Visual recognition: peeking inside computations in the brain

Gabriel Kreiman

Gabriel.kreiman@tch.harvard.edu

klab.tch.harvard.edu

Center for Brains, Minds and Machines

Camille Gomez

Richard Born

Jojo Nassi

Laura Groomes

Hanlin Tang

Thomas Miconi

Bill Lotter

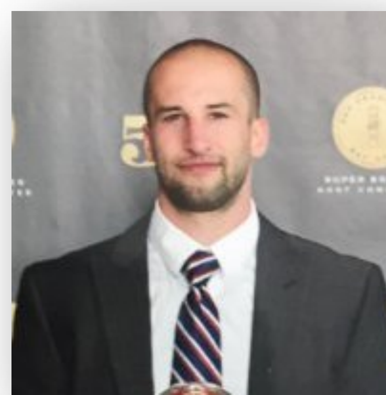

David Cox

Biologically-inspired computations are powerful

Over millions of years of evolution, "interesting" solutions to difficult problems have emerged through changes in neuronal circuits

- Hardware and software that work for many decades
- Parallel computation (with serial bottlenecks)
- Reprogrammable architecture
- Low power
- Single-shot learning
- "Discover" structure in data
- Fault tolerance
- Robustness to sensory transformations
- Component interaction and integration of sensory modalities

Algorithms, solutions [redacted]

Bottom-up models of object recognition

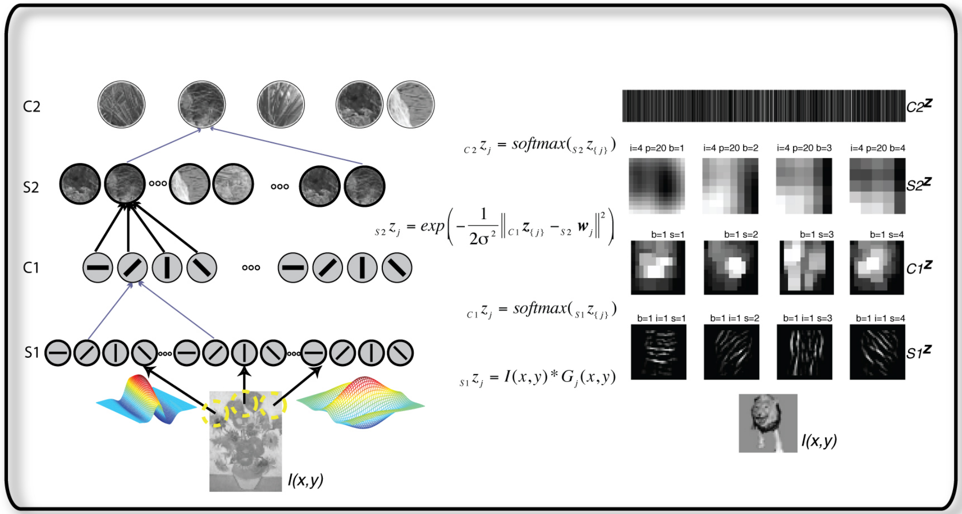

Fukushima, Mel, Olshausen, LeCun, Riesenhuber, Rolls, DiCarlo, …

Serre et al 2007

Deep convolutional networks

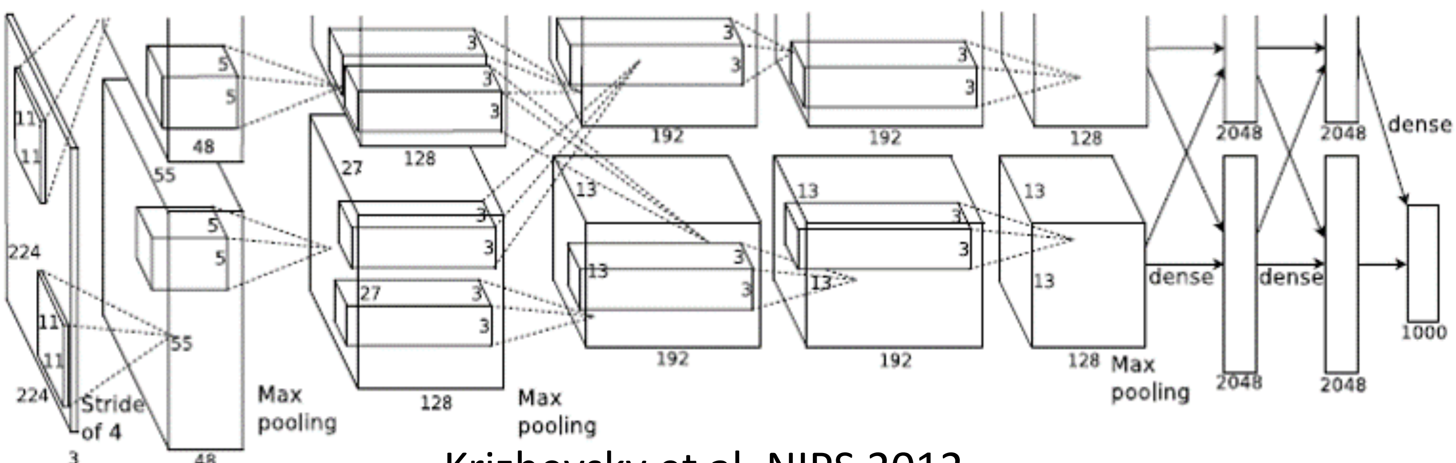

Krizhevsky et al, NIPS 2012

There is more. Much more.

A schematic diagram of visual cortex connections in macaque monkeys

This is a major oversimplification …

Each box contains a bewildering and magnificent world of computations

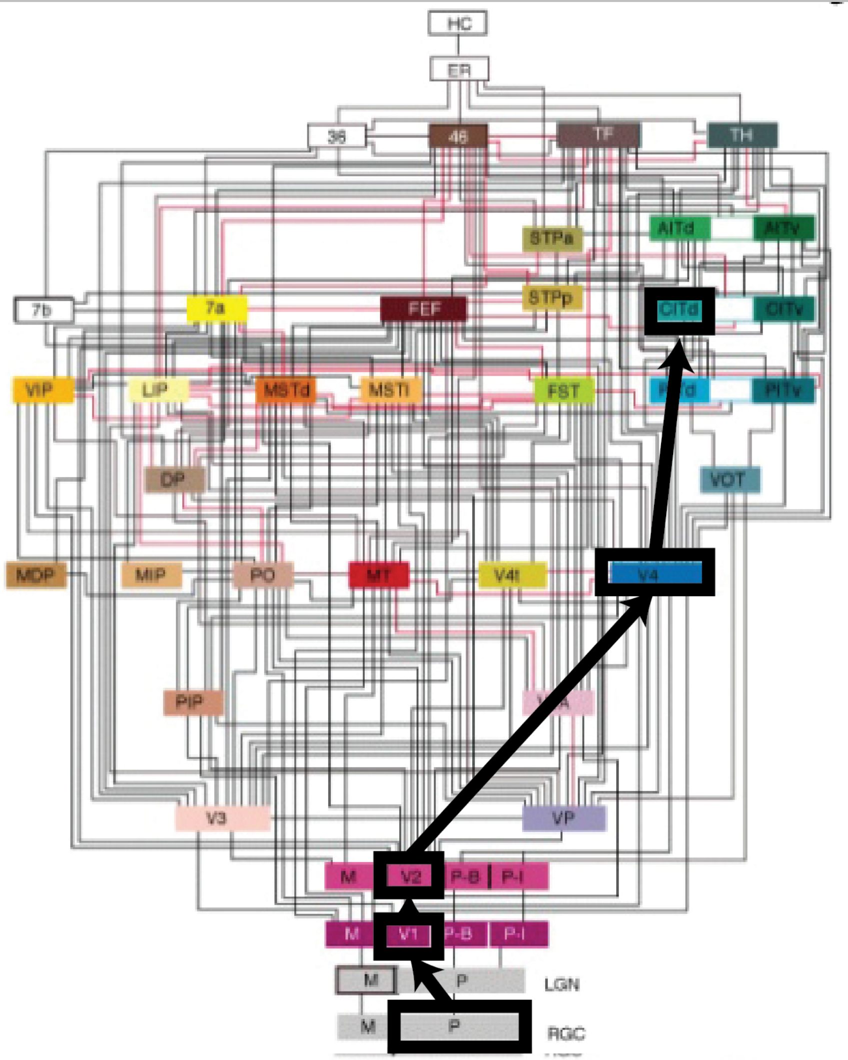

Felleman and Van Essen 1991

Why are there so many feedback and recurrent connections?

There are more horizontal + top-down projections than bottom-up ones (e.g. Douglas 2004, Callaway 2004)

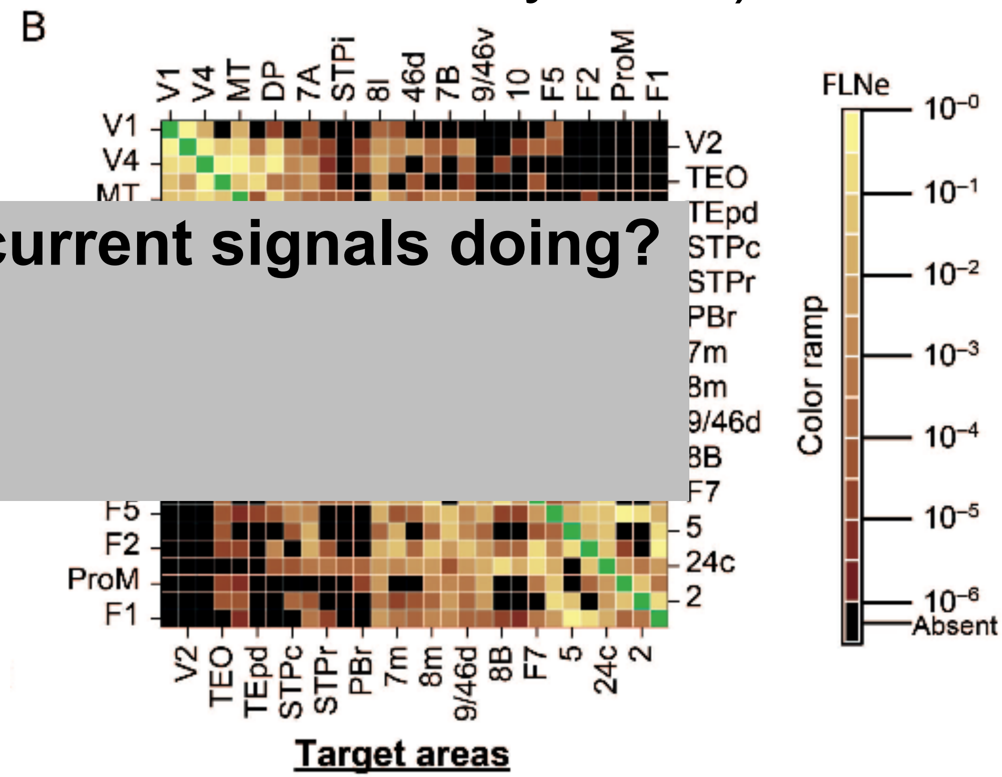

What are feedback/recurrent signals doing?
When?
Why?
How?

Markov et al 2014

Computational roles of recurrent/feedback signals

1. Pattern completion (recurrent computations)
2. Predictive coding (feedback computations)

Image by Hanlin Tang

Pattern completion as a hallmark of intelligence

A, C, E, G,	**I**
1, 2, 3, 5, 7, 11,	**13**
V-s-a- R-c-g-i-i-n	**Visual Recognition**
Even though it was raining heavily, Jonathan decided to go out without an	**Umbrella**

Also:
Other sensory modalities
Music
Social interactions

Objects can be recognized from partial information

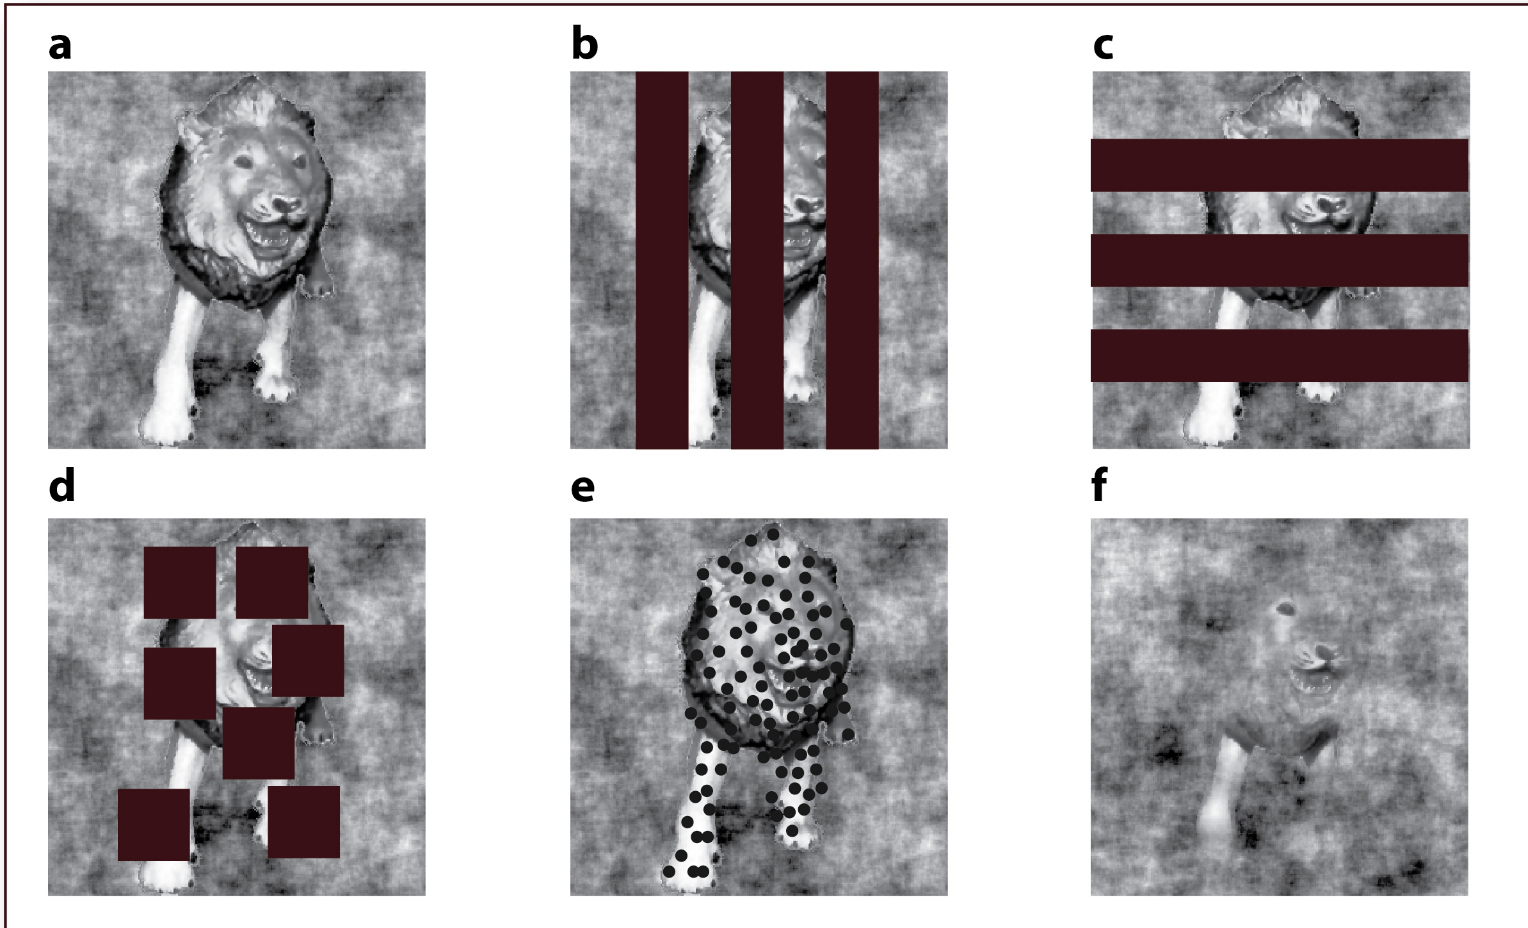

4 bubbles

Evaluating pattern completion

20 bubbles

10 bubbles

6 bubbles

4 bubbles

A

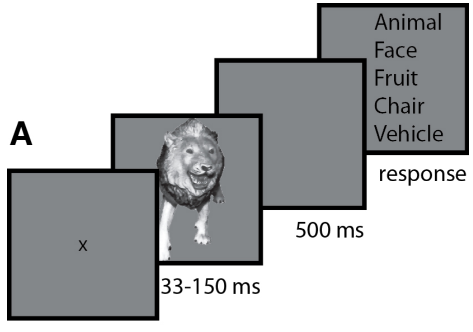

C

Whole

Partial

D

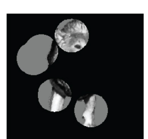

Occluded

Hanlin Tang, Bill Lotter, Martin Schrimpf

Strong robustness to limited visibility

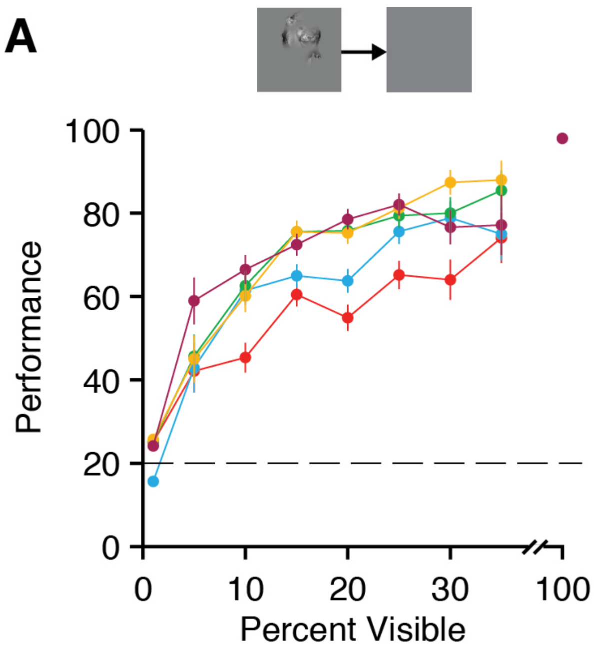

Backward masking interrupts processing (presumably of feedback/recurrent computations)

<u>Models</u>:

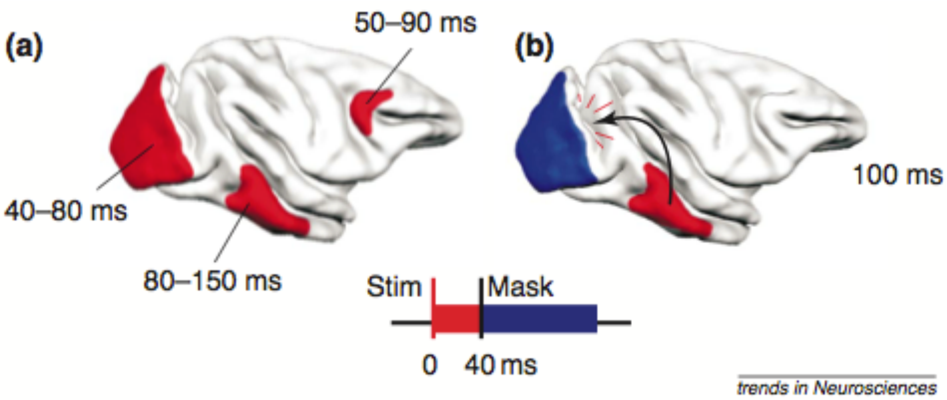

Lamme V, Roelfsema P (2000)

<u>Masks</u>:

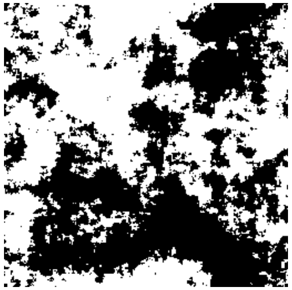

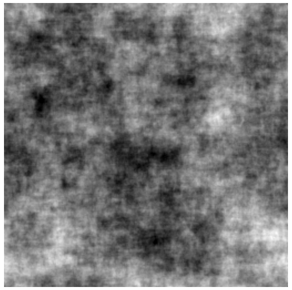

- Short delays (SOA<20ms): mask reduces visibility
- Longer delays: mask is purported to disrupt recurrent/top-down processing

V1: Bridgeman 1980, Maknik and Livinsgtone 1998, Lamme et al 2002
IT: Kovacs et al 1995, Rolls et al 1999

Evaluating pattern completion abilities

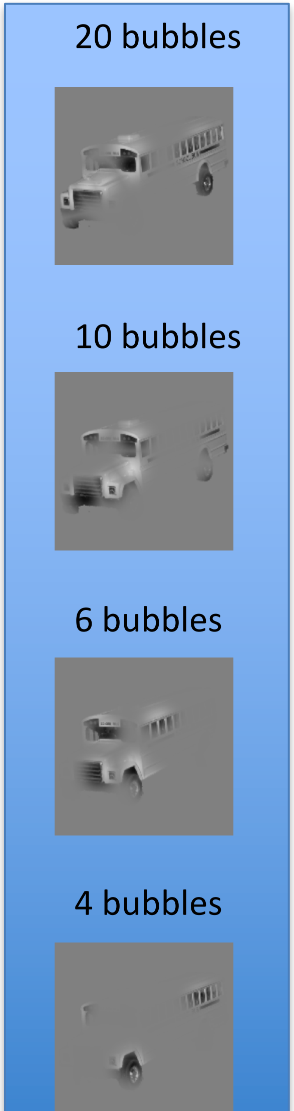

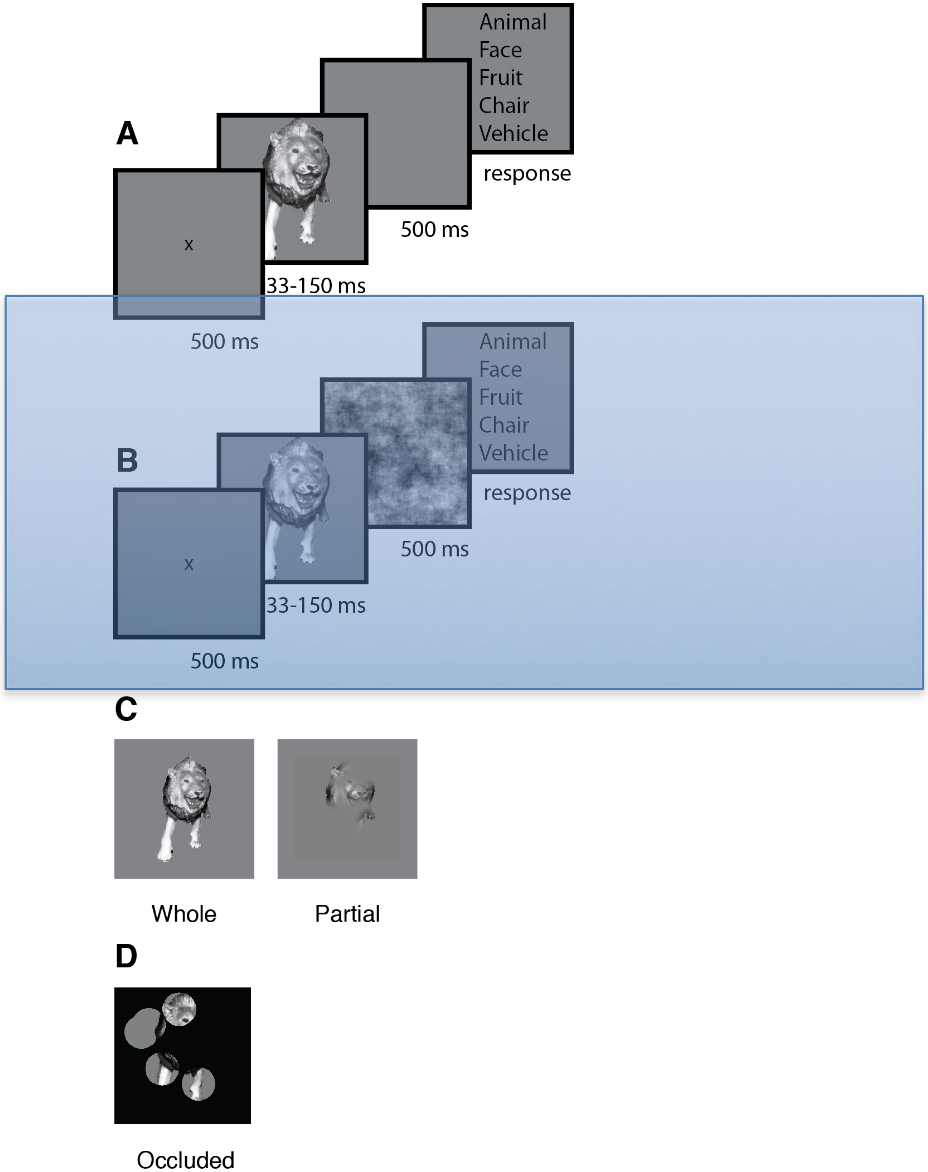

Backward masking disrupts pattern completion

E

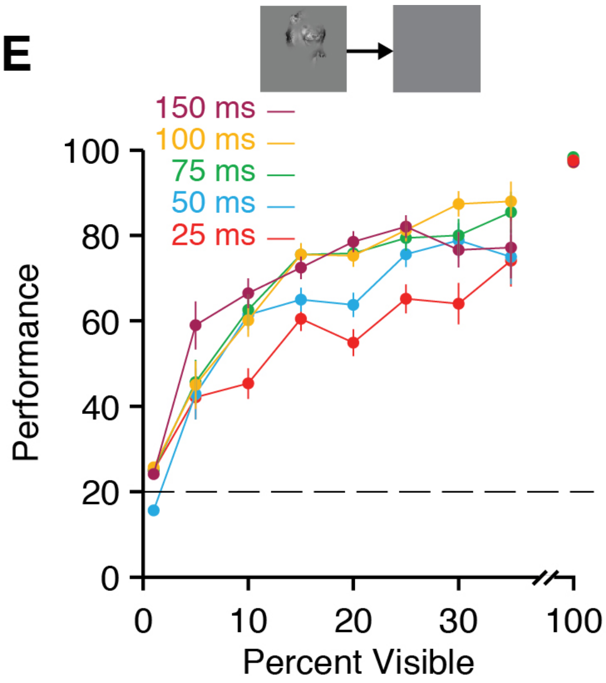

Peeking inside the human brain

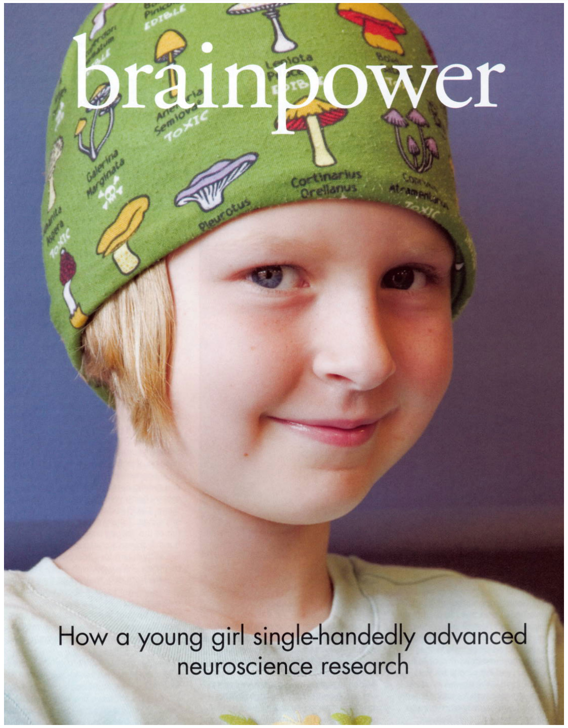

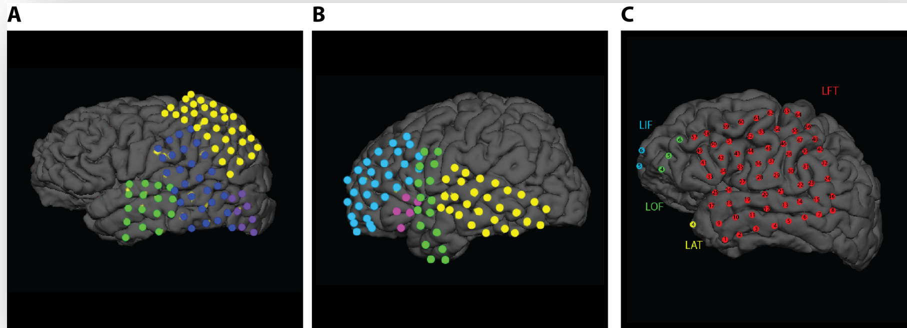

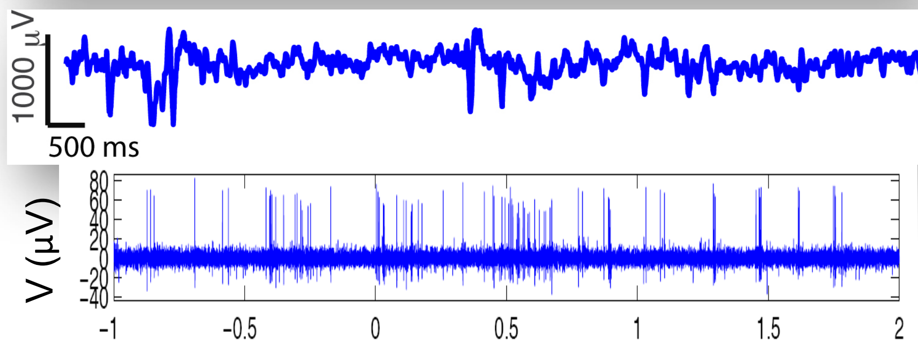

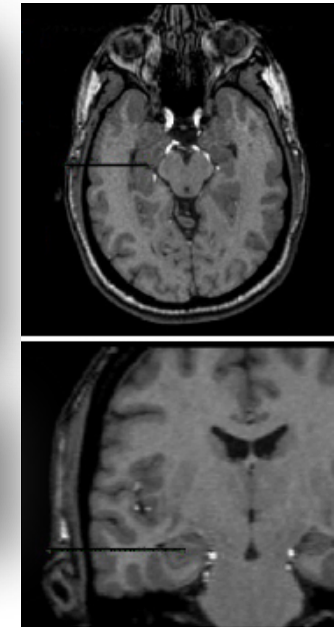

- Patients with pharmacologically intractable epilepsy
- Multiple electrodes implanted to localize seizure focus
- Patients stay in the hospital for about 7-10 days
- All experiments are approved by the Institutional Review Boards
- All testing is performed with the subjects’ consent

Neurosurgeons: **William Anderson, Joseph Madsen, Itzhak Fried**

Reliable, selective and rapid responses in human inferior temporal cortex

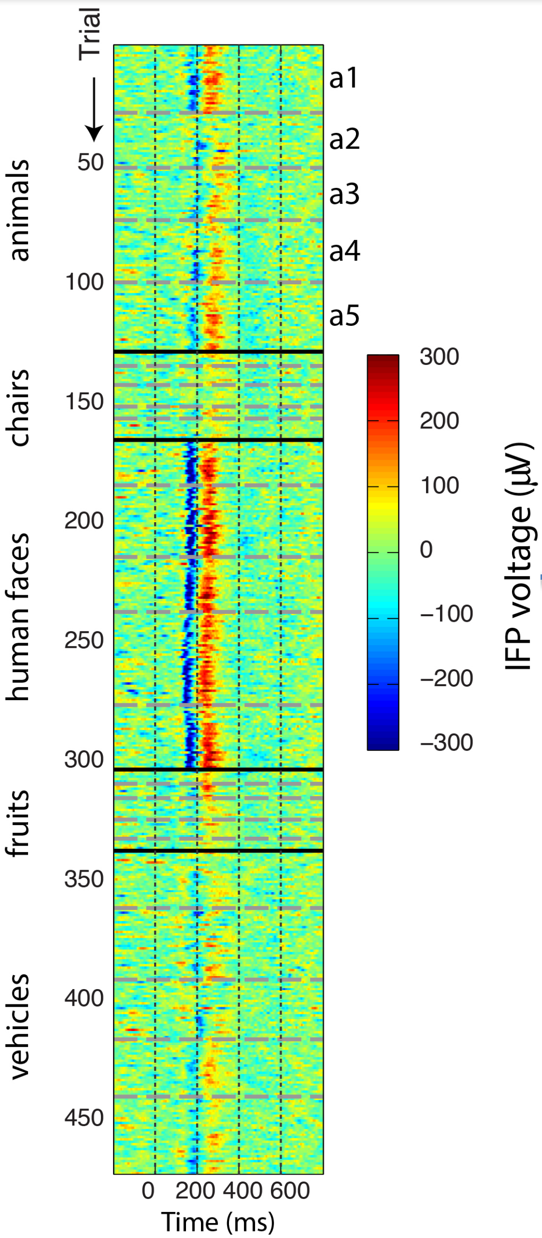

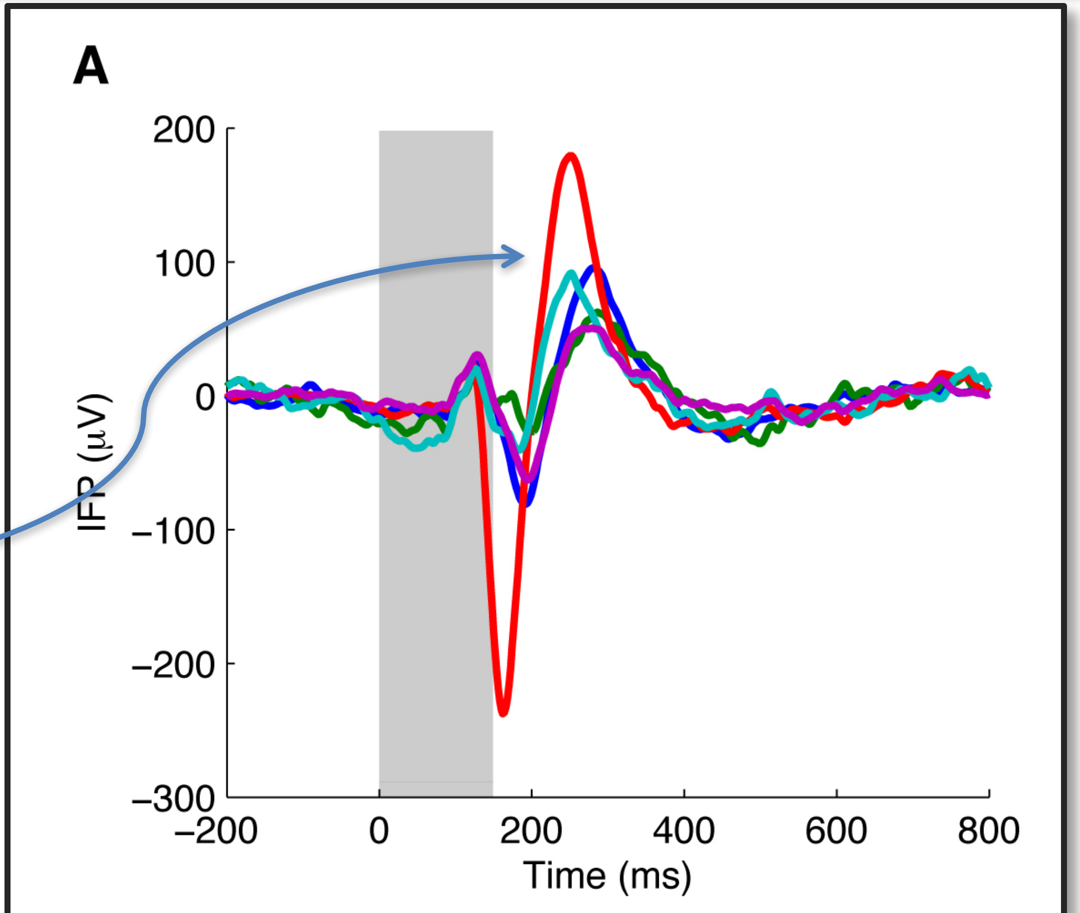

Inferior temporal gyrus

Tang et al 2014

Visual selectivity along the human ventral visual cortex

2205 electrodes
27 subjects

Main areas showing visual selectivity
Inferior-occipital gyrus
Fusiform gyrus
Medial temporal gyrus
Inferior temporal gyrus
Temporal pole

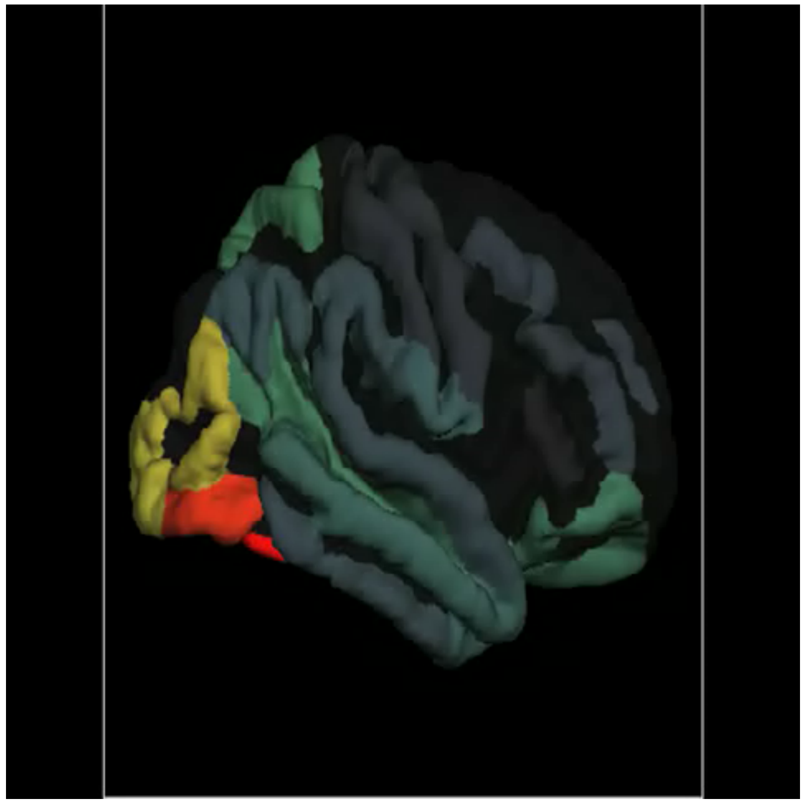

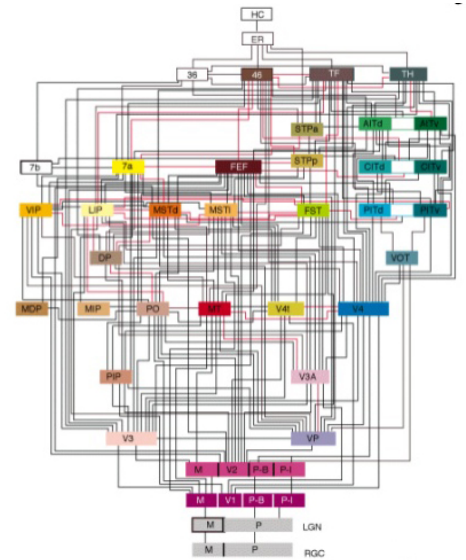

Agam, Liu, Singer, Tang

Gross, Desimone, Logothetis, Richmond, Tanaka, DiCarlo, Rolls, Connors, Ito, Vogels

Example responses during object completion

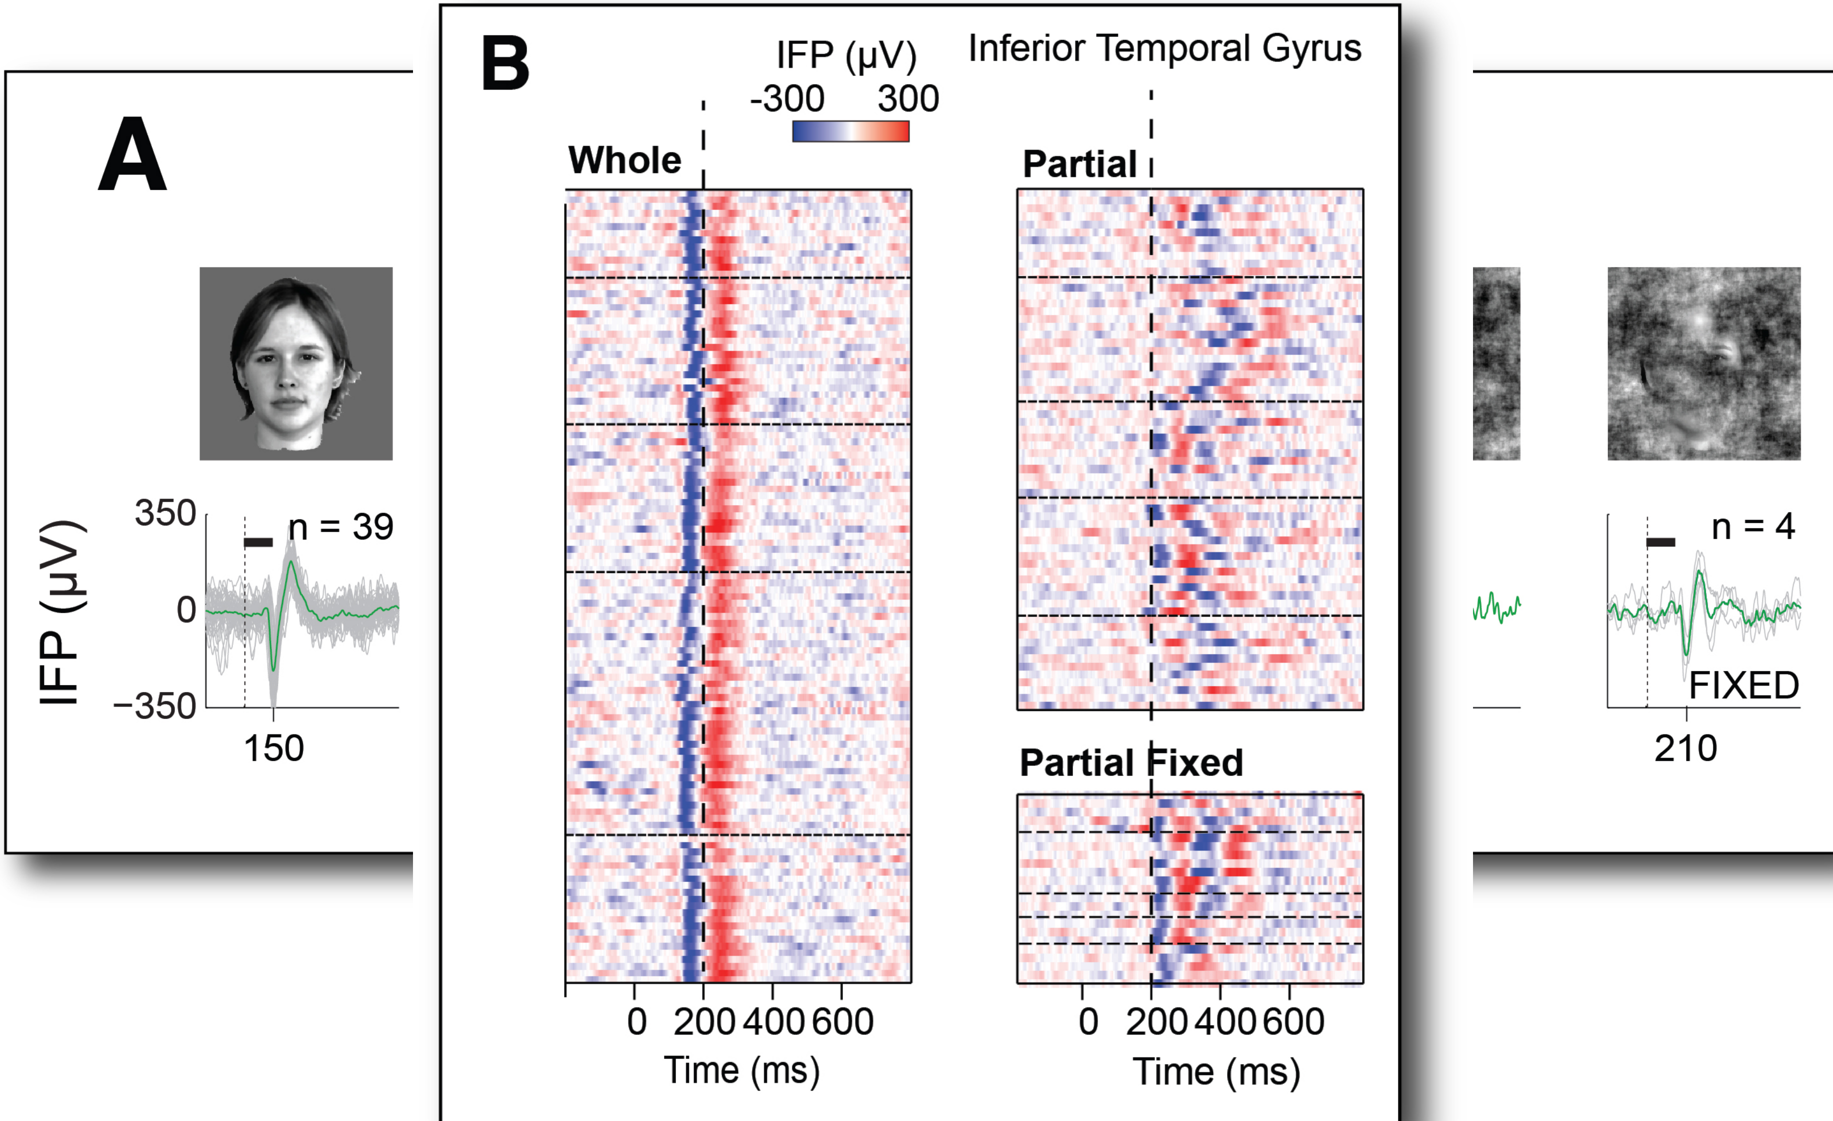

Inferior Temporal Gyrus

Tang et al, 2014

The behavioral effect of masking correlated with the neural response latency on an image-by-image basis

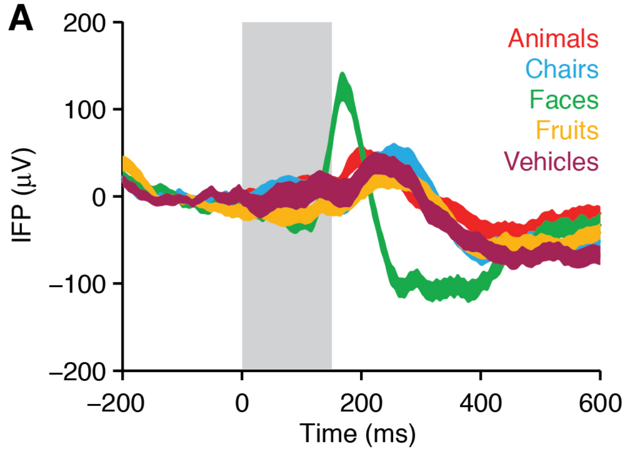

Bottom-up models significantly underperform in recognition of partial images

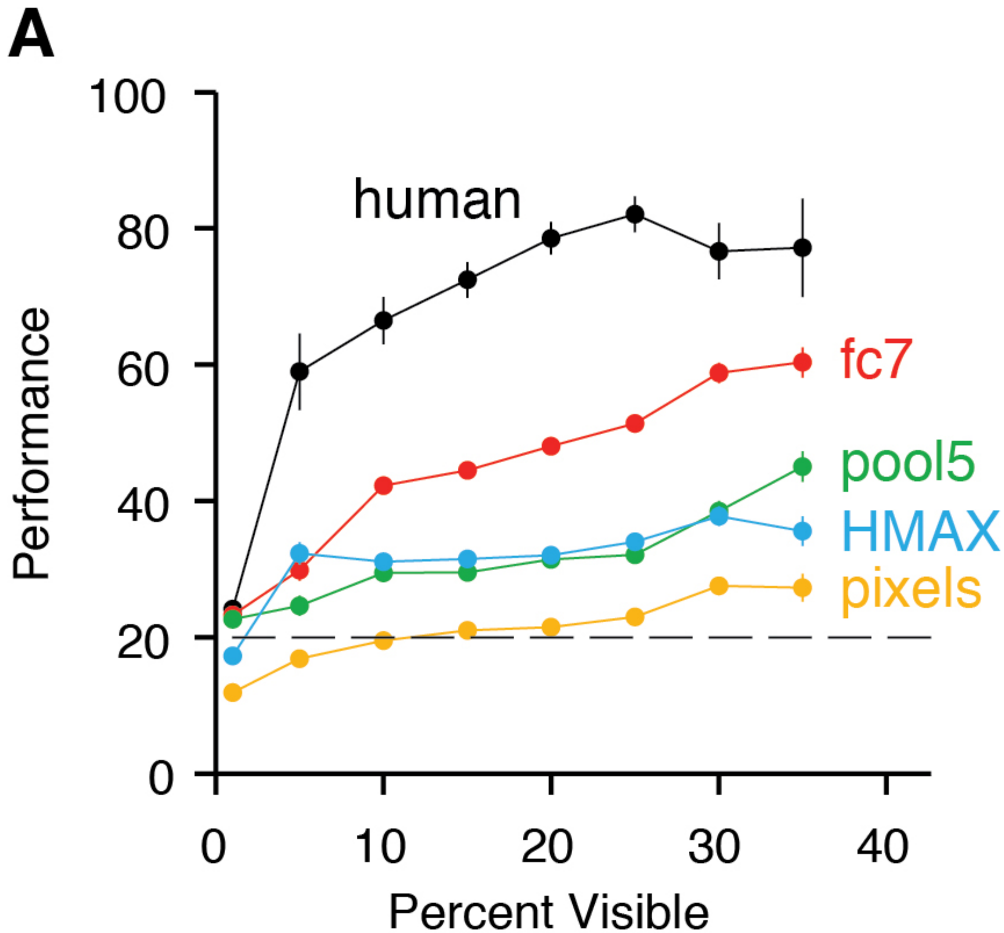

See also Pepik et al 2015, Wyatte et al 2012

2D object representation at the top of the model hierarchy is not robust to occlusion

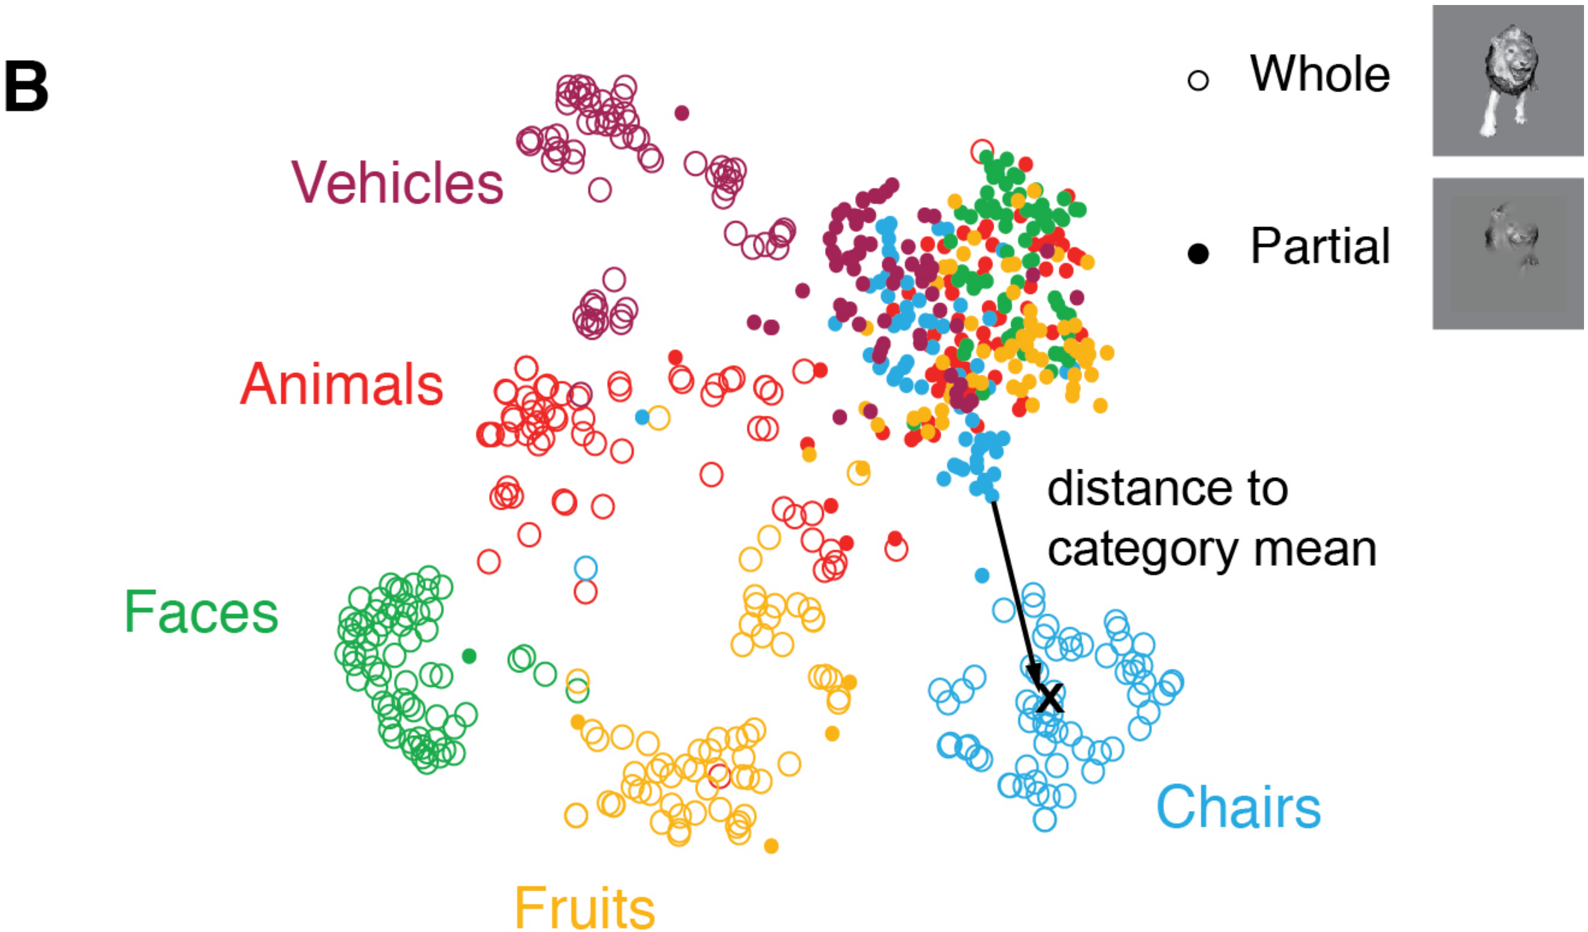

Stochastic neighborhood embedding. Van der Maaten 2008

The neural latency for each image was correlated with the distance to category center

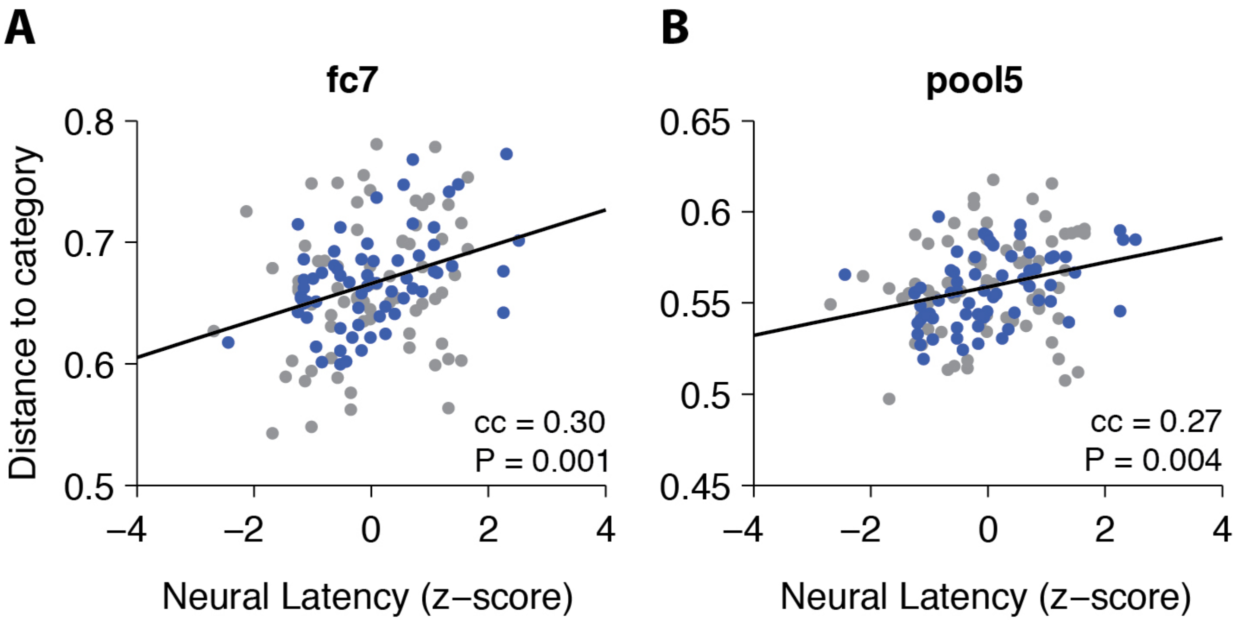

Hopfield network with binary neurons

Each neuron i has two states: V_i=0 or V_i=1

Ensemble: $\mathbf{V} = [V_1, V_2, ..., V_N]$ Note: $\mathbf{V}$=$\mathbf{V}$(t)

Synaptic strength: T_{ij}

If two neurons are not connected: $T_{ij}=0$

No self connections: $T_{ii}=0$

Update rule: $V_i(t)=1$ iff $\sum_j T_{ij} V_j(t) > 0$

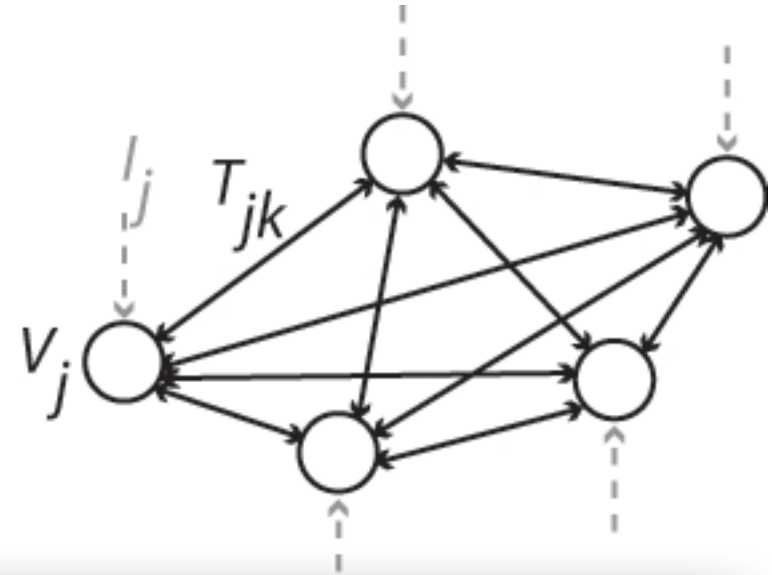

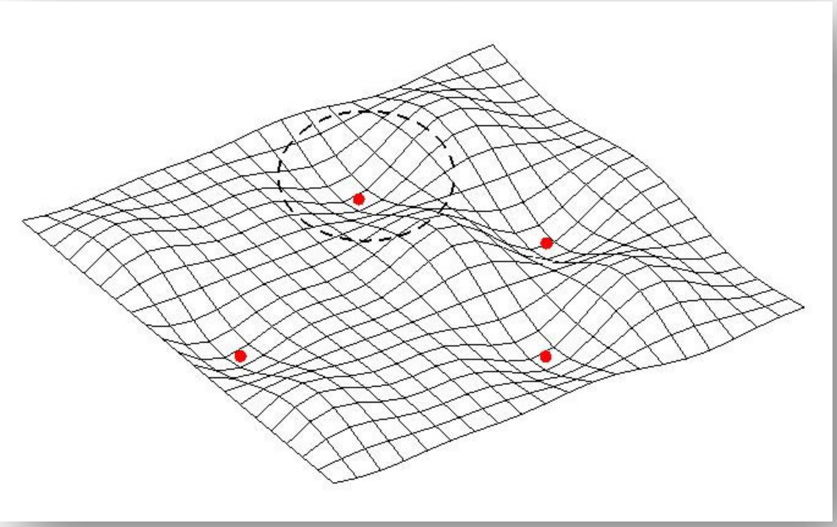

Hopfield, 1982

A recurrent network may ameliorate the problem of missing information

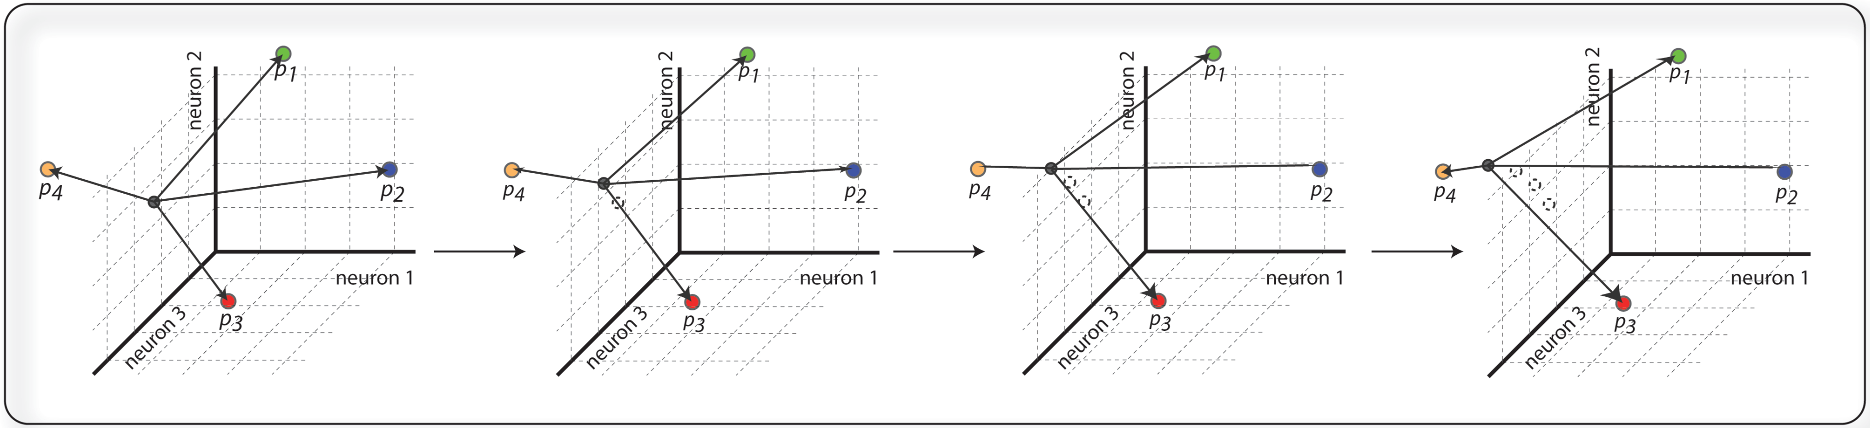

$\boldsymbol{p}$ = prototypes (fixed)

Recurrent Hopfield network (RNN_h) improves recognition performance for partial images

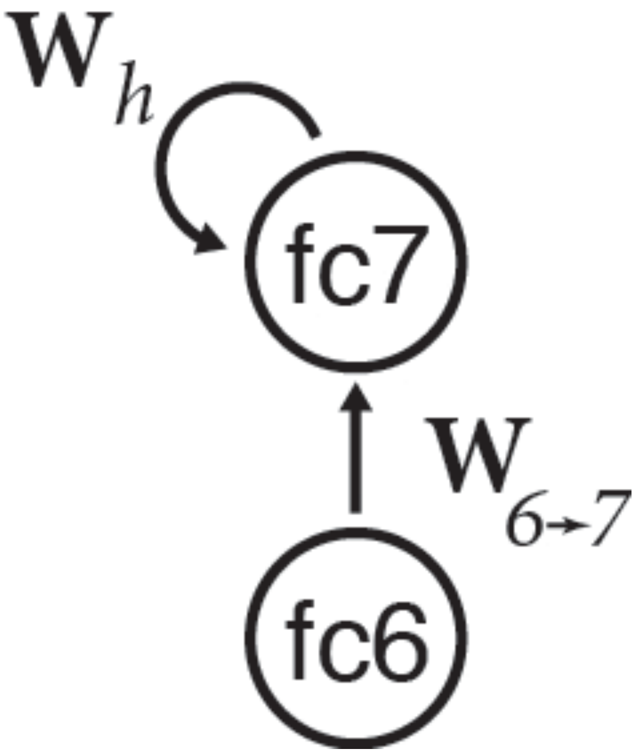

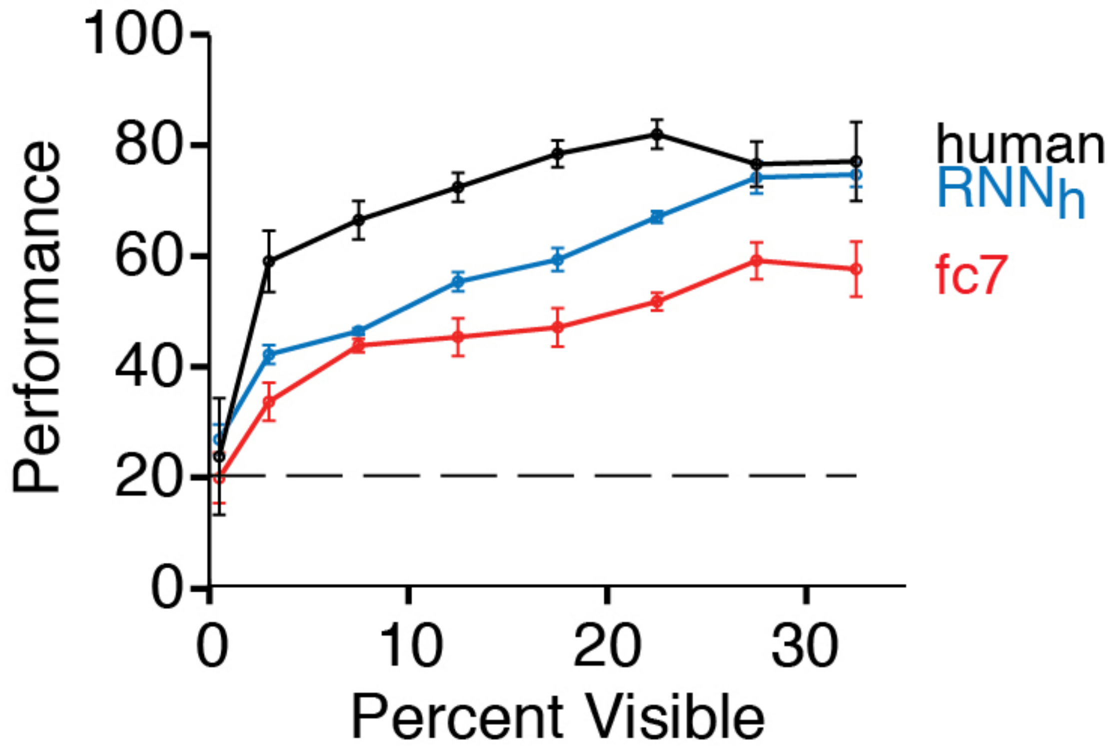

NOTE: 0 free parameters

Training with occluded objects leads to matching human performance in pattern completion

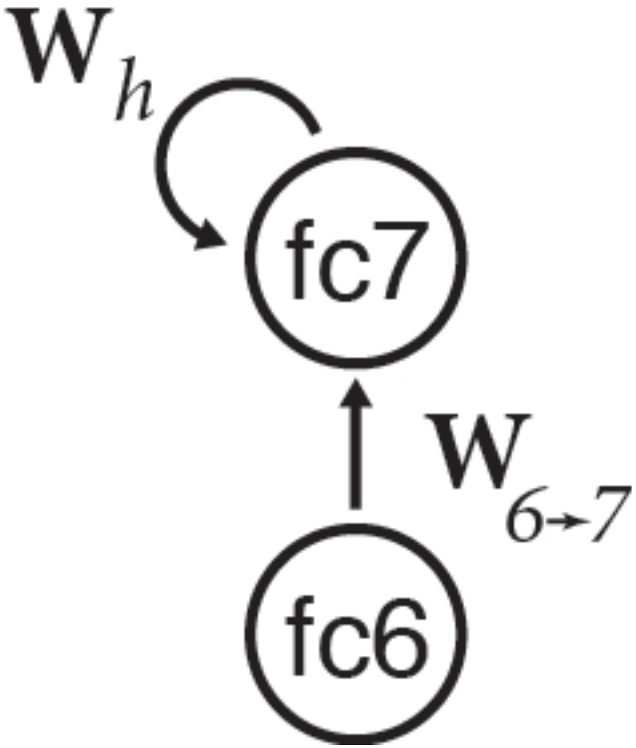

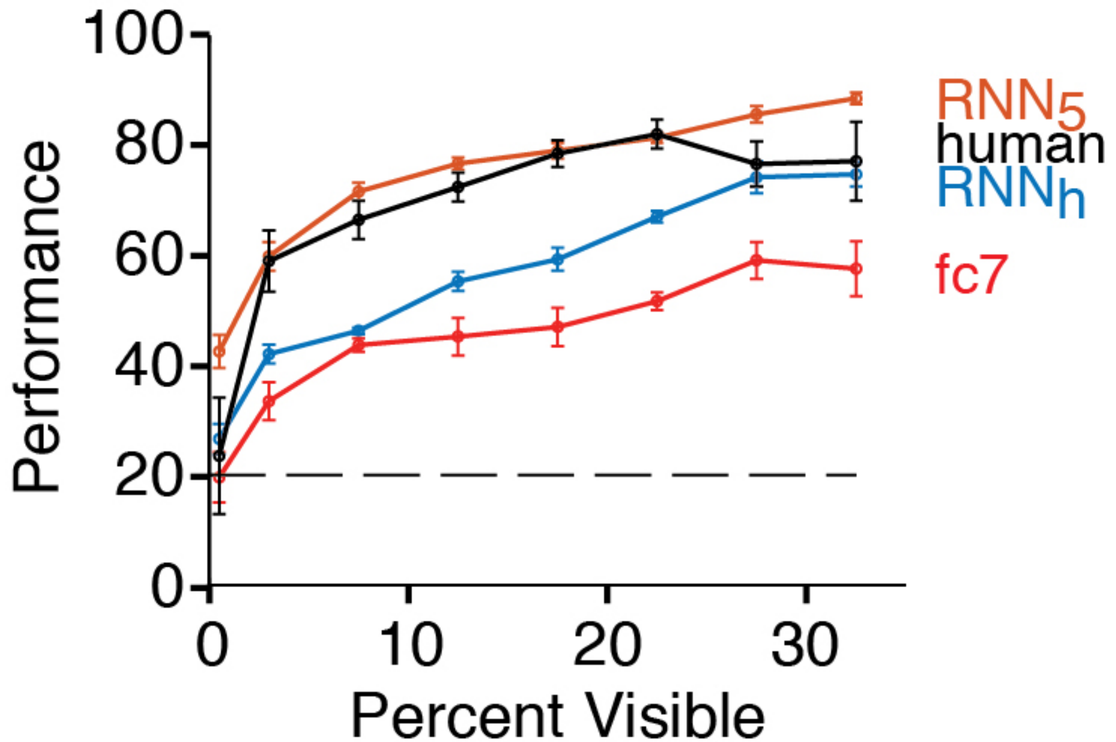

Temporal evolution in recurrent networks

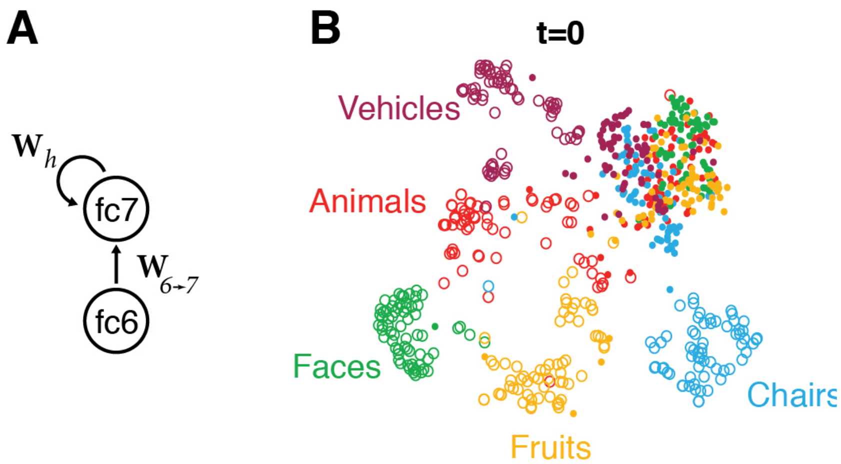

Correlation between RNN models and human performance for individual objects

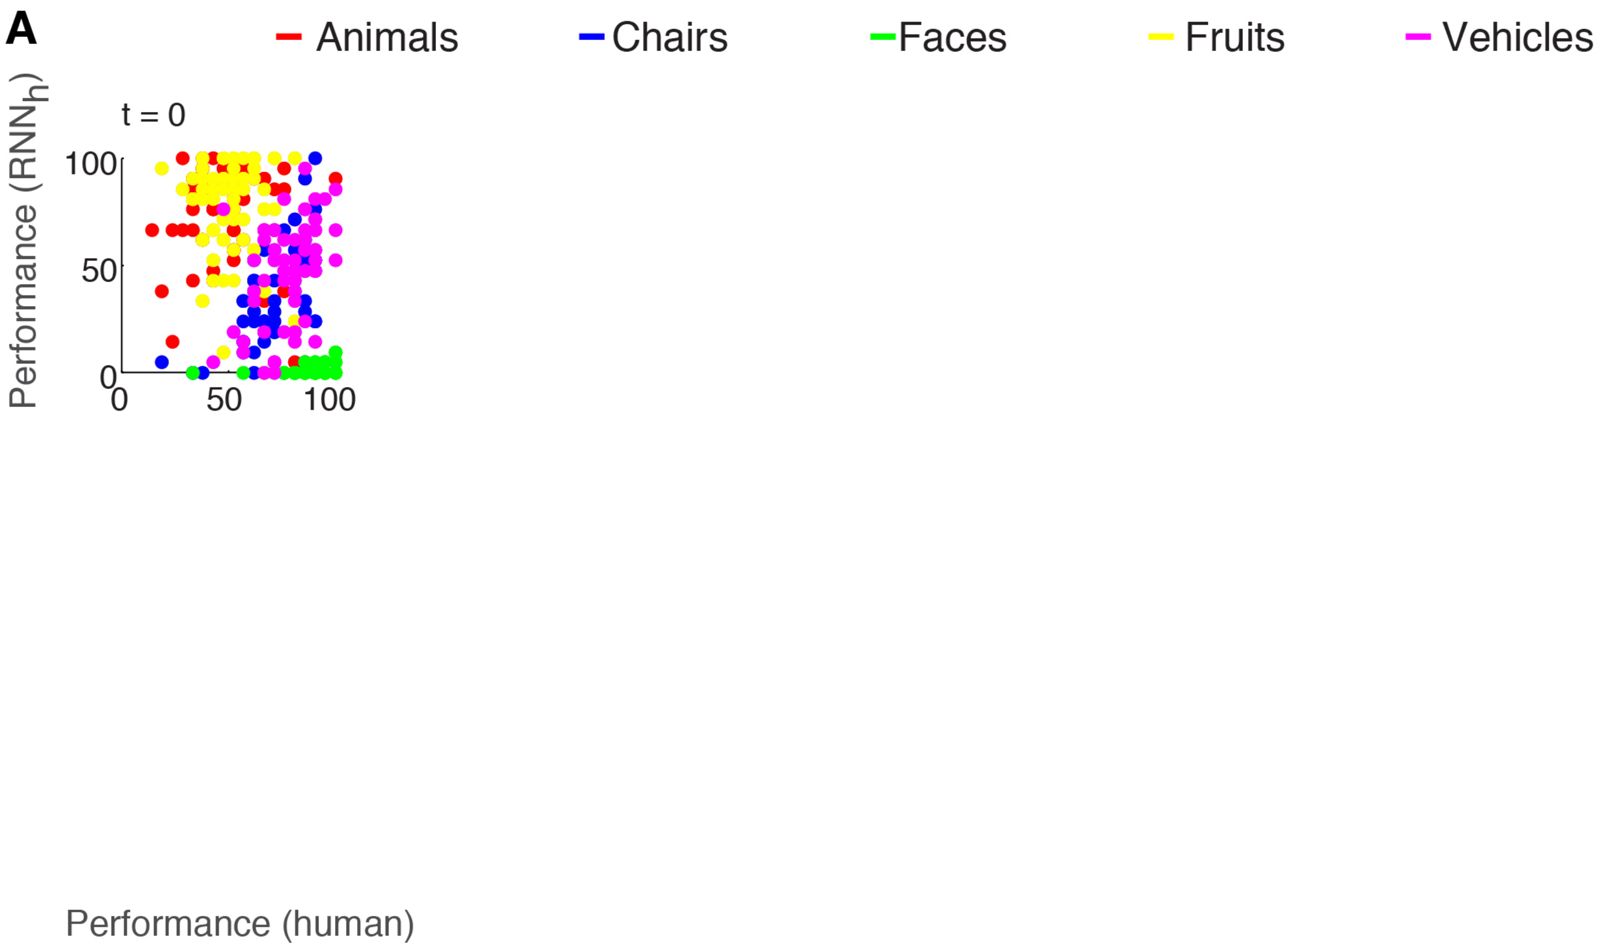

Recurrent neural networks match human performance in pattern completion

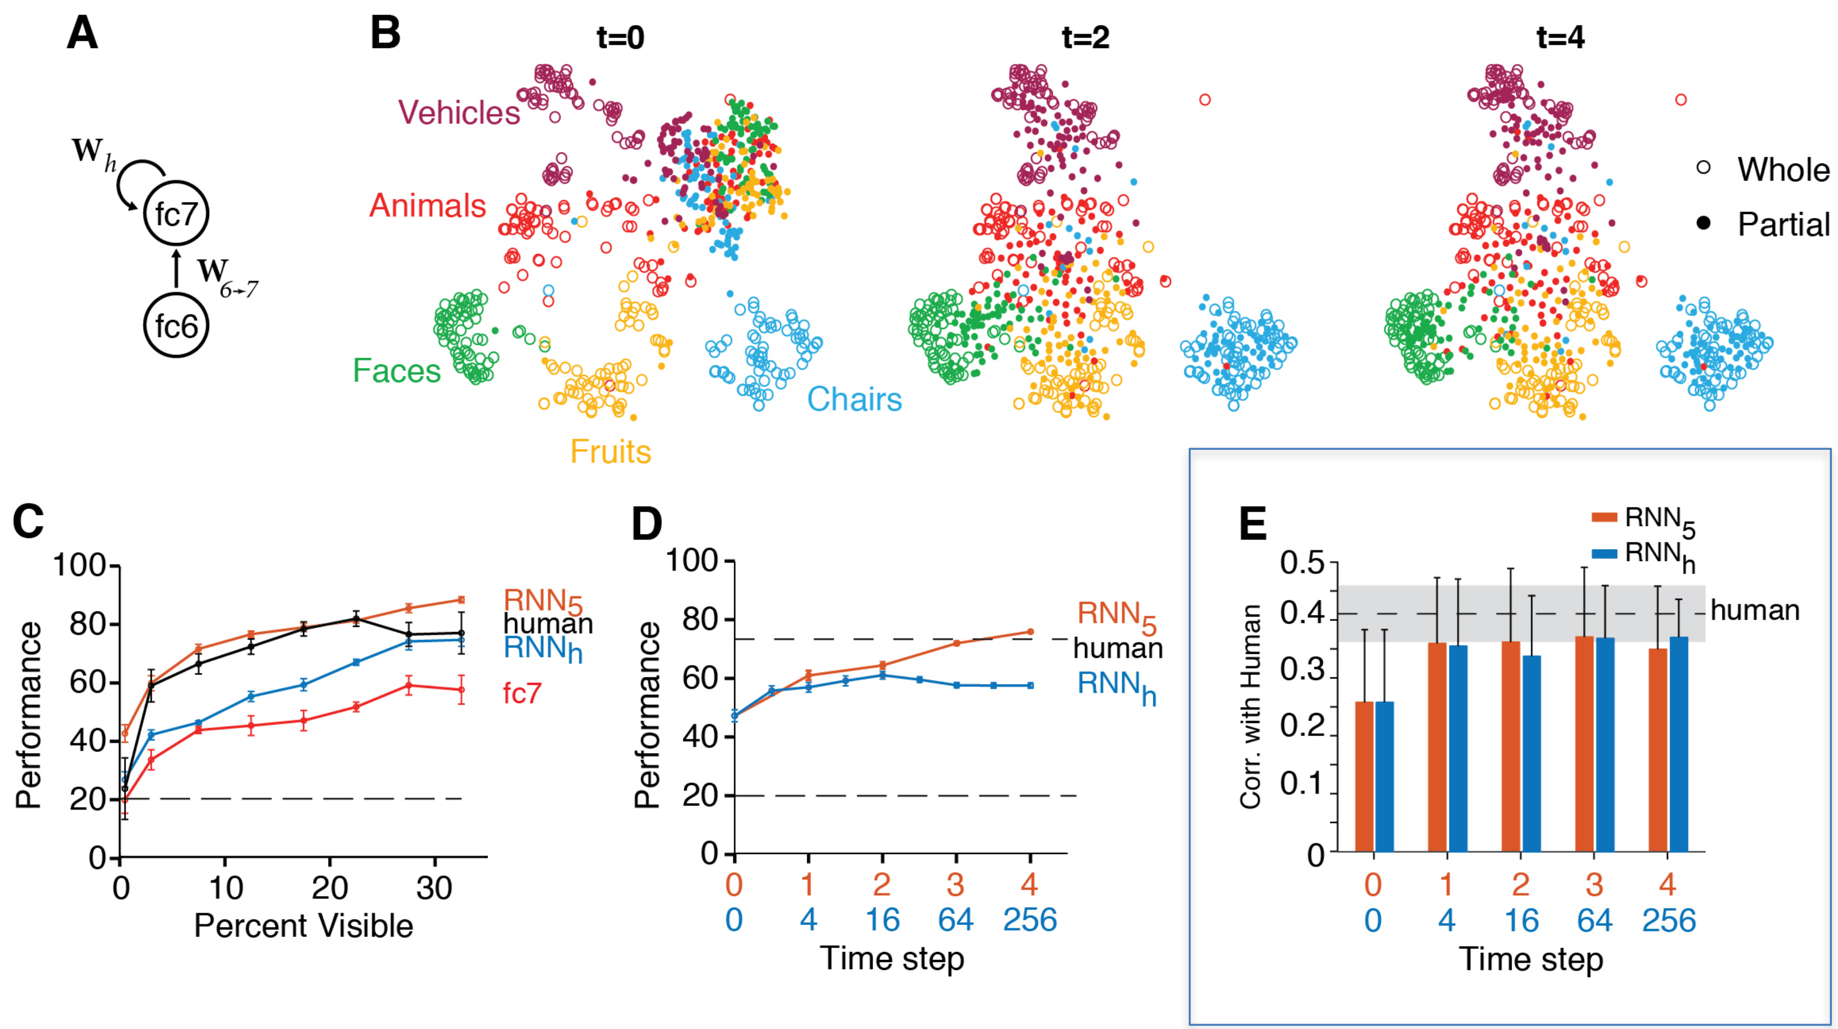

Computational roles of recurrent/feedback signals

1. Pattern completion (recurrent computations)
2. Predictive coding (feedback computations)

Image by Hanlin Tang

Neurophysiology led the way to basic filters

Orientation selectivity

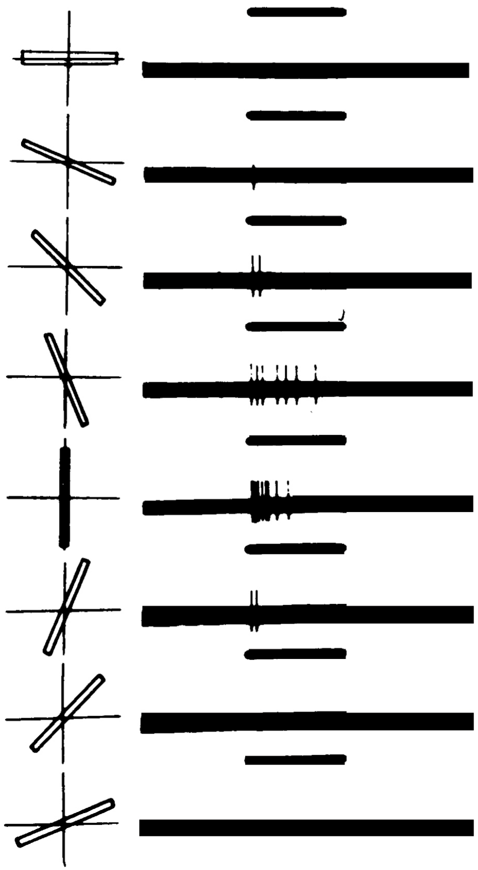

Gabor function

$$D(x,y)=\frac{1}{2\pi\sigma_x\sigma_y}\exp\left[-\frac{x^2}{2\sigma_x^2}-\frac{y^2}{2\sigma_y^2}\right]\cos(kx-\phi)$$

Hubel and Wiesel 1968

Hubel – Nobel Lecture

Area summation curve in V1

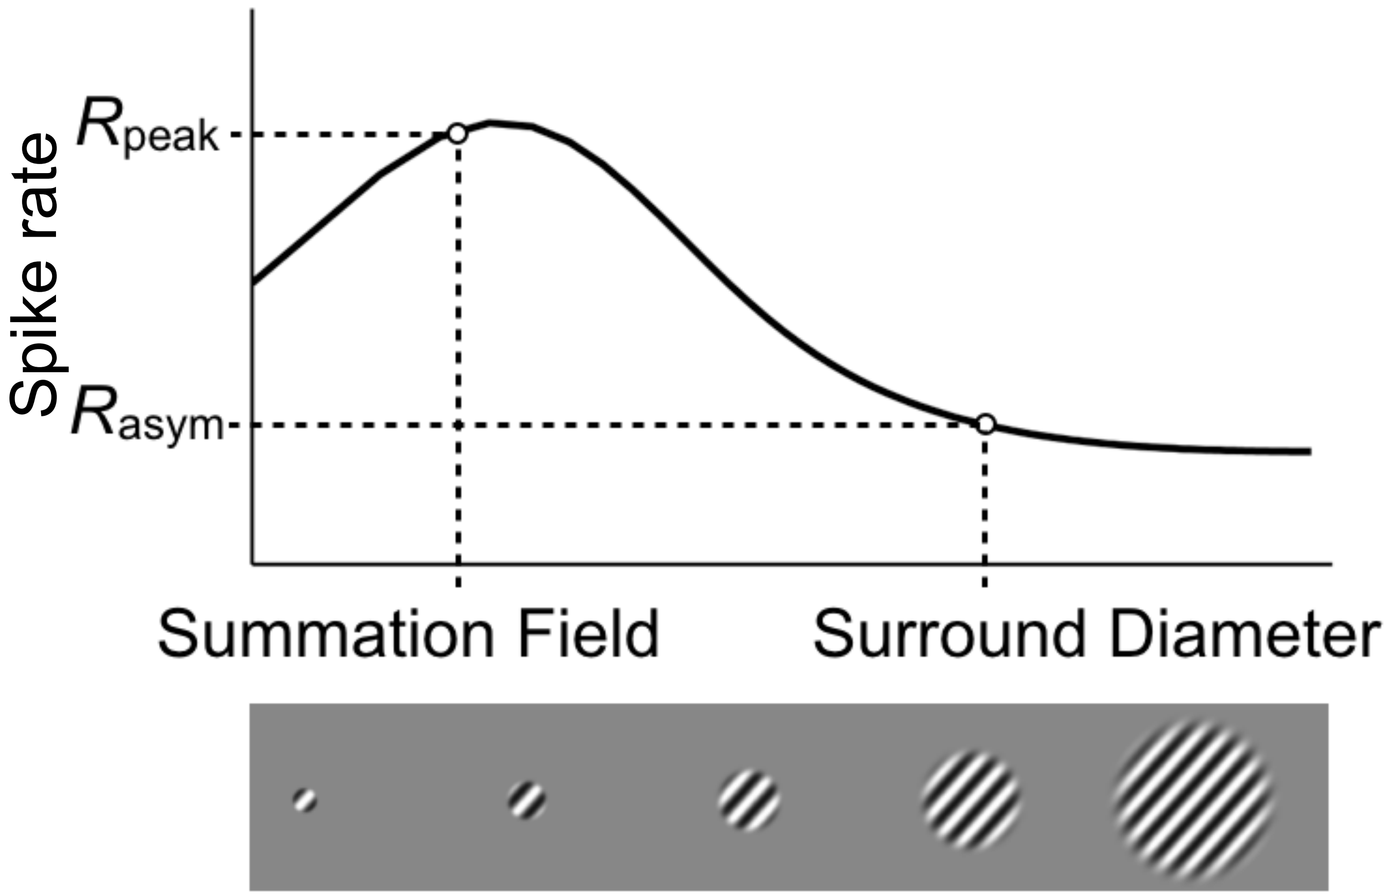

Reversible inactivation of feedback signals (from V2/V3 to V1)

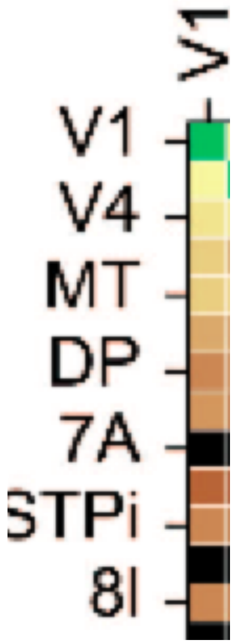

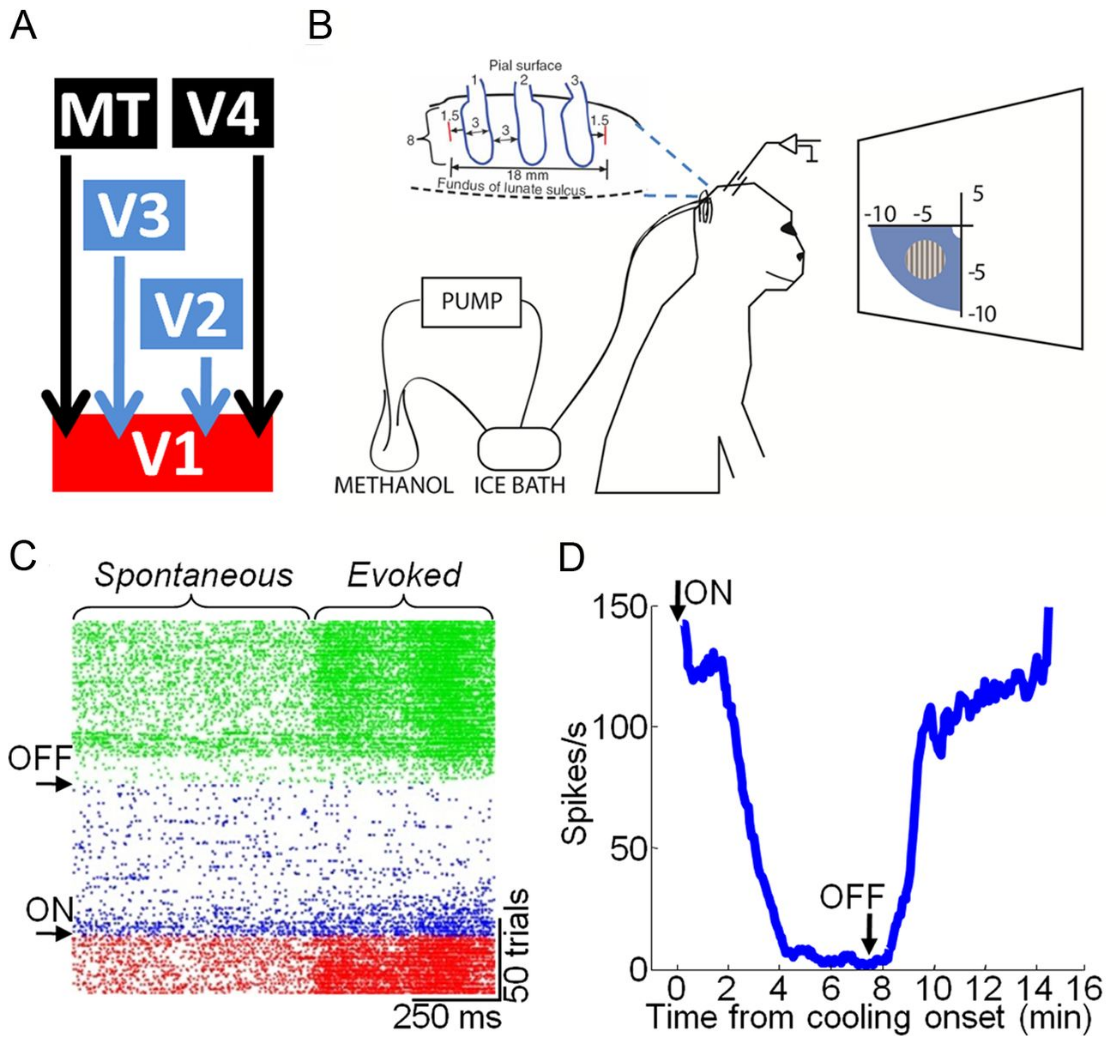

Camille Gomez-Laberge, JoJo Nassi and Richard Born

Feedback inactivation does not change orientation or direction selectivity

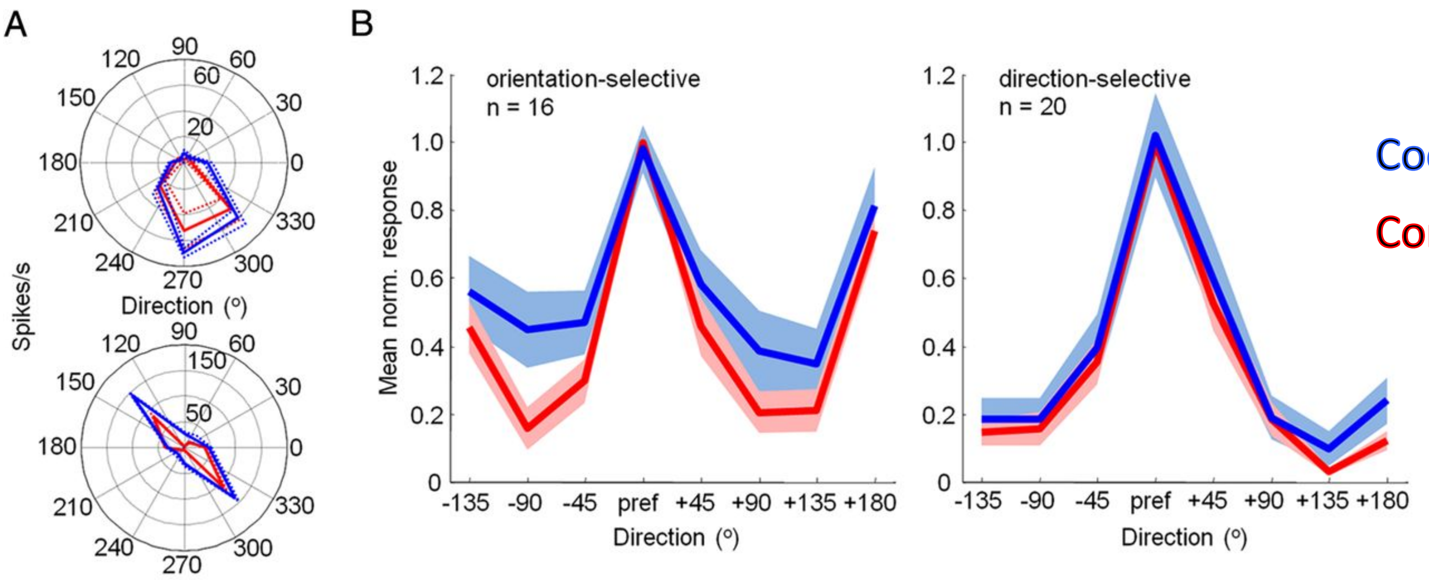

Cooling V2/V3

Control

Nassi et al 2013, Gomez-Laberge et al 2014, 2016

Feedback inactivation leads to reduced surround suppression

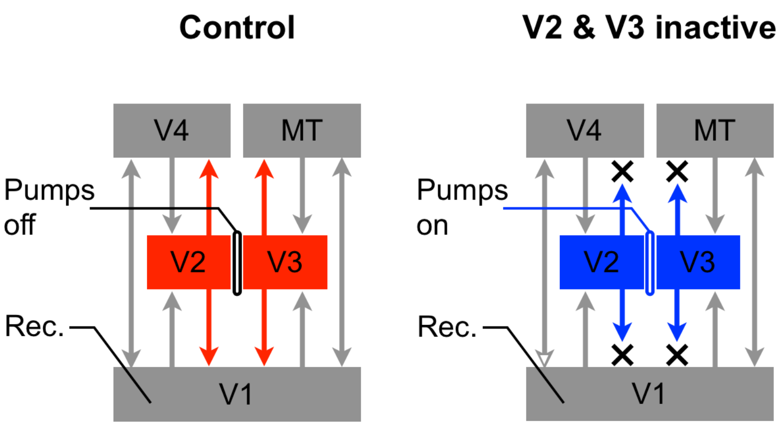

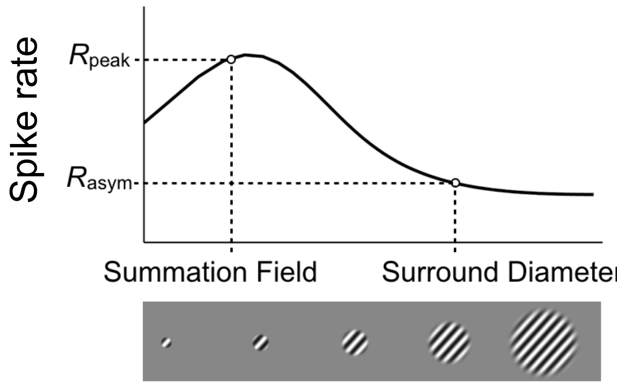

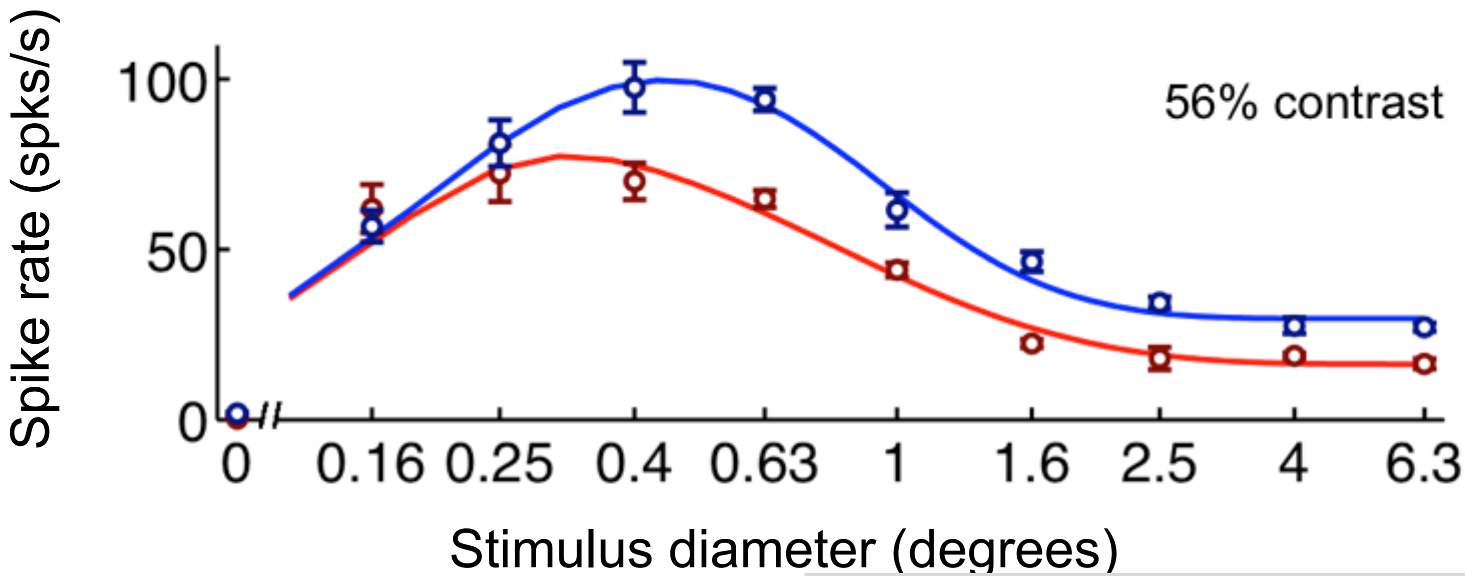

Nassi et al 2013, Gomez-Laberge et al 2014, 2016

Recurrent connections also contribute to surround suppression (Adesnik et al 2012)

A simple normalization model to explain area summation curves

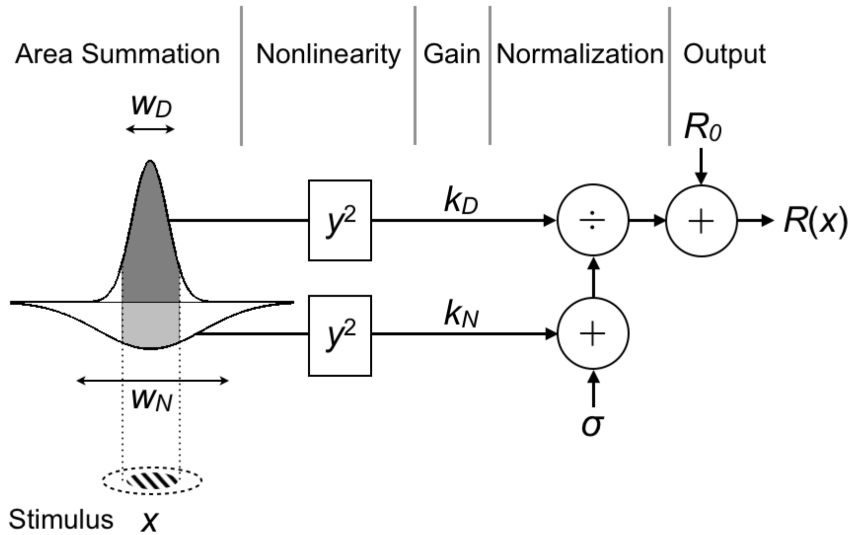

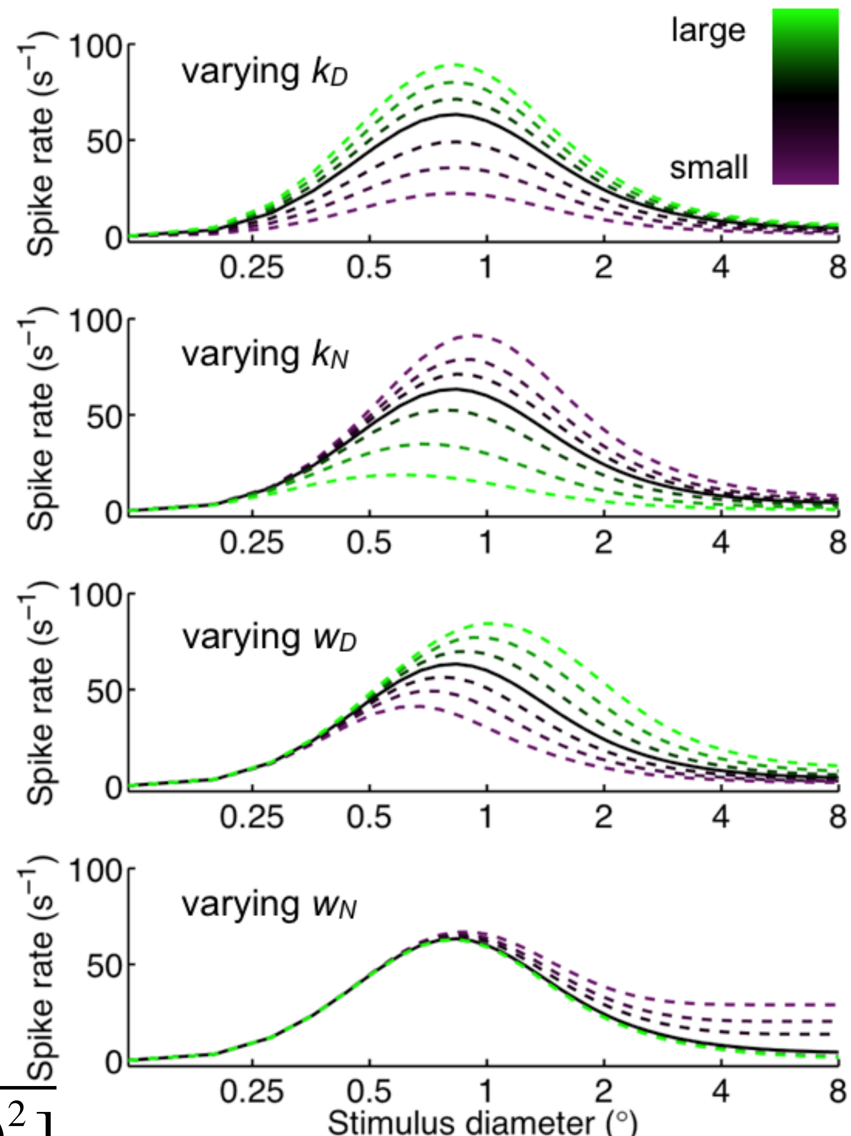

$$R_{ROG}(x) = R_0 + \frac{D(x)}{\sigma + N(x)}$$

$$R_{ROG}(x) = R_0 + \frac{k_D[w_D erf(x/2w_D)^2]}{\sigma + k_N[w_N erf(x/2w_N)^2]}$$

Predictive coding in visual cortex

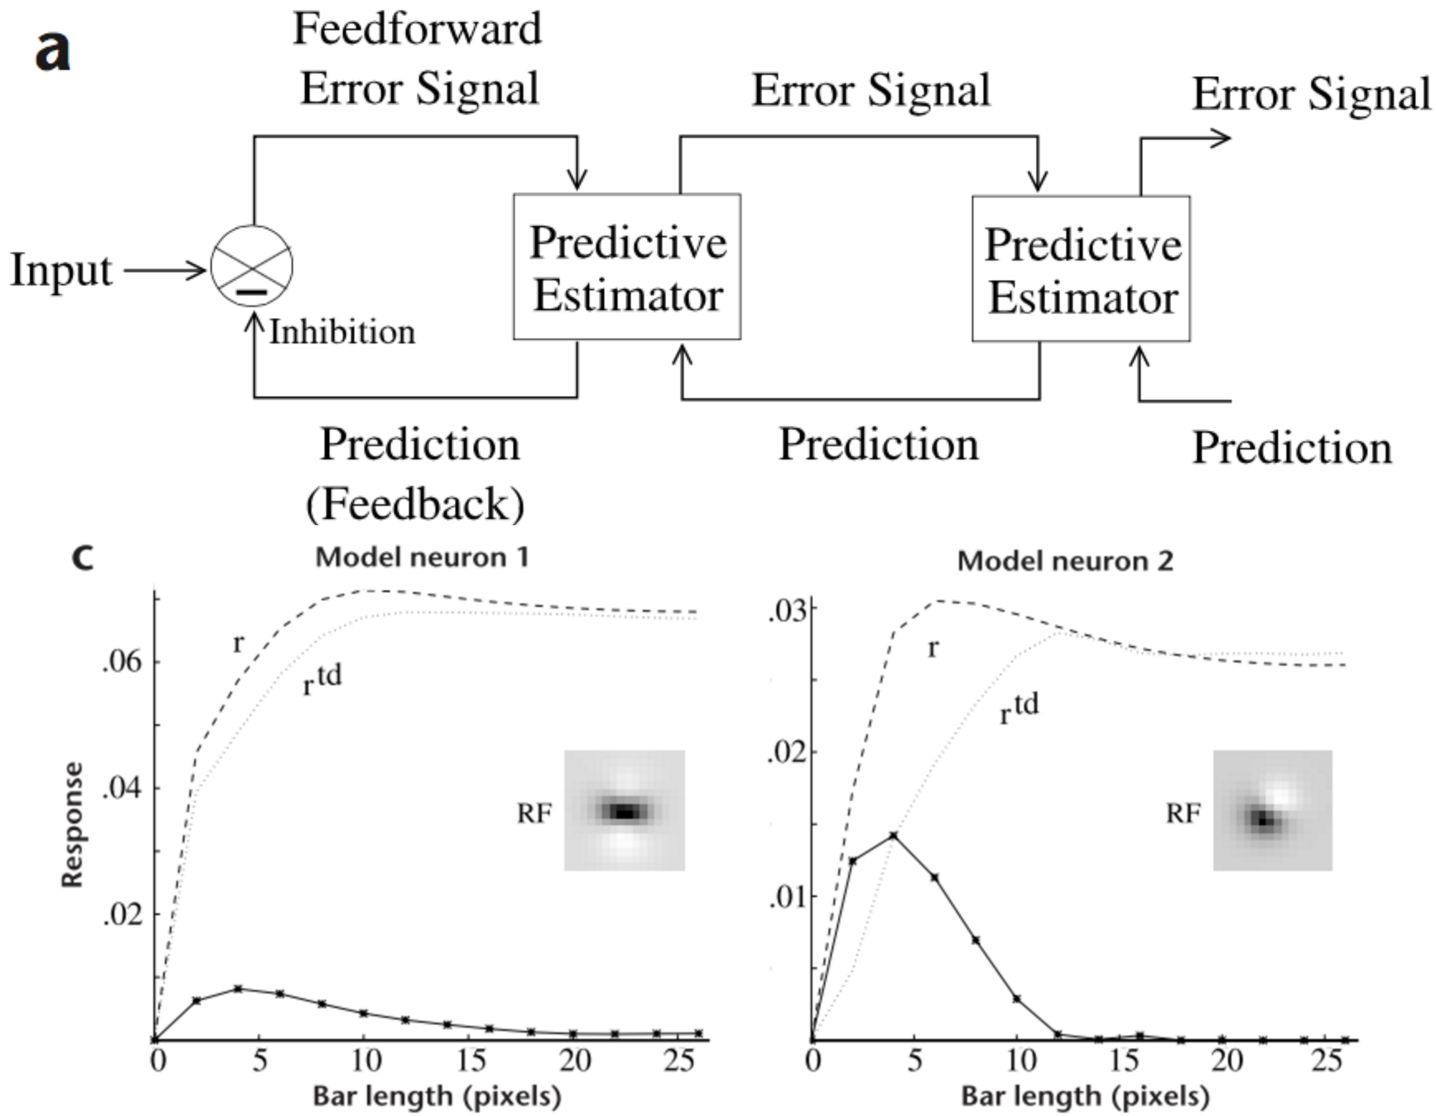

Rao and Ballard, 1999

Deep Learning Implementation of Predictive Coding

Essential elements:

- "**Representation**" neurons: hold "state of world"
- **Predictions**
- **Targets**
- "**Error**" neurons

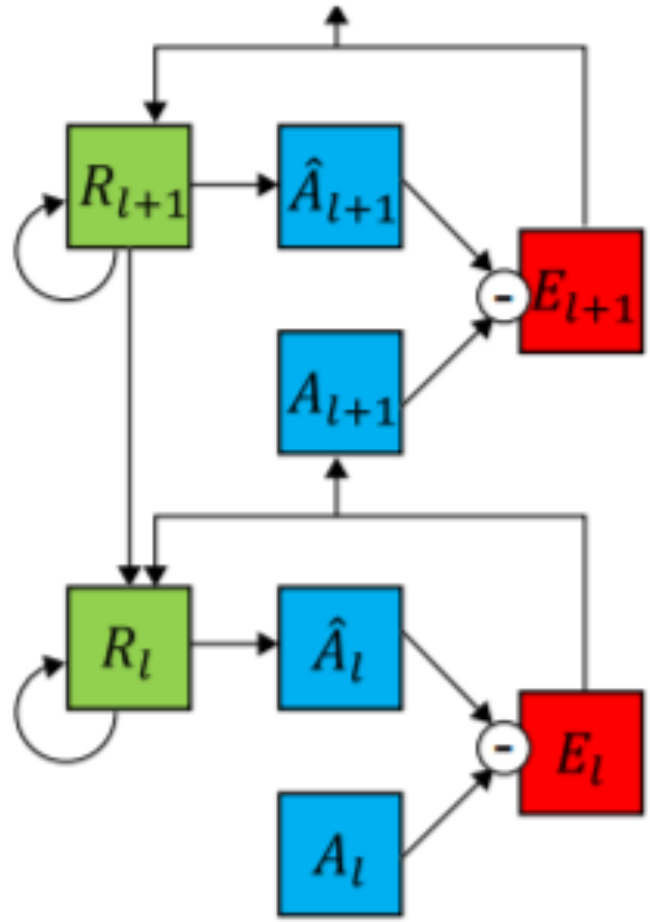

Bill Lotter and David Cox

“PredNet” Details

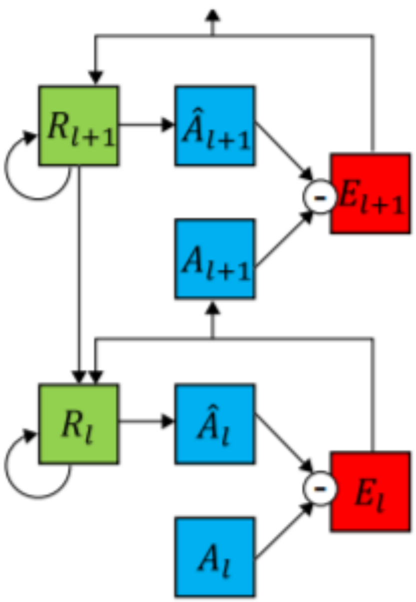

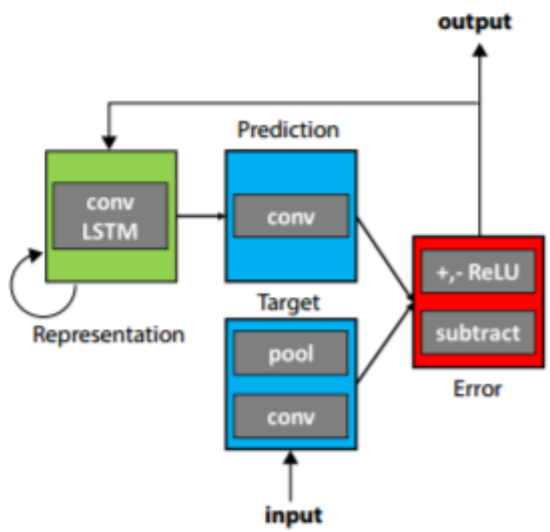

$$A_l^t = \begin{cases} x_t & \text{l=0} \\ \mathrm{MAXPOOL}(\mathrm{RELU}(\mathrm{CONV}(E_{l-1}^t))) & \text{l>0} \end{cases}$$

$$\hat{A}_l^t = \mathrm{RELU}(\mathrm{CONV}(R_l^t))$$

$$E_l^t = [\mathrm{RELU}(A_l^t - \hat{A}_l^t); \mathrm{RELU}(\hat{A}_l^t - A_l^t)]$$

$$R_l^t = \mathrm{CONVLSTM}(E_l^{t-1}, R_l^{t-1}, \mathrm{UPSAMPLE}(R_{l+1}^t))$$

$$L = \sum_{t=0}^{T} \lambda_t \sum_{l=0}^{N_l} \lambda_l / n_l \sum_{i=1}^{n_l} E_l^t(i)$$

Testing the model with rotating faces

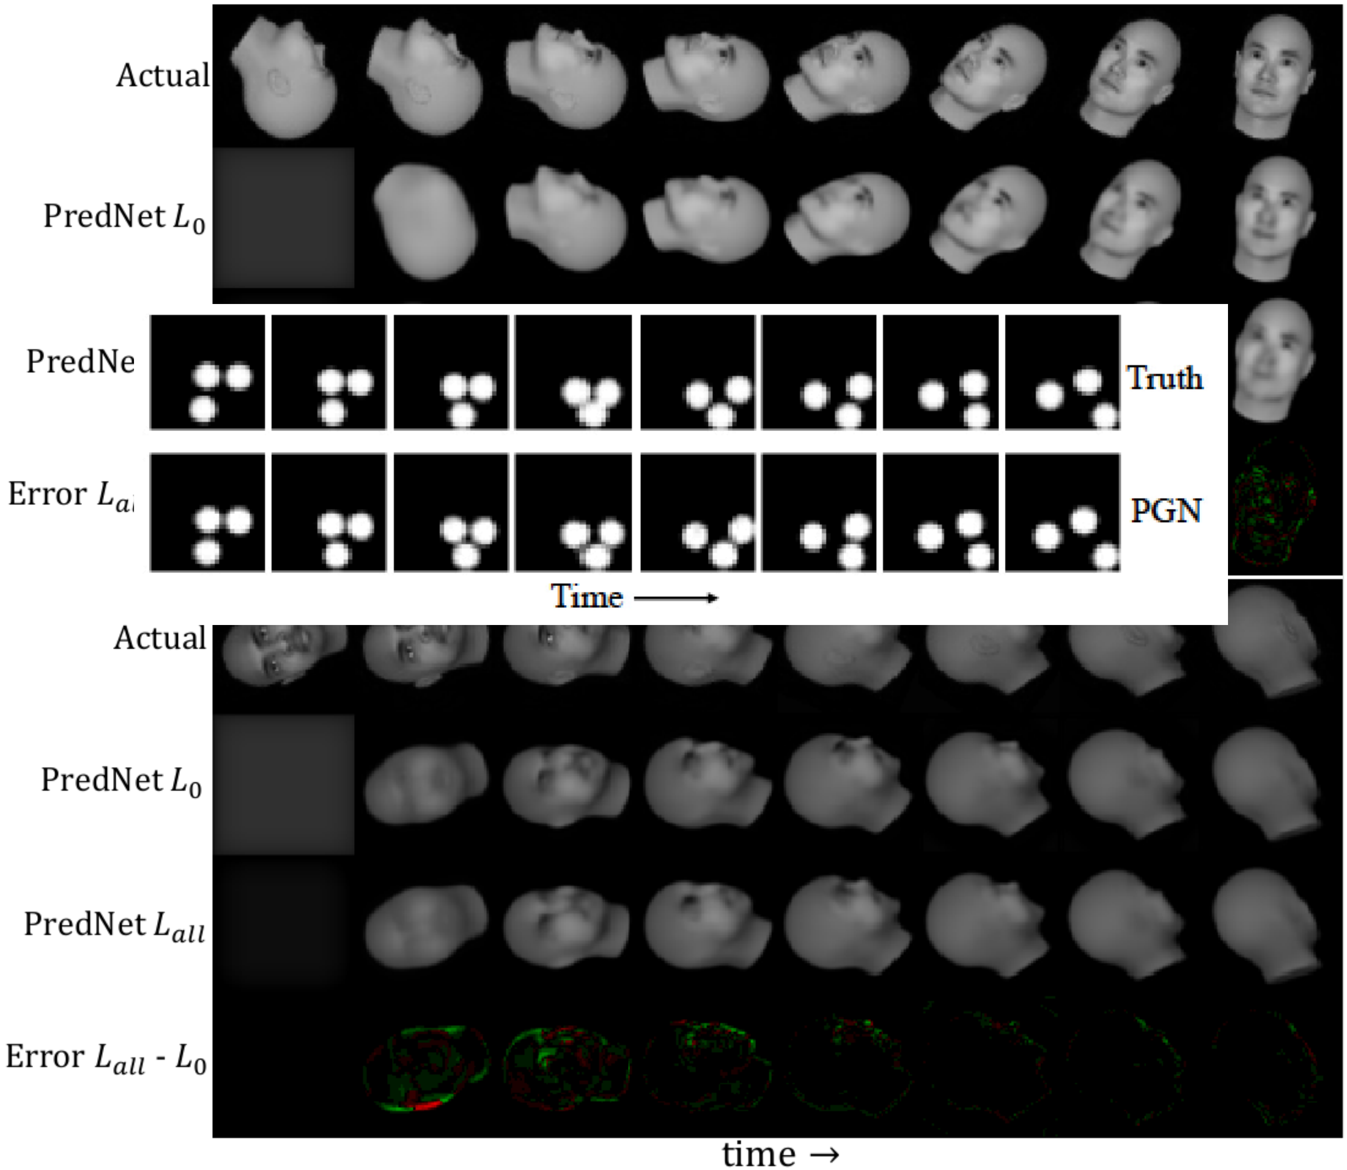

Training for prediction → successful image classification

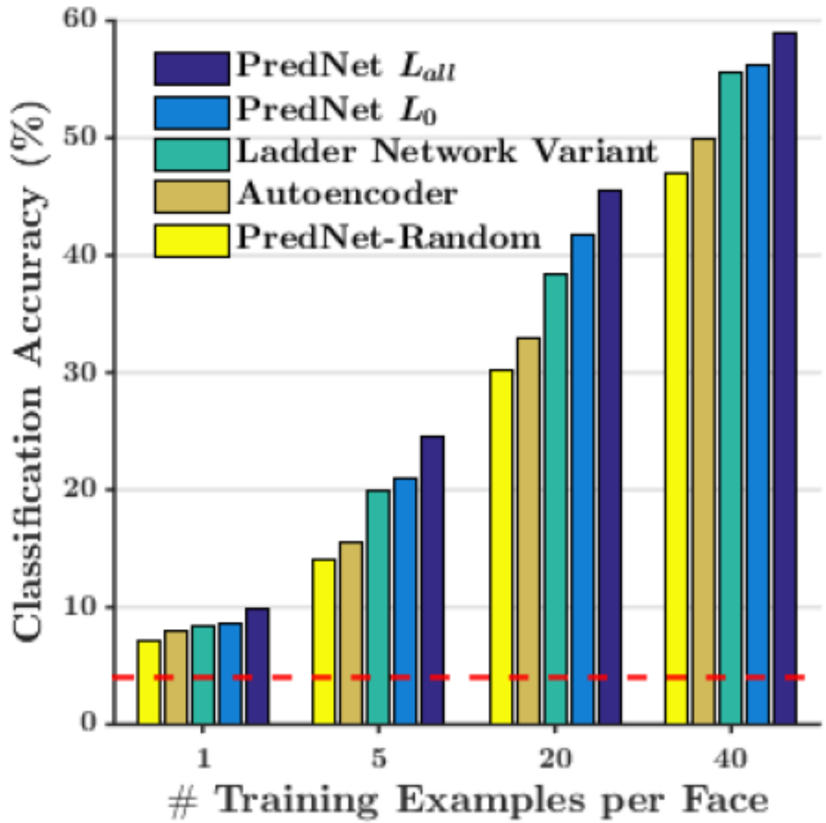

Feedback signals in visual search

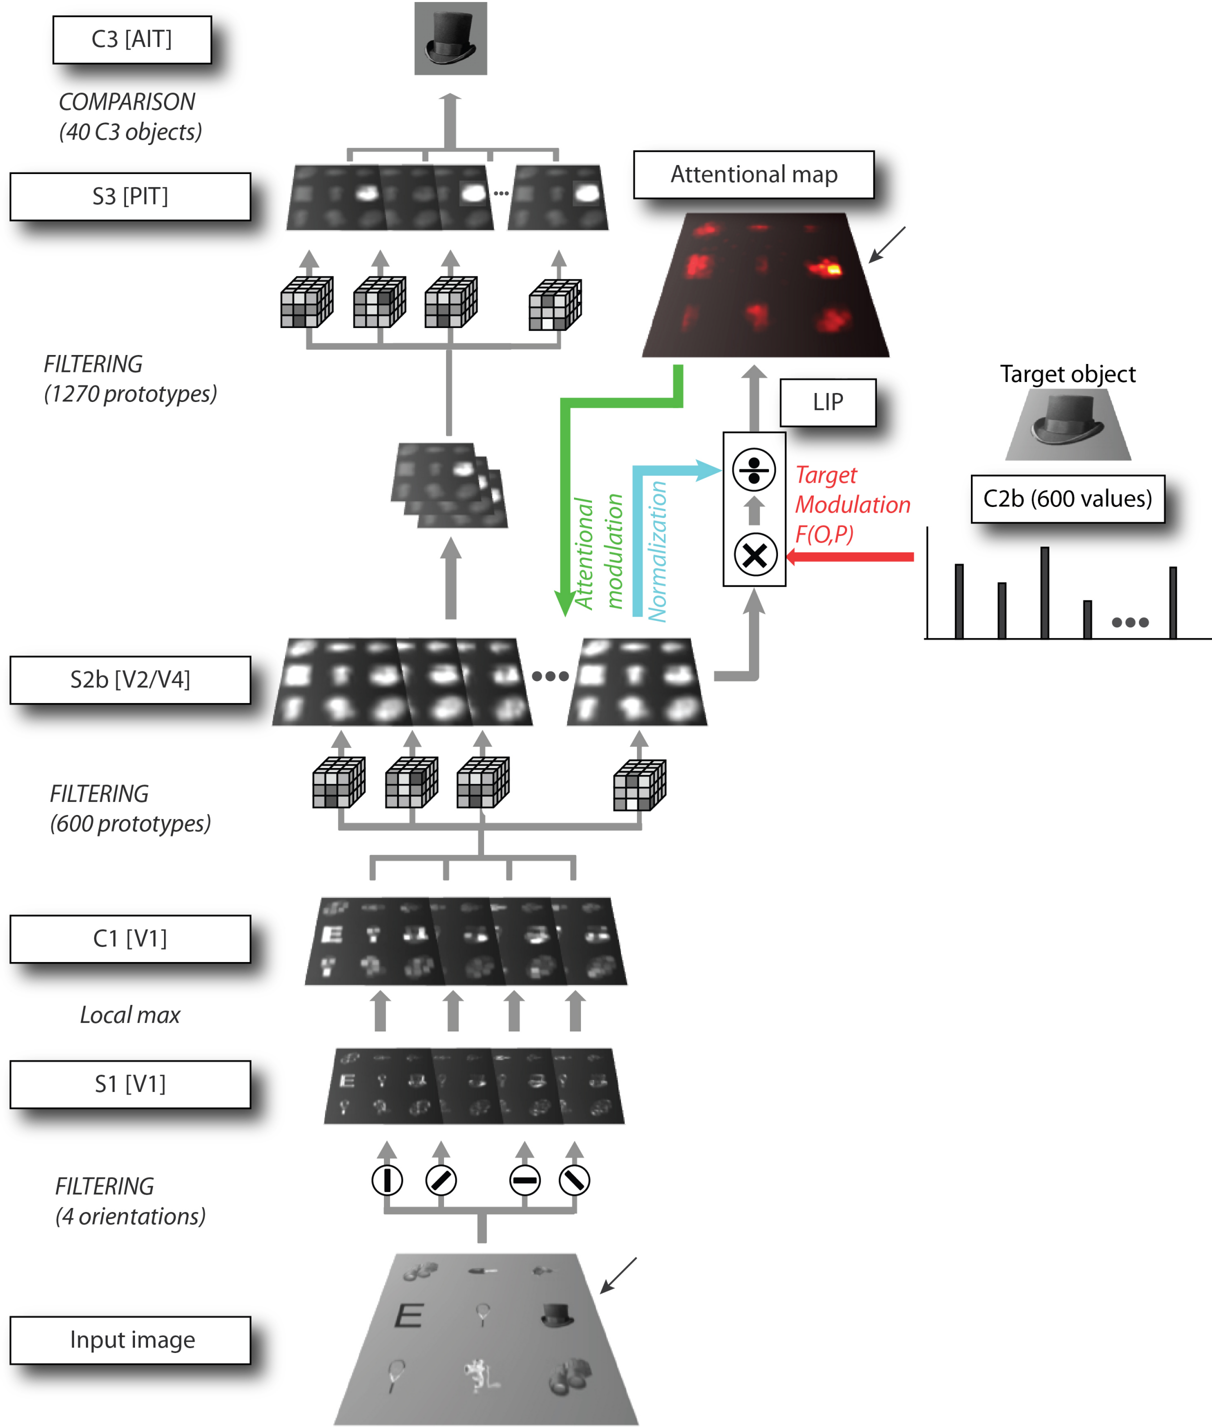

Miconi et al, 2016

Summary

Pattern completion: Recurrent connections can help recognize heavily occluded objects and pattern completion

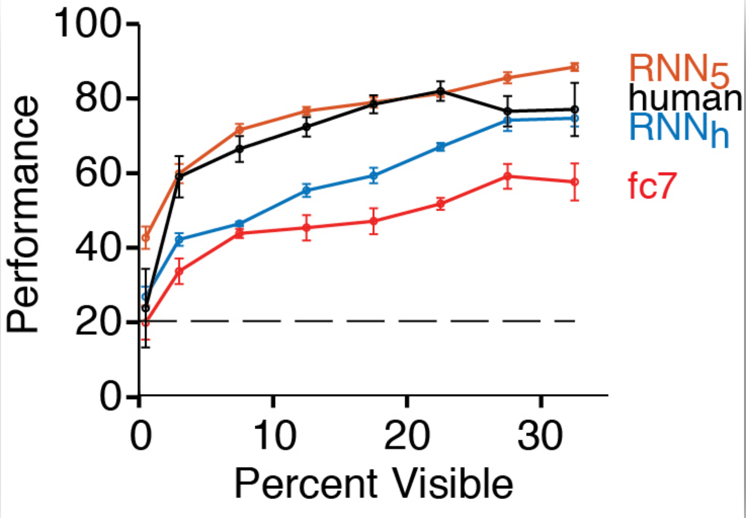

Feedback signals enhance surround suppression and may provide a signal for predictive coding that can help in unsupervised learning

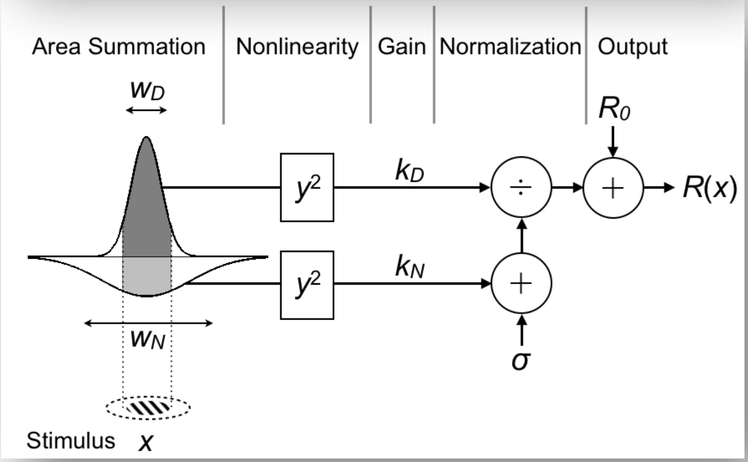

From biological codes to computational codes

Biological codes

Computer codes

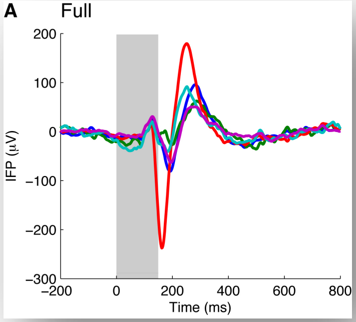

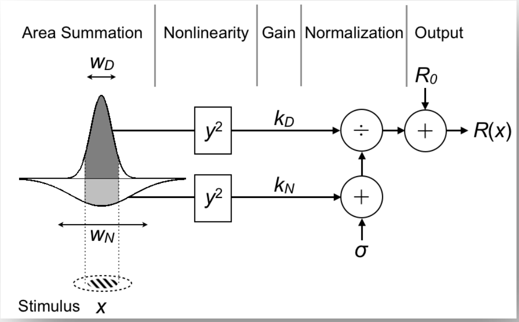

Visual recognition: peeking inside computations in the brain

Gabriel Kreiman

gabriel.kreiman@tch.harvard.edu

klab.tch.harvard.edu

Center for Brains, Minds and Machines

Camille Gomez

Richard Born

Jojo Nassi

Laura Groomes

Hanlin Tang

Thomas Miconi

Bill Lotter

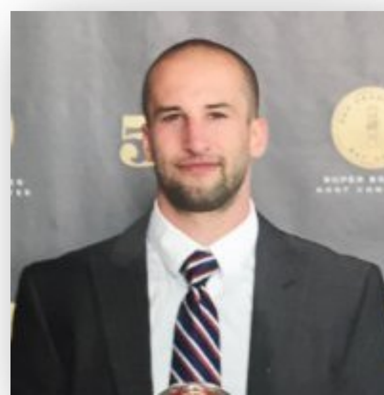

David Cox