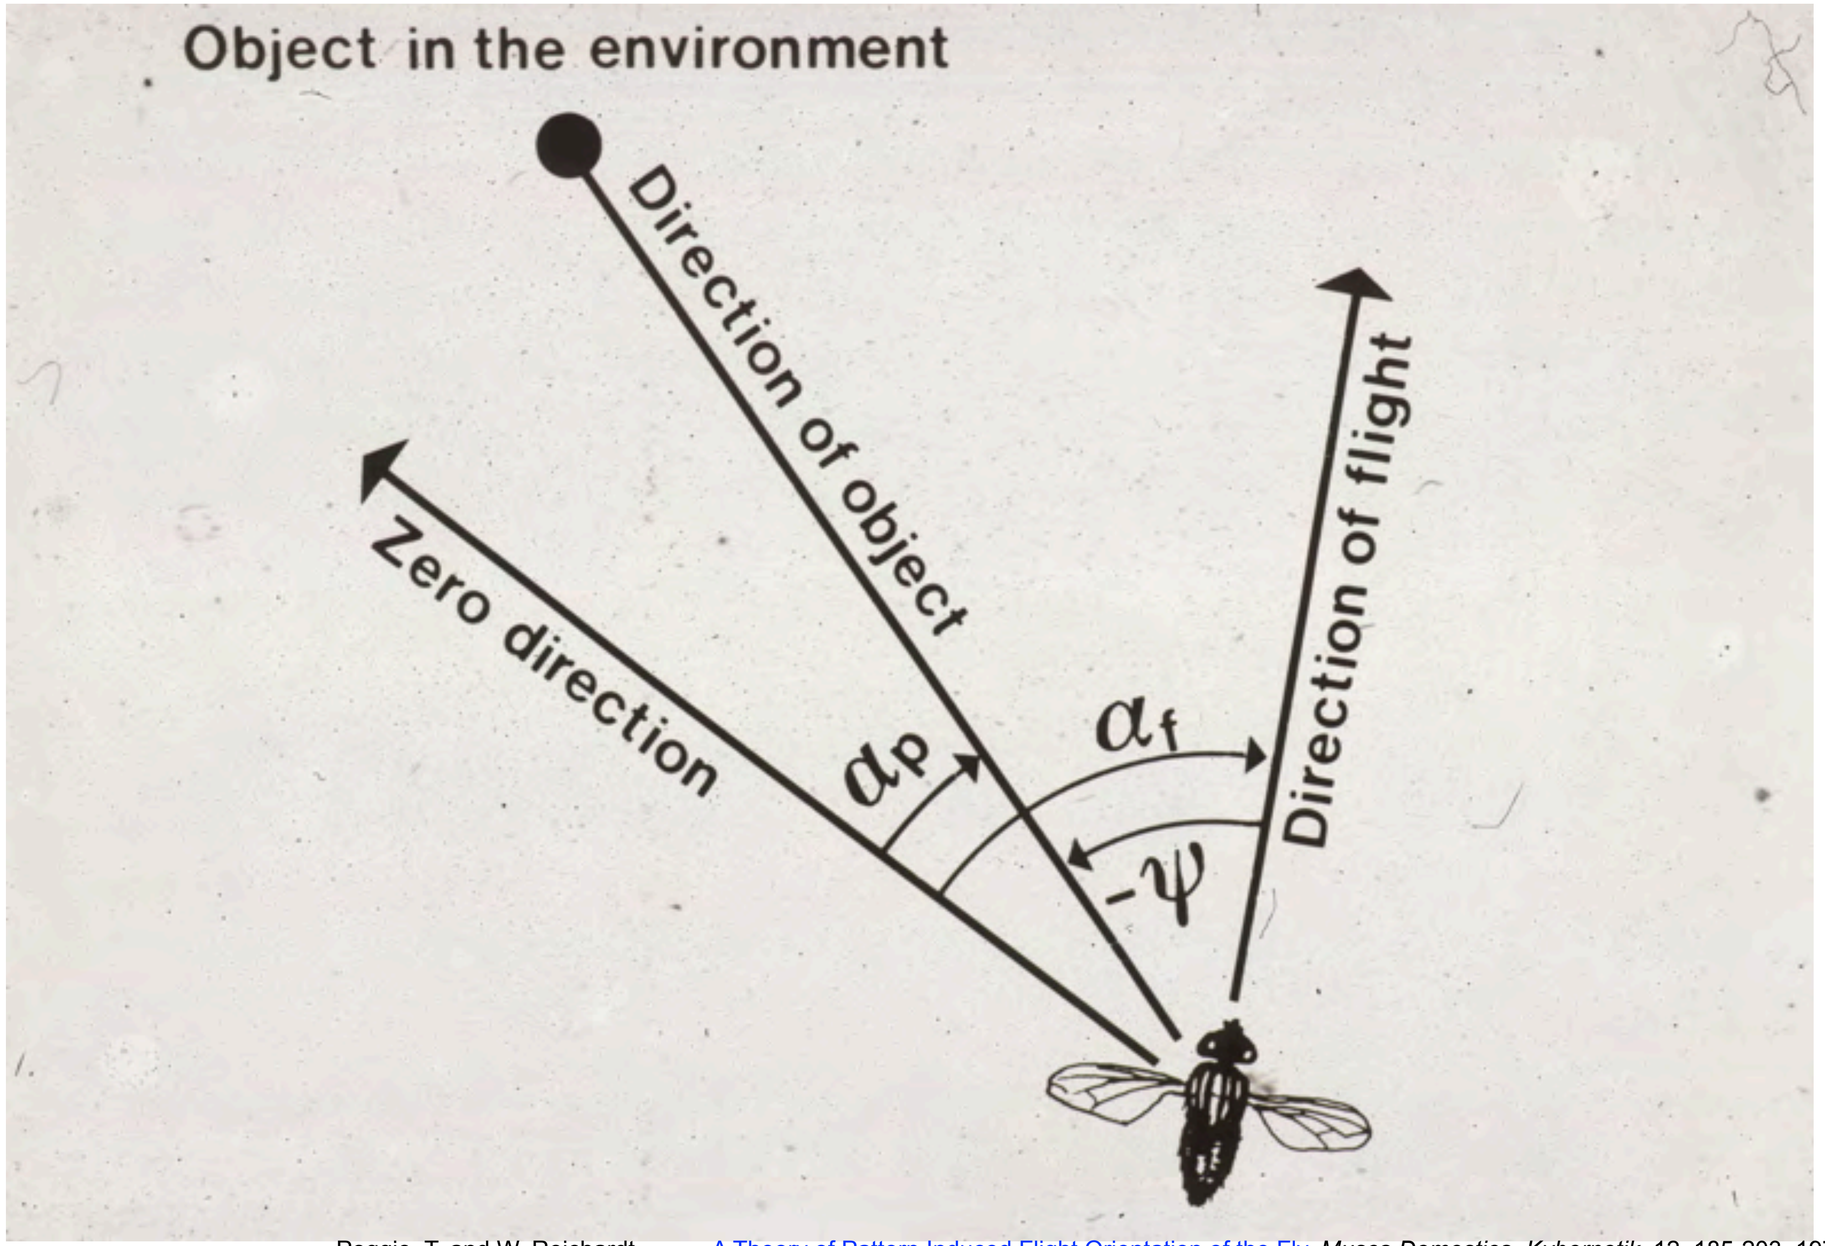
Work at 3 levels

- Fixation and tracking behavior of the fly
- Motion algorithms and circuits: the beetle (and the fly); relative motion algorithms and circuits
- Biophysics of computation

Fixation and tracking behavior:
Reichardt's closed loop flight simulator



Fixation and tracking behavior



Poggio, T. and W. Reichardt. A Theory of Pattern Induced Flight Orientation of the Fly, *Musca Domestica*, *Kybernetik*, 12, 185-203, 1972.

Cognition in flies: probabilistic theories then (coming only now to humans)

The basic equation is

$$\text{(H)}\quad \ddot{\psi} + K\dot{\psi} = N(t) - D(\psi(t)) - \tau(\psi)\dot{\psi}(t) + S(t)$$

If N is a "white noise" process the corresponding Kolmogoroff eq. is

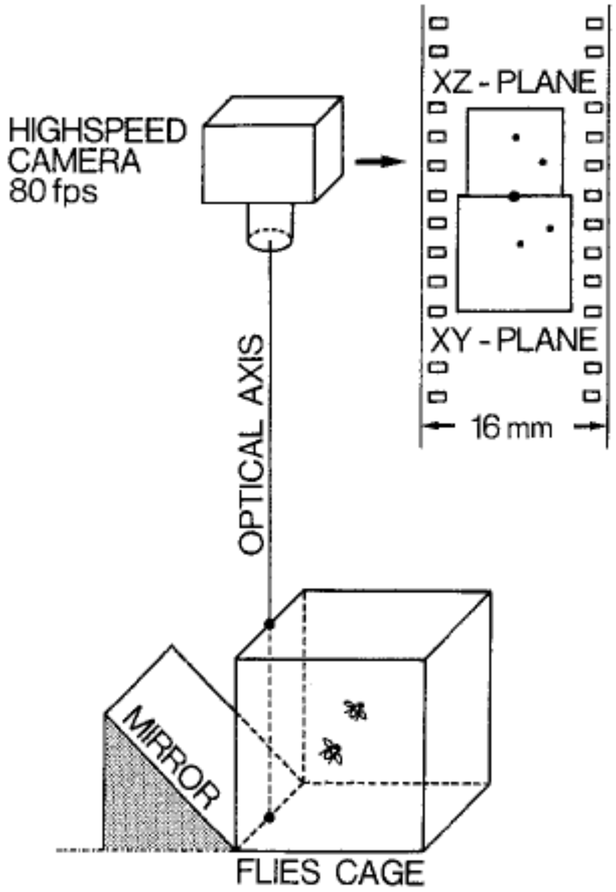
$$\frac{\partial}{\partial t} P(\psi,t) = -\frac{\partial}{\partial \psi}\left[\frac{1}{K+\tau}\frac{\partial U(\psi)}{\partial \psi} P(\psi,t)\right] + \frac{c}{(K+\tau)^2}\frac{\partial^2}{\partial \psi^2} P(\psi,t)$$

The beginning of untethered flight analysis
Bülthoff, Poggio & Wehrhahn Z. Naturforsch. 35c, 811-815 (1980)

- most behavioral fly research was done with the Götz torque meter
- in 1976, based on this recording technology, Reichardt & Poggio developed their theory for: *Visual control of orientation behaviour in the fly, Part I +II.* Quart. Rev. Biophysics ***9(3)***, 311-375

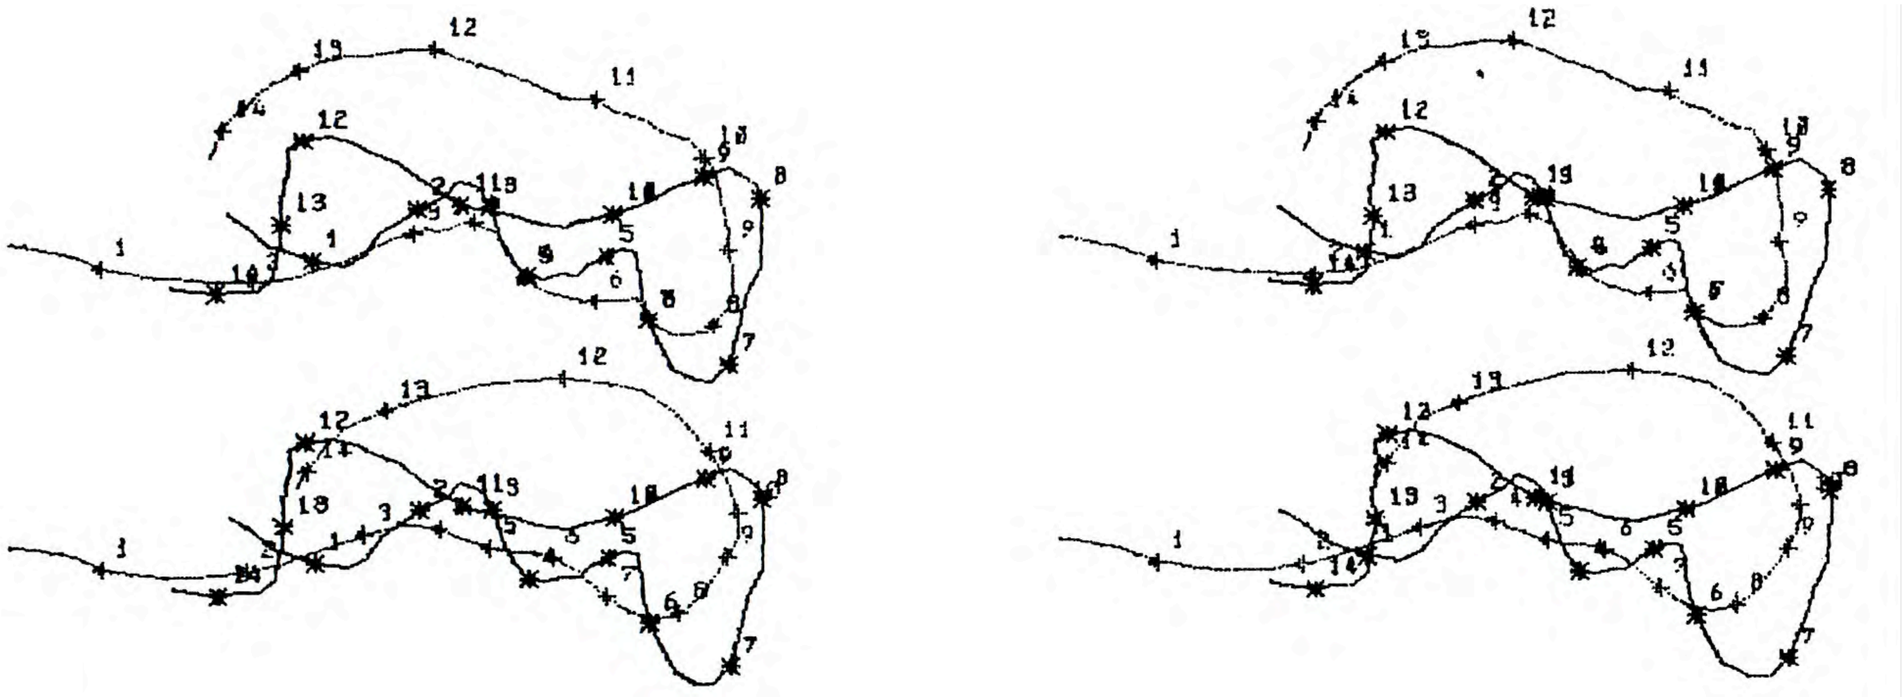
- open question: how well does this theory describe fly behavior of natural flight

- in 1980 Wehrhan started high-speed film recording of flies chasing each other
- single frame analysis
- 3D stereo reconstruction



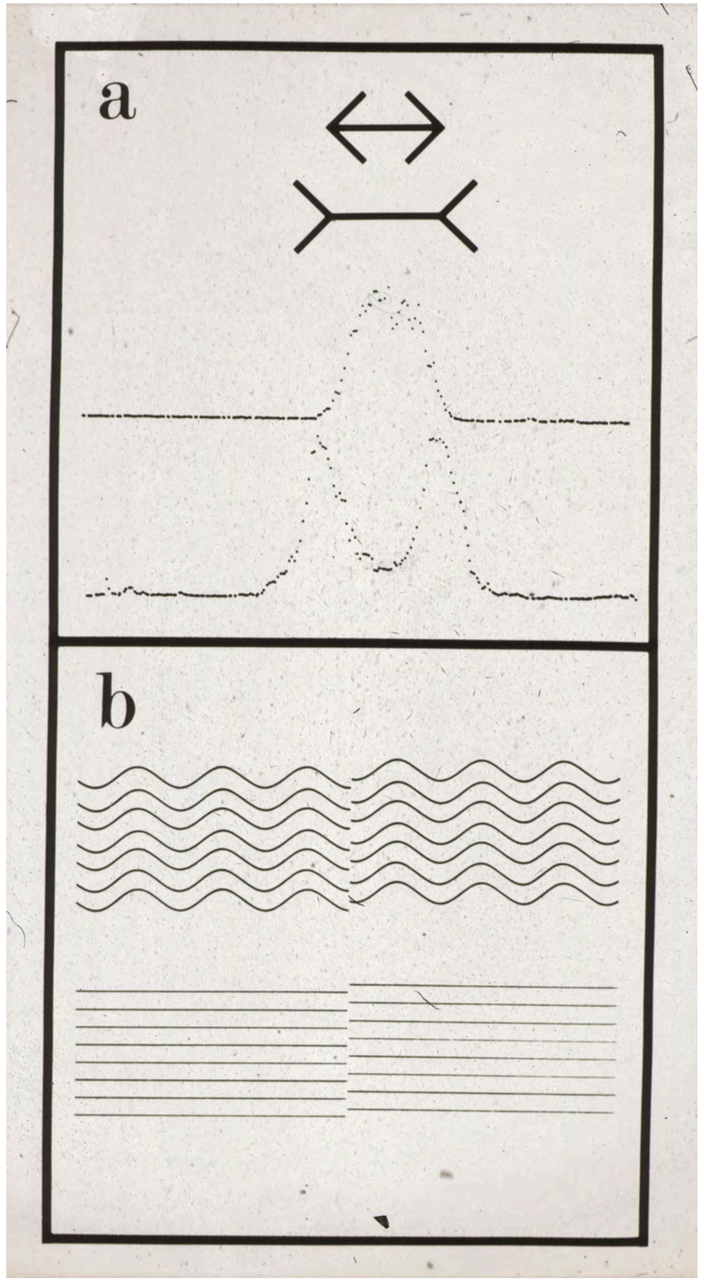
Cognitive theory of basic fly instincts predicts trajectory of chasing fly ...

$$\Theta\ddot{\psi} + [k + r(\psi)]\,\dot{\psi} + \frac{\partial U(\psi)}{\partial \psi} = N(t) + S(t),$$



Wehrhahn, C., T. Poggio and H. Bülthoff, *Biological Cybernetics*, 45, 123-130, 1982.

Cognition in flies

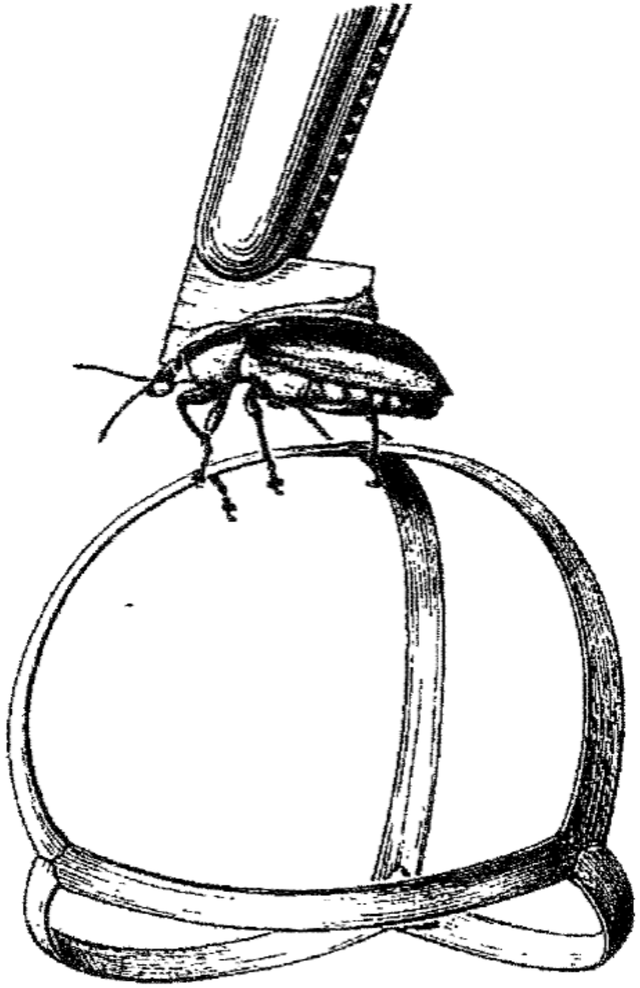


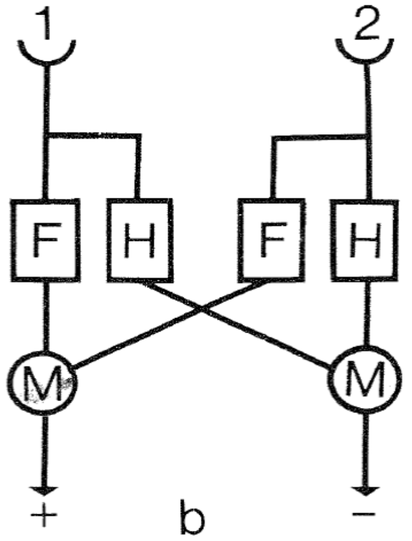
Geiger, G. and T. Poggio. The Muller-Lyer Figure and the Fly, *Science*, 190, 479-480, 1975.

Work at 3 levels

- Fixation and tracking behavior of the fly (cognition in the fly…similar to Bayesian approach to cognition in humans…no neurons!)
- Motion algorithms and circuits: the beetle (and the fly); relative motion algorithms and circuits
- Biophysics of computation

Motion algorithm: the beetle Clorophanus and Reichardt's motion detector

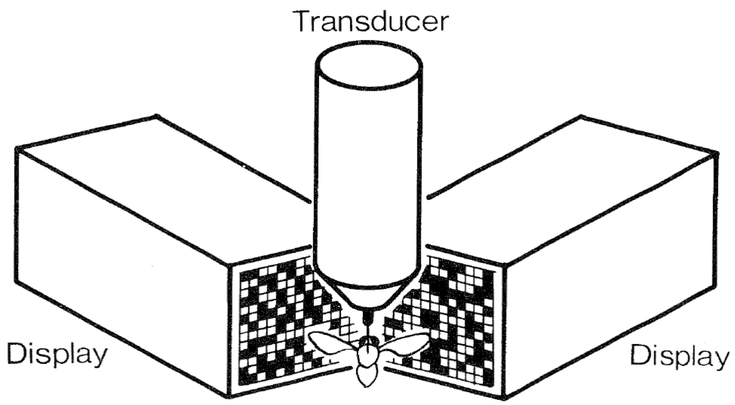


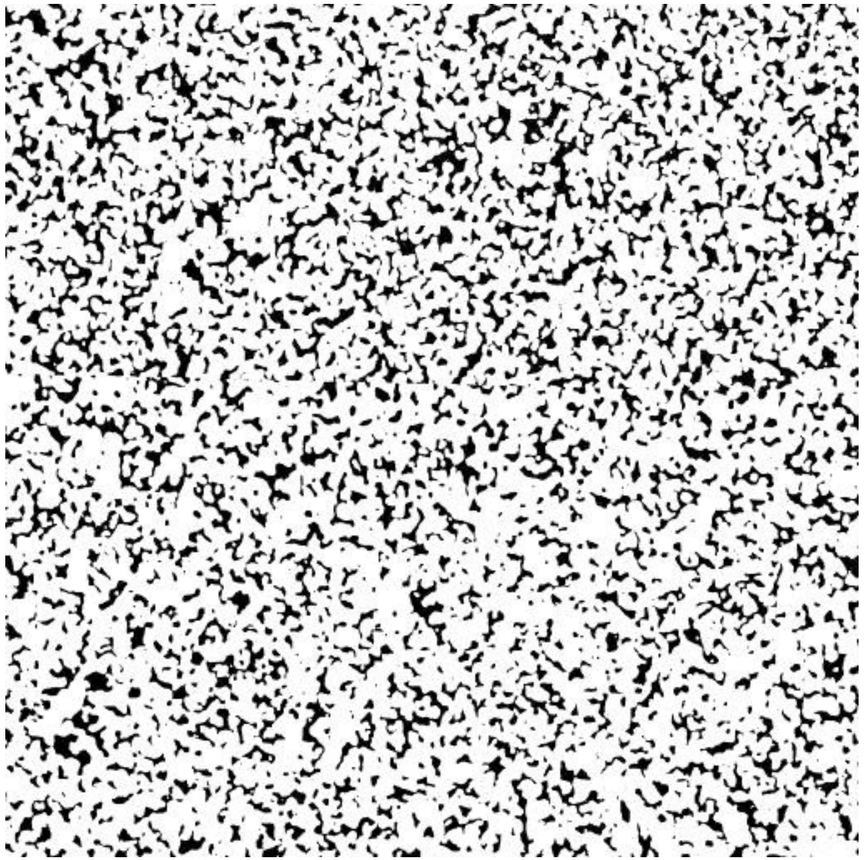


Motion algorithm: the beetle and the fly

- The beetle follows the motion
- Each photoreceptor sees only an alternation of dark and light: how is motion computed?
- Reichardt and Hassenstein (and Peter Kunze) found the rules used by neural circuits. The algorithm (refined by D. Varju) explained many data : *Reichardt detector.*
- The same model describes motion perception in flies: beautiful papers on anatomy, optics and organization of motion perception by Braitenberg, Kirschfeld, Goetz.
- An equivalent (“energy”) model (Adelson) describes motion cells in primate cortex.
- A form of it has been used by Matsushita in the first chips to stabilize videocameras (see also Buelthoff, Little and Poggio, Nature, 1989)

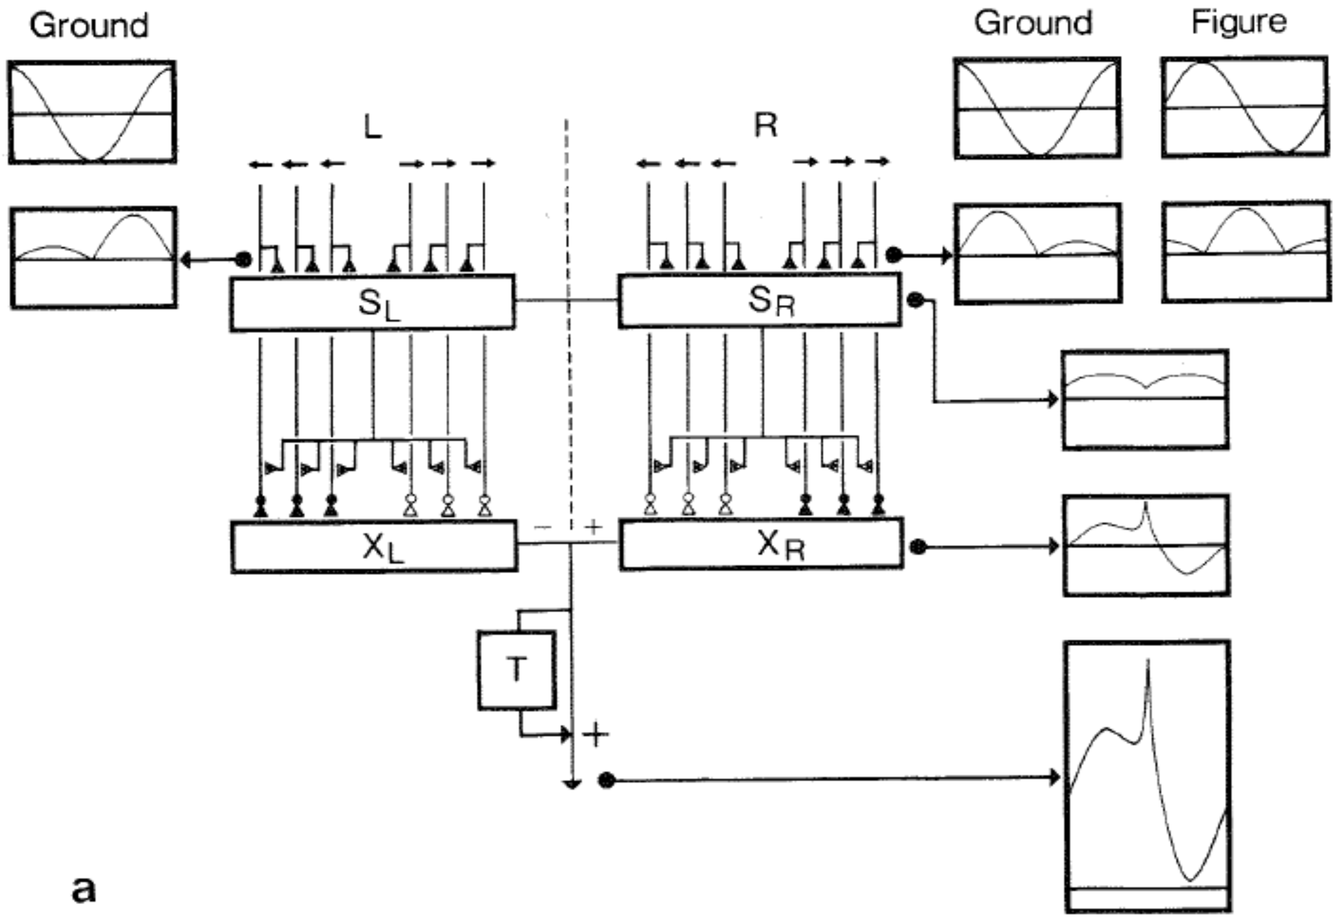
Relative motion and figure-ground discrimination: the fly





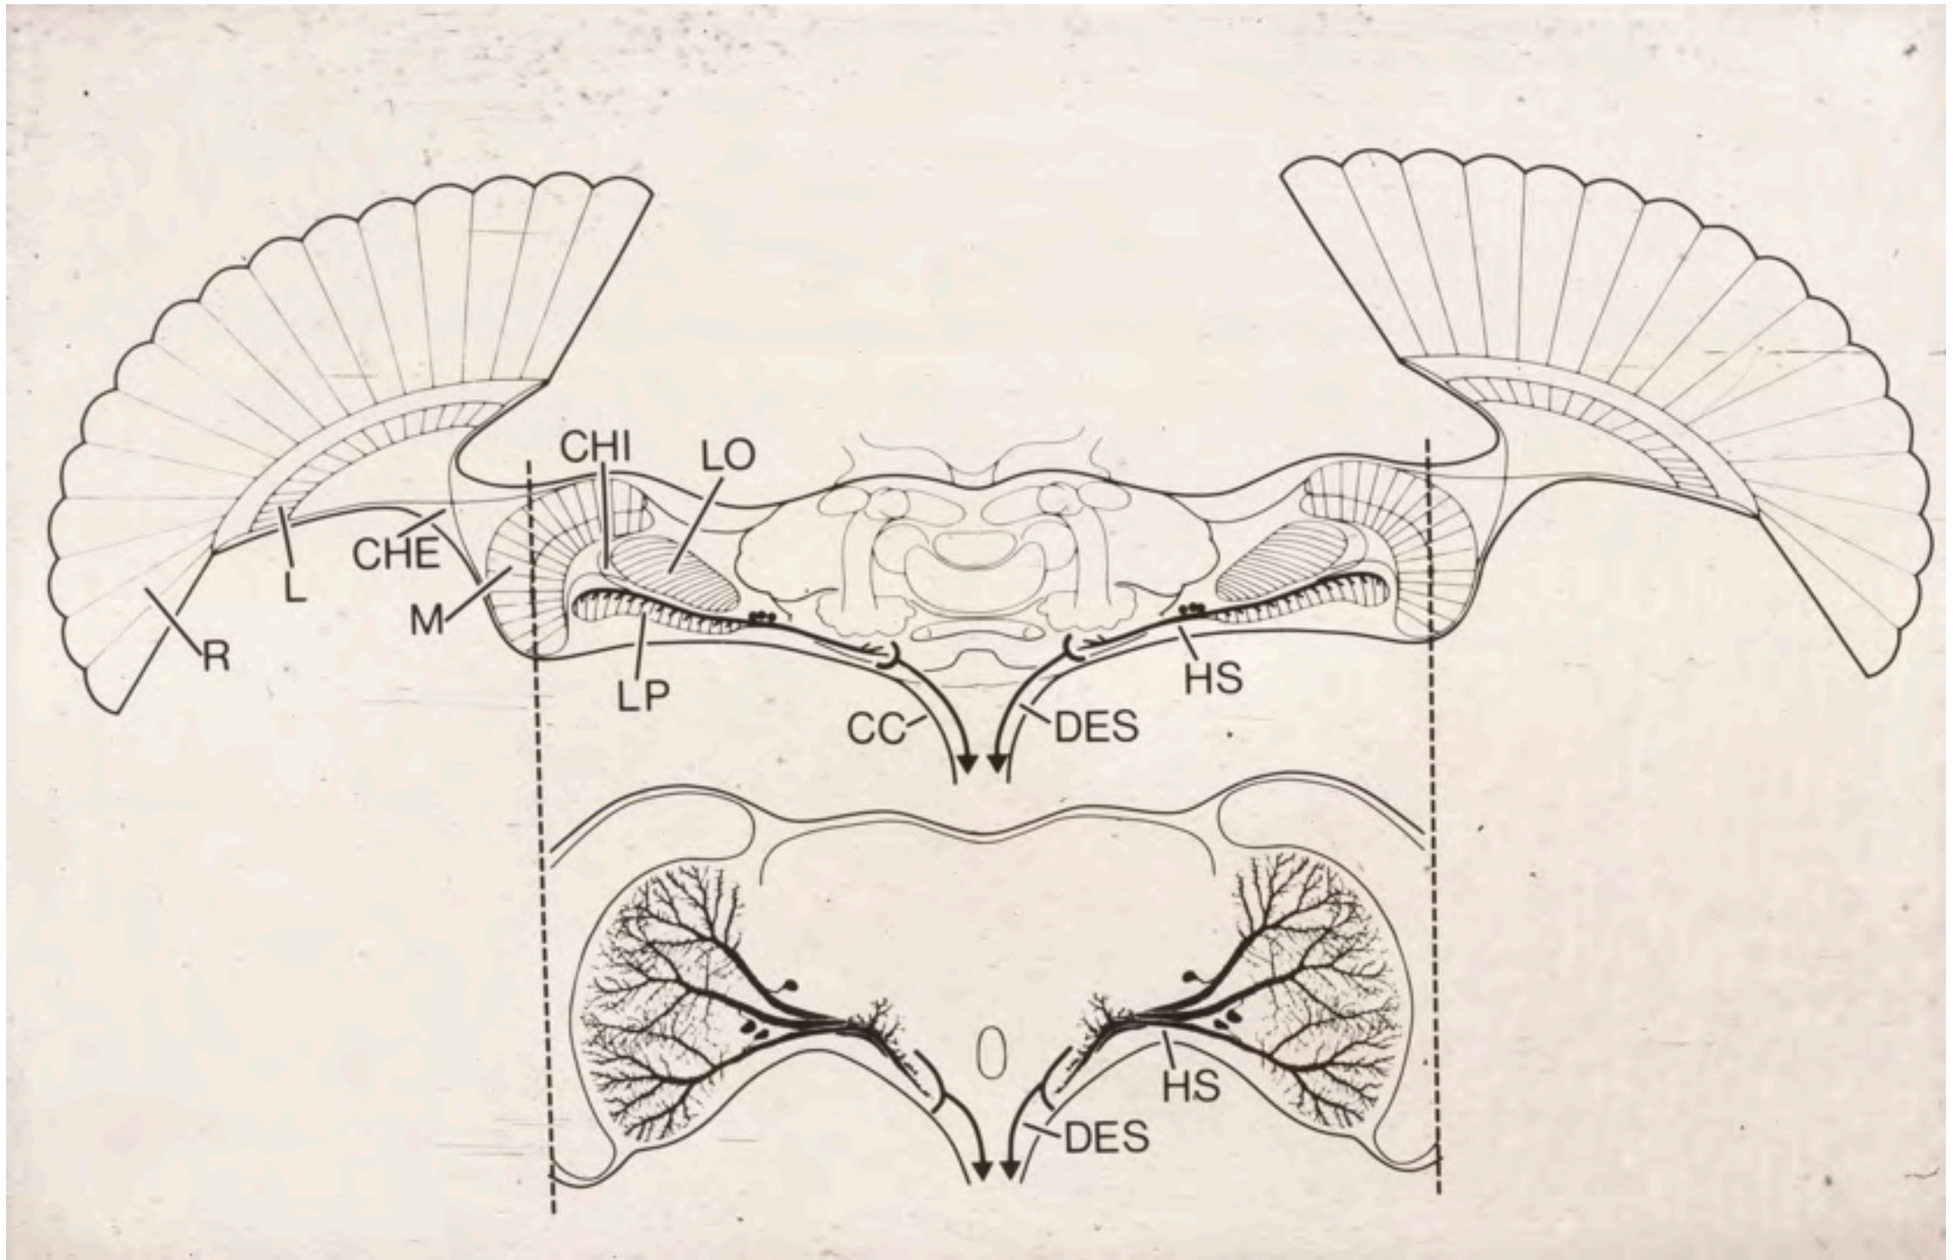
Work by Werner Reichardt
(with Poggio and Hausen and later with M. Egelhaaf and A. Borst)

Motion discontinuities and figure-ground discrimination: neural circuitry

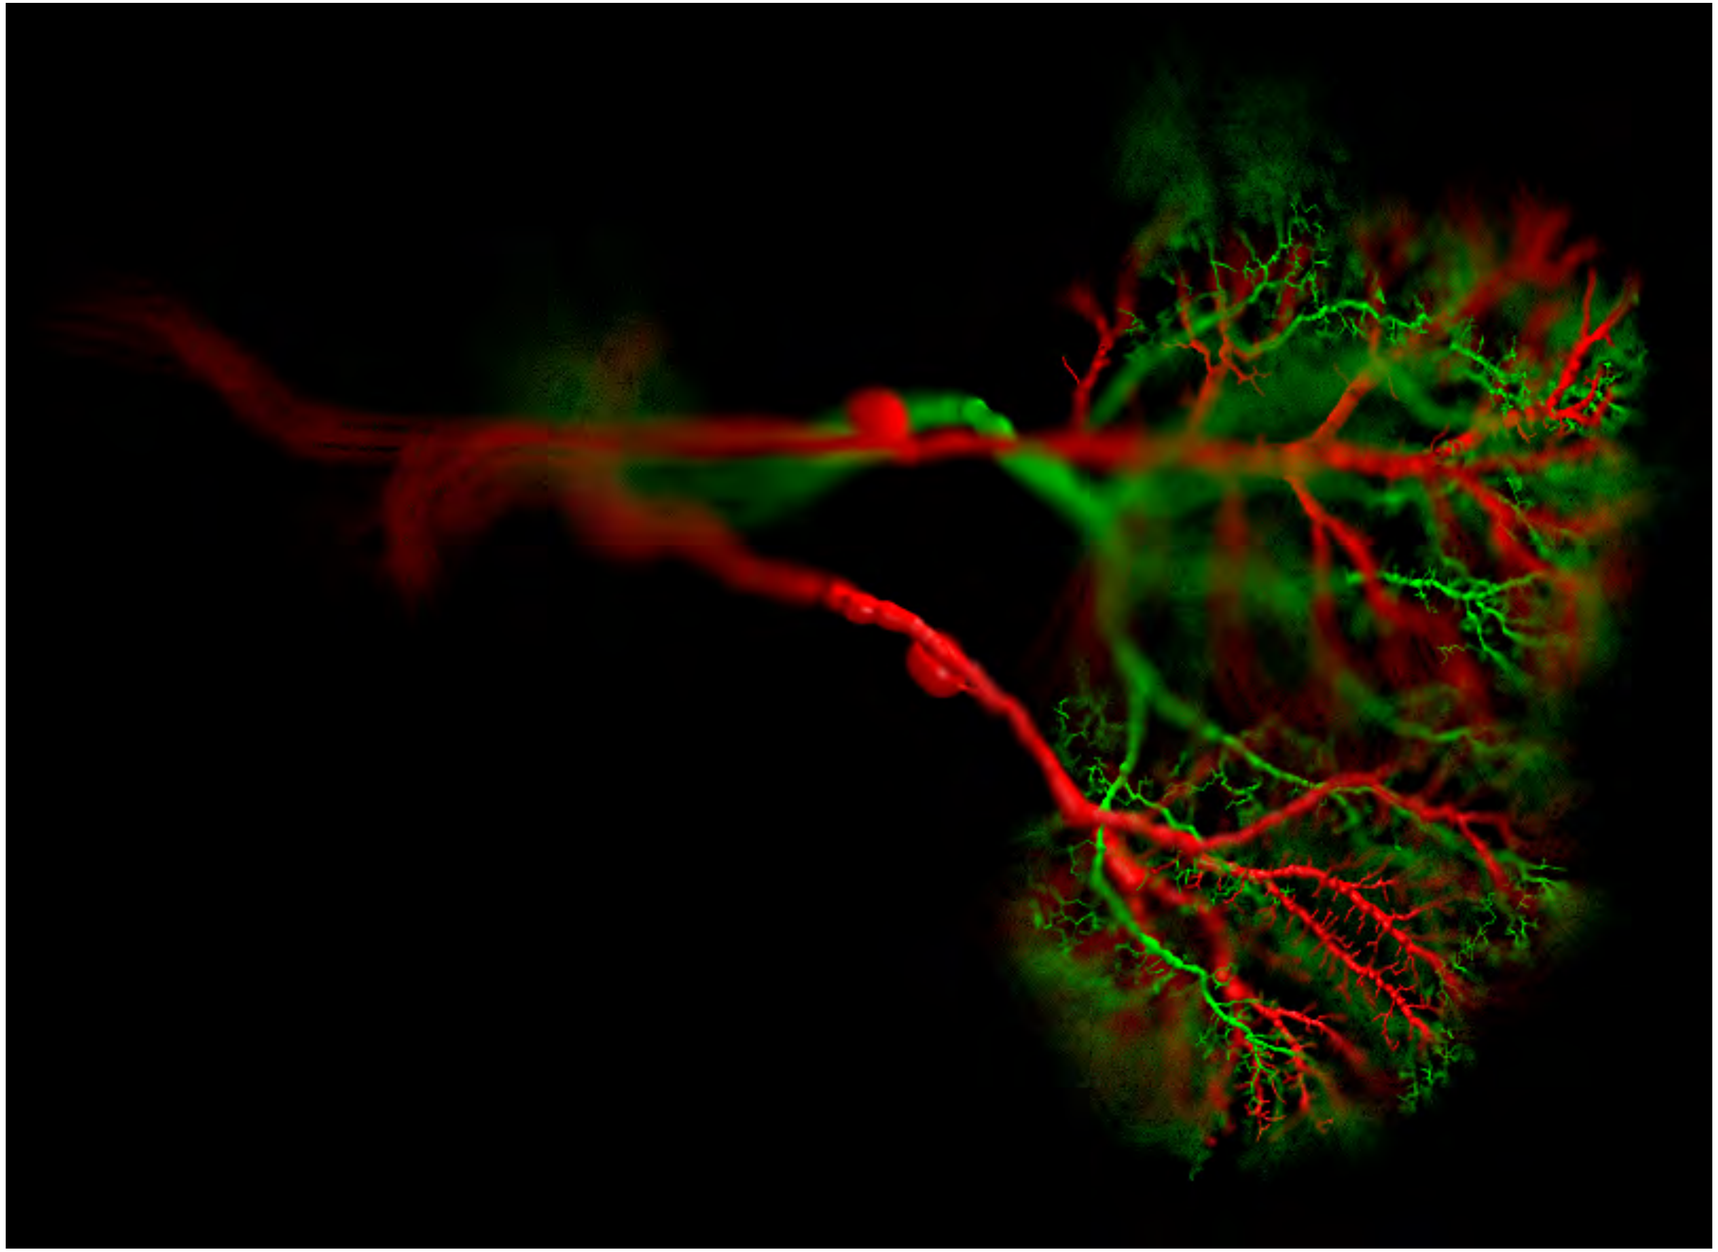


Towards the neural circuitry, Reichardt, Poggio, Hausen, 1983

Relative motion



Two of the neurons....



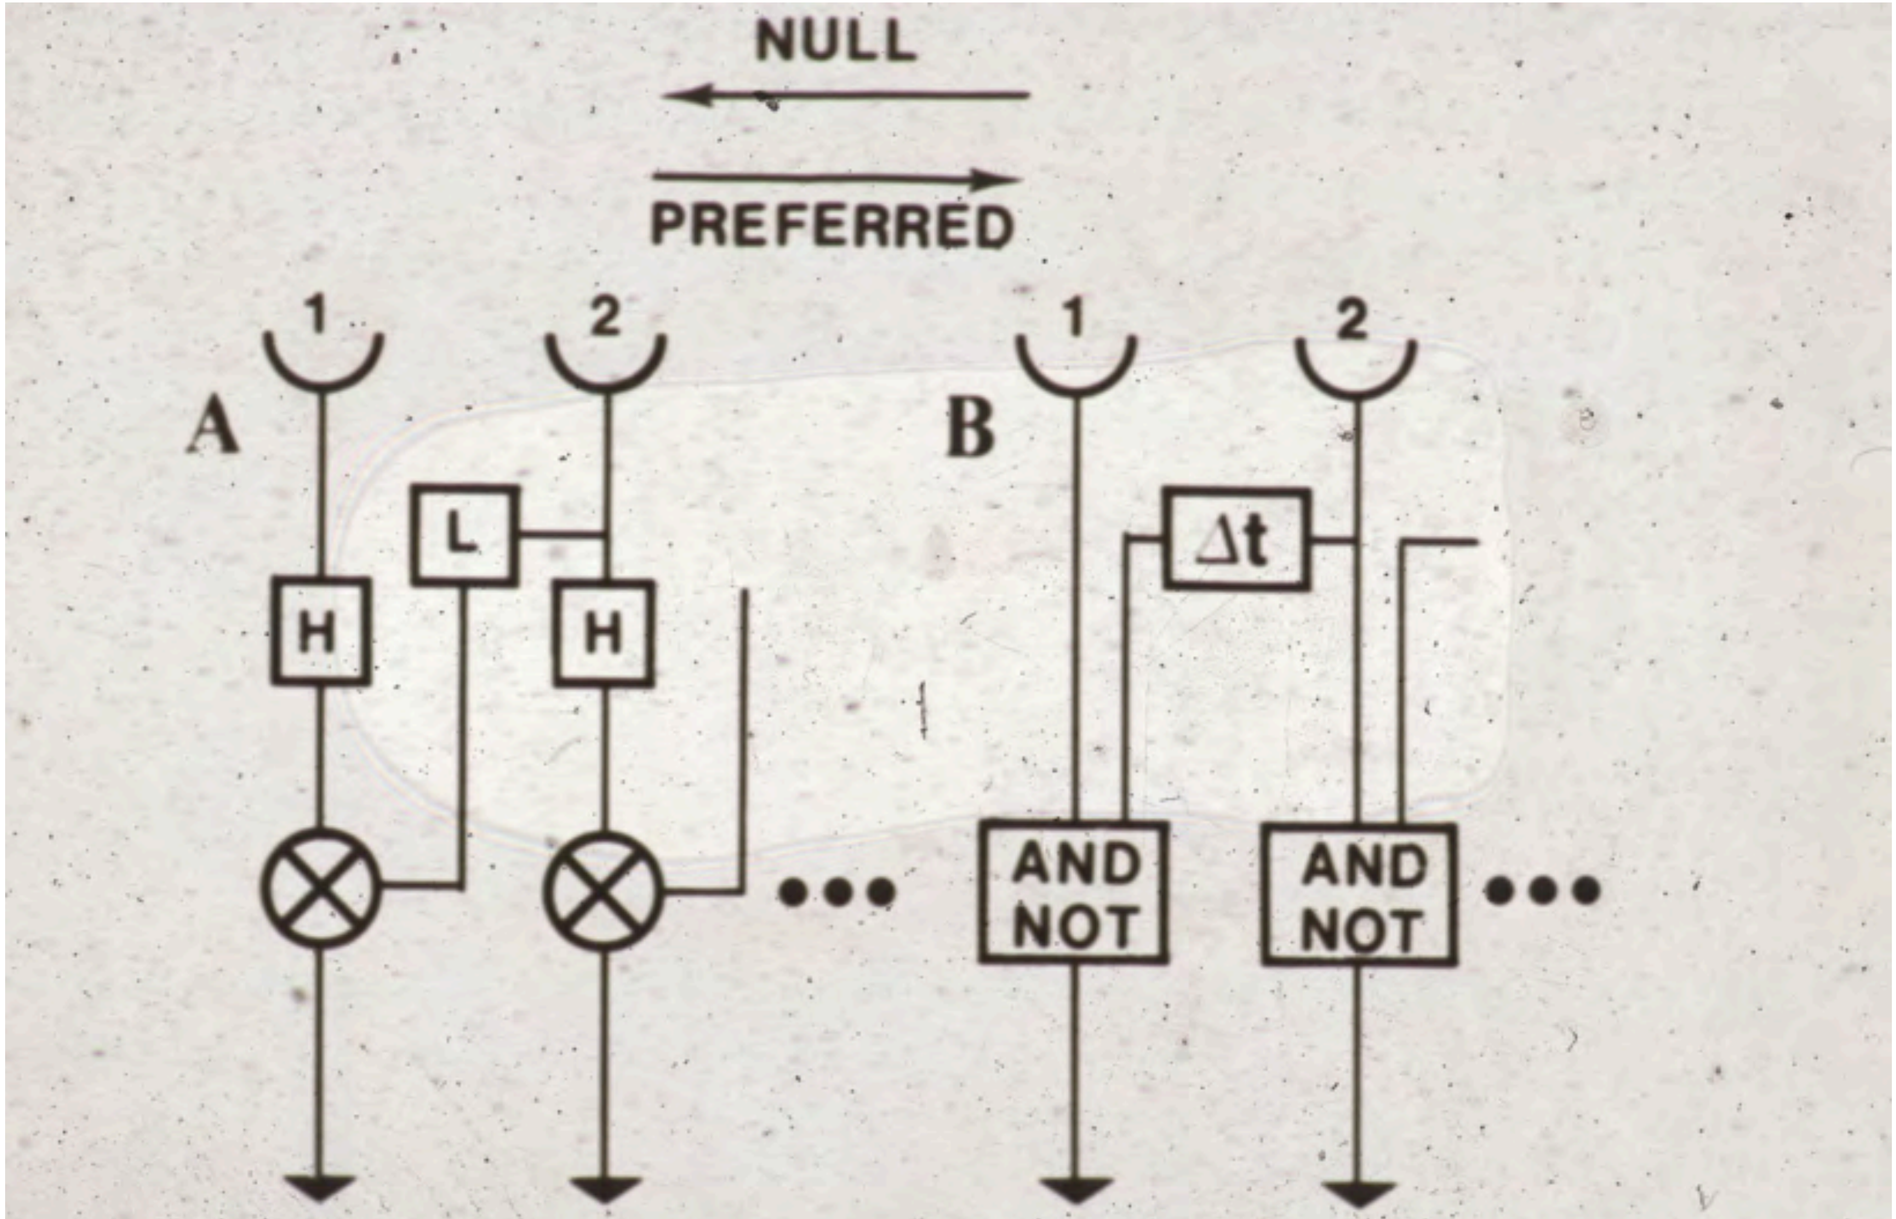
Hermann Cuntz, Ju¨rgen Haag, and Alexander Borst, 2003

Work at 3 levels

- Fixation and tracking behavior of the fly

- Motion algorithms and circuits: the beetle (and the fly); relative motion algorithms and circuits (similar in spirit to HMAX and DiCarlo)
- Biophysics of computation

Biophysics of computation (motion detection)



Biophysics of Computation



314 REVIEW ARTICLE NATURE VOL. 317 26 SEPTEMBER 1985

Computational vision and regularization theory

Tomaso Poggio, Vincent Torre* & Christof Koch

Artificial Intelligence Laboratory and Center for Biological Information Processing, Massachusetts Institute of Technology, 545 Technology Square, Cambridge, Massachusetts 02193, USA
* Istituto di Fisica, Universita di Genova, Genova, Italy



Proc. R. Soc. Lond. B. **202**, 409–416 (1978)
Printed in Great Britain

A synaptic mechanism possibly underlying directional selectivity to motion

By V. Torre† and T. Poggio‡

† *Università di Genova, Istituto di Fisica, Genoa, Italy*
‡ *Max-Planck-Institut für biologische Kybernetik, Tübingen, Germany*

(*Communicated by B. B. Boycott, F.R.S. – Received 1 February* 1978)

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

and

CENTER FOR BIOLOGICAL INFORMATION PROCESSING
WHITAKER COLLEGE

A.I. Memo 795
C.B.I.P. Paper 008

October, 1984

Biophysics of Computation: Neurons, Synapses and Membranes

Christof Koch and Tomaso Poggio

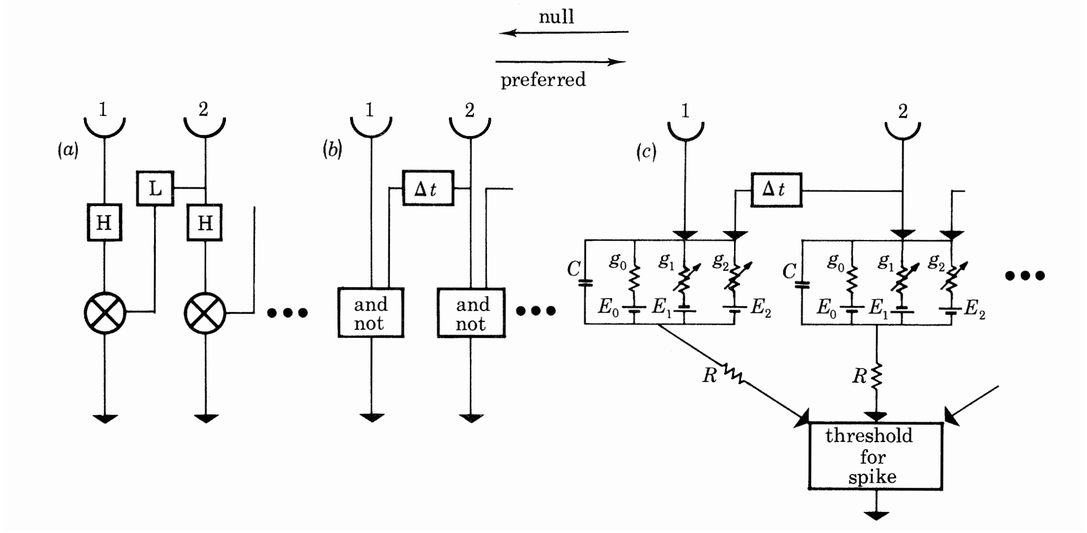
Proc. R. Soc. Lond. B. **202**, 409–416 (1978)
Printed in Great Britain

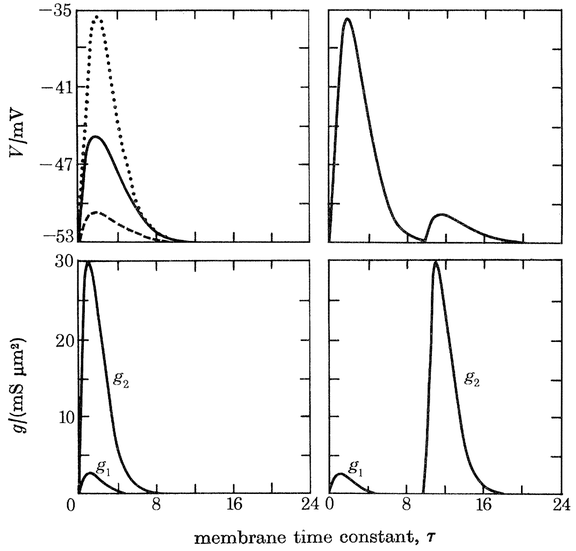
A synaptic mechanism possibly underlying directional selectivity to motion

By V. Torre† and T. Poggio‡

† *Università di Genova, Istituto di Fisica, Genoa, Italy*
‡ *Max-Planck-Institut für biologische Kybernetik, Tübingen, Germany*

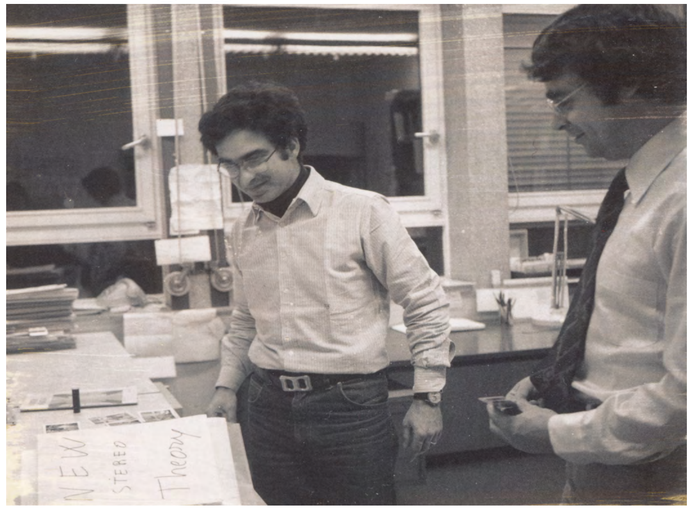
(*Communicated by B. B. Boycott, F.R.S. – Received 1 February* 1978)





Cooperative neural network for stereo

Science, New Series, Vol. 194, No. 4262. (Oct. 15, 1976), pp. 283-287.



~ 1979 , T. Poggio and D. Marr, MPI, Tuebingen

Cooperative Computation of Stereo Disparity

A cooperative algorithm is derived for extracting disparity information from stereo image pairs.

D. Marr and T. Poggio

Perhaps one of the most striking differences between a brain and today's com-

able hardware (*3*). Nevertheless, one can expect a nervous system and a digital

which local operatior
ate in forming globa
regulated manner. C
ena are well known i
it has beer
an import
as well (*6*–
tions alon
lesz (*11*),
scopic fus
His model
dipole ma
tips of adja
gestive m
its biologi
stereoscop
and yet ur
mation pr

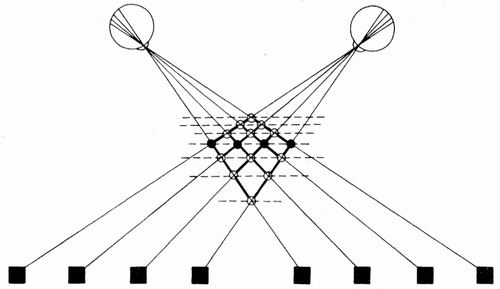
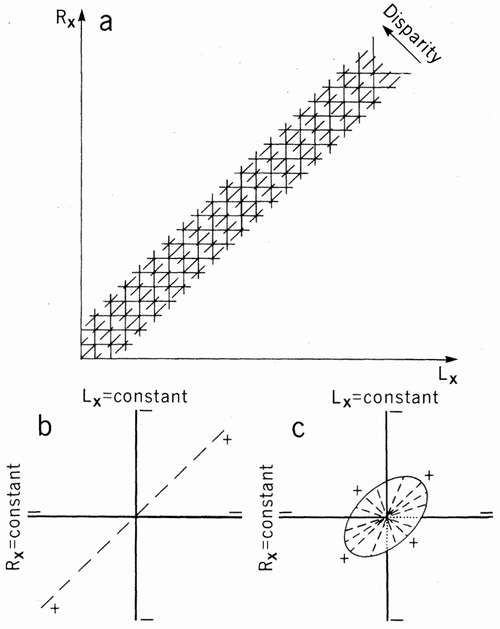


Fig. 1 (left). Ambiguity in the correspondence between the two retinal projections. In this figure, each of the four points in one eye's view could match any of the four projections in the other eye's view. Of the 16 possible matchings only four are correct (closed circles), while the remaining 12 are "false targets" (open circles). It is assumed here that the targets (closed squares) correspond to "matchable" descriptive elements obtained from the left and right images. Without further constraints based on global considerations, such ambiguities cannot be resolved. Redrawn from Julesz (*11*, figure 4.5-1). Fig. 2 (right). The explicit structure of the two rules R1 and R2 for the case of a one-dimensional image is represented in (a), which also shows the structure of a network for implementing the algorithm described by Eq. 2. Solid lines represent "inhibitory" interactions, and dotted lines represent "excitatory" ones. The local structure at each node of the network in (a) is given in (b). This algorithm may be extended to two-dimensional images, in which case each node in the corresponding network has the local structure shown in (c). Such a network was used to solve the stereograms exhibited in Figs. 3 to 6.



Vision:
what is where



The MIT Press

Vision

A Computational Investigation into the Human Representation and Processing of Visual Information

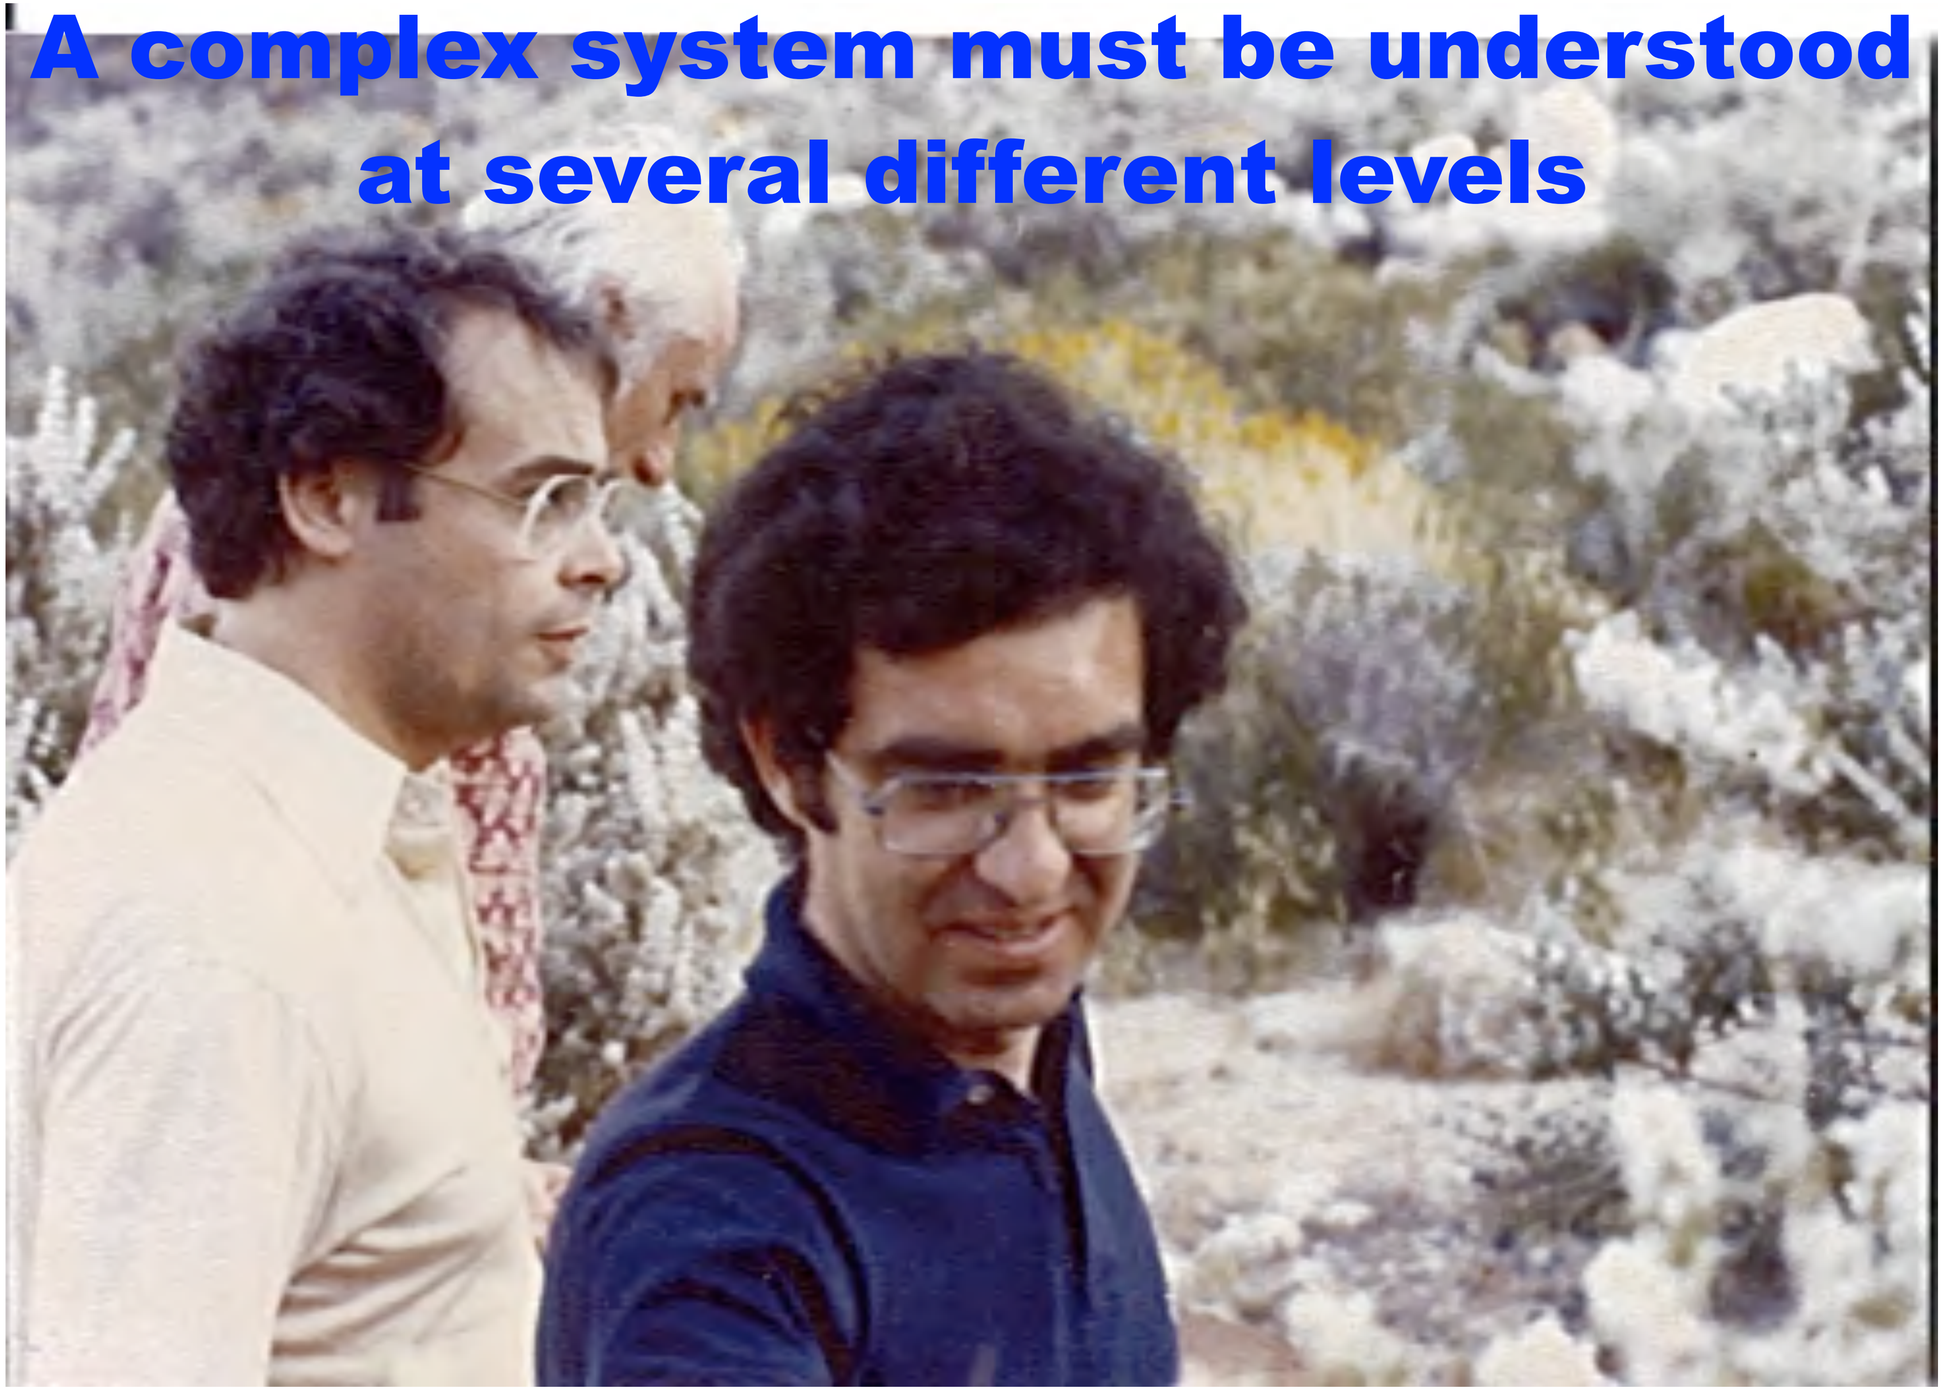
Foreword by
Afterword by Tomaso Poggio

David Marr's posthumously published *Vision* (1982) influenced a generation of brain and cognitive scientists, inspiring many to enter the field. In *Vision*, Marr describes a general framework for understanding visual perception and touches on broader questions about how the brain and its functions can be studied and understood. Researchers from a range of brain and cognitive sciences have long valued Marr's creativity, intellectual power, and ability to integrate insights and data from neuroscience, psychology, and computation. This MIT Press edition makes Marr's influential work available to a new generation of students and scientists.

In Marr's framework, the process of vision constructs a set of representations, starting from a description of the input image and culminating with a description of three-dimensional objects in the surrounding environment. A central theme, and one that has had far-reaching influence in both neuroscience and cognitive science, is the notion of different levels of analysis—in Marr's framework, the computational level, the algorithmic level, and the hardware implementation level.

Now, thirty years later, the main problems that occupied Marr remain fundamental open problems in the study of perception. *Vision* provides inspiration for the continui

A complex system must be understood at several different levels





Werner Reichardt's scientific legacy: Integrative Neuroscience

- Marr's book *Vision* (Marr, 1982) had a great impact on computational neuroscience: a system as complex as the brain must be understood at several different levels:
  - — computation
  - — algorithms
  - — biophysics and circuits

- The argument came from "From Understanding Computation to Understanding Neural Circuits", Marr and Poggio, 1977…

- …part of which comes from Reichardt and Poggio, 1976 (Q. Rev. Biophysics, Part I)…

- …which is a follow-up of Werner's argument for starting the Max-Planck-Institute fuer Biologische Kybernetik!

MIT, (1981-)

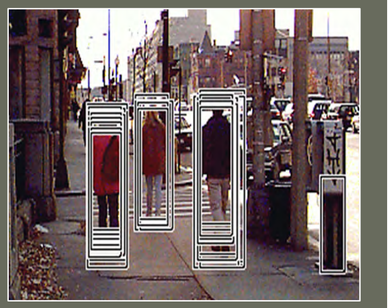
43rd Stated Meeting of the NRP Associates, March 14-17, 1982

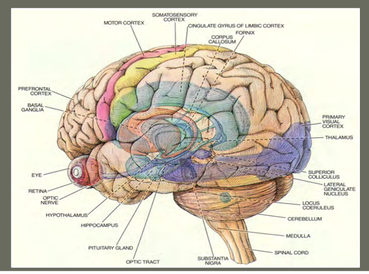


MIT, (1981-)

$$\min_{f \in H} \left[\frac{1}{\ell} \sum_{i=1}^{\ell} V(y_i, f(x_i)) + \mu \|f\|_K^2 \right]$$

$$f(x) = \sum_{i=1}^{l} c_i K(\mathbf{x}_i, \mathbf{x})$$





Learning theory + algorithms

Predictive regularization algorithms

Theorems on foundations of learning:

ENGINEERING APPLICATIONS

- Bioinformatics
- Computer vision
- Computer graphics, speech synthesis, creating a virtual actor

Computational Neuroscience: models+experiments

How visual cortex works – and how it may suggest better computer vision systems

BULLETIN (New Series) OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 39, Number 1, Pages 1–49
S 0273-0979(01)00923-5
Article electronically published on October 5, 2001

ON THE MATHEMATICAL FOUNDATIONS OF LEARNING

FELIPE CUCKER AND STEVE SMALE

The problem of learning is arguably at the very core of the problem of intelligence, both biological and artificial.

T. Poggio and C.R. Shelton

Introduction

(1) A main theme of this report is the relationship of approximation to learning and the primary role of sampling (inductive inference). We try to emphasize relations of the theory of learning to the mainstream of mathematics. In particular, there are large roles for probability theory, for algorithms such as *least squares*, and for tools and ideas from linear algebra and linear analysis. An advantage of doing this

General conditions for predictivity in learning theory

Tomaso Poggio[1], Ryan Rifkin[1,4], Sayan Mukherjee[1,3] & Partha Niyogi[2]

[1]*Center for Biological and Computational Learning, McGovern Institute Computer Science Artificial Intelligence Laboratory, Brain Sciences Department, MIT, Cambridge, Massachusetts 02139, USA*
[2]*Departments of Computer Science and Statistics, University of Chicago, Chicago, Illinois 60637, USA*
[3]*Cancer Genomics Group, Center for Genome Research/Whitehead Institute, MIT, Cambridge, Massachusetts 02139, USA*
[4]*Honda Research Institute USA Inc., Boston, Massachusetts 02111, USA*

Developing theoretical foundations for learning is a key step towards understanding intelligence. 'Learning from examples' is a paradigm in which systems (natural or artificial) learn a functional relationship from a training set of examples. Within this paradigm, a learning algorithm is a map from the space of training sets to the hypothesis space of possible functional solutions. A central question for the theory is to determine conditions under which a learning algorithm will generalize from its finite training set to novel examples. A milestone in

What we assume in the above examples is a machine that is trained, instead of programmed, to perform a task, given data of the form $S = (x_i, y_i)_{i=1}^{n}$. Training means synthesizing a function that best represents the relation between the inputs x_i and the corresponding outputs y_i.

The basic requirement for any learning algorithm is generalization: the performance on the training examples (empirical error) must be a good indicator of the performance on future examples (expected error), that is, the difference between the two must be 'small' (see Box 1 for definitions; see also Fig. 1).

Probably the most natural learning algorithm is ERM: the algorithm 'looks' at the training set S, and selects as the estimated function the one that minimizes the empirical error (training error) over the functions contained in a hypothesis space of candidate

Box 1
Formal definitions in supervised learning

Convergence in probability. A sequence of random variables $\{X_n\}$ converges in probability to a random variable X (for example, $\lim_{n\to\infty} |X_n - X| = 0$ in probability) if and only if for every $\epsilon > 0$, $\lim_{n\to\infty} \mathbb{P}(|X_n - X| > \epsilon) = 0$.

Why do hierarchical architectures work?

THE COMPUTATIONAL MAGIC OF THE VENTRAL STREAM: TOWARDS A THEORY

Tomaso Poggio[⋆,†] (section 4 with Jim Mutch[⋆]; appendix 7.2 with Joel Leibo[⋆] and appendix 7.9 with Lorenzo Rosasco†)

⋆ *CBCL, McGovern Institute, Massachusetts Institute of Technology, Cambridge, MA, USA*

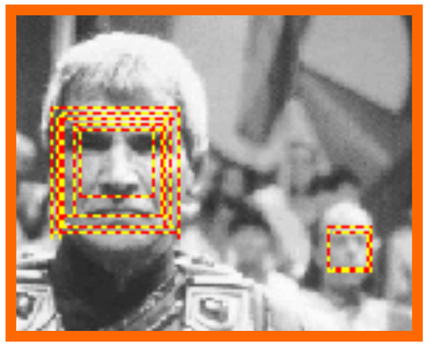
† Istituto Italiano di Tecnologia, Genova, Italy

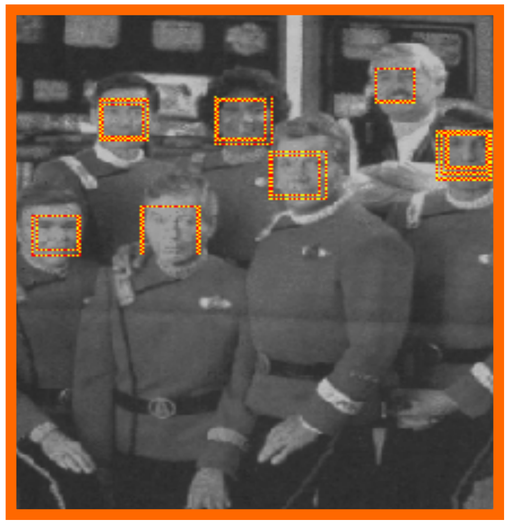
~15 year old CBCL computer vision research:

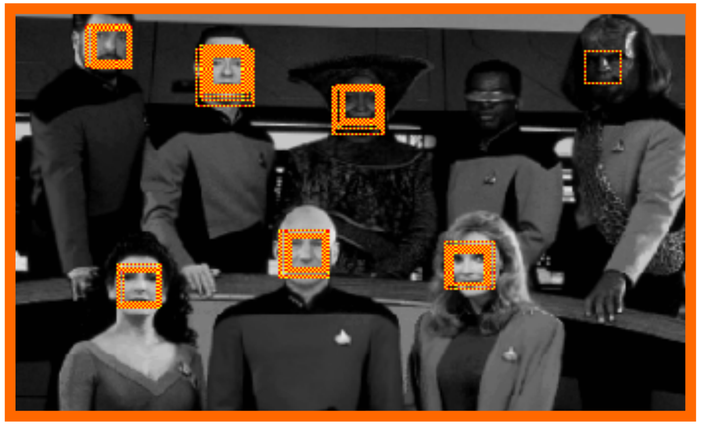
face detection;

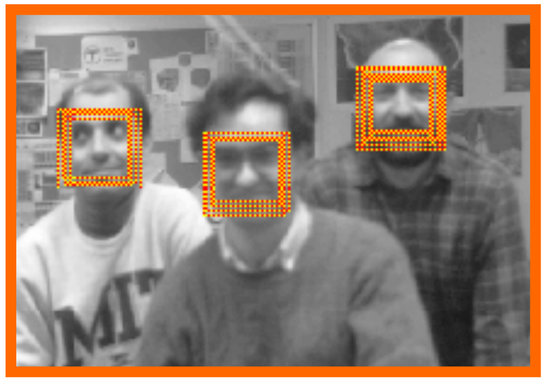
since 2006 on the market (digital cameras…)

- Training Database
- 1000+ Real, 3000+ VIRTUAL
- 50,0000+ Non-Face Pattern









Sung & Poggio 1995

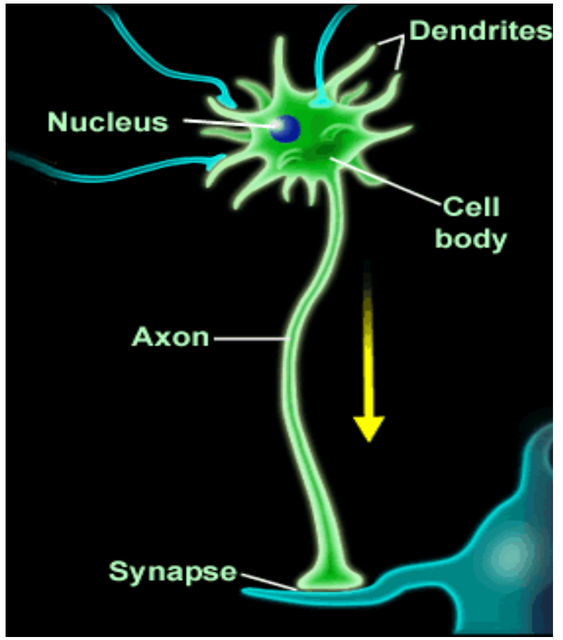

Image
Output
13

Moore-like law for ML (1995-2018)

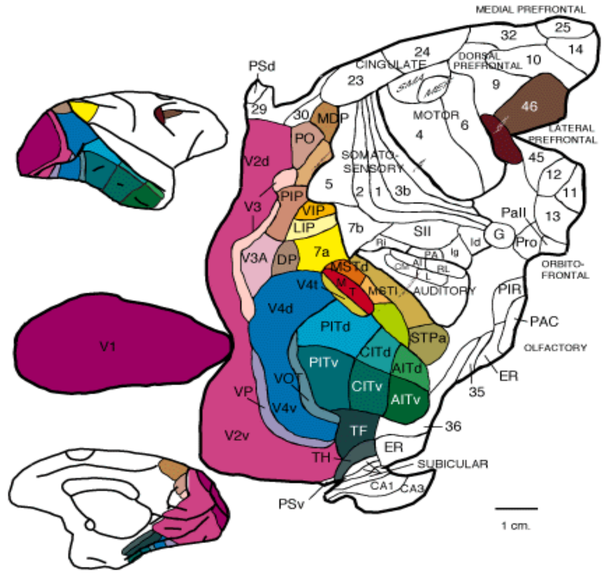


Vision: what is where

- Human Brain
  - 10^{10}-10^{11} neurons (~1 million flies)
  - 10^{14}- 10^{15} synapses

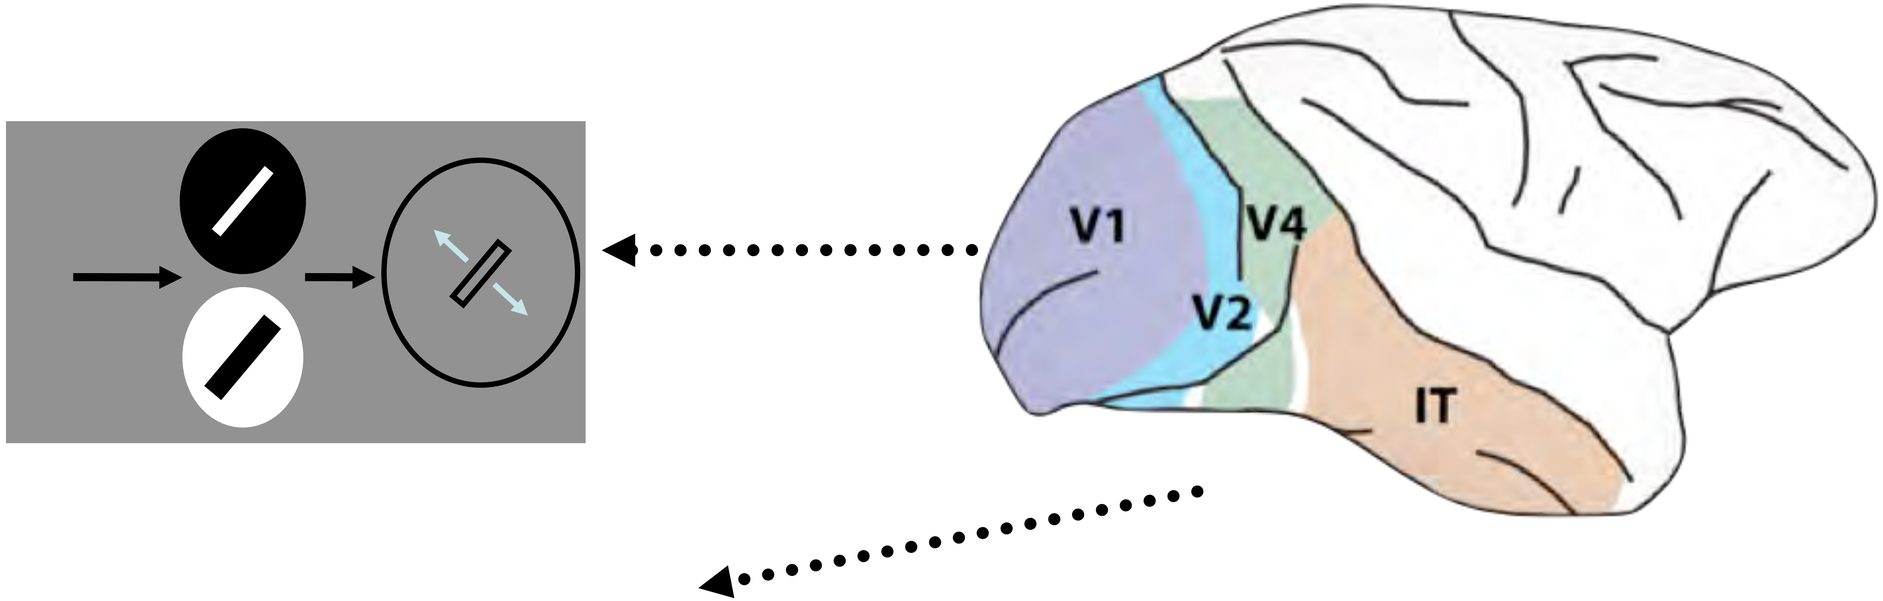


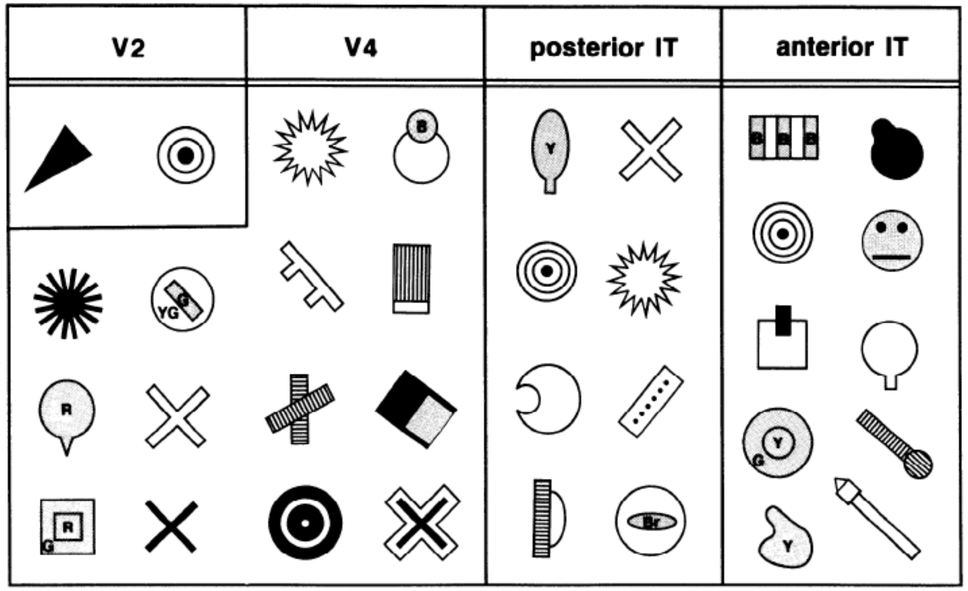
- Ventral stream in rhesus monkey
  - ~10^{9} neurons in the ventral stream (350 10^{6} in each emisphere)
  - ~15 10^{6} neurons in AIT (Anterior InferoTemporal) cortex



Van Essen & Anderson, 1990

Vision: ventral stream





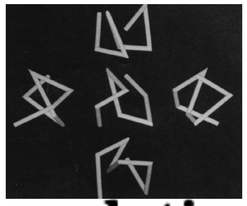
The ventral stream hierarchy: V1, V2, V4, IT

A gradual increase in the receptive field size, in the "**complexity**" of the preferred stimulus, in **"invariance"** to position and scale changes

Kobatake & Tanaka, 1994

Cognition in people

Proc. Natl. Acad. Sci. USA
Vol. 89, pp. 60–64, January 1992
Psychology



Psychophysical support for a two-dimensional view interpolation theory of object recognition

(regularization networks/computer graphics psychophysics/generalization)

HEINRICH H. BÜLTHOFF*† AND SHIMON EDELMAN‡

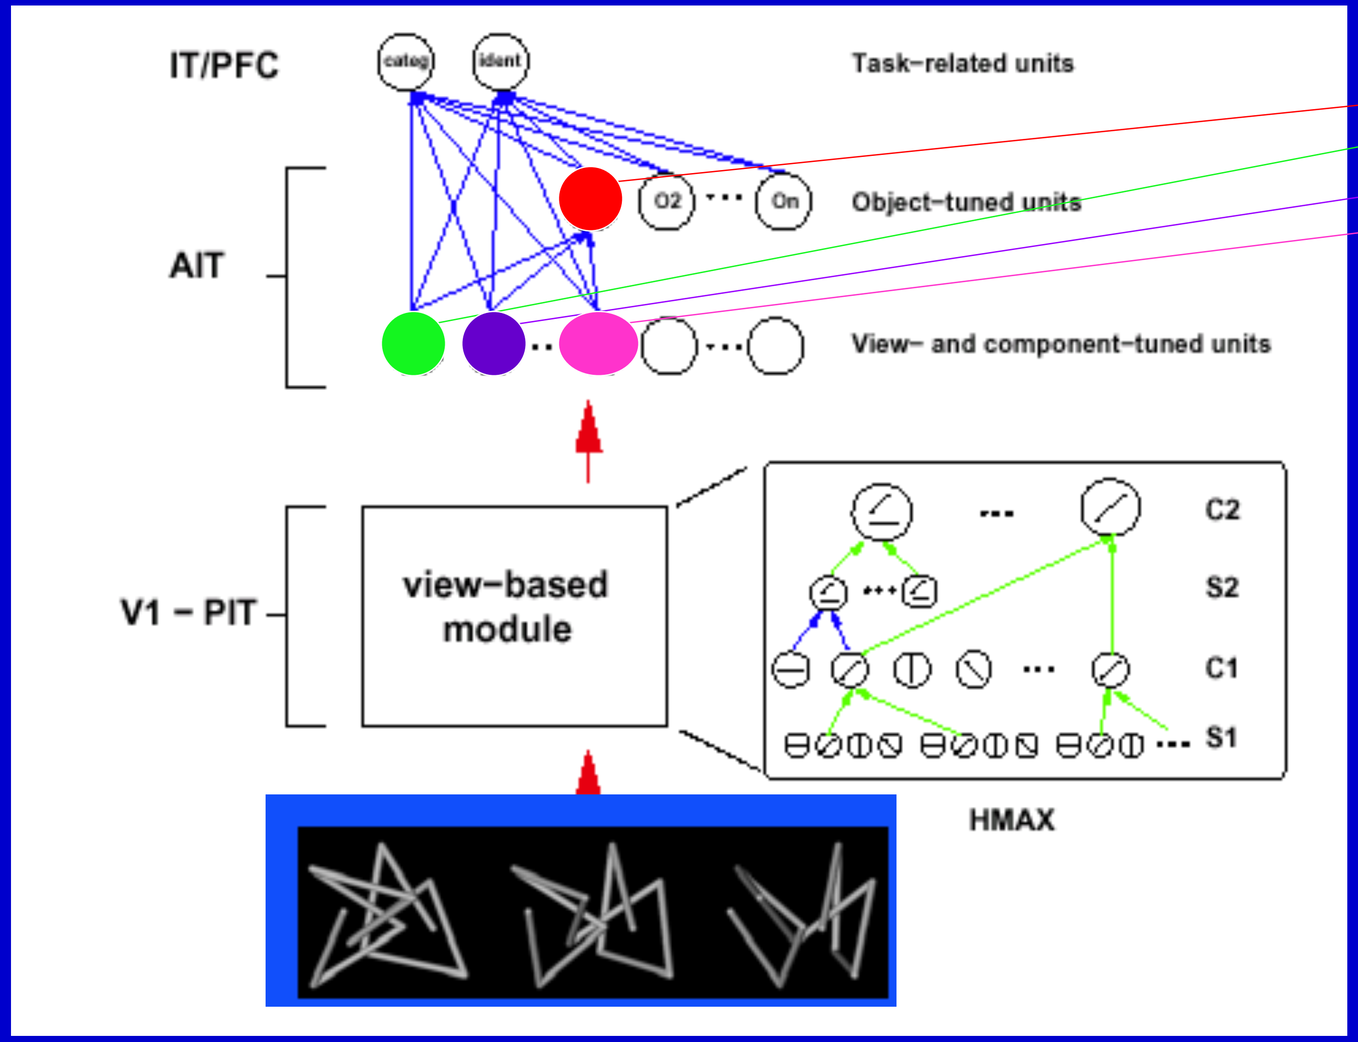
Shape representation in the inferior temporal cortex of monkeys

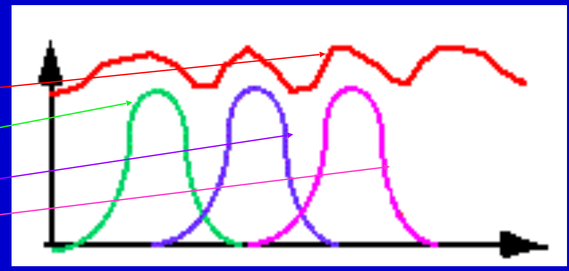
Nikos K. Logothetis*, Jon Pauls* and Tomaso Poggio†

*Division of Neuroscience, Baylor College of Medicine, One Baylor Plaza, Houston, Texas 77030, USA. †Center for Computational and Biological Learning, and Department of Brain Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA.



Model's early predictions: neurons become view-tuned during recognition



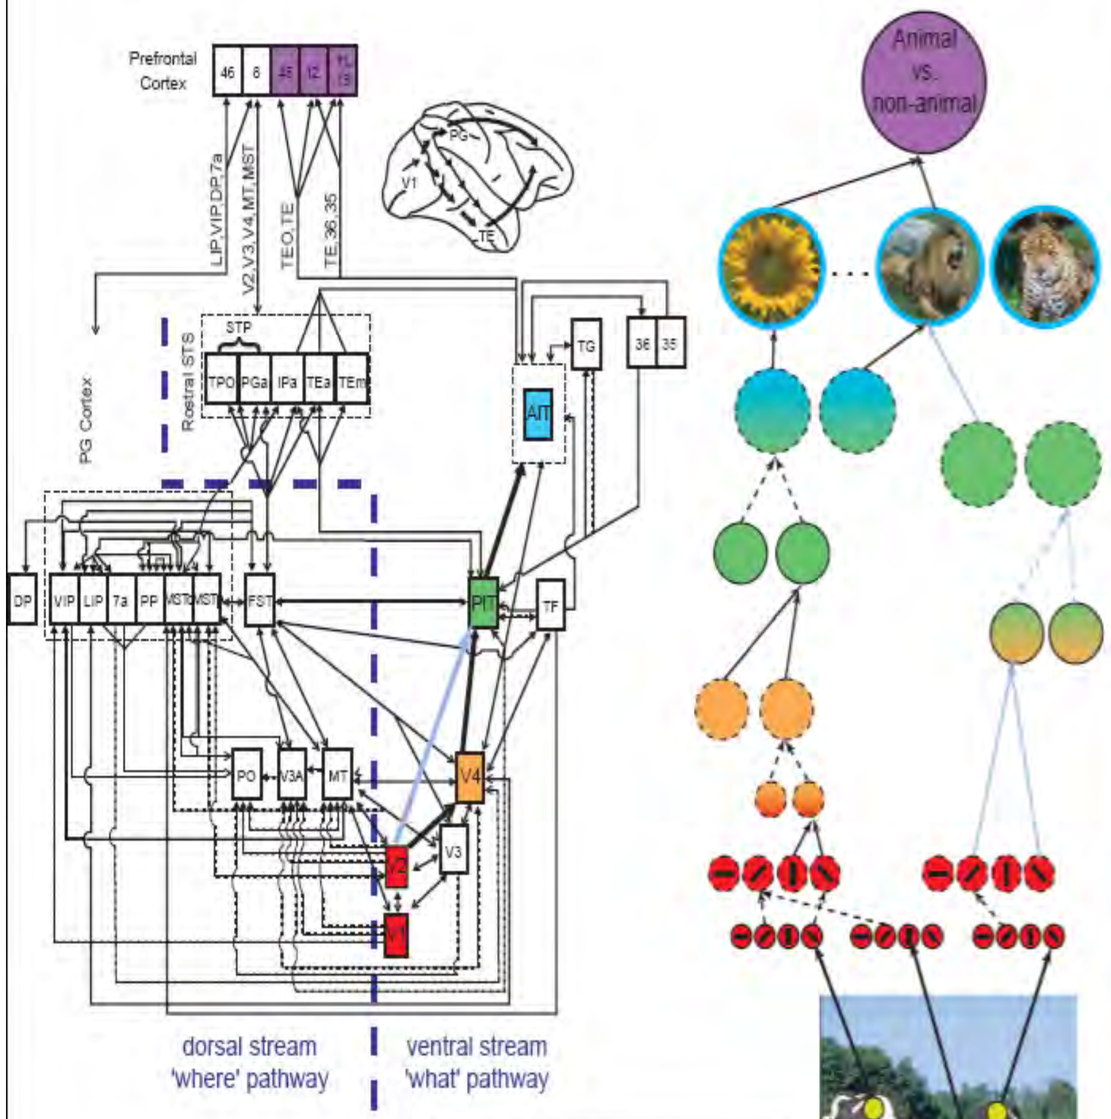


Logothetis, Pauls, and Poggio, 1995;
Logothetis, Pauls, 1995

Poggio, Edelman, Riesenhuber (1990, 2000)

A model of the ventral stream in visual cortex

(starting from work with Buelthoff and Logothetis)





Model layers	Corresponding brain area (tentative)	RF sizes	Number units	
classifier	PFC		$1.0\ 10^0$	
S4	AIT	>4.4°	$1.5\ 10^2$	~ 5,000 subunits
C3	PIT - AIT	>4.4°	$2.5\ 10^3$	
C2b	PIT	>4.4°	$2.5\ 10^3$	
S3	PIT	1.2°- 3.2°	$7.4\ 10^4$	~ 100 subunits
S2b	V4 - PIT	0.9°- 4.4°	$1.0\ 10^7$	~ 100 subunits
C2	V4	1.1°- 3.0°	$2.8\ 10^5$	
S2	V2 - V4	0.6°- 2.4°	$1.0\ 10^7$	~ 10 subunits
C1	V1 - V2	0.4°- 1.6°	$1.2\ 10^4$	
S1	V1 - V2	0.2°- 1.1°	$1.6\ 10^6$	

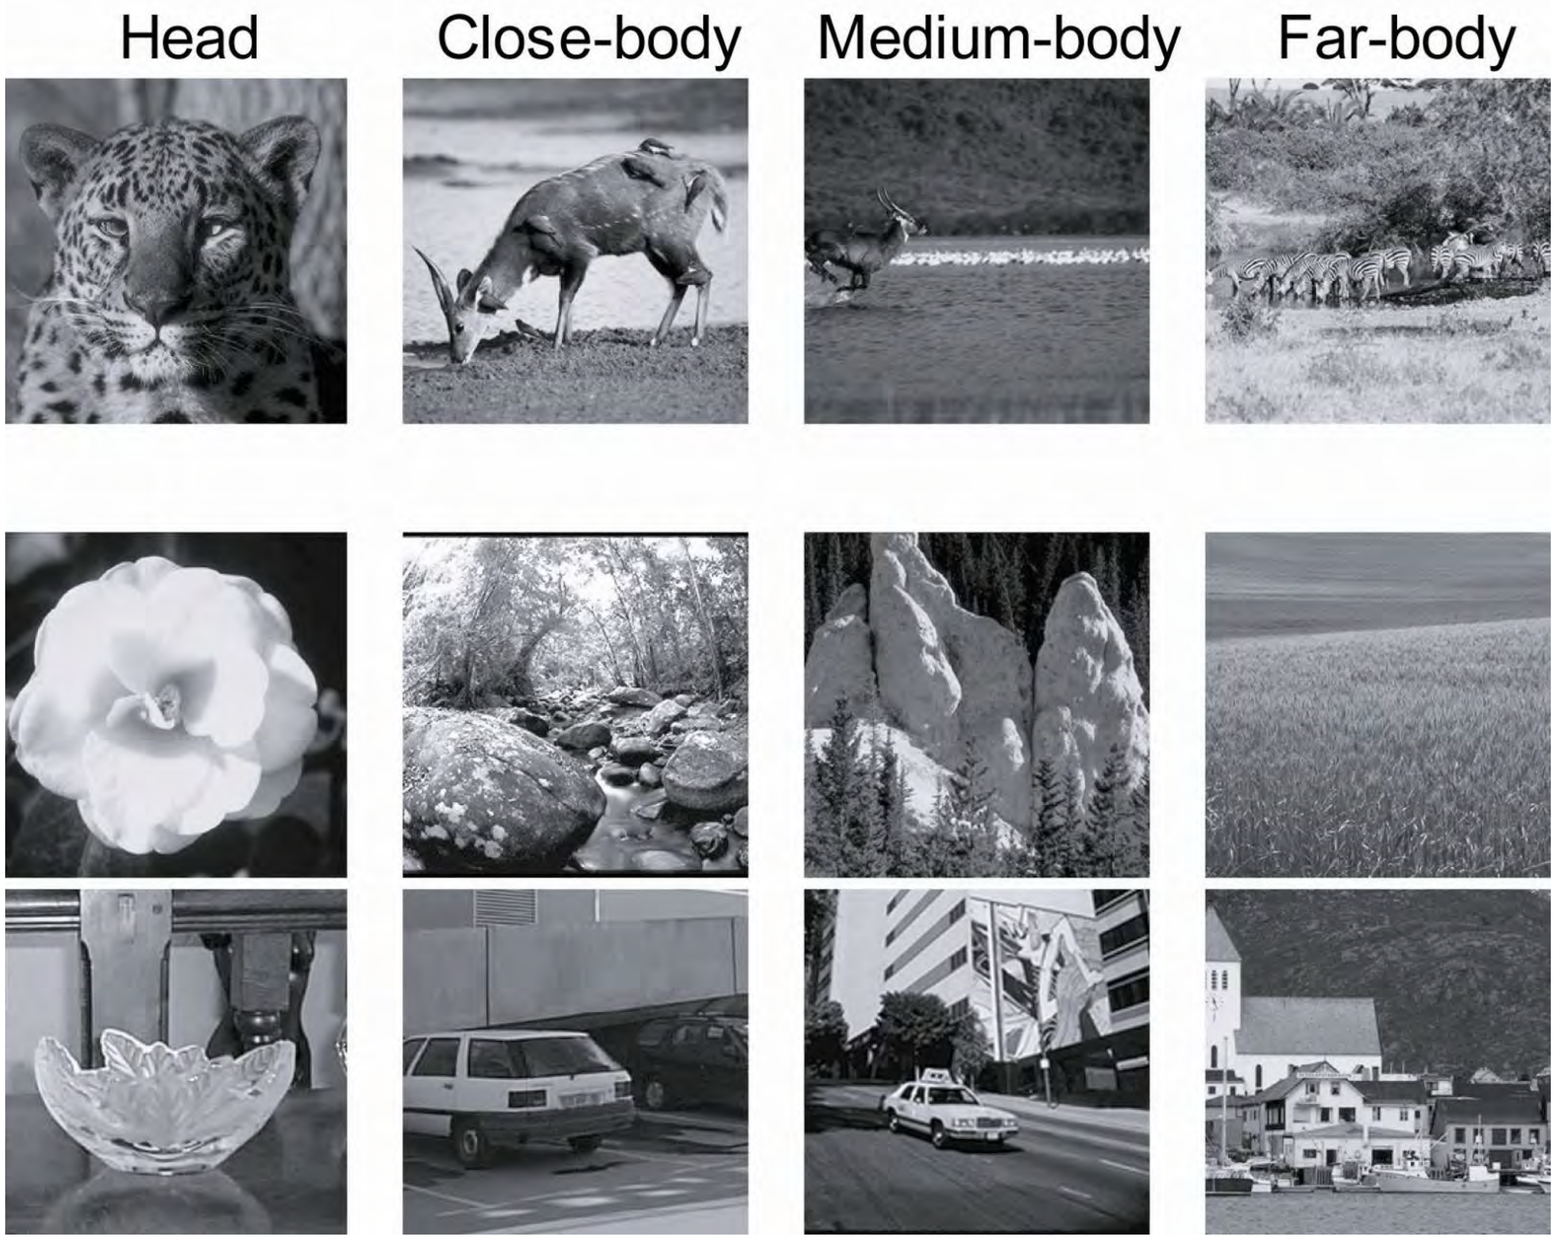
Supervised task-dependent learning

Unsupervised task-independent learning

increase in complexity (number of subunits), RF size and invariance

Riesenhuber & Poggio 1999, 2000;
Serre Kouh Cadieu Knoblich Kreiman & Poggio 2005; Serre Oliva Poggio 2007

Psychophysics of rapid categorization



Database collected by Oliva & Torralba

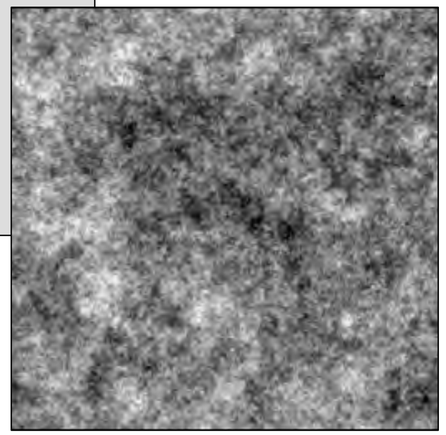
Rapid categorization task (with mask to test feedforward model)



Image

Interval

Image-Mask



Mask

1/f noise

20 ms

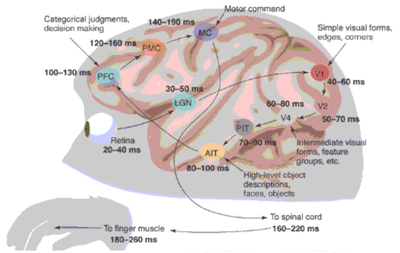
30 ms ISI

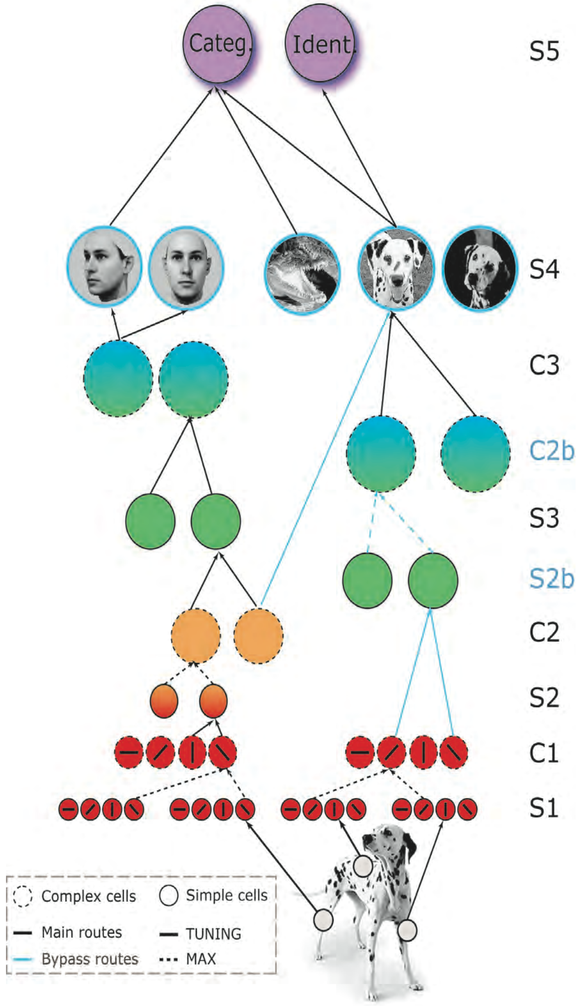
80 ms

Animal present or not ?

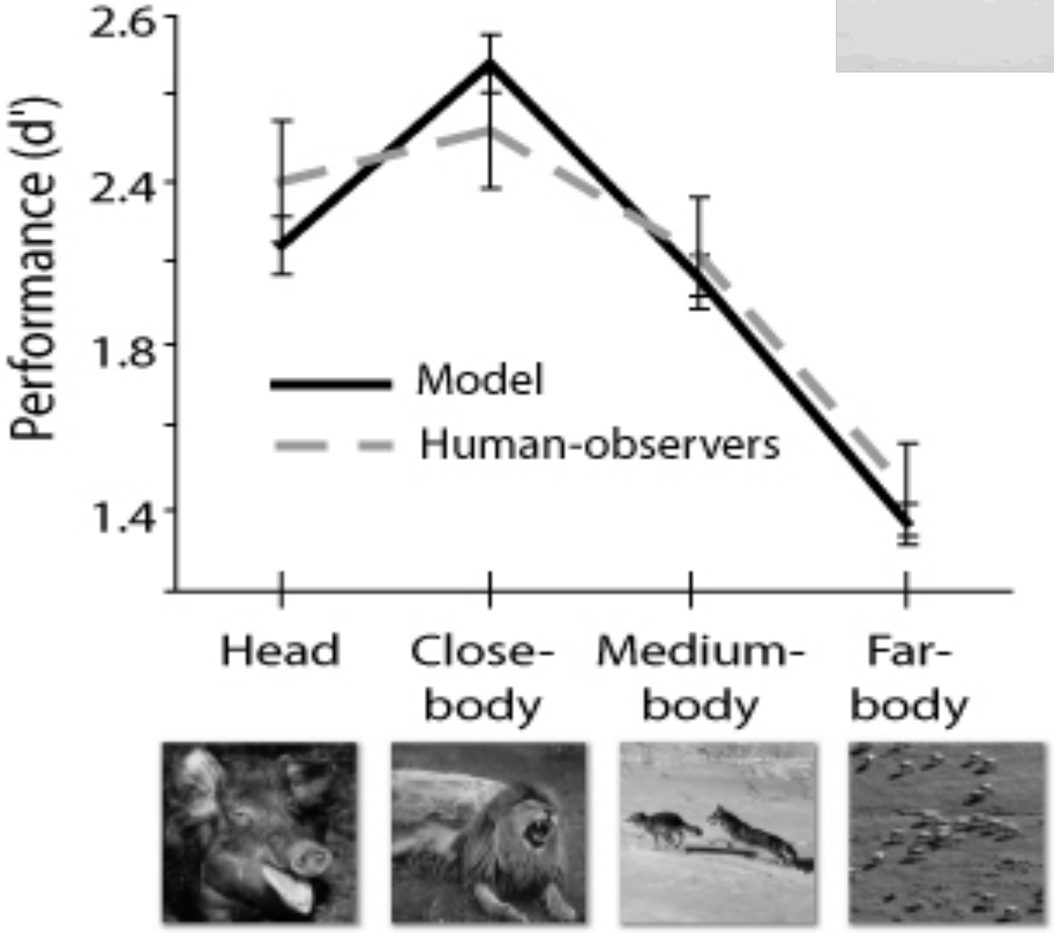
Thorpe et al 1996; Van Rullen & Koch 2003; Bacon-Mace et al 2005

Hierarchical feedforward models of the ventral stream



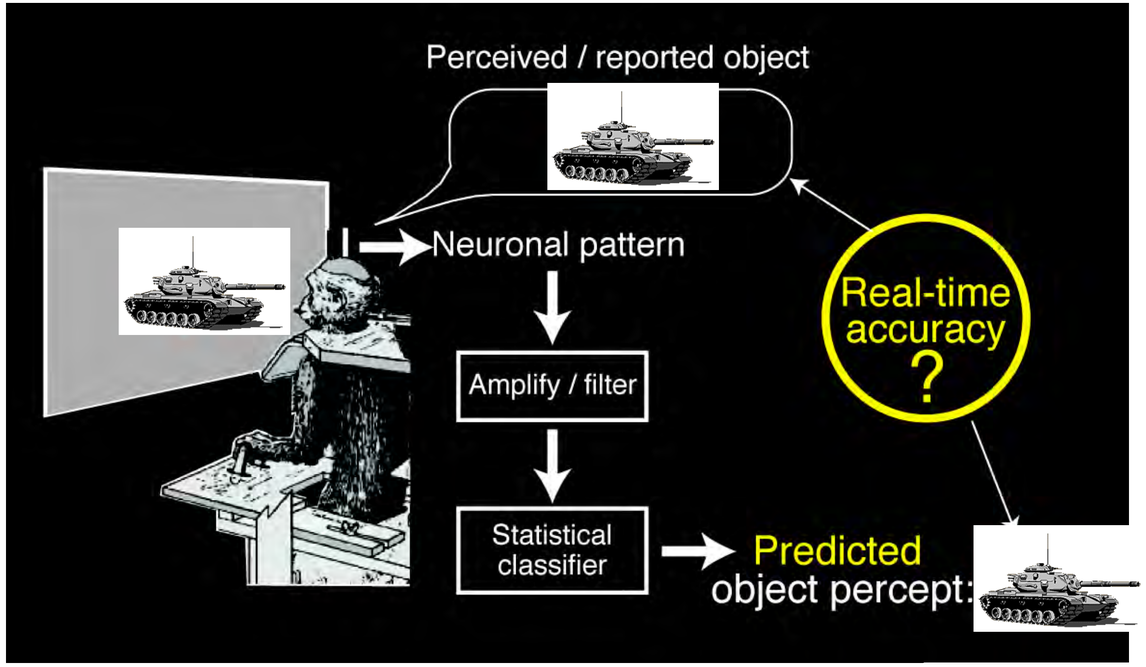


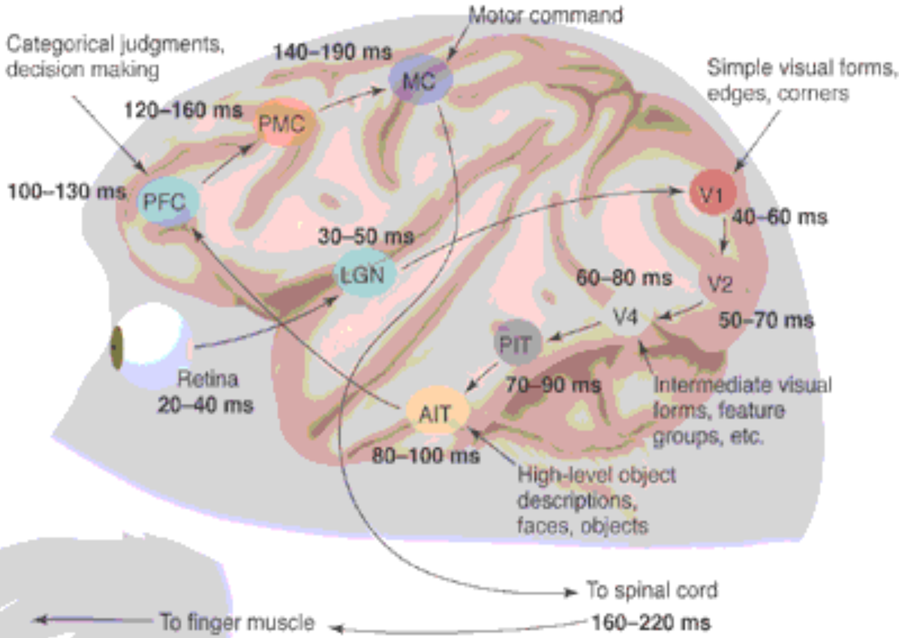
Feedforward Models:
"predict" rapid categorization
(82% model vs. 80% humans)





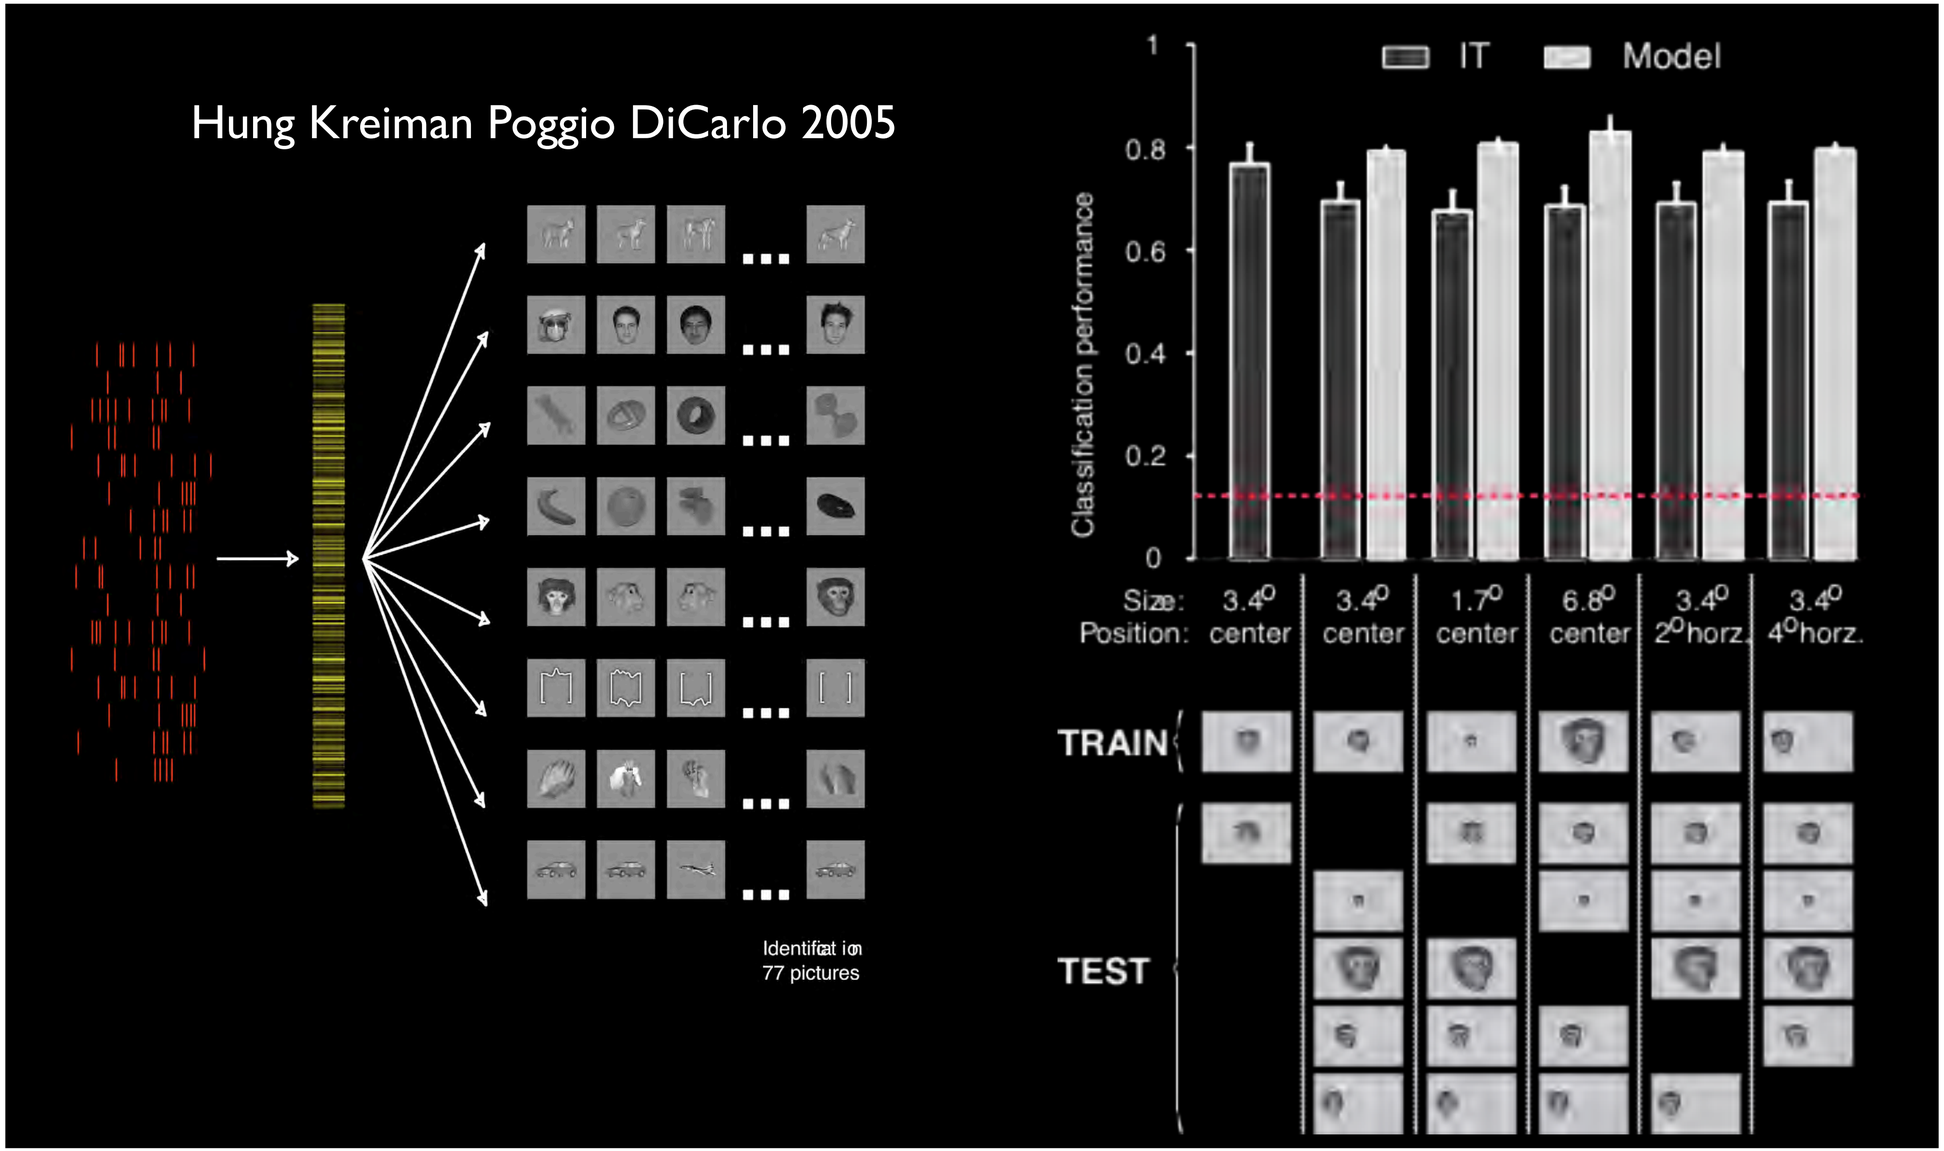
Decoding the neural code: Matrix-like read-out from the brain





Agreement of model w| IT Readout data

Reading out category and identity invariant to position and scale



Serre Kouh Cadieu Knoblich Kreiman & Poggio 2005

.....



in 2013.....