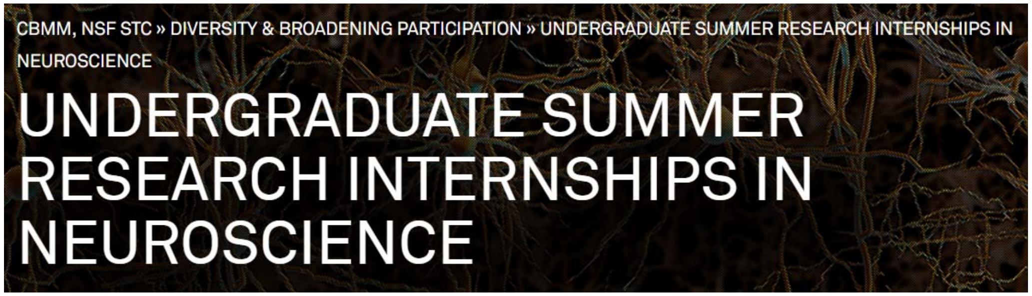

The visual alphabet 2.0

Carlos R. Ponce, M.D., Ph.D.
Assistant Professor
Department of Neuroscience

Washington University in St. Louis
School of Medicine

we study the primate visual system - current focus on **object recognition**

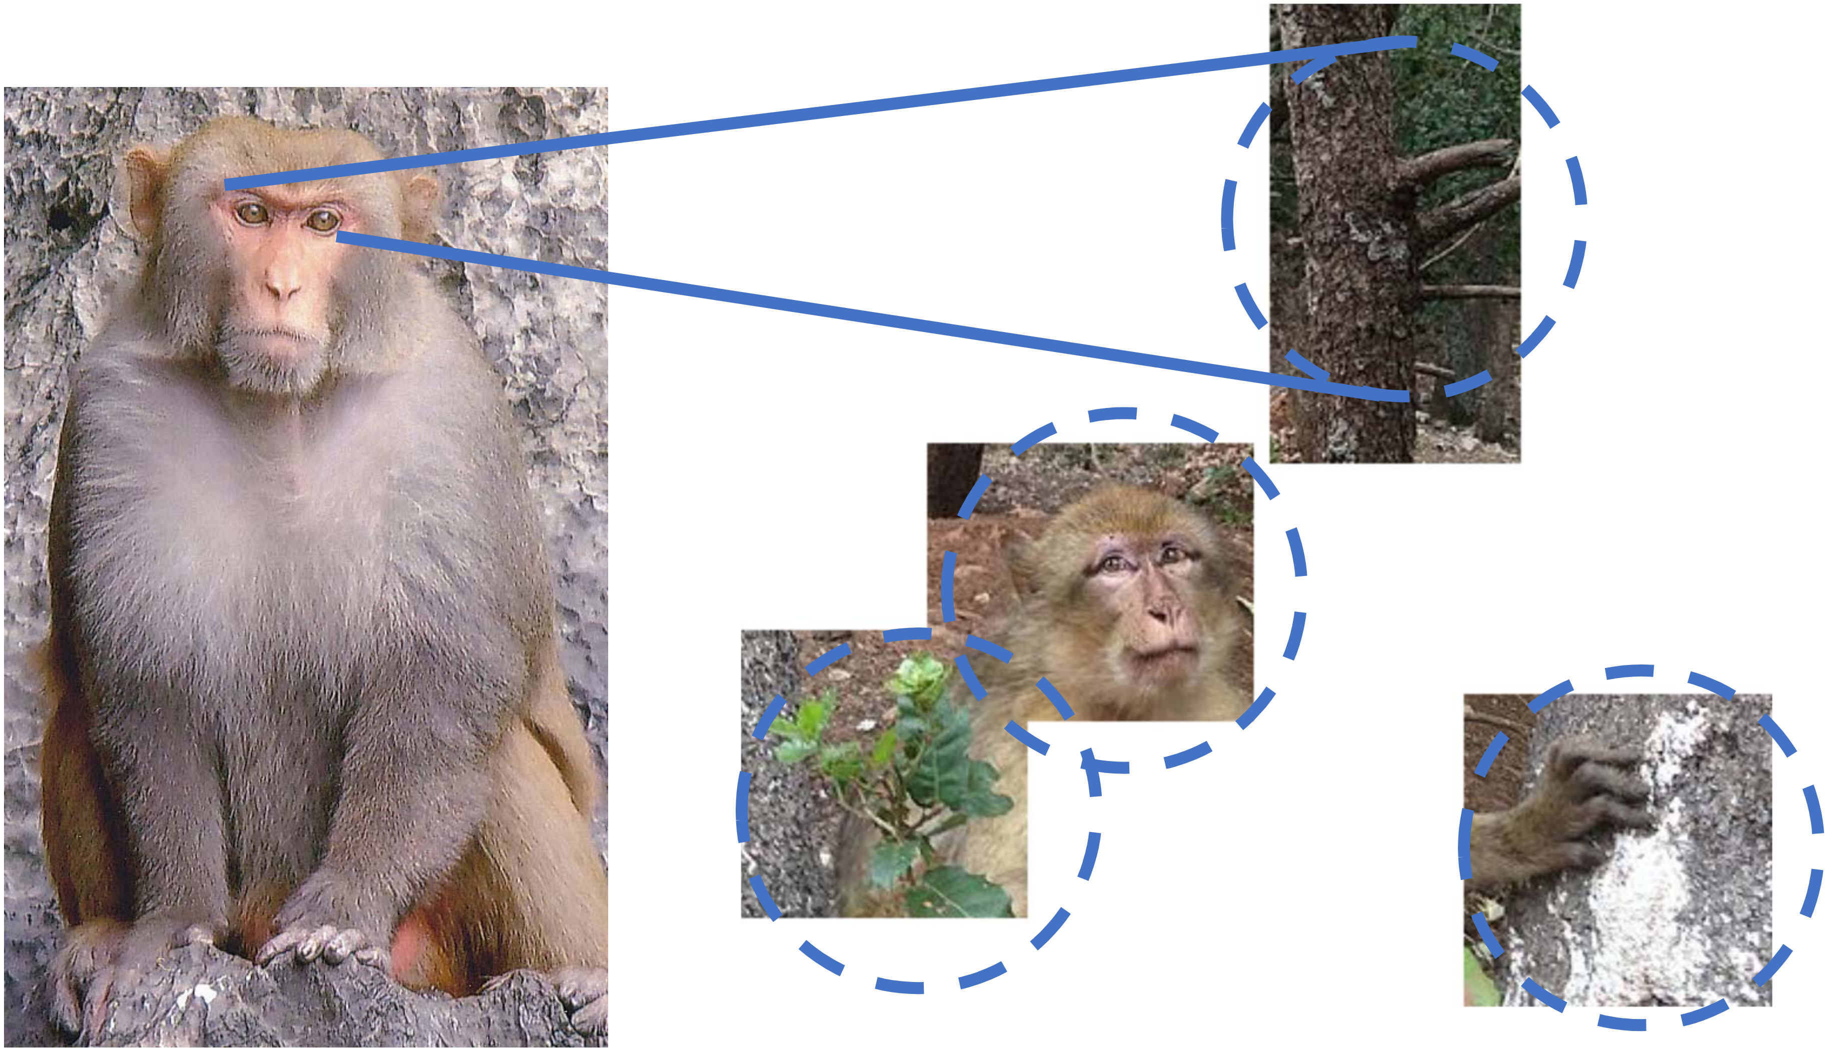

object recognition realized by
a stream of cortical areas

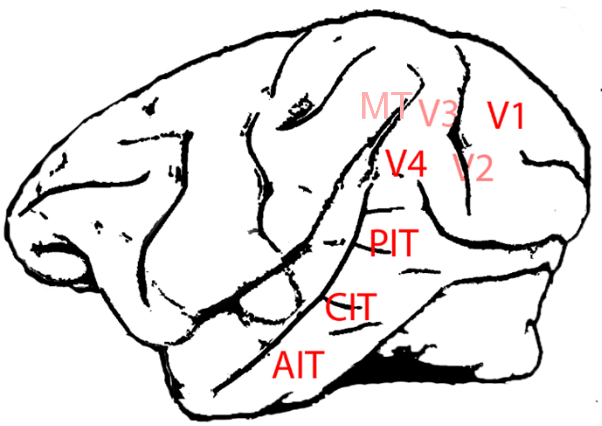

V1, V2, V3, V4

IT = inferotemporal cortex
(posterior, central and anterior)

“ventral stream”

object recognition realized by a set of cortical areas

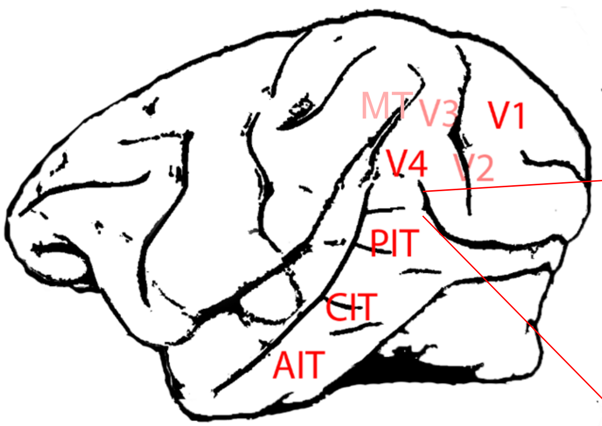

V1, V2, V3, V4

IT = inferotemporal cortex (posterior, central and anterior)

“ventral stream”

containing neurons that respond to specific images (by emitting action potentials)

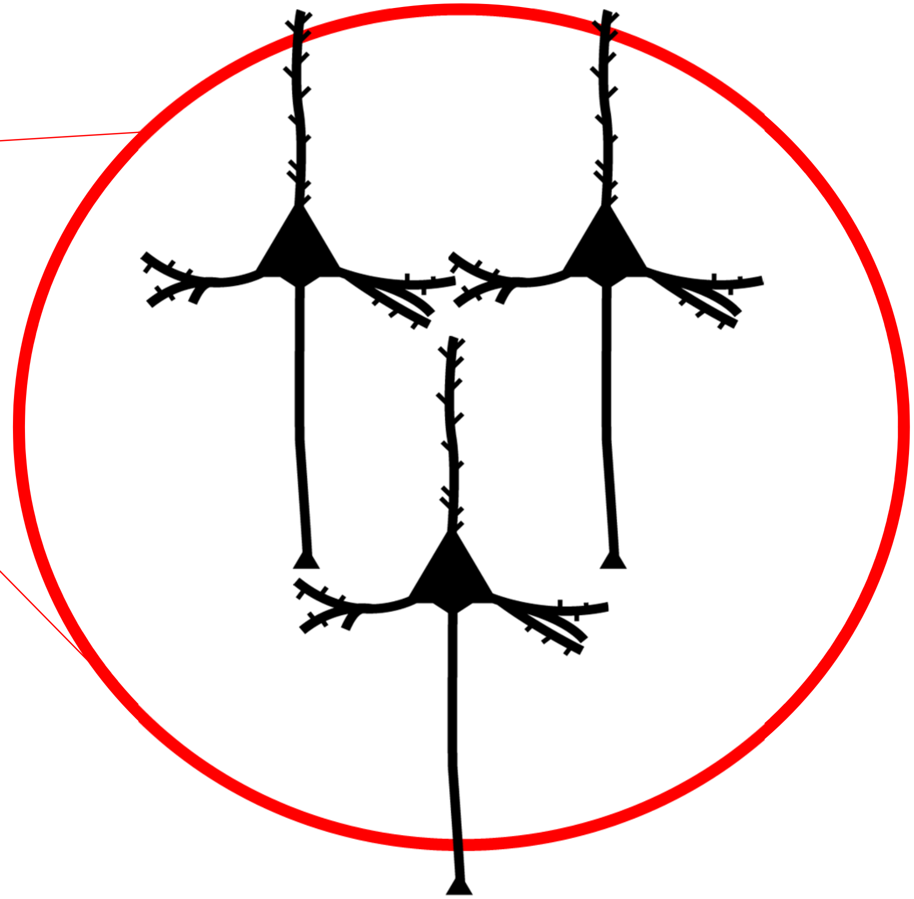

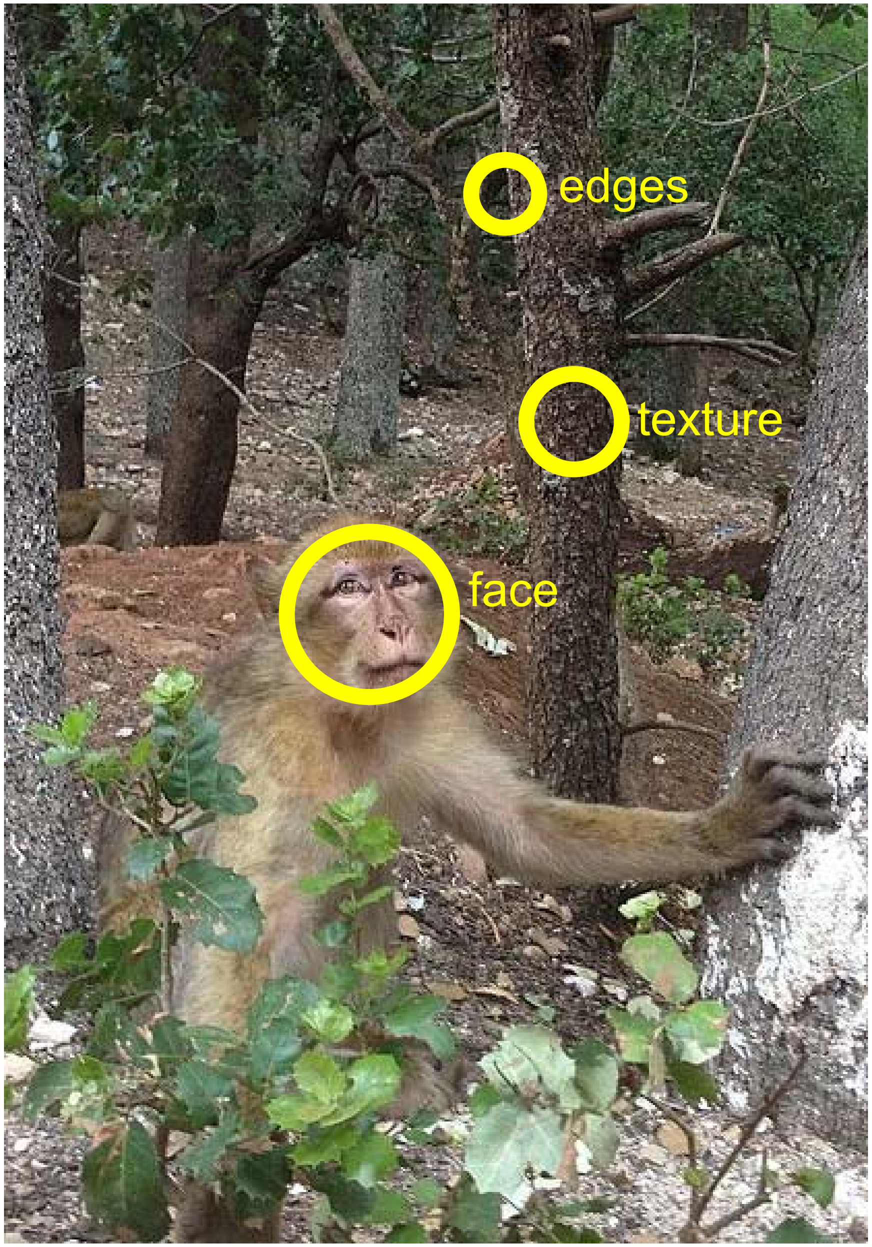

responses signal presence of (*represent*) special features in the world

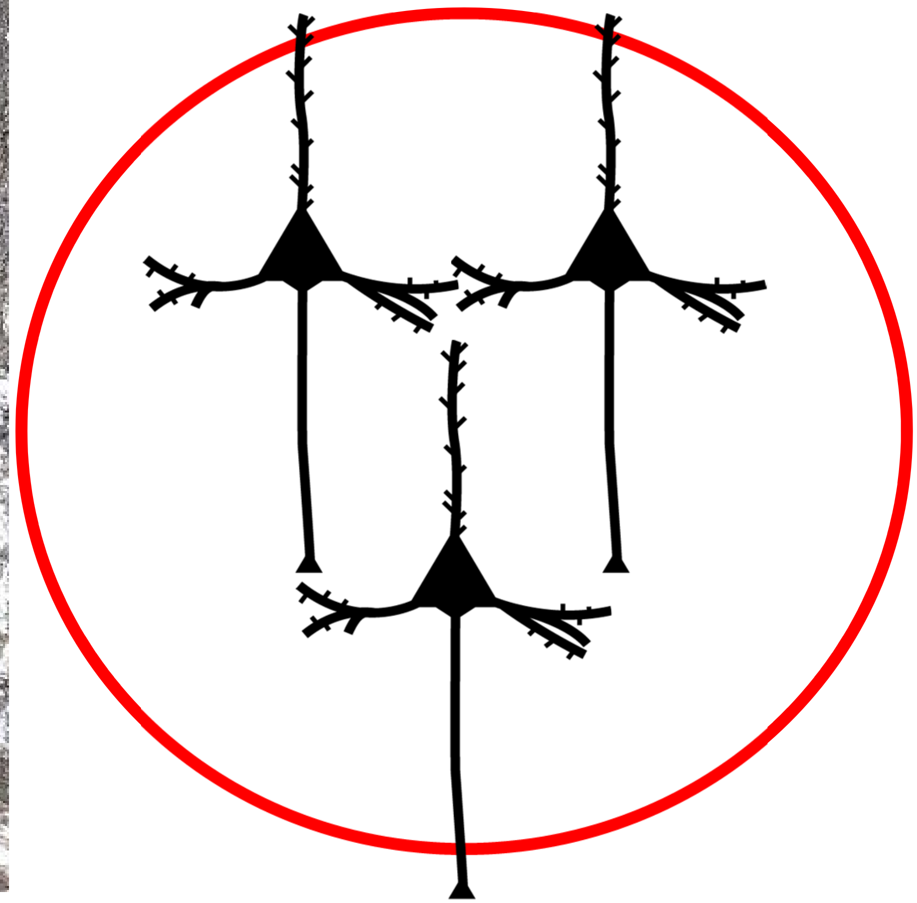

special because too any
features in the world and many
fewer neurons

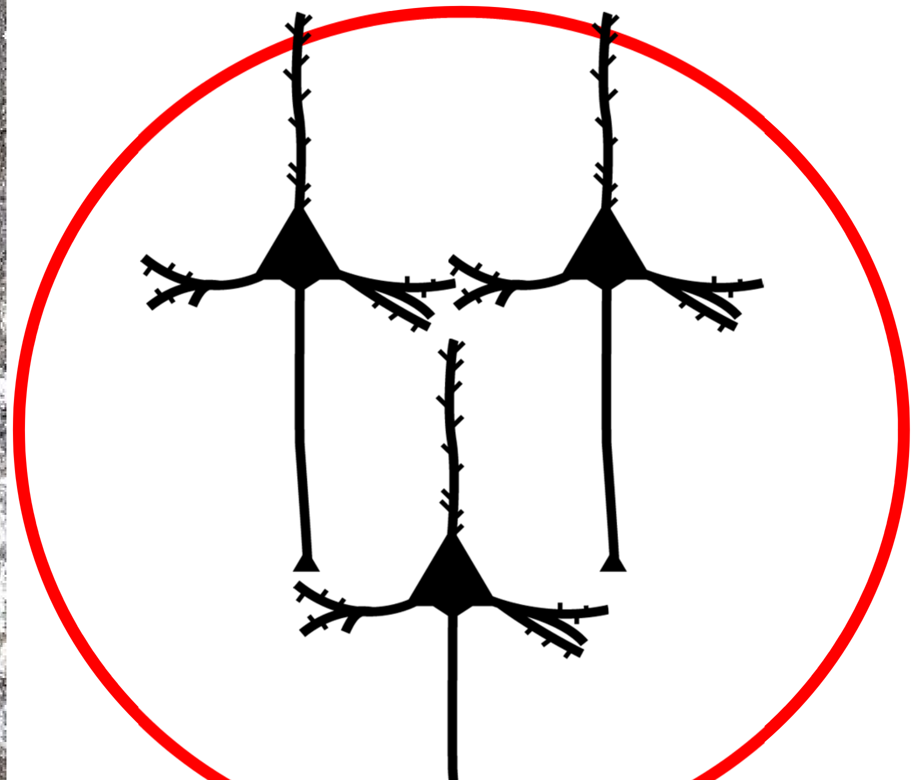

the visual system must be efficient
in allocating its neurons

Uncovering the visual “alphabet”

Leslie G. Ungerleider *, Andrew H. Bell

Vision Research 51 (2011) 782–799

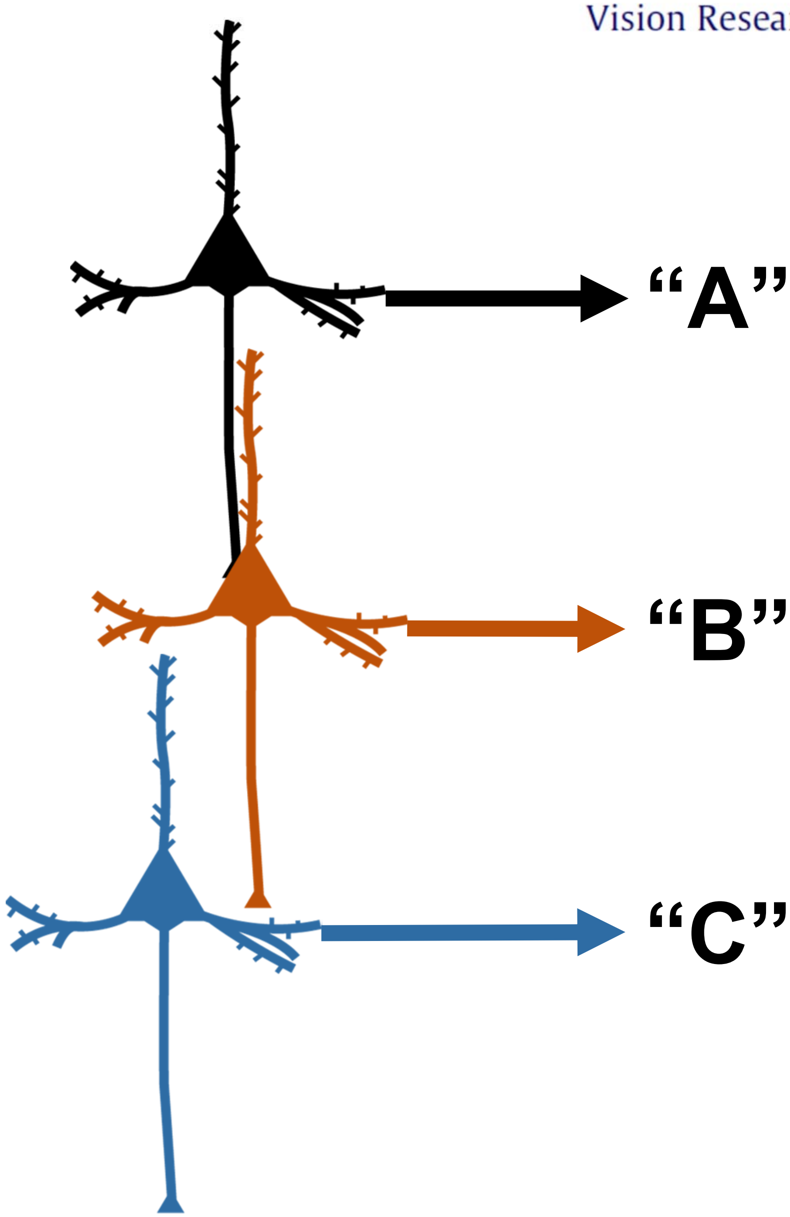

in English, 26 letters can act *compositionally* to express almost any thought

a subset of neuronal *representations* may allow the decoding of any visual scene progress in defining a visual alphabet

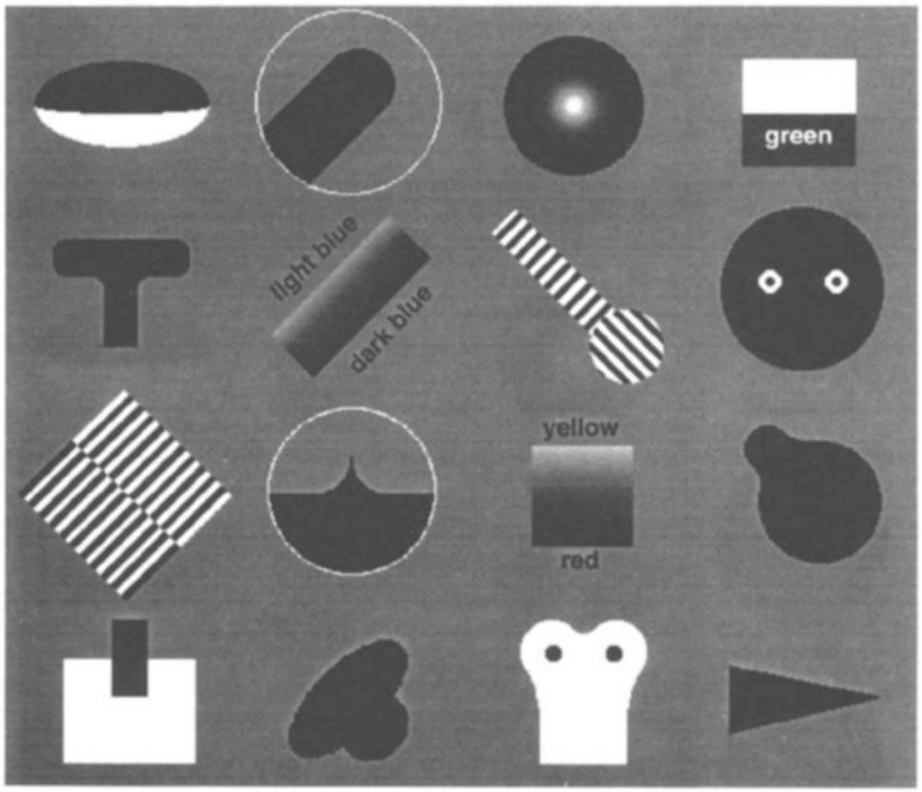

Figure 3 **Sixteen examples of the critical features of cells in TE. They are moderately complex.**

Keiji Tanaka, 1996

Uncovering the visual "alphabet"

Leslie G. Ungerleider *, Andrew H. Bell

Vision Research 51 (2011) 782–799

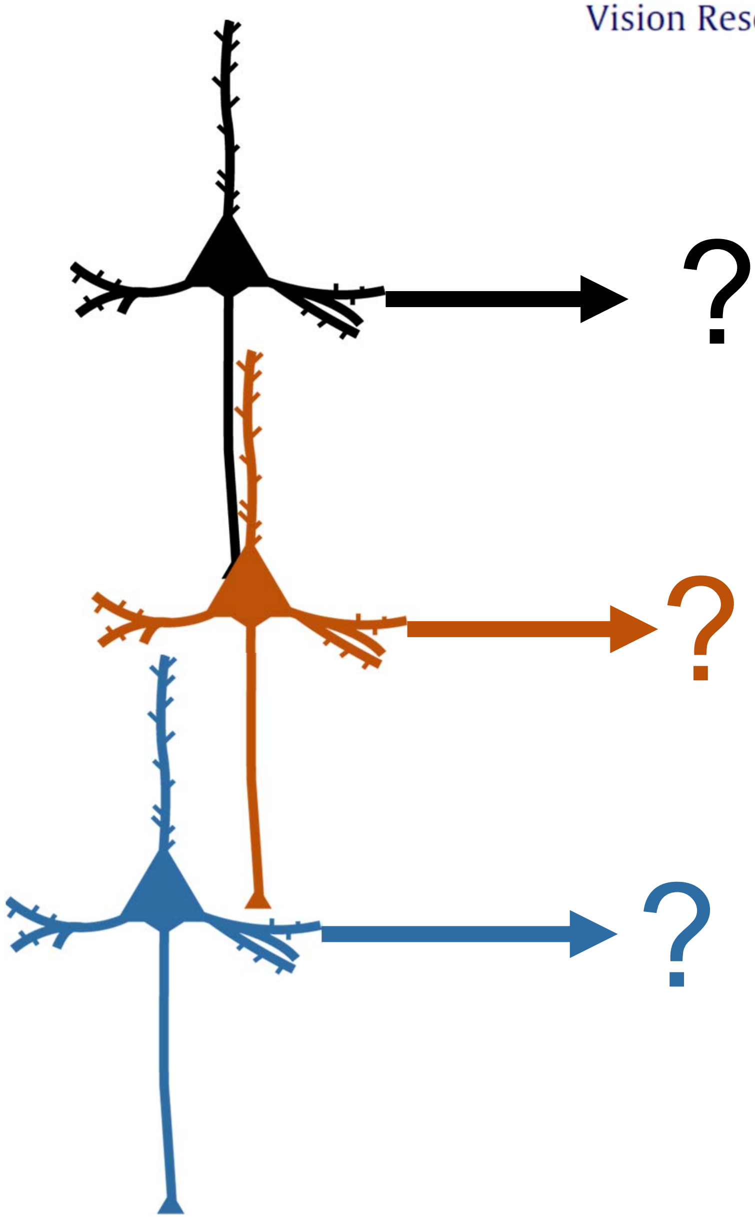

how do we know
what any given neuron
represents about the world?

the canonical approach to visual selectivity

David Hubel and Torsten Wiesel

the canonical approach to visual selectivity

David Hubel and Torsten Wiesel

introduced a tungsten microelectrode into primary visual cortex (V1)

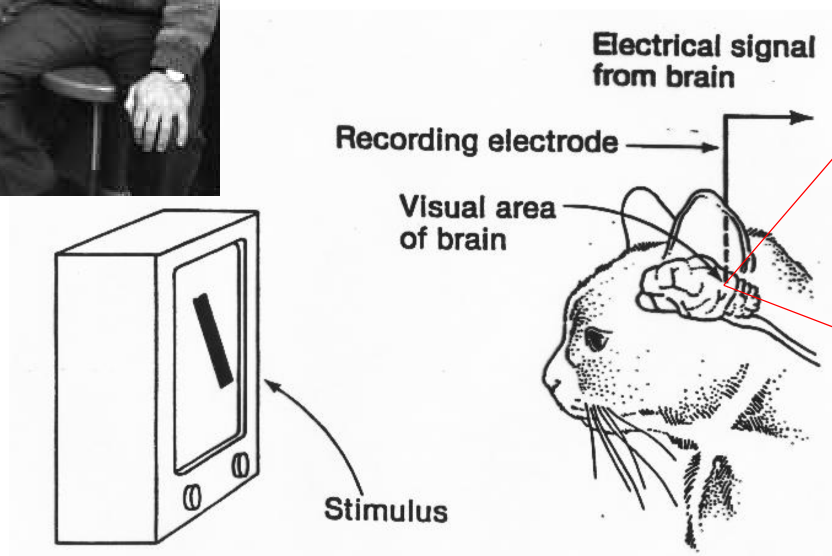

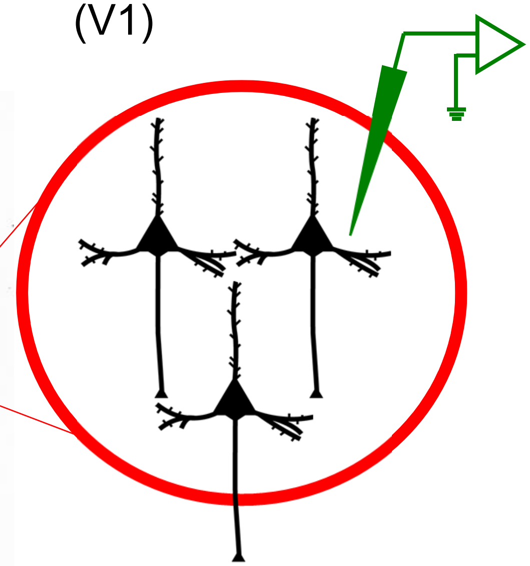

the canonical approach to visual selectivity

David Hubel and Torsten Wiesel

introduced a tungsten microelectrode into primary visual cortex (V1)

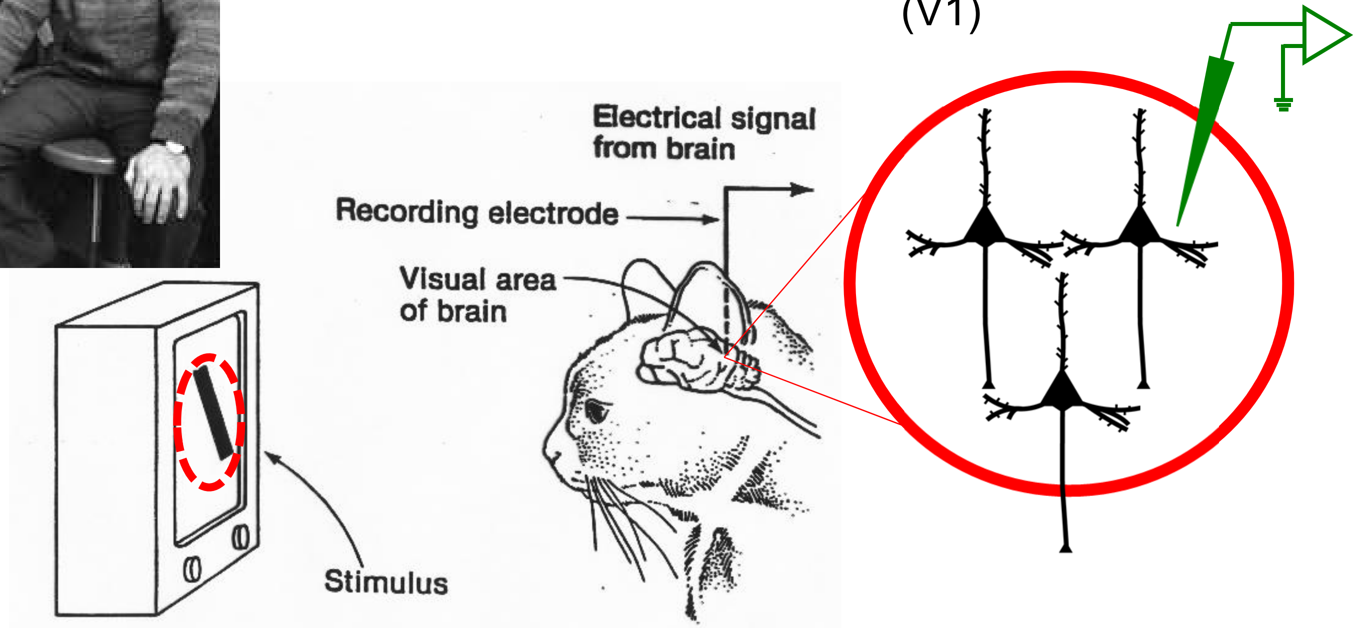

visual region to which the neuron responds (“receptive field”)

David Hubel Stephen Kuffler

http://braintour.harvard.edu

DISCHARGE PATTERNS AND FUNCTIONAL ORGANIZATION OF MAMMALIAN RETINA*

STEPHEN W. KUFFLER
The Wilmer Institute, Johns Hopkins Hospital and University Baltimore, Maryland

(1951)

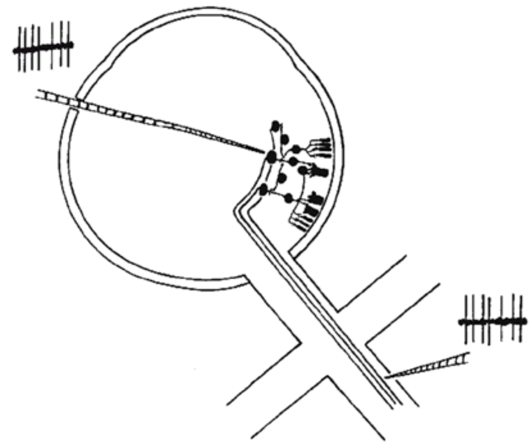

Enroth-Cugell and Robson, 1984

David Hubel Stephen Kuffler

http://braintour.harvard.edu

DISCHARGE PATTERNS AND FUNCTIONAL ORGANIZATION OF MAMMALIAN RETINA*

STEPHEN W. KUFFLER

The Wilmer Institute, Johns Hopkins Hospital and University
Baltimore, Maryland

(1951)

simple light spots are effective stimuli
for retinal ganglion cells

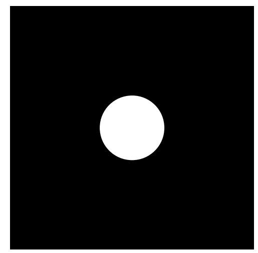

action potentials (“spikes”)

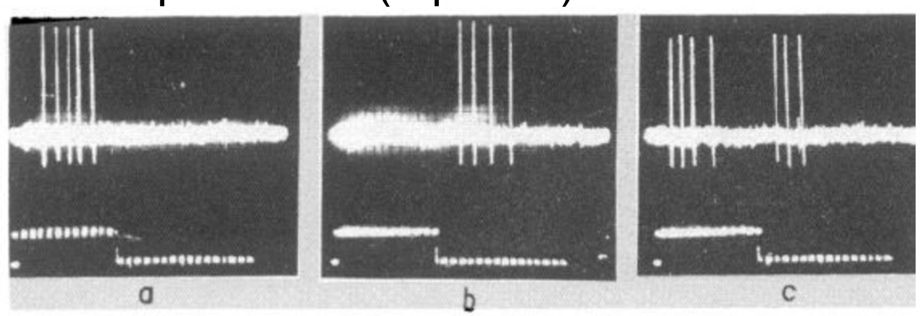

FIG. 4. Specific regions within receptive field. 0.2 mm. diameter light spot

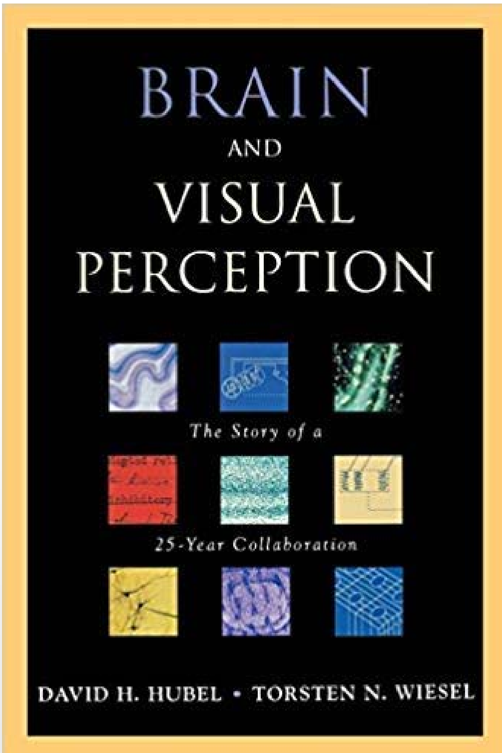

“in our very first experiments, we used circular spots…because these had served Stephen Kuffler so well”

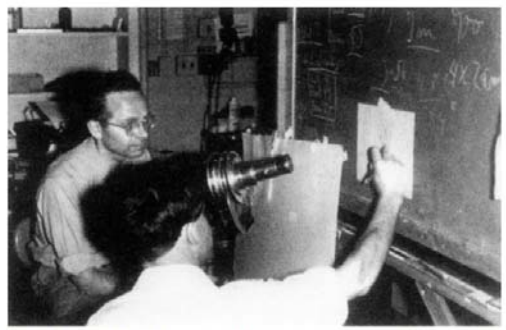

Figure 2. Hubel and Wiesel mapping a receptive field in cat visual cortex using a 'crude projector and screen'. (Photo source: Harvard Medical Library in the Francis A. Countway Library of Medicine.)

modified ophthalmoscope with spots painted on glass microscope slides

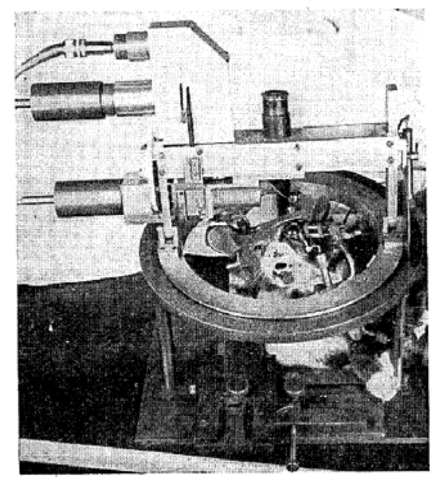

FIG. 1. Assembly for retinal studies from the unopened cat's eye. Cat in position (for details see text).

“early failures…were a matter of finding the right stimulus”

“We worked away, in shifts.

“Suddenly, as we inserted one of our glass slides into the ophthalmoscope, the cell seemed to come to life…

the cell was responding to the fine moving shadow cast by the edge of the glass slide” the neuron was selective for an edge at a given orientation

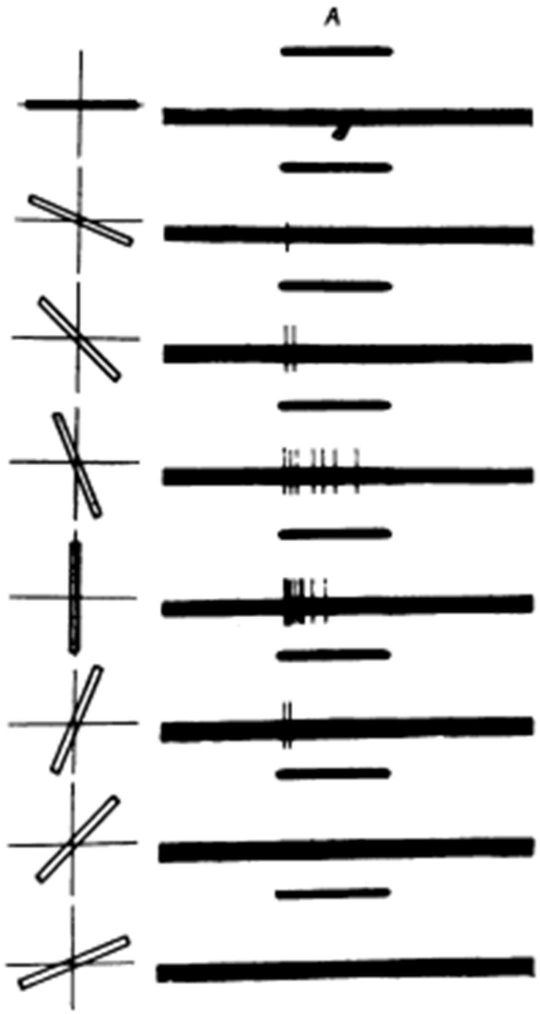

Hubel and Wiesel (1959)

the neuron was selective for an edge at a given orientation

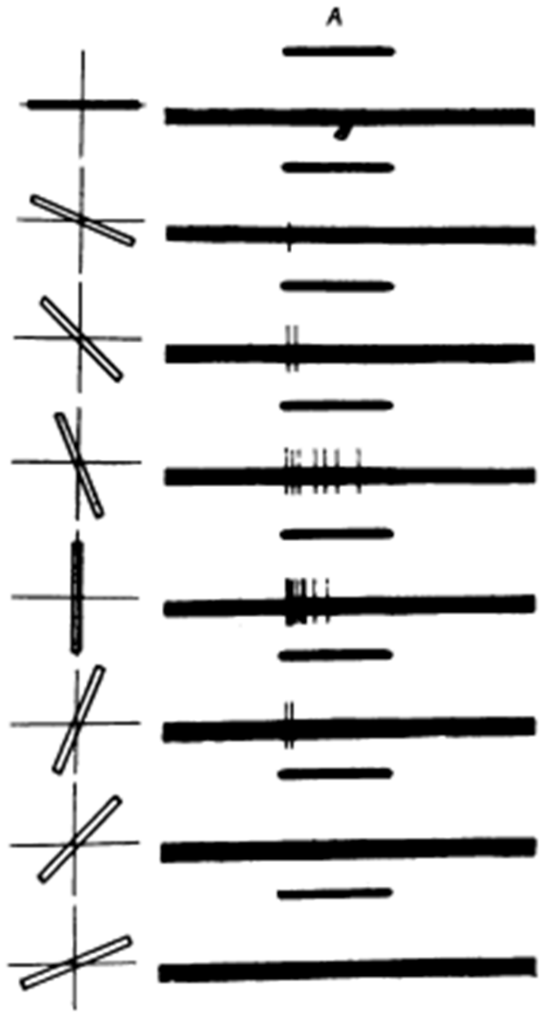

Hubel and Wiesel (1959)

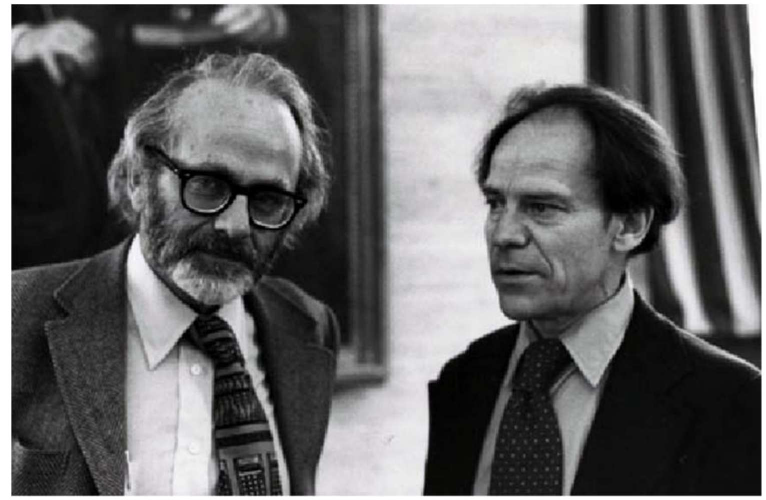

Orientation selectivity:

V1 neurons respond to lines/edges:

thus they *represent* orientation values present in the retinal scene

A wonderful foundation for visual neuroscience:

- their scientific mission required exploration
  (not enough known about the visual system)

- they relied on previous successes (Kuffler's spots)

- they succeeded because of a little luck and "bullheaded persistence" (Hubel, 2005)

They also highlighted an interesting challenge to their successors:

as we explore the rest of the ventral stream,
how do we not miss the proverbial *slide's edge*?

this problem grows as we move along the ventral stream

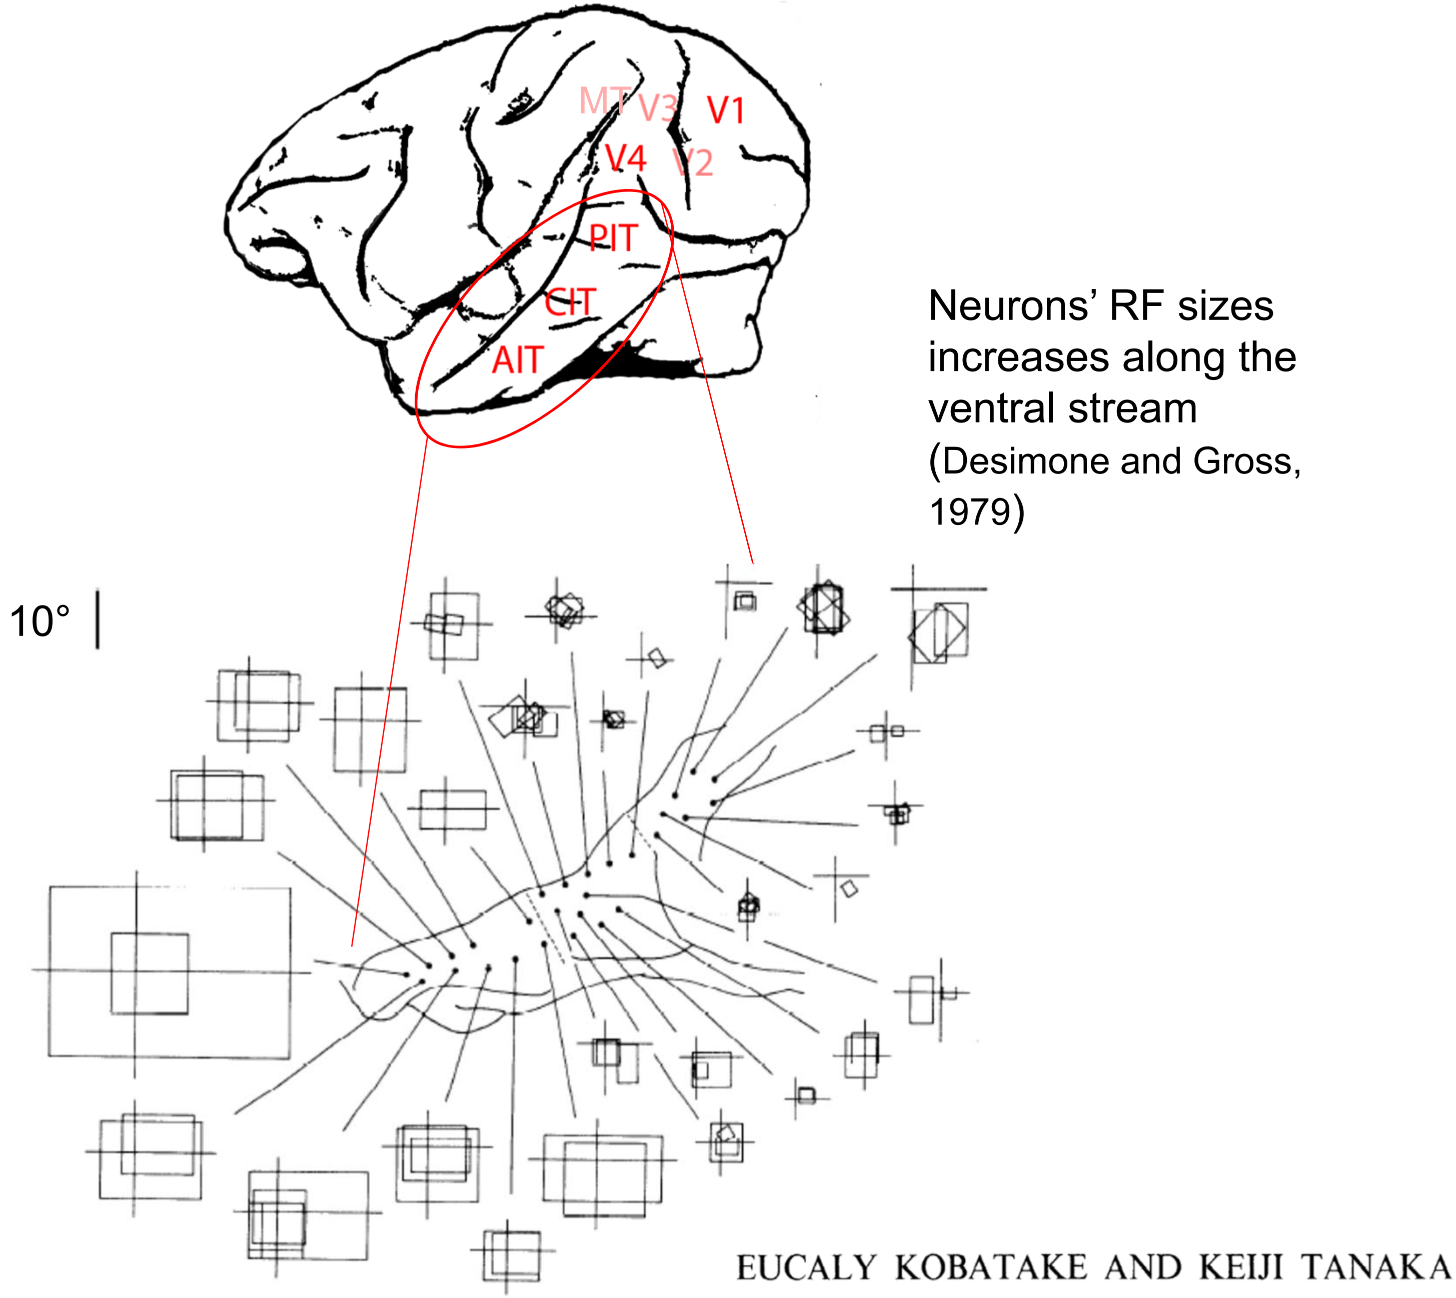

Neurons' RF sizes increases along the ventral stream (Desimone and Gross, 1979)

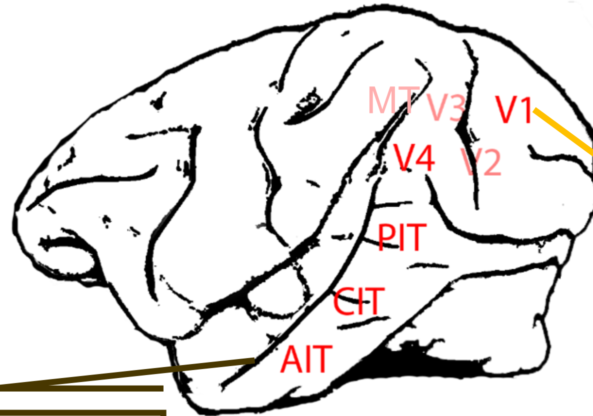

IT

neurons with larger RFs can respond to much more complex stimuli –

Hubel and Wiesel (1959)

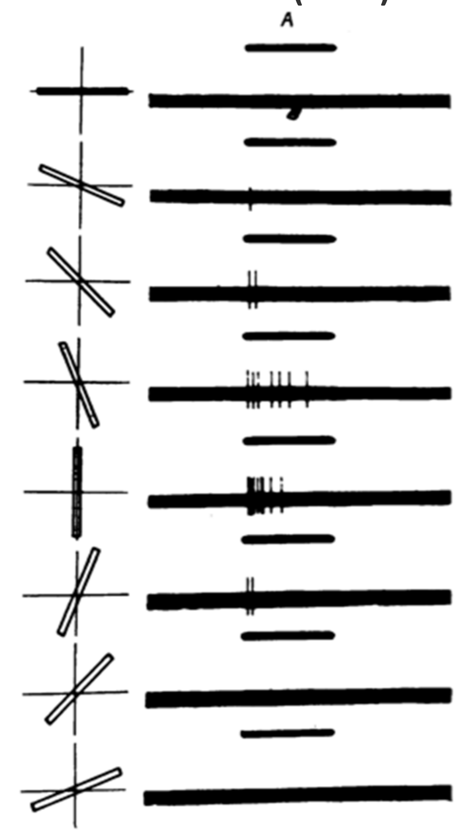

in IT, neurons respond to pictures of monkey and human faces, hands, places, artificial objects…

STIMULUS-SELECTIVE PROPERTIES OF INFERIOR TEMPORAL NEURONS IN THE MACAQUE[1]

ROBERT DESIMONE,[*,2] THOMAS D. ALBRIGHT,‡ CHARLES G. GROSS,‡ AND CHARLES BRUCE§

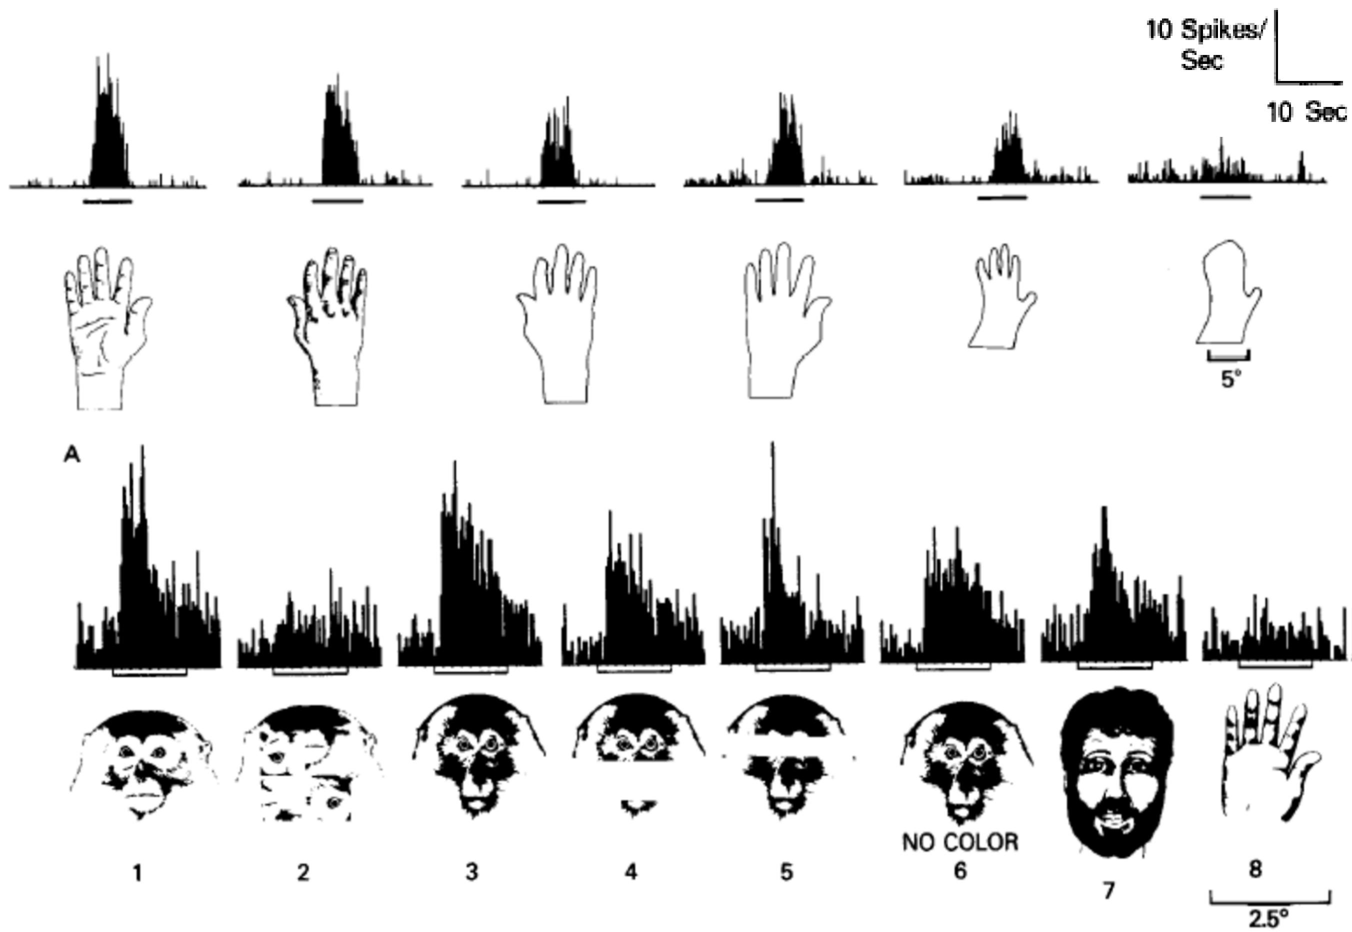

edges

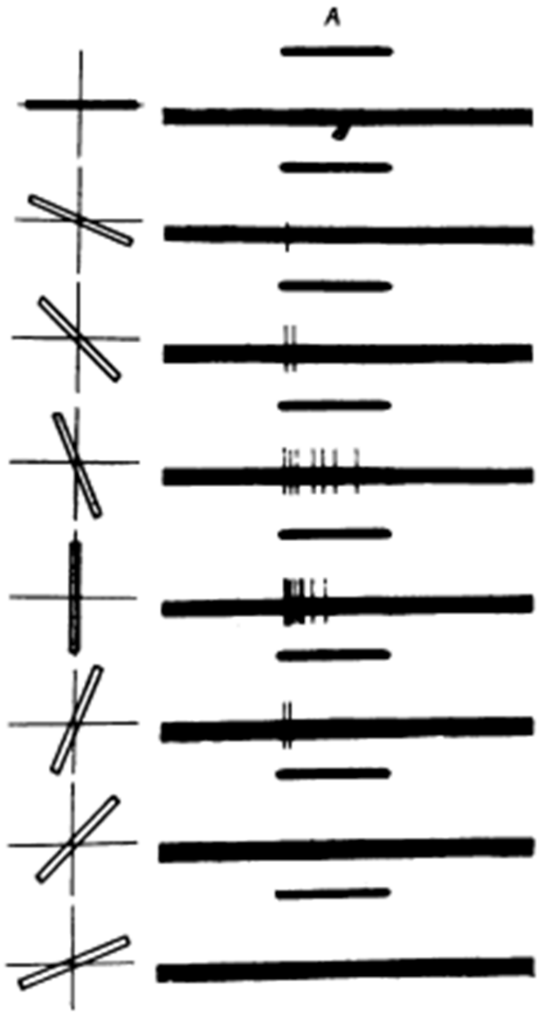

Orientation selectivity

also spatial frequency
wavelength/color

natural images

much more difficult to
parameterize!

thus we interpret neuronal tuning according to colloquial categories:

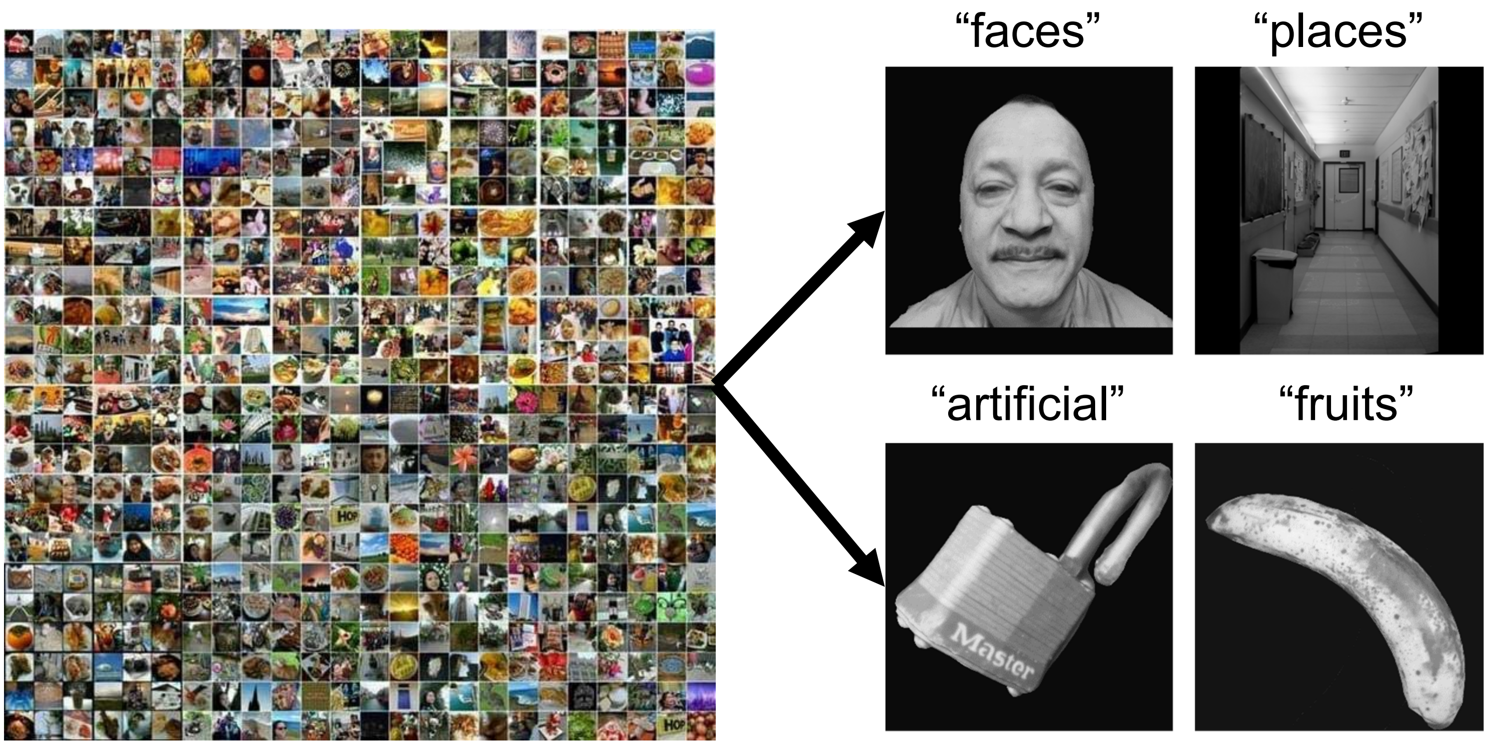

no theoretically rigorous way to

1) define what the complete set of category labels *should be*
   (every lab does its own thing)

2) know how to assign any given image to a category

3) Neurons do not respect boundaries: face cells can respond to fruits, place cells to bodies, …

“a matter of finding the right stimulus”

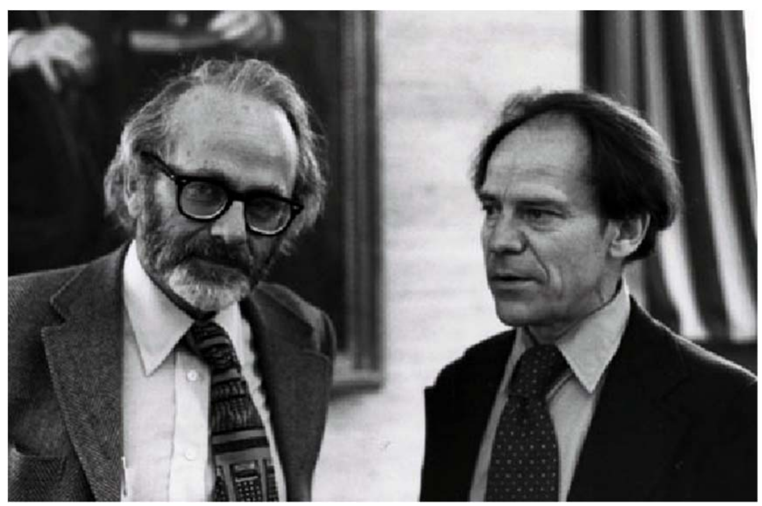

using evolutionary search algorithms with IT neurons

Yamane, Carlson, Bowman, Wang and Connor (2008)

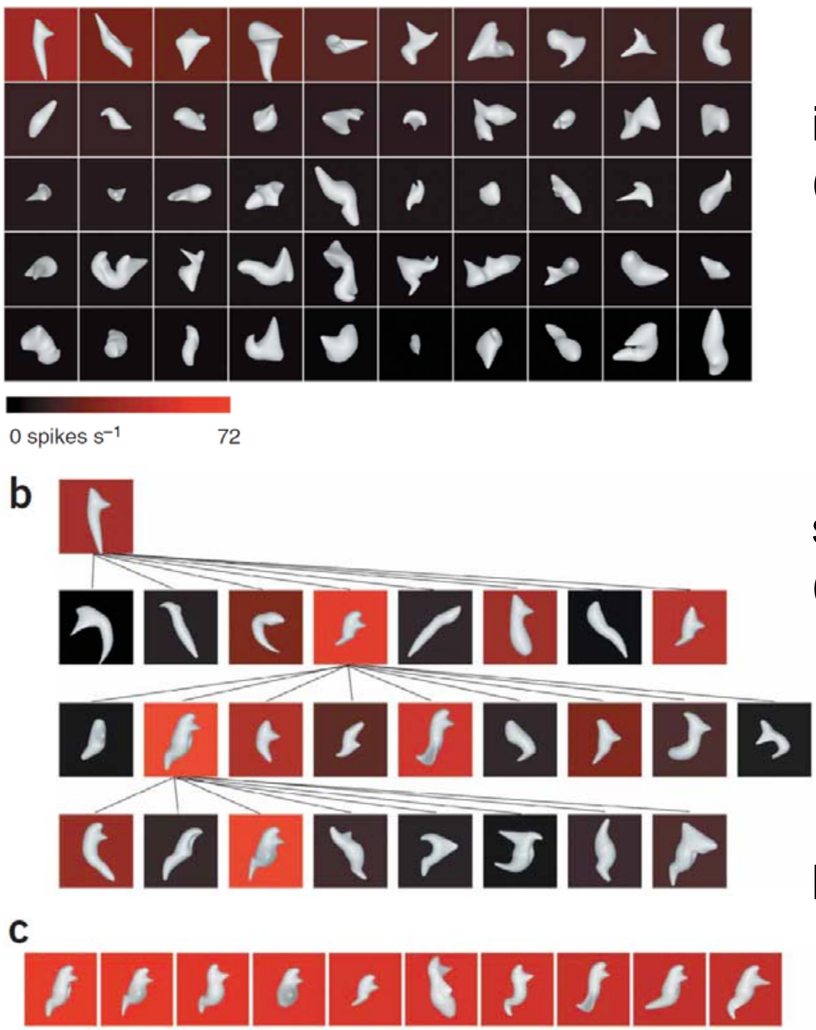

initial set of random 3-D shapes
(non-uniform rational basis splines, via OpenGL)

stimuli evolve under guidance of neuronal activity
(using an evolutionary search algorithm)

preferred 3-D shapes as directed by the neuron evolutionary algorithms with a larger feature space

analogy: a map of the observable universe

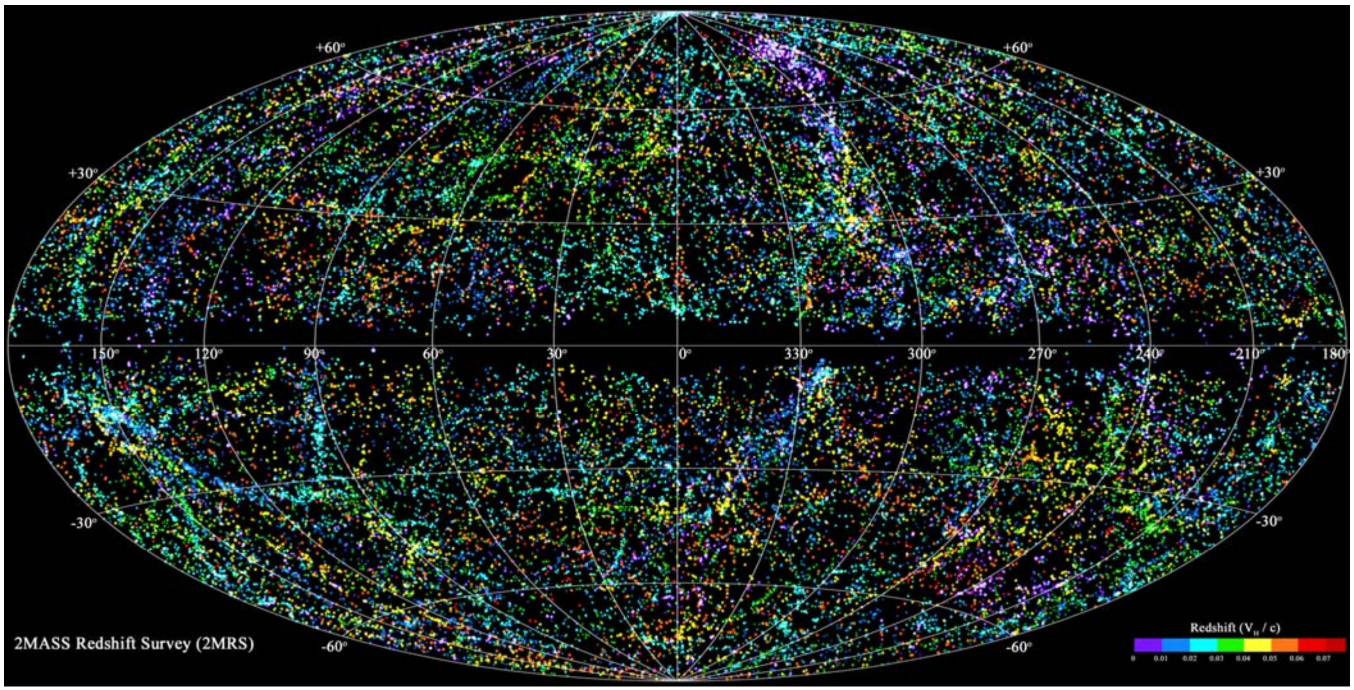

evolutionary algorithms with a larger feature space

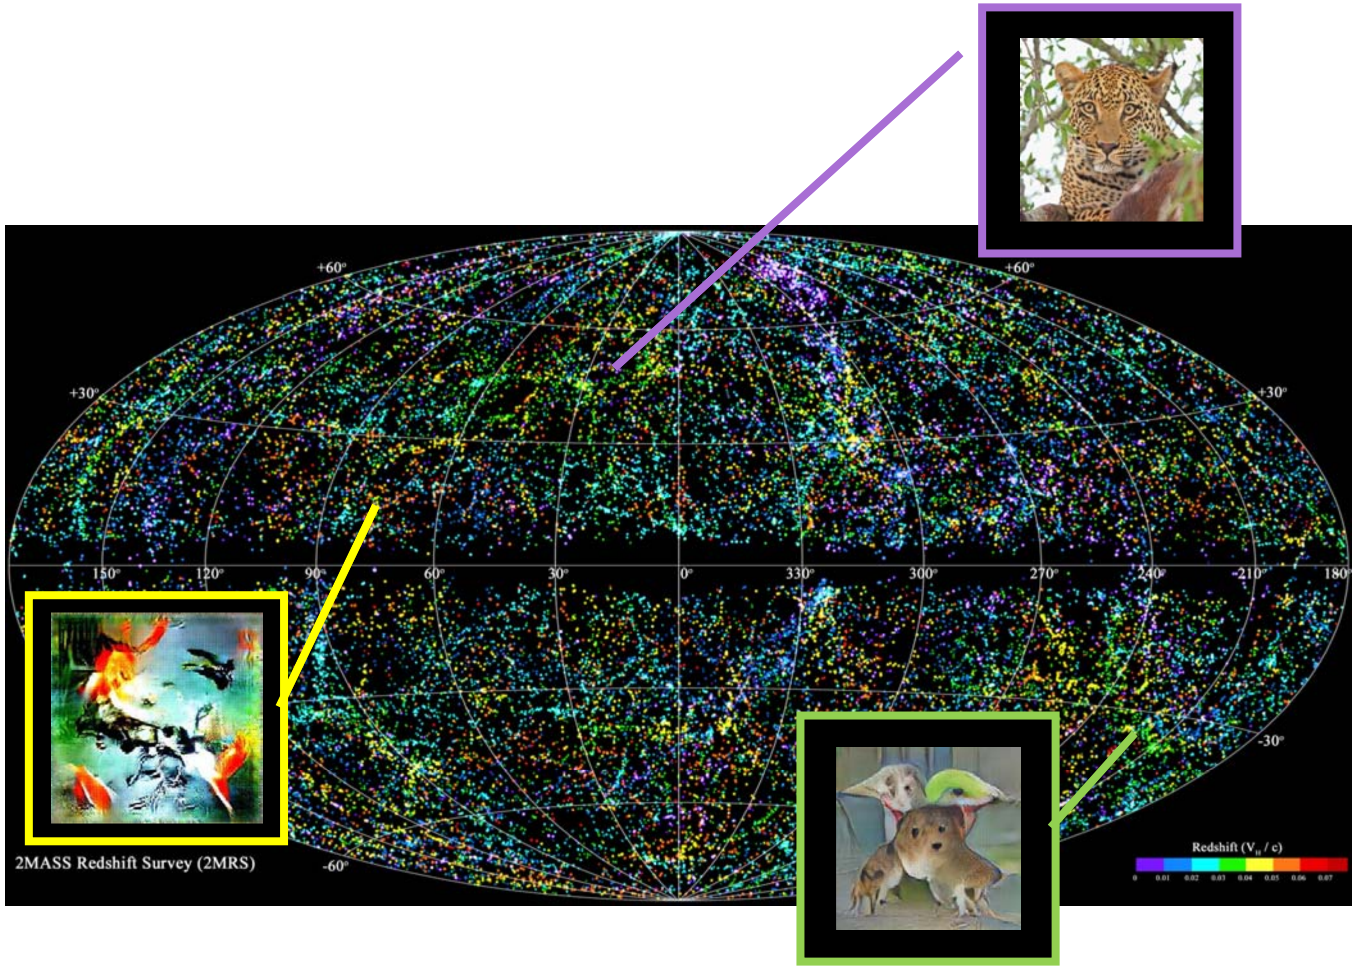

evolutionary algorithms with a larger feature space:
generative adversarial networks

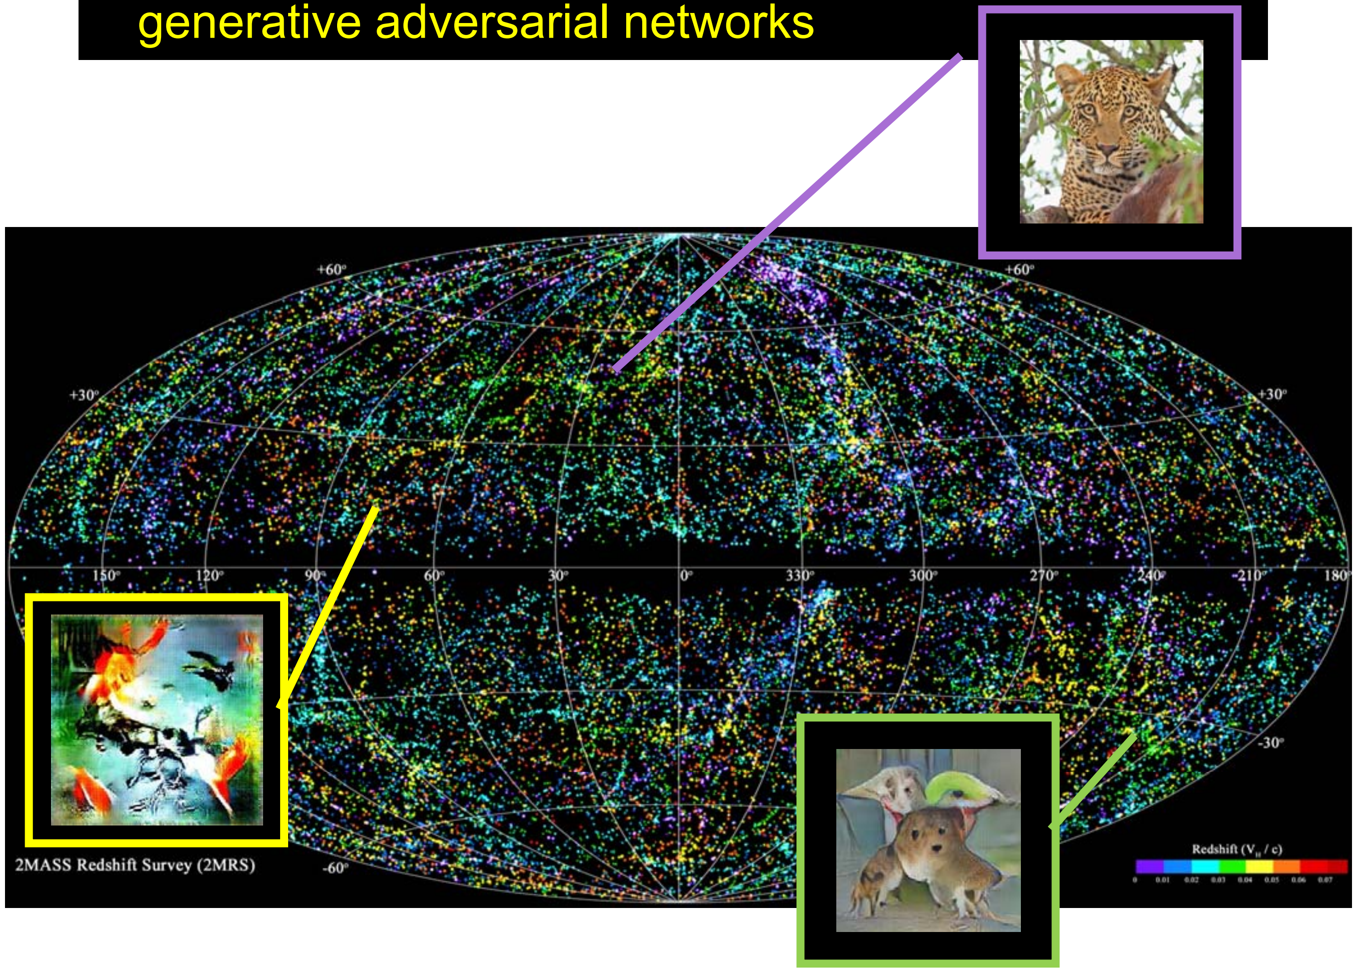

what are generative adversarial networks?

meanwhile in machine learning…

ImageNet Large Scale Visual Recognition Challenge

error rate

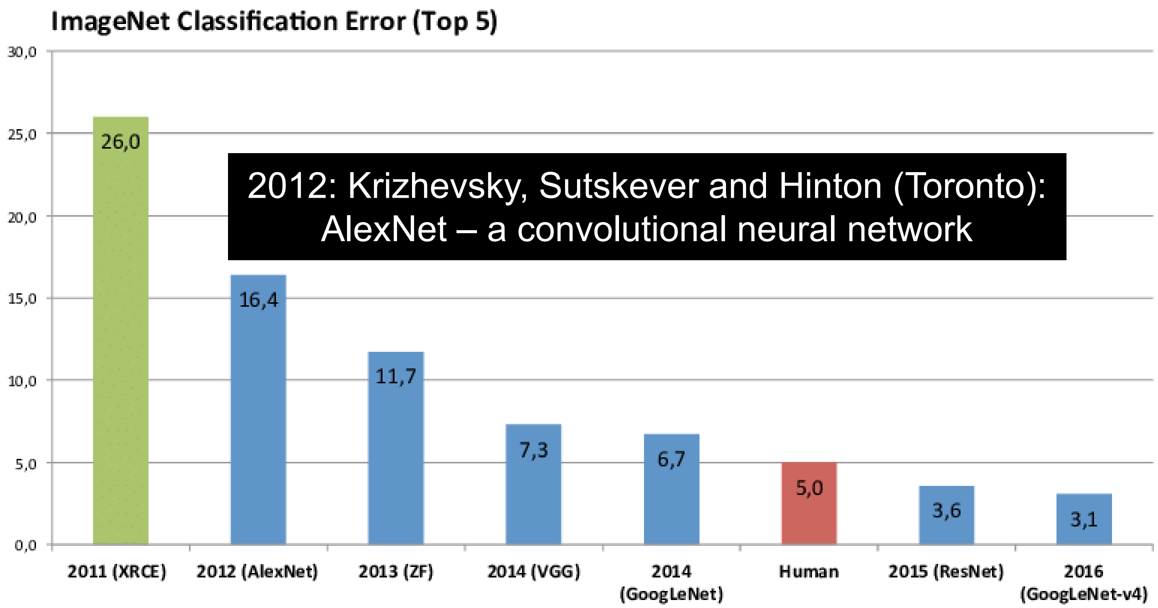

Gustav von Zitzewitz

convolutional neural networks ("convnets")

image as input

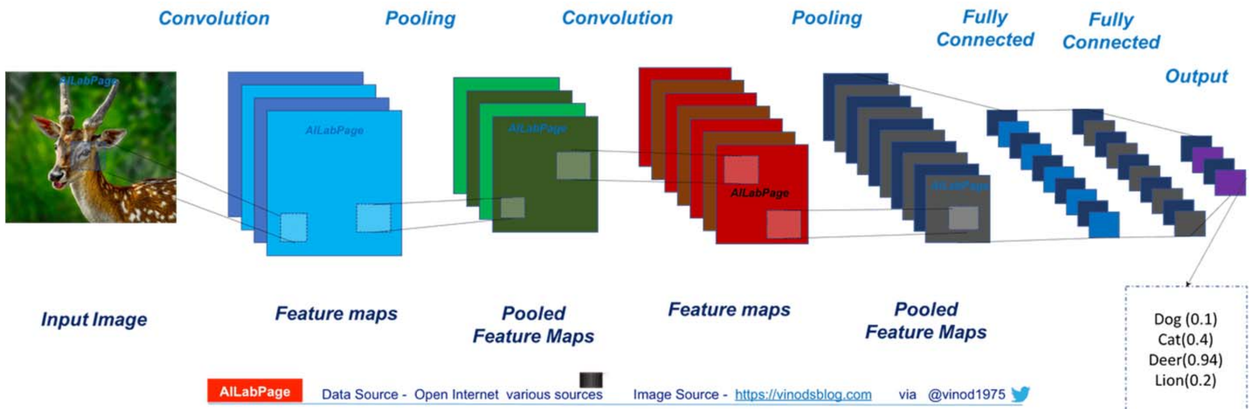

outputs a vector (compact representation of a complex visual feature, e.g. deer or cat present?)

convolutional neural networks ("convnets")

image as input

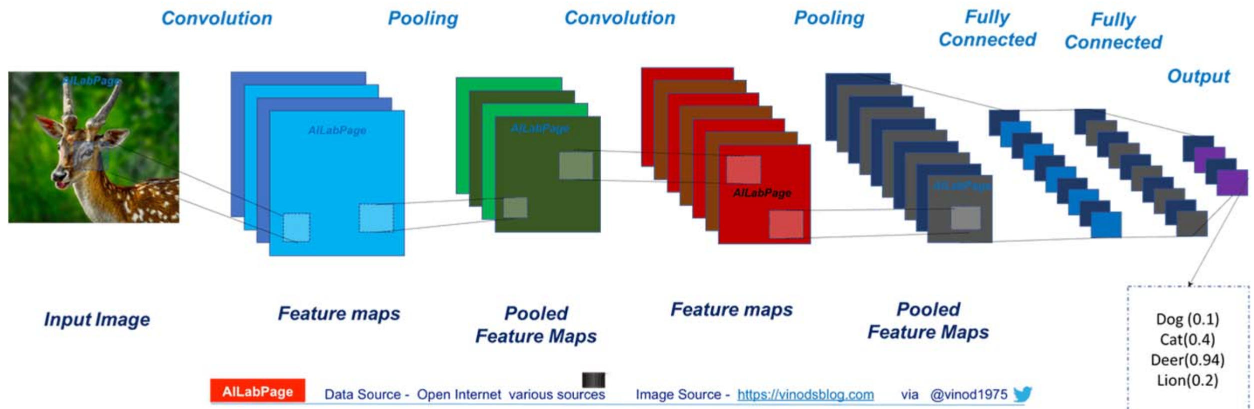

outputs a vector (compact representation of a complex visual feature, e.g. deer or cat present?)

AlexNet (Toronto)

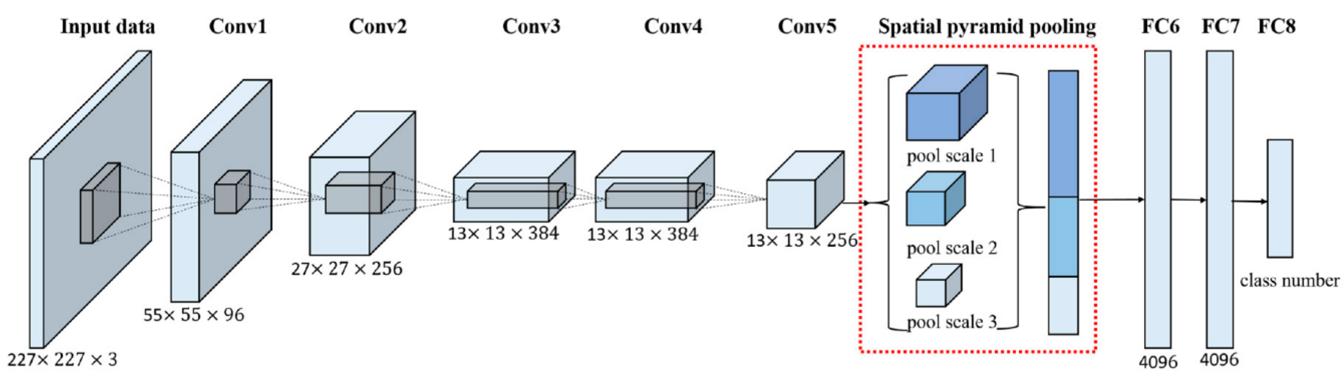

convolutional neural networks (“convnets”)

image as input

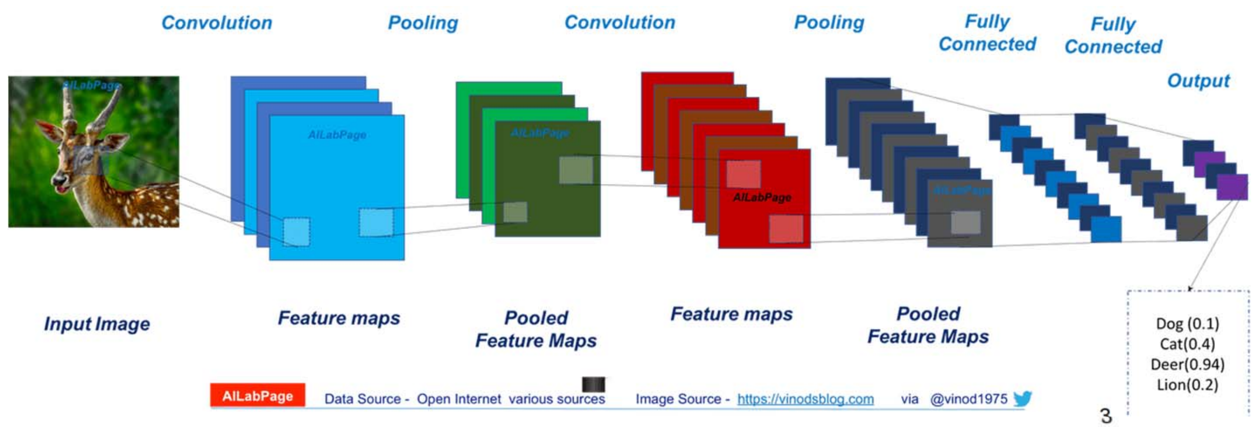

outputs a vector (compact representation of a complex visual feature, e.g. deer or cat present?)

generative adversarial networks (“GANs”)

vector as input

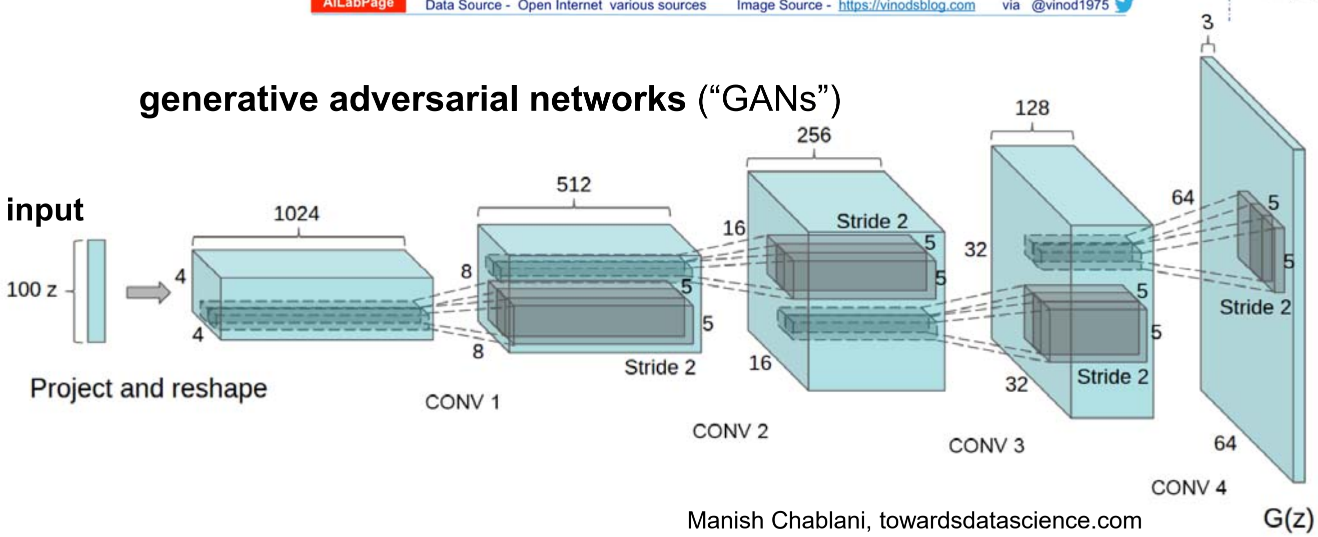

Manish Chablani, towardsdatascience.com

image as output

GANs learn to copy distributions:

features

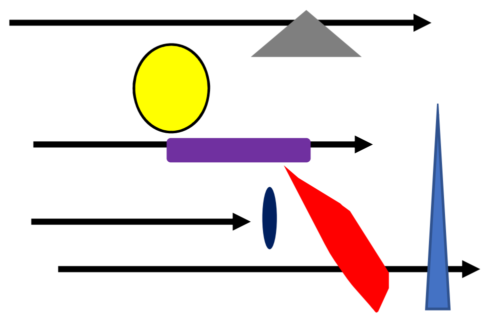

probability distribution

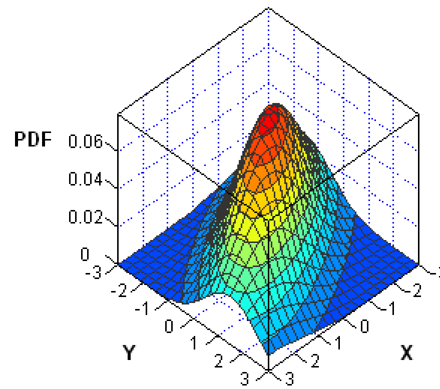

Karras et al, 2019, NVIDIA

What are they good for?

Uses being explored in the videogame, fashion, security industries, creating data for training self-driving cars

pix2pix

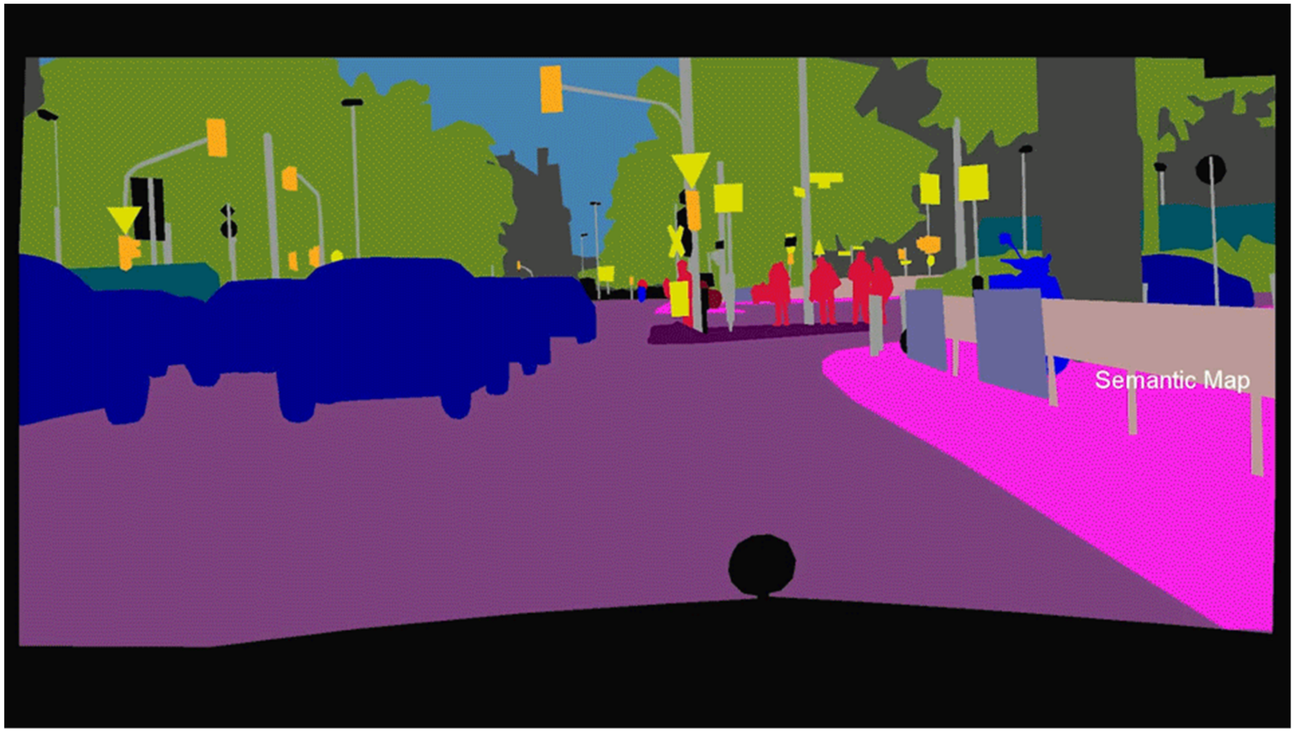

Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. "High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs", in CVPR, 2018.

BigGAN: Brock et al (2019), Google

@memotv

can we link a neuron in a monkey's brain to a GAN
–
and let it "build" its preferred complex image?

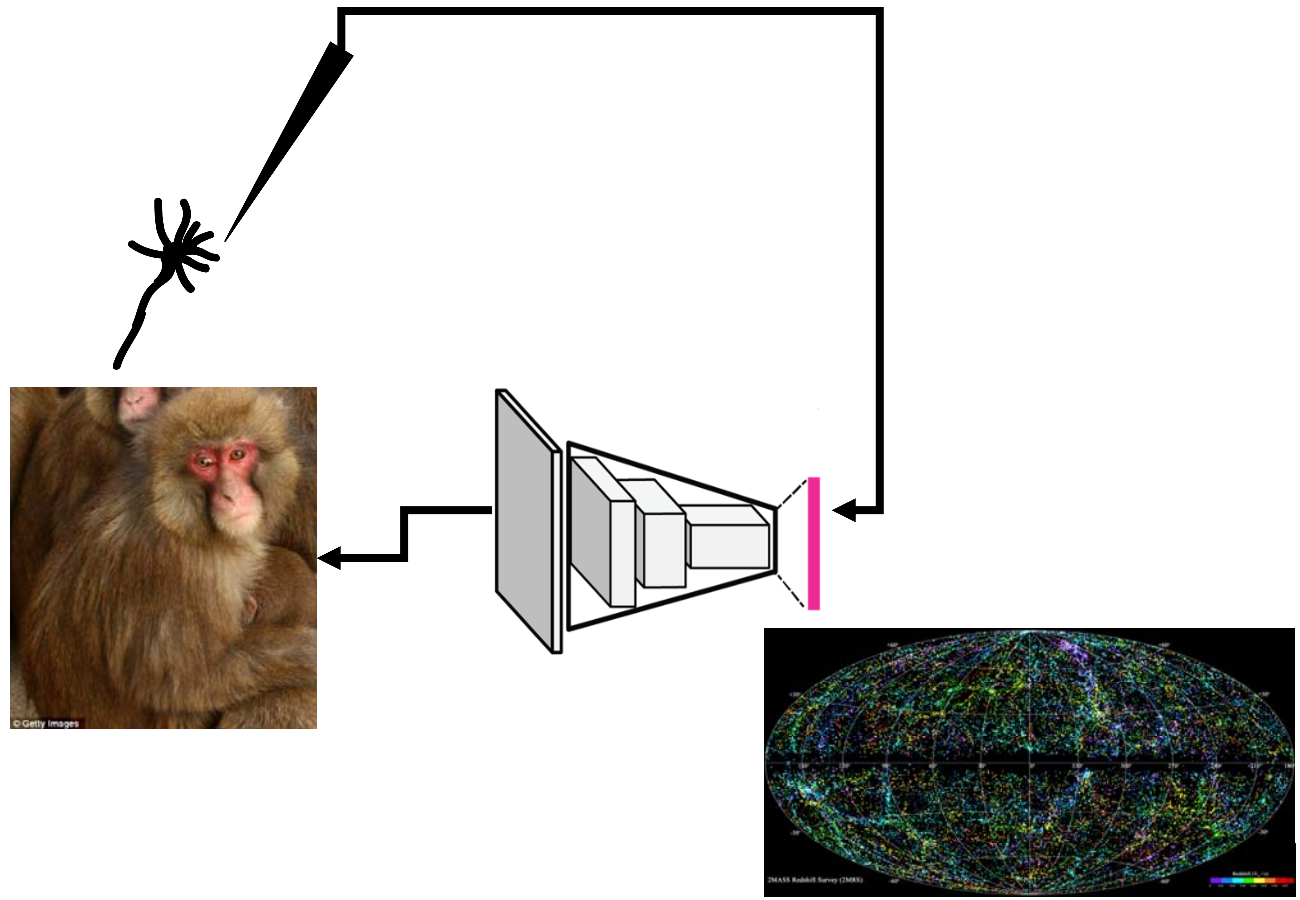

Will Xiao

Peter Schade

Till Hartmann

Gabriel Kreiman

Margaret Livingstone

Cell

Volume 177, Issue 4, 2 May 2019, Pages 999-1009.e10

Article

Evolving Images for Visual Neurons Using a Deep Generative Network Reveals Coding Principles and Neuronal Preferences

Carlos R. Ponce [1, 4, 5], Will Xiao [2, 5], Peter F. Schade [1, 5], Till S. Hartmann [1], Gabriel Kreiman [3], Margaret S. Livingstone [1, 6]

[1] Department of Neurobiology, Harvard Medical School, Boston, MA 02115, USA

[2] Department of Molecular and Cellular Biology, Harvard University, Cambridge, MA 02138, USA

[3] Department of Ophthalmology, Boston Children's Hospital, Harvard Medical School, Boston, MA 02115, USA

[4] Department of Neuroscience, Washington University School of Medicine, St. Louis, MO 63110, USA

Received 6 August 2018, Revised 5 November 2018, Accepted 2 April 2019, Available online 2 May 2019.

Synthesizing the preferred inputs for neurons in neural networks via deep generator networks

Anh Nguyen, Alexey Dosovitskiy, Jason Yosinski, Thomas Brox, Jeff Clune. 2016

Input: 4096-element vector

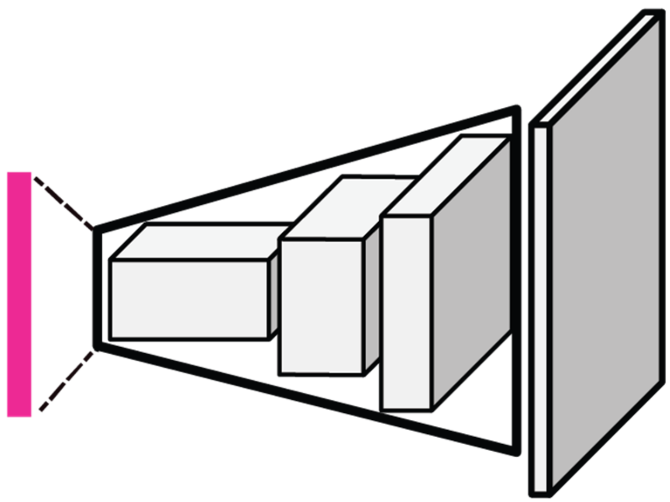

Output:
images
(256 x 256 x 3)

trained to invert
representations in **AlexNet**
layer fc6

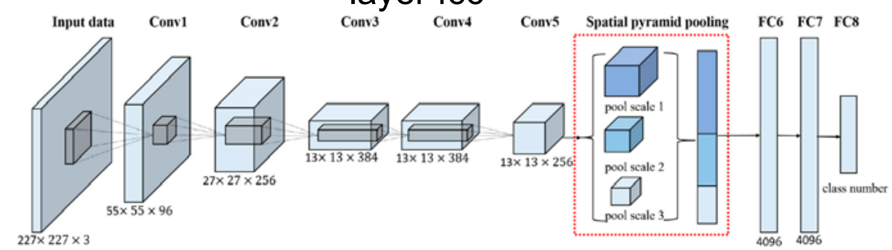

Synthesizing the preferred inputs for neurons in neural networks via deep generator networks

Nguyen *et al.* 2016

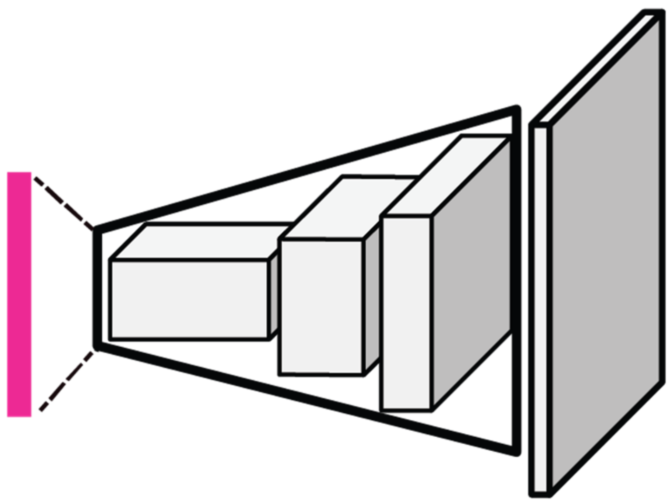

random vectors

non-random vectors

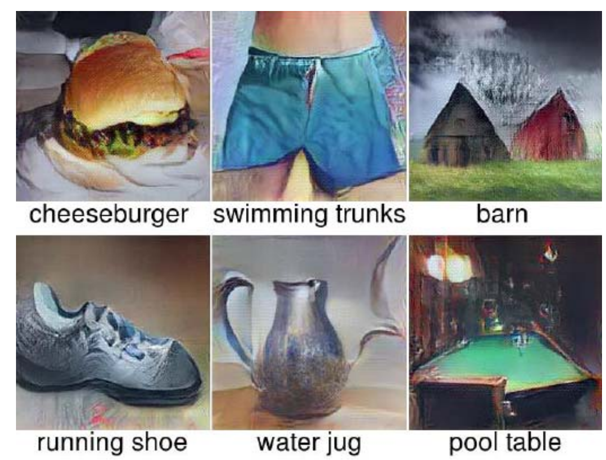

we recorded from six monkeys
with chronically implanted arrays

Posterior/central inferotemporal cortex (one in primary visual cortex)

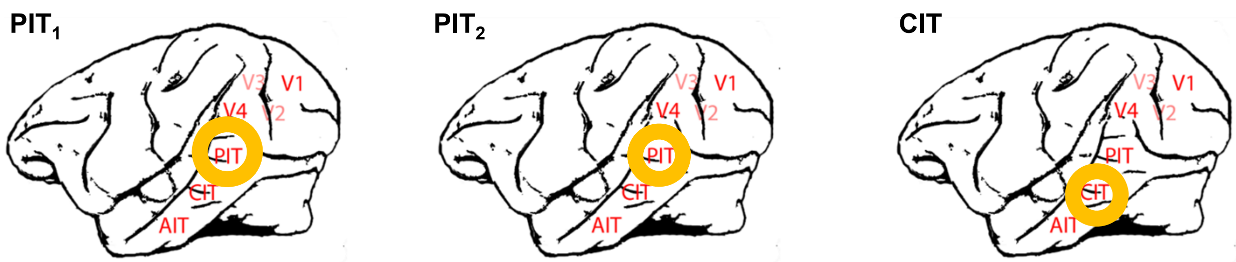

in my lab, now along the full ventral stream

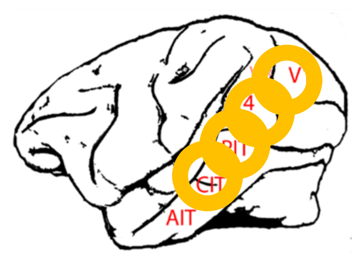

To let the neuron search through the vector space, we used a genetic algorithm (similar to those in the Connor lab)

Will Xiao

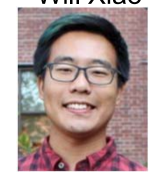

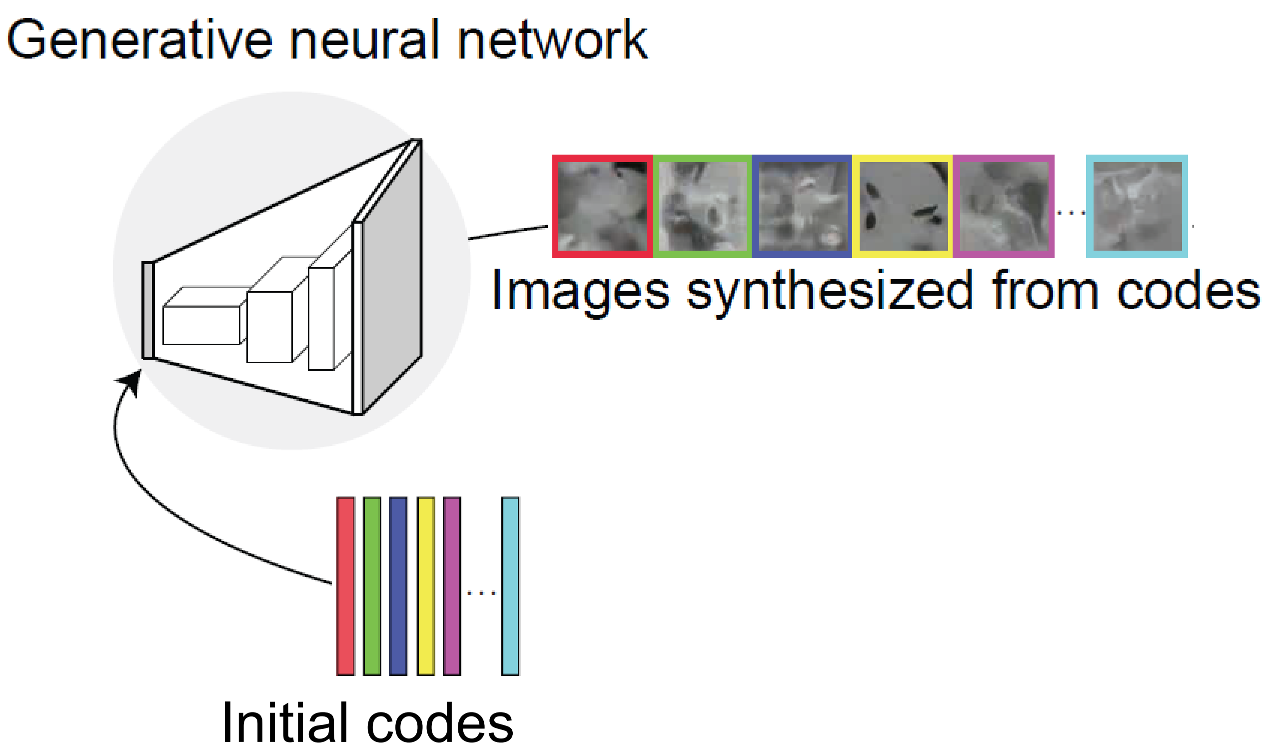

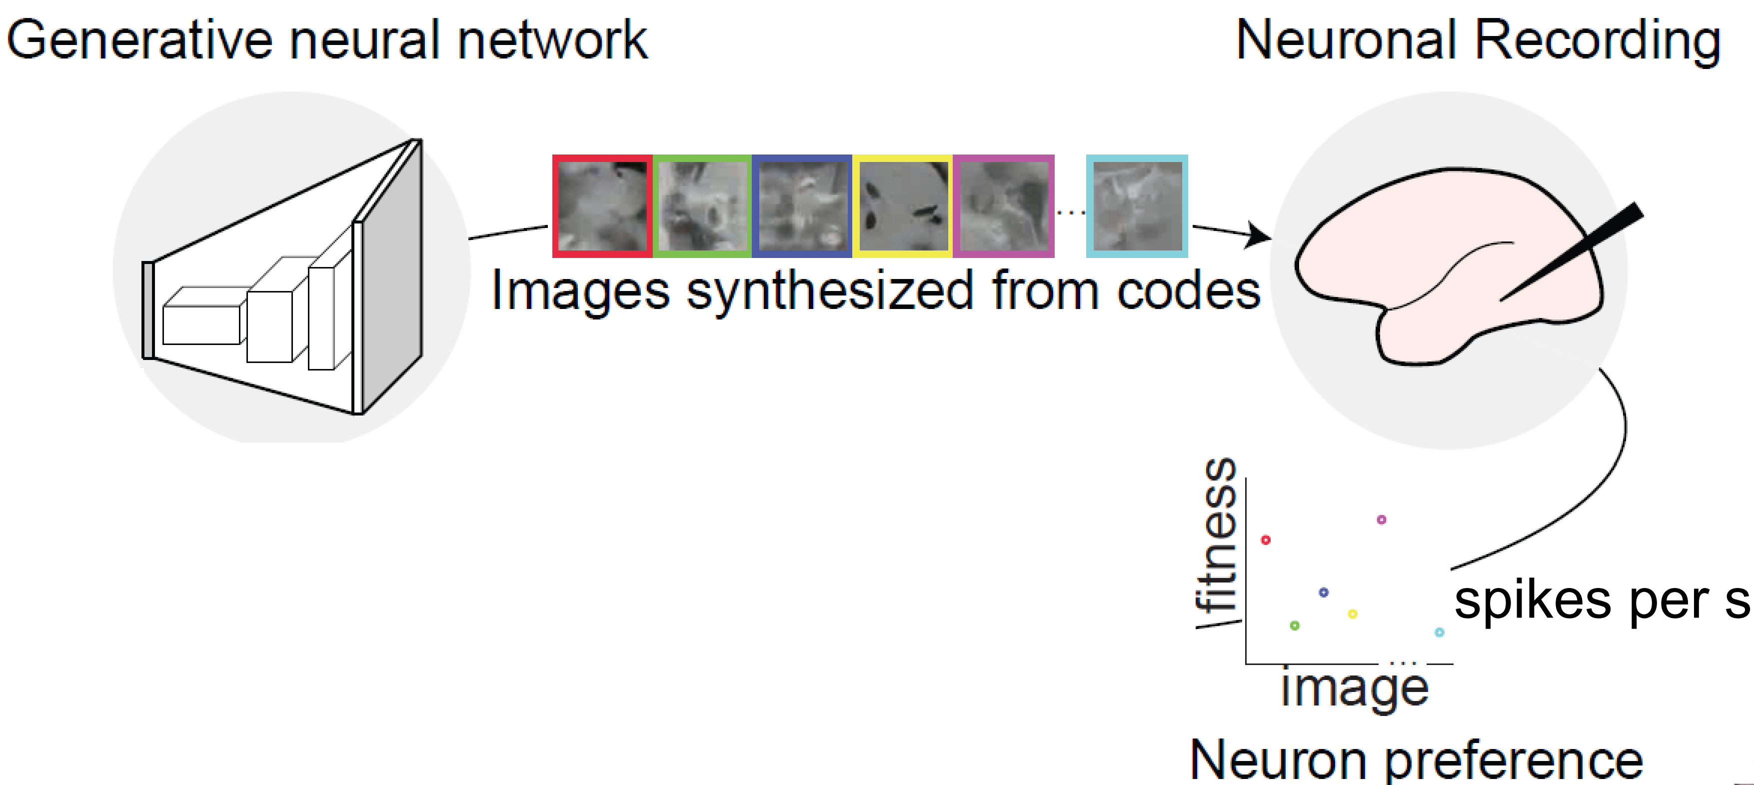

Generative neural network
Neuronal Recording
Images synthesized from codes
fitness
image
spikes per s
Neuron preference

Will Xiao

Generative neural network

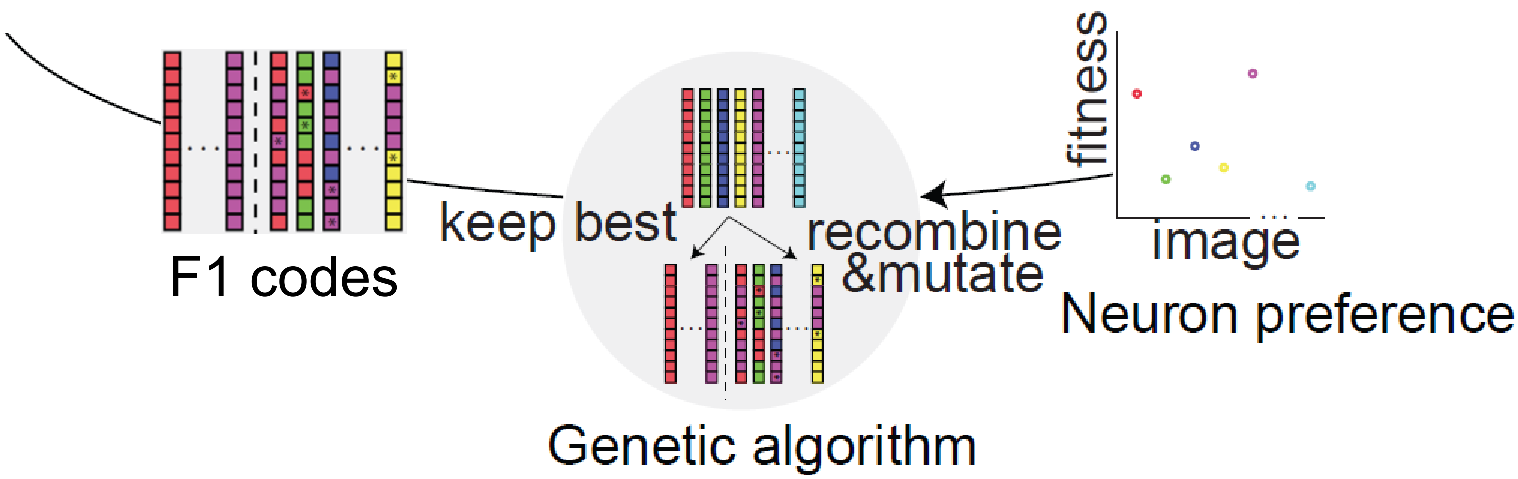

Will Xiao

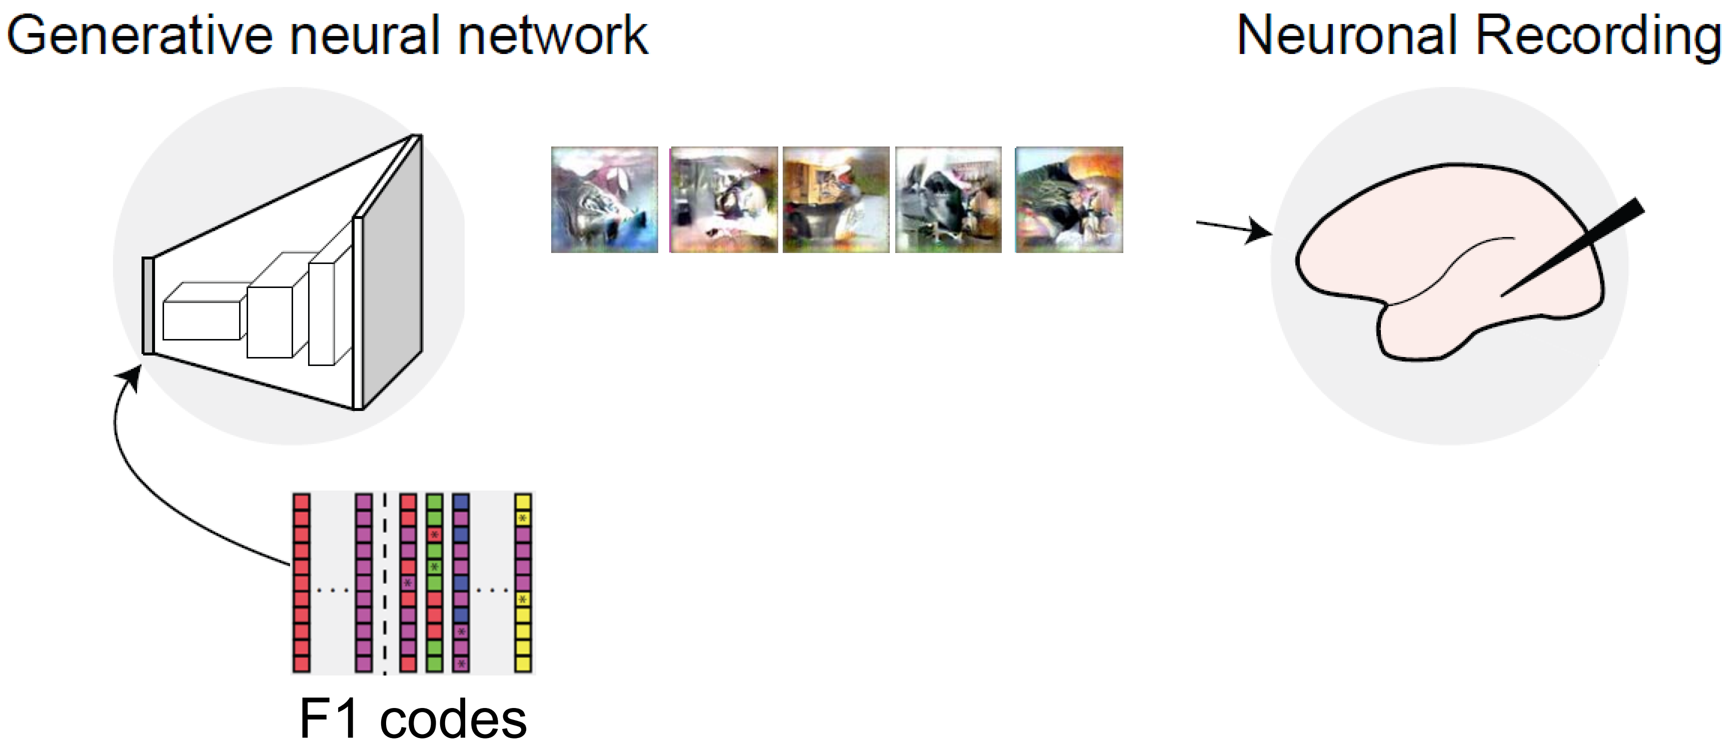

Will Xiao

Behavioral task

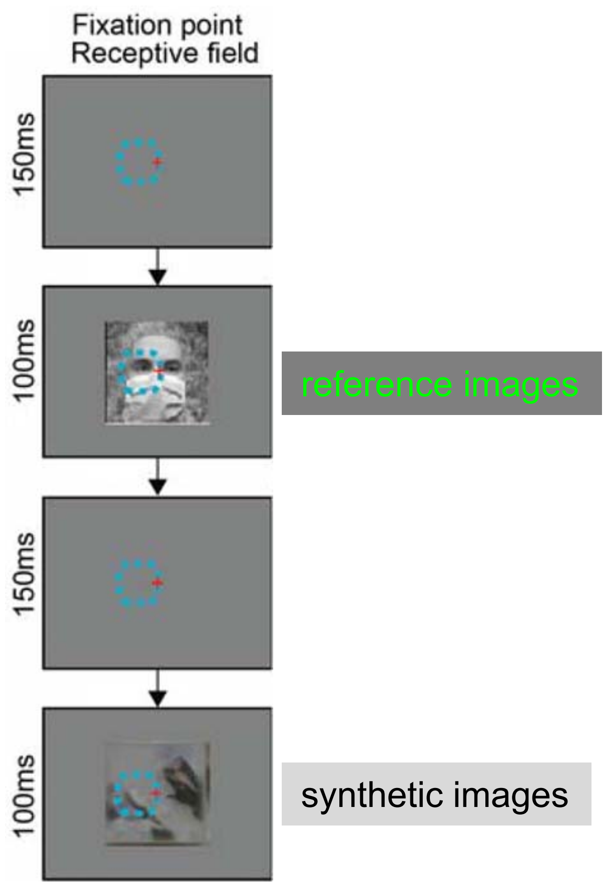

Behavioral task

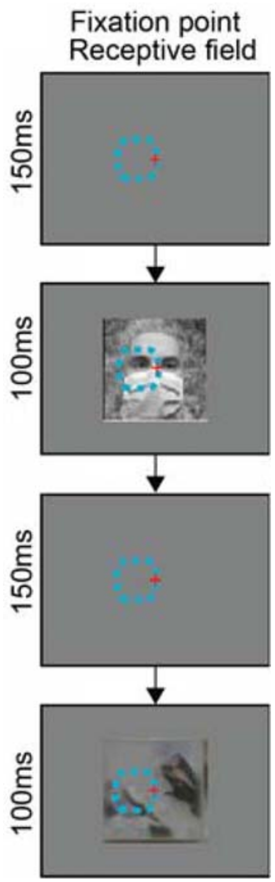

reference images

synthetic images

neurons showed increases in firing rate during the evolution of new images

PIT single unit: mean firing rate per generation

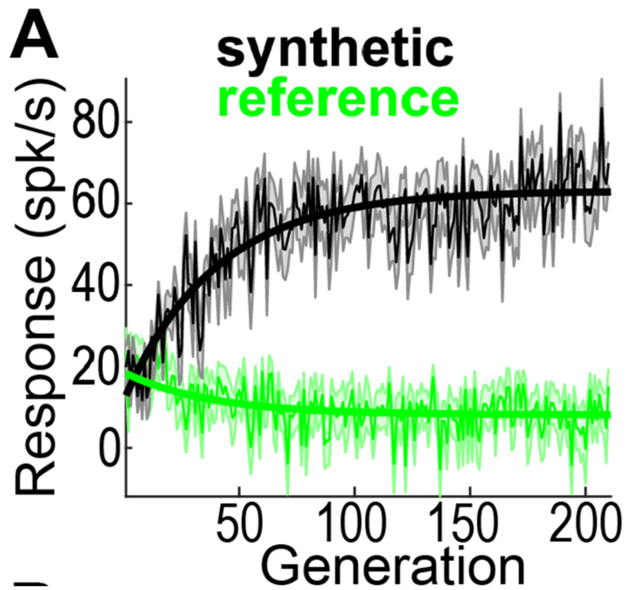

We replicated this effect in 6 animals, 46 experiments

change in response to evolving synthetic images (spikes/s)

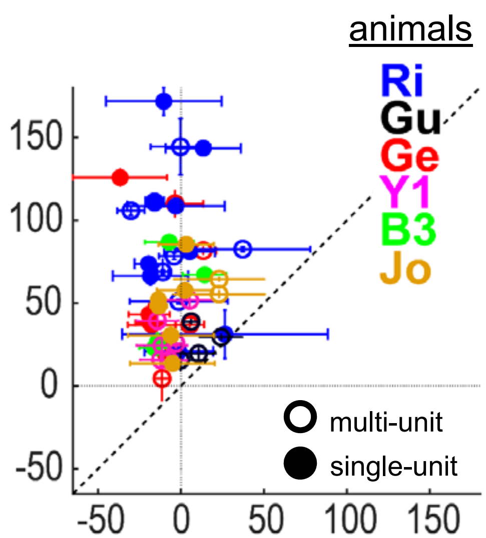

change in response to fixed natural images (spikes/s)

PIT single unit: mean firing rate per generation

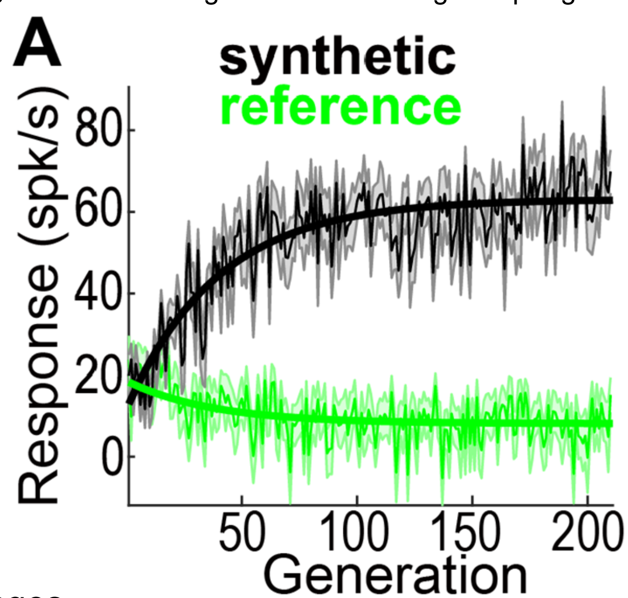

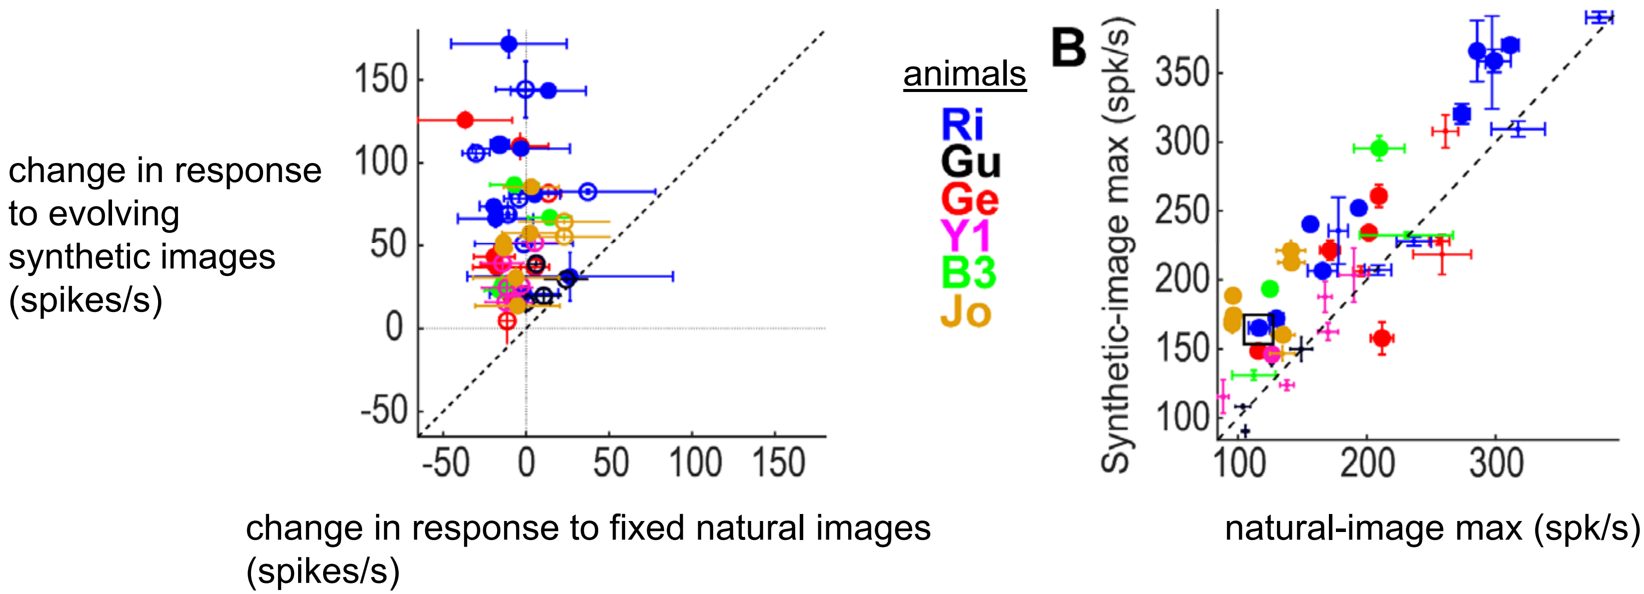
change in response to evolving synthetic images (spikes/s)
150
100
50
0
-50
-50
0
50
100
150
change in response to fixed natural images (spikes/s)
animals
Ri
Gu
Ge
Y1
B3
Jo
B
Synthetic-image max (spk/s)
350
300
250
200
150
100
100
200
300
natural-image max (spk/s)

so what's happening here?

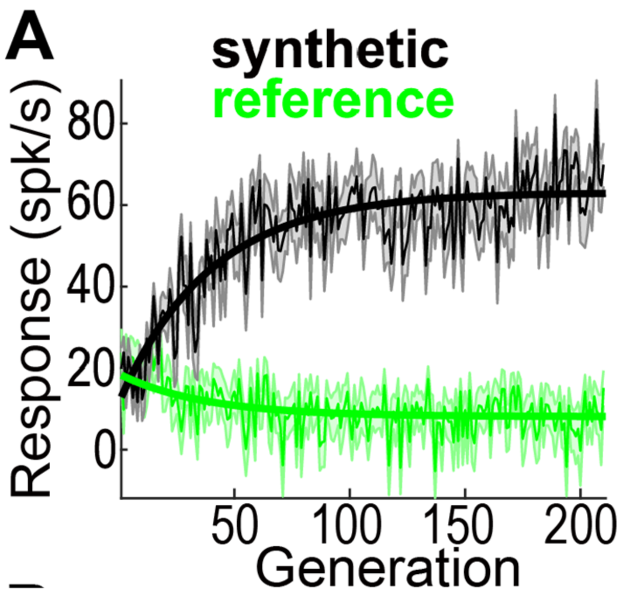

Initial generation: 30-40 Simoncelli and Portilla textures

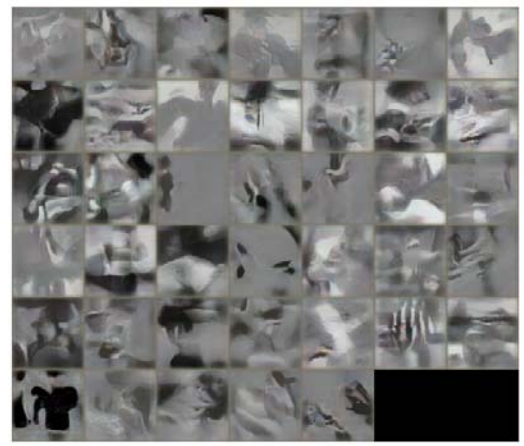

next slide: one complete experiment,
mean synthetic image (top 5 images) per generation

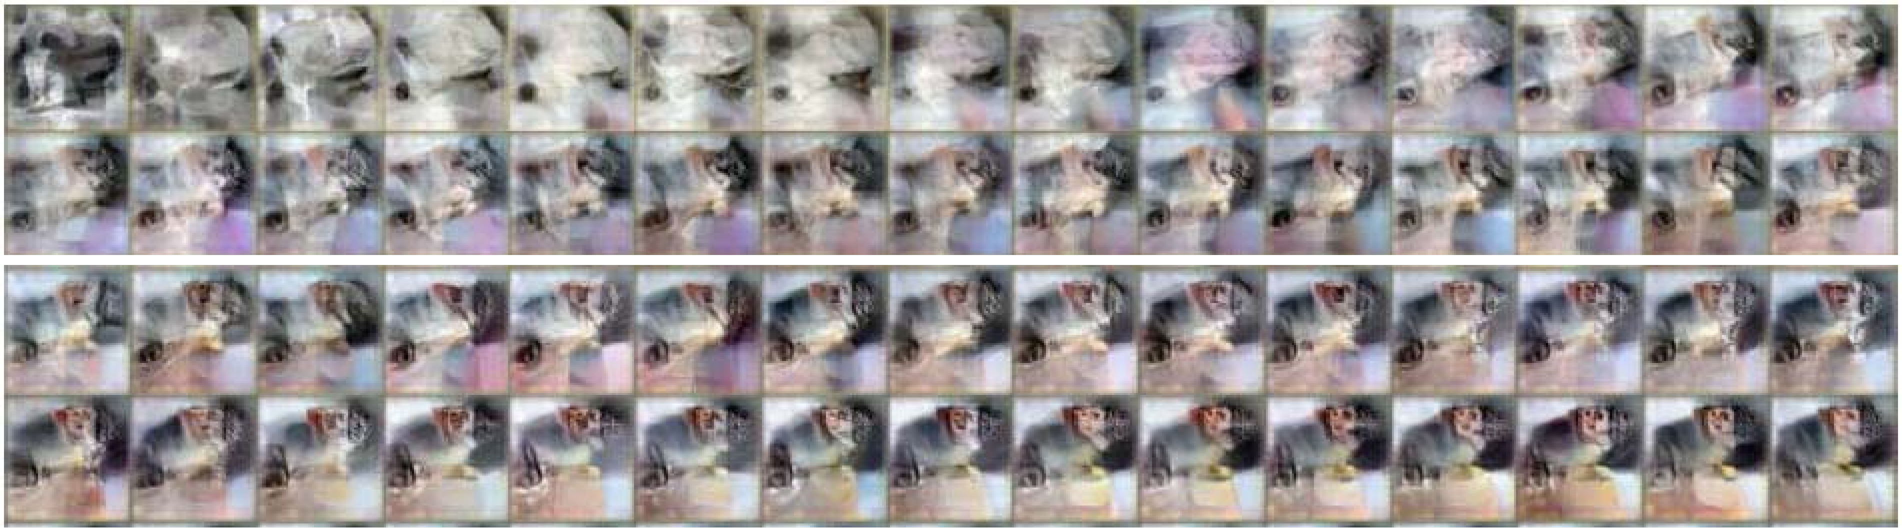

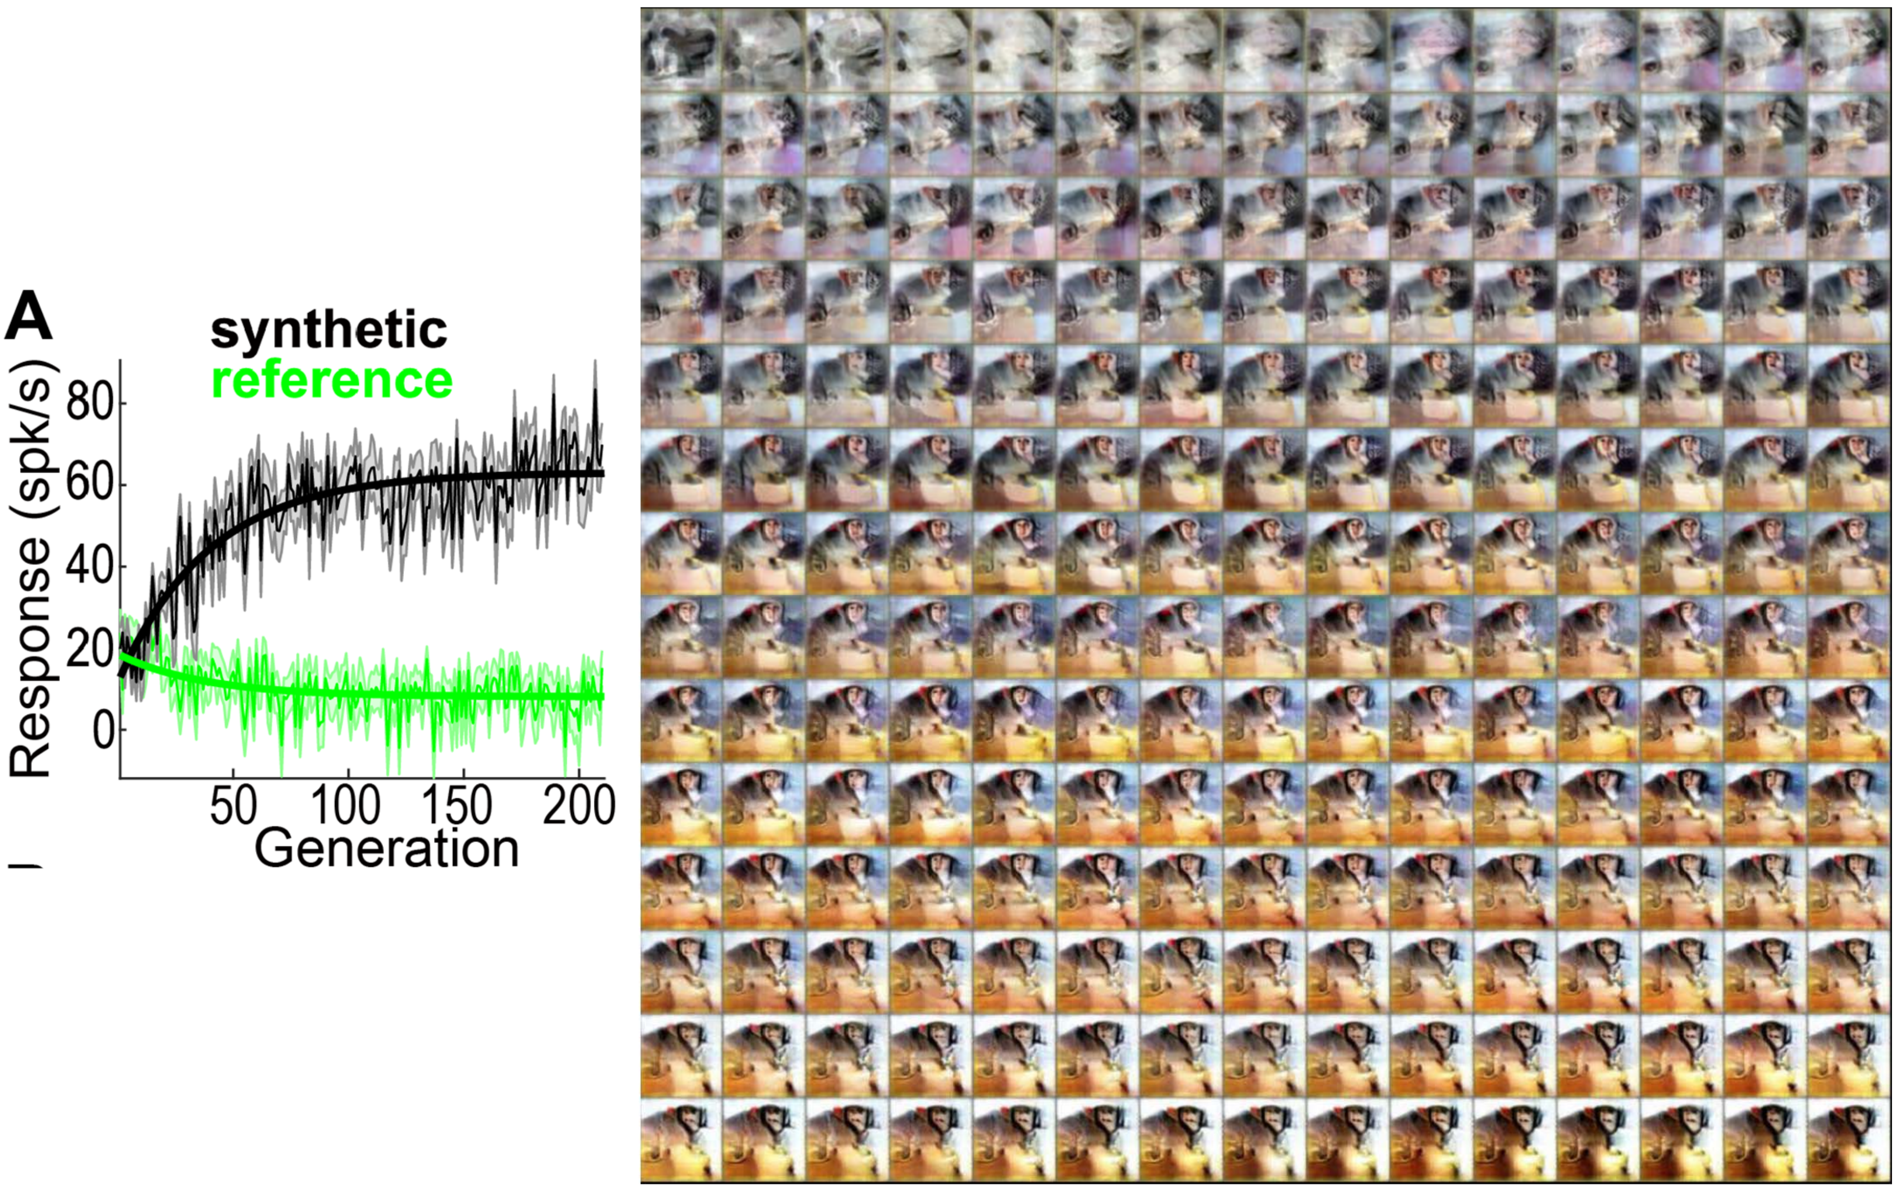
A
synthetic
reference
Response (spk/s)
0
20
40
60
80
50
100
150
200
Generation

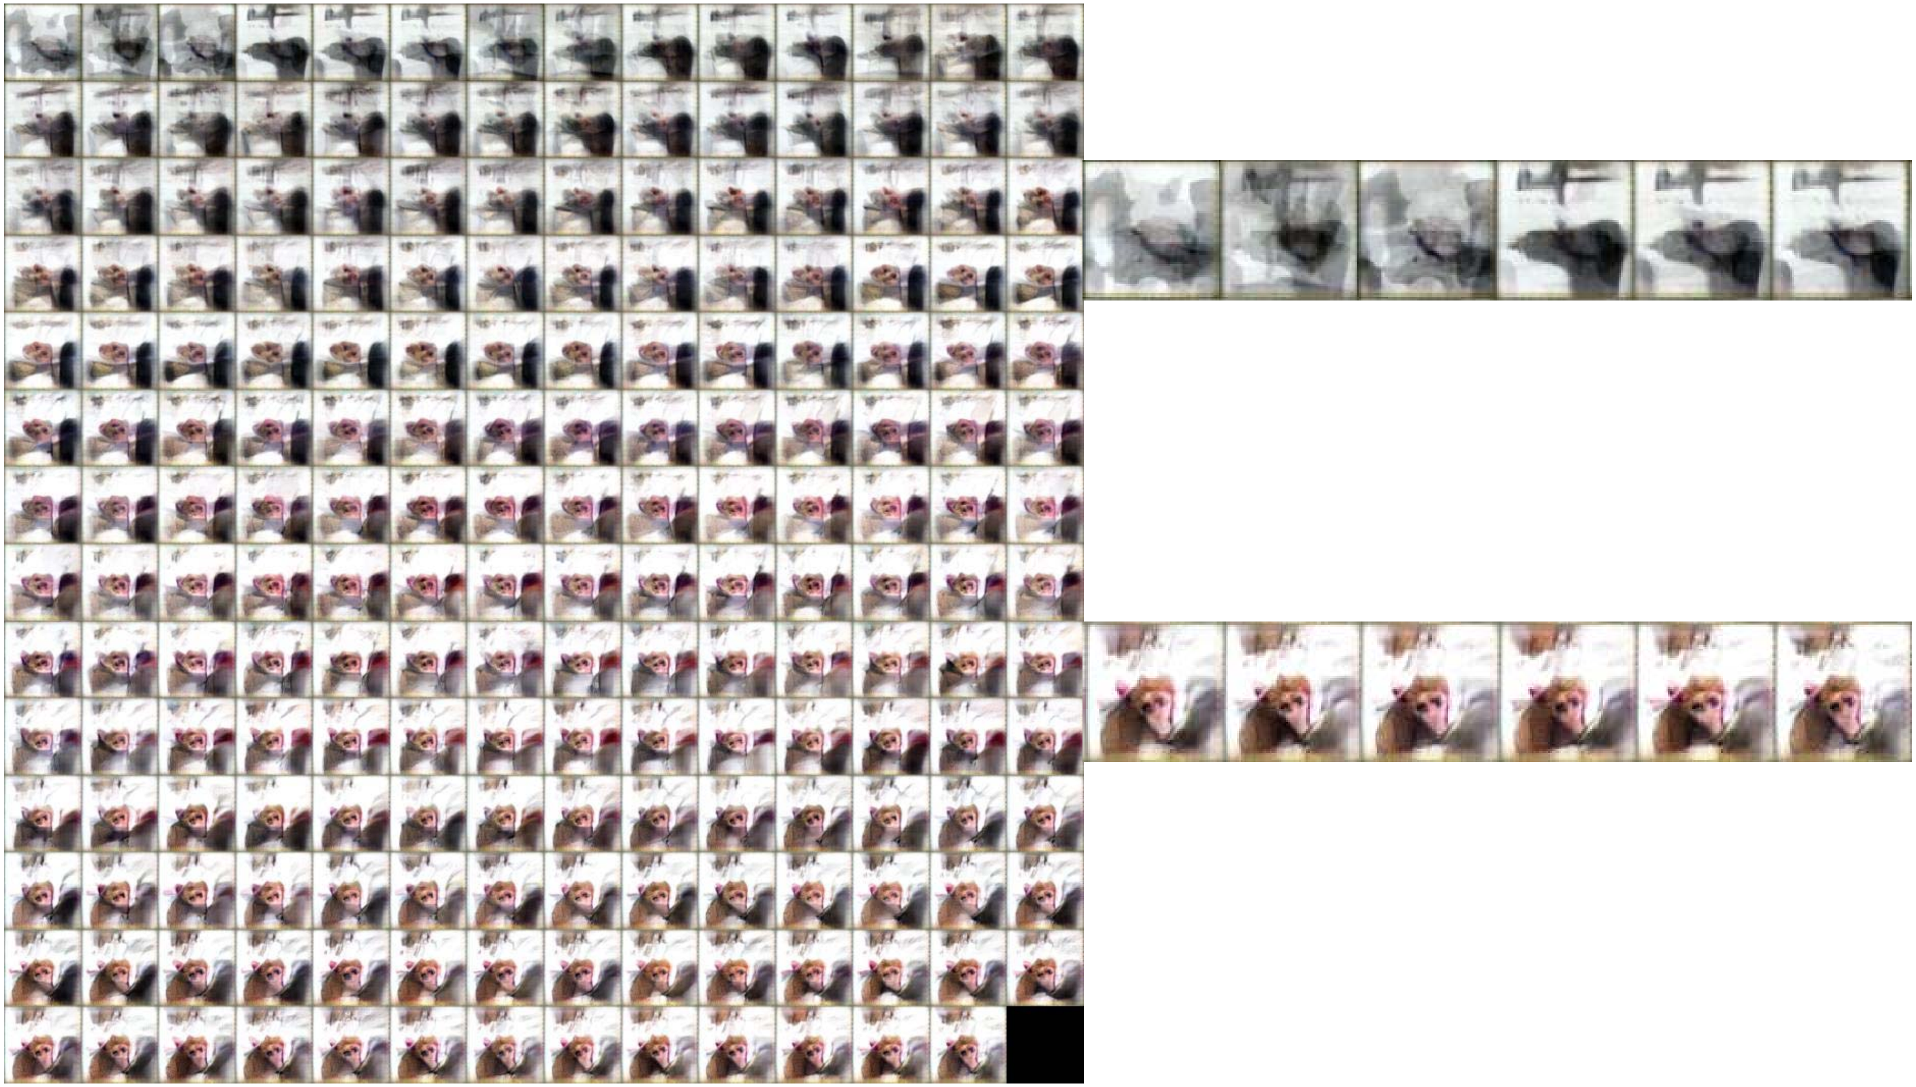

What can we say about these images?

we showed >2500 unique natural images to same cell

Synthesized image

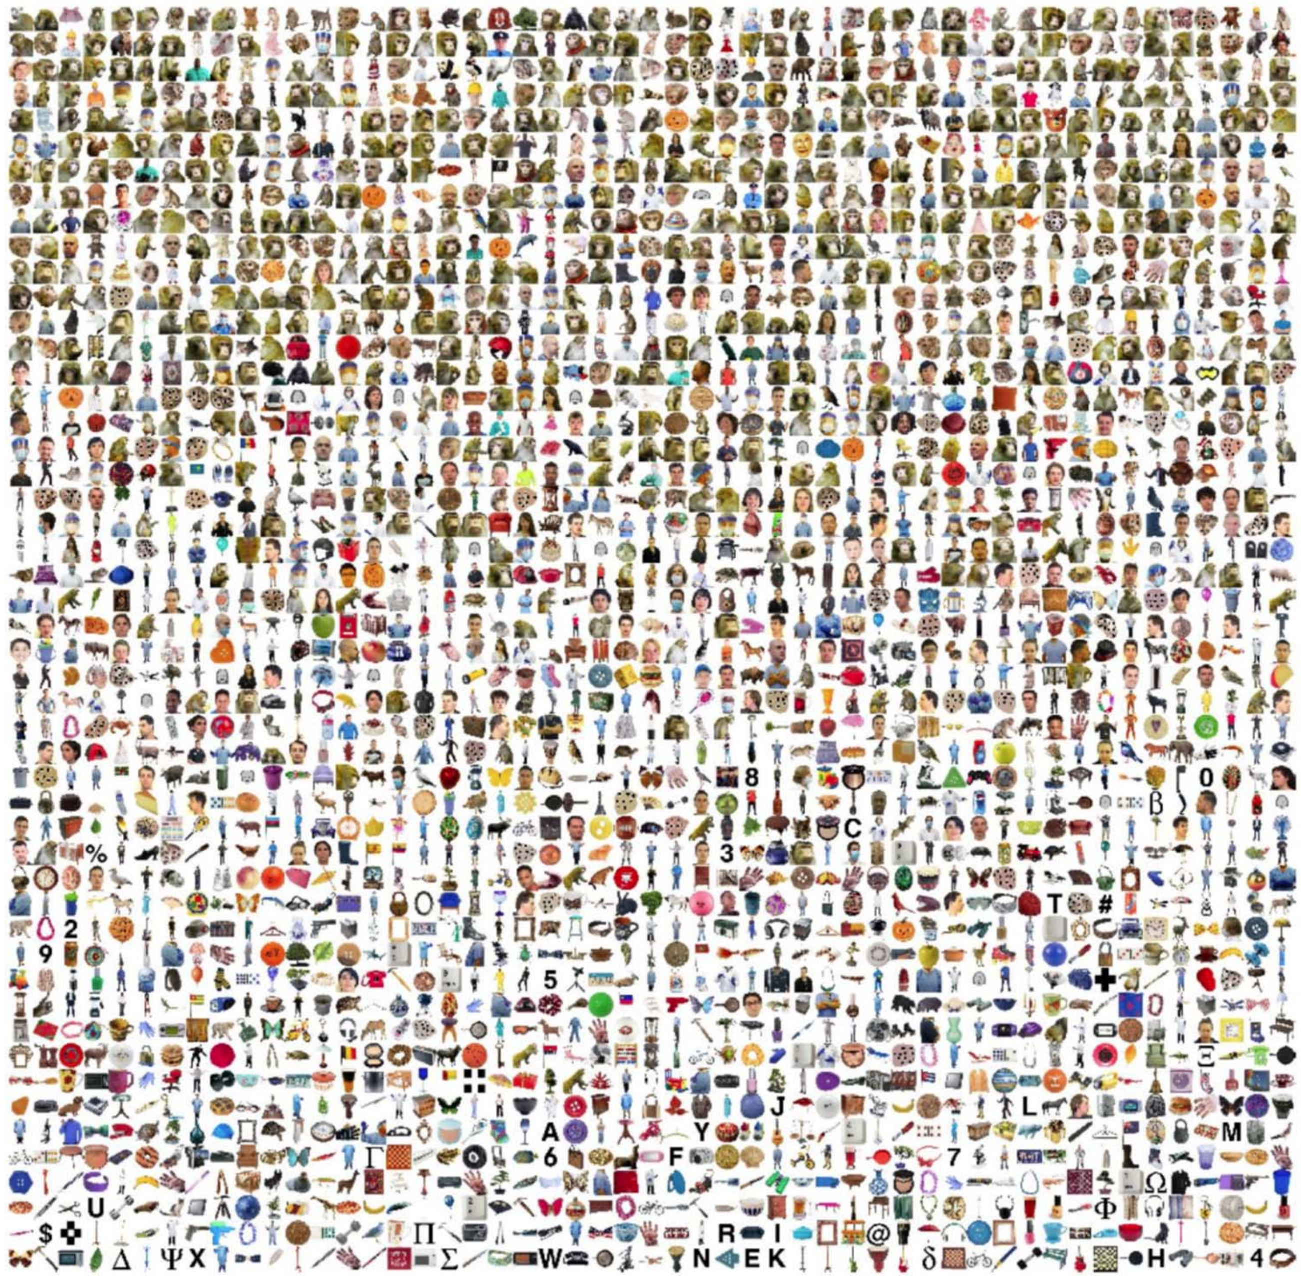

Neurons encode multidimensional features

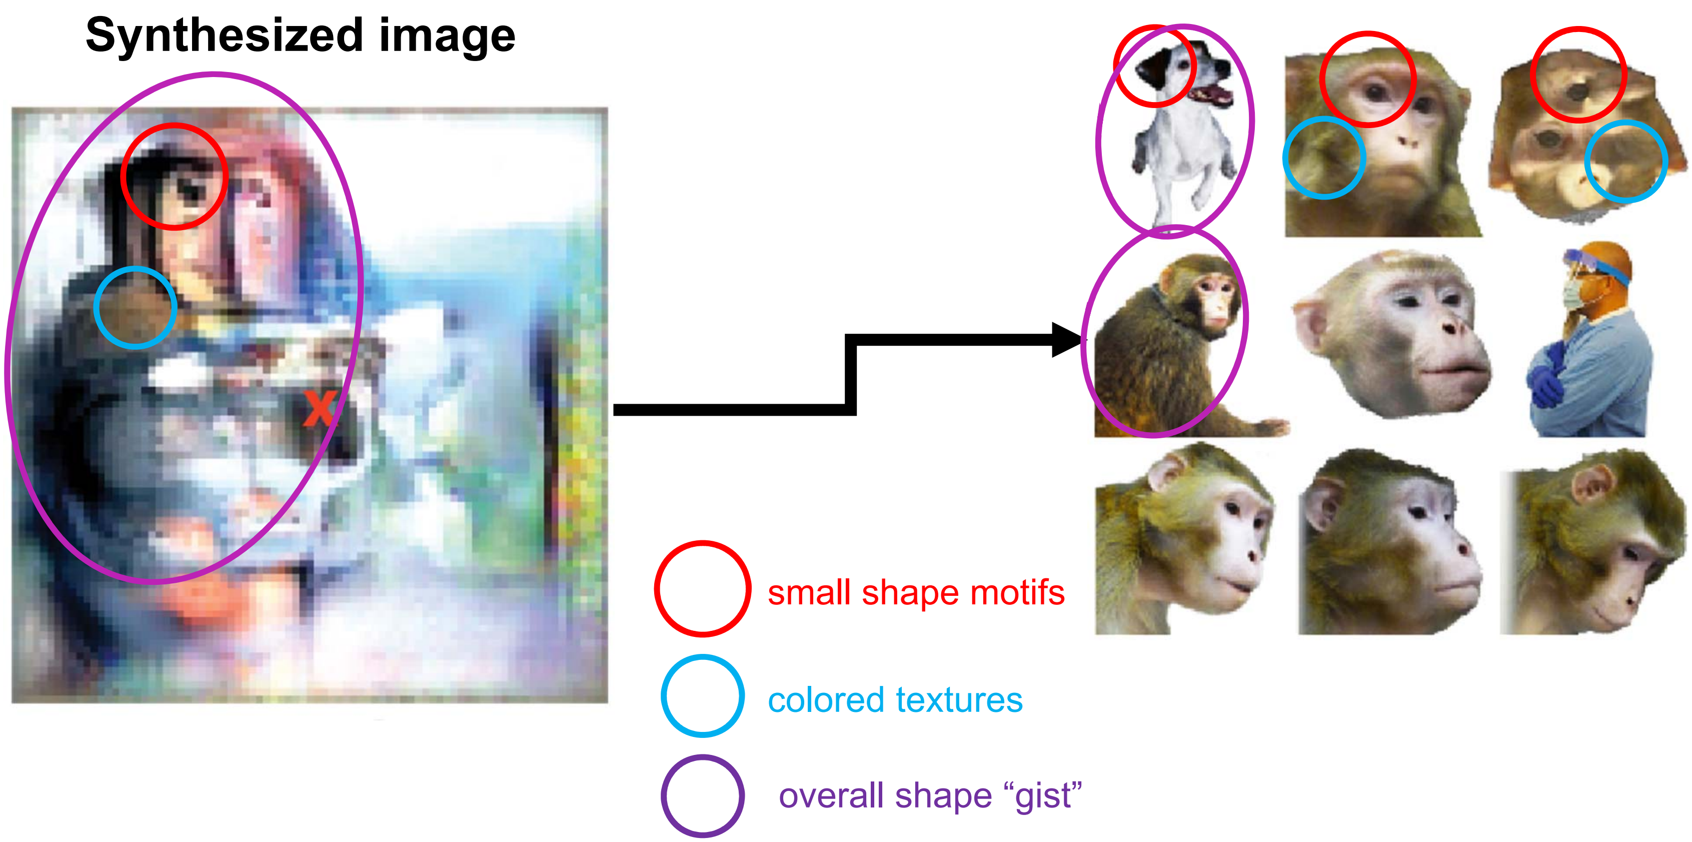

Synthesized image

GAN trained to invert representations in **AlexNet** layer fc6

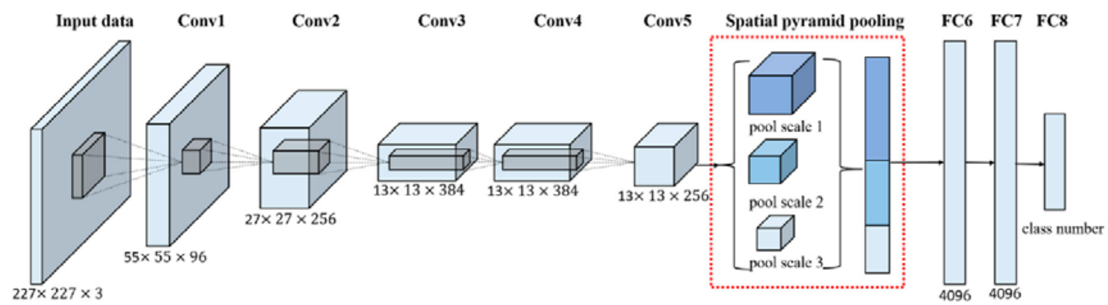

Synthesized image

evolved

GAN trained to invert representations in **AlexNet** layer fc6

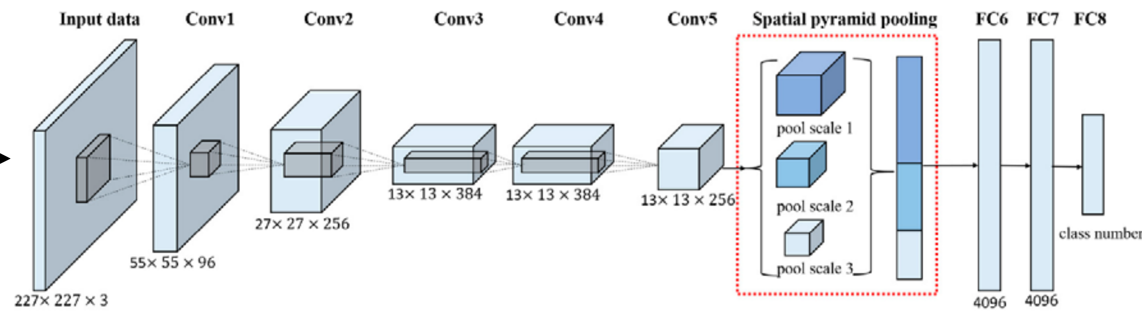

100,000 images from **ImageNet**, a labeled-image database

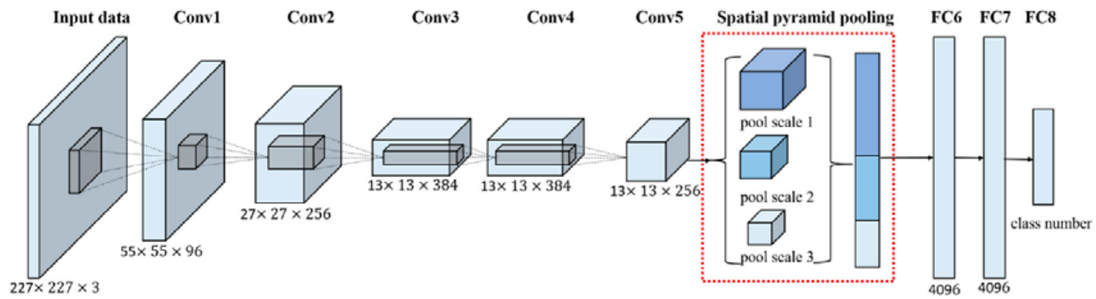

Synthesized image

most similar images
according to AlexNet fc6

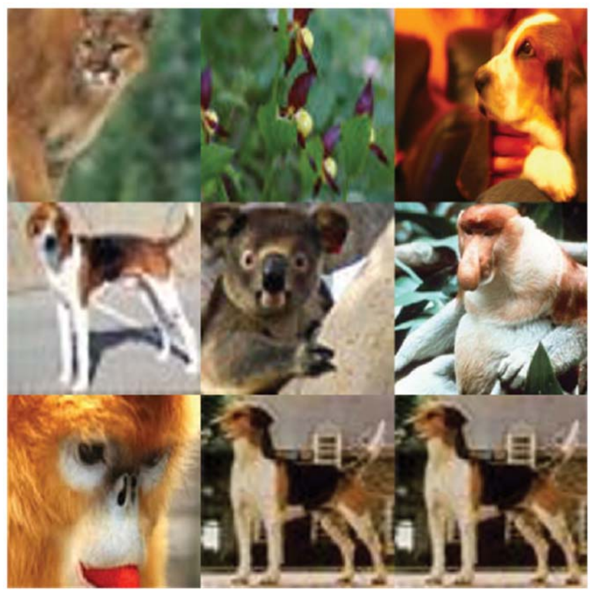

because ImageNet pictures
are *labeled*, we can
quantify those words

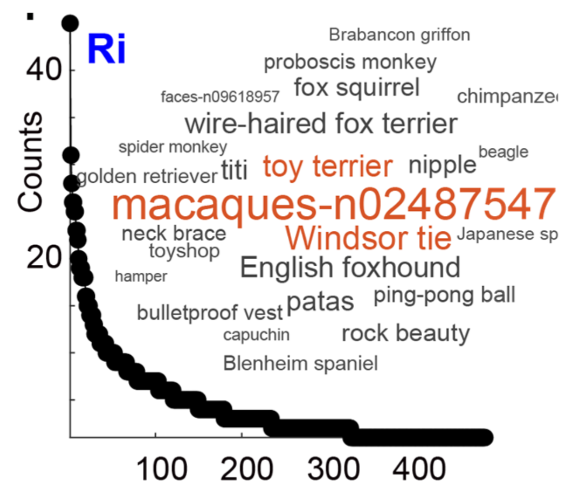

we used the AlexNet interpretations to predict responses to novel images

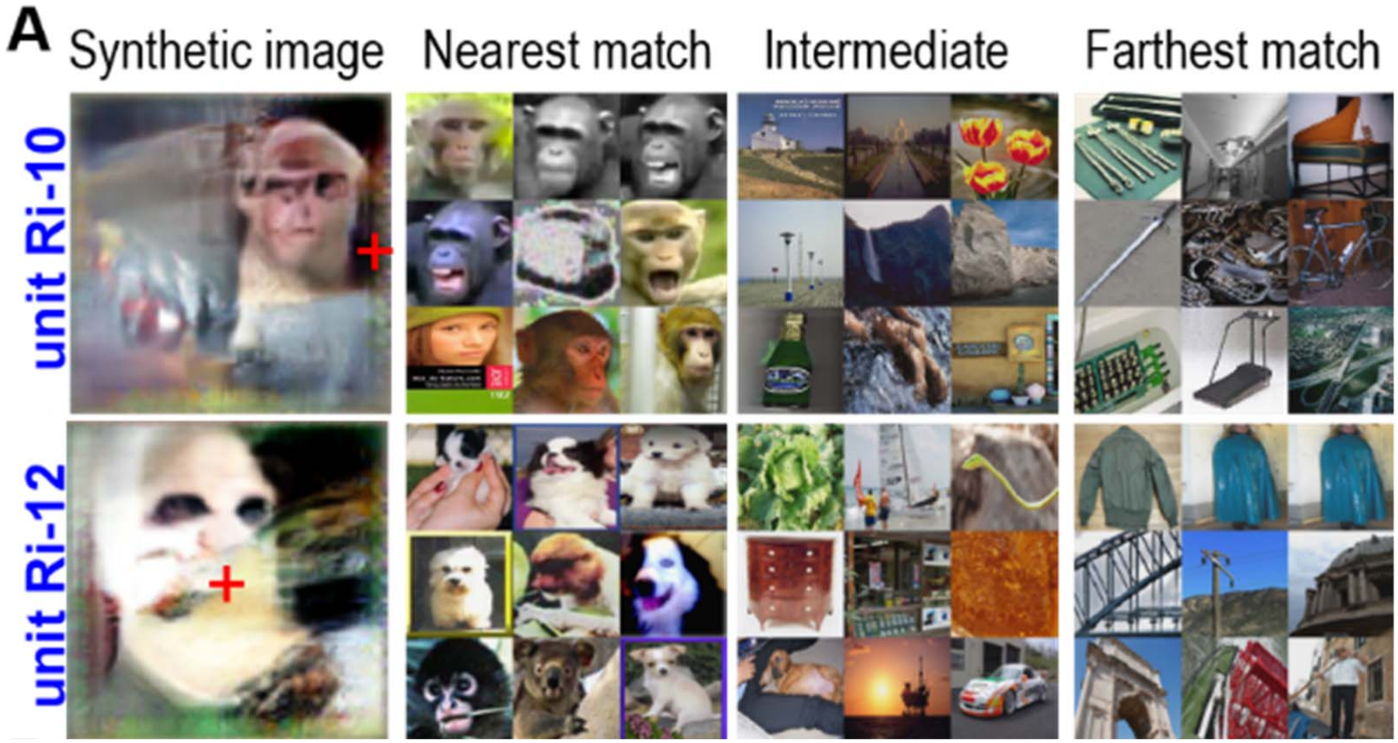

we used the AlexNet interpretations to predict responses to novel images

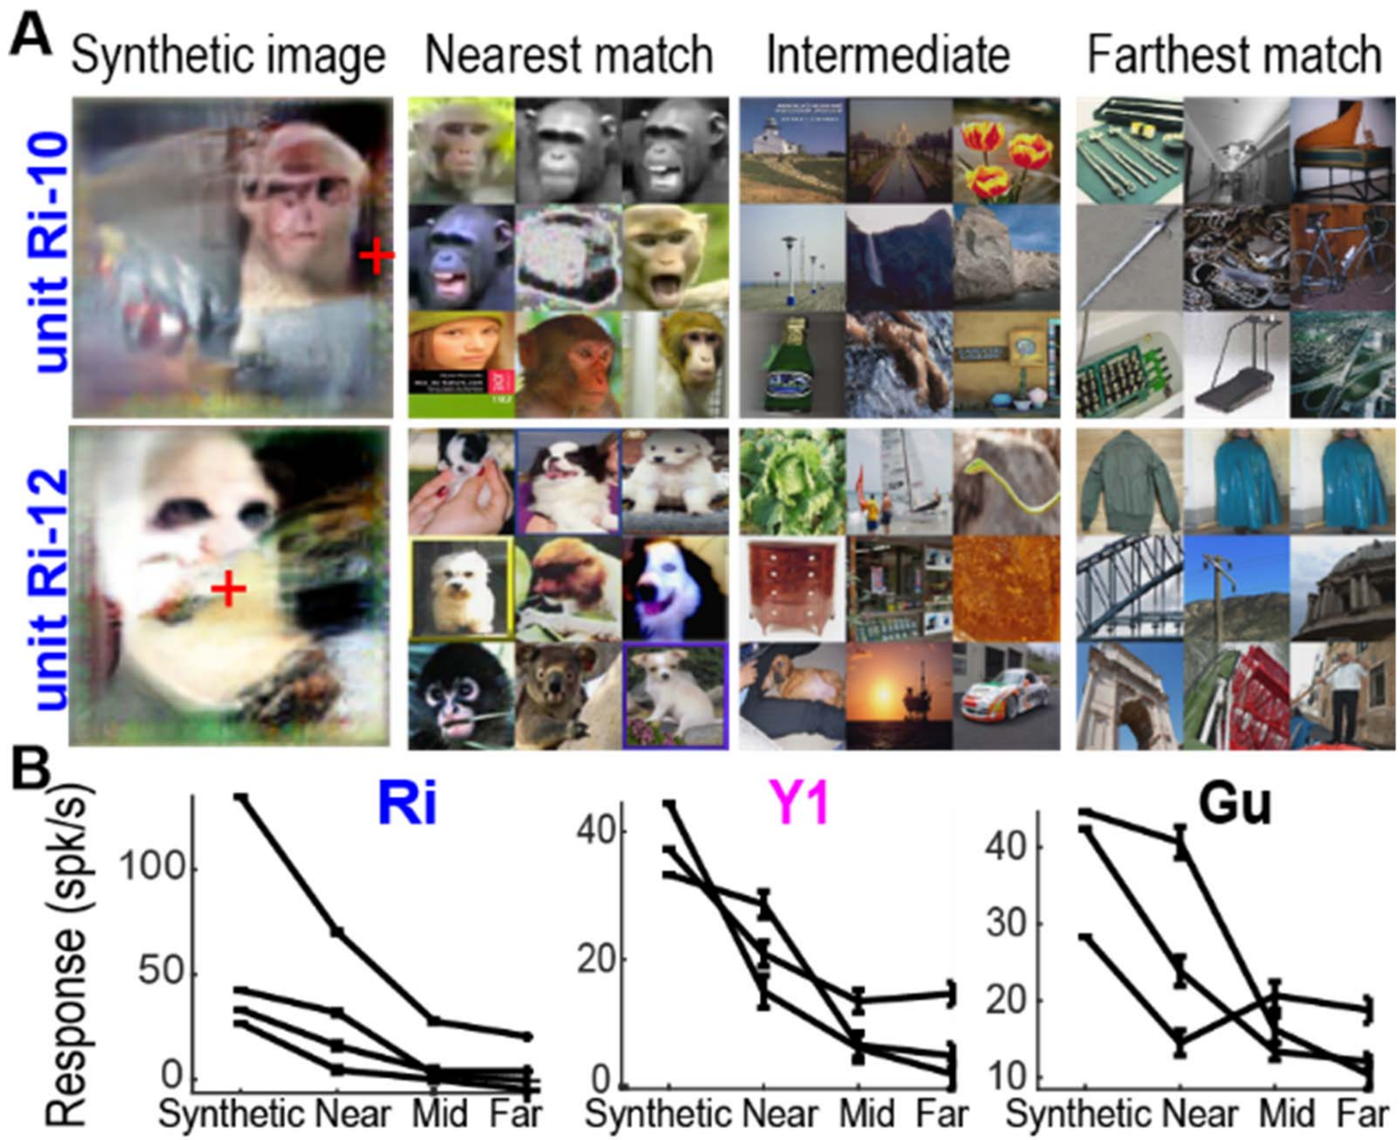

where did these representations come from?

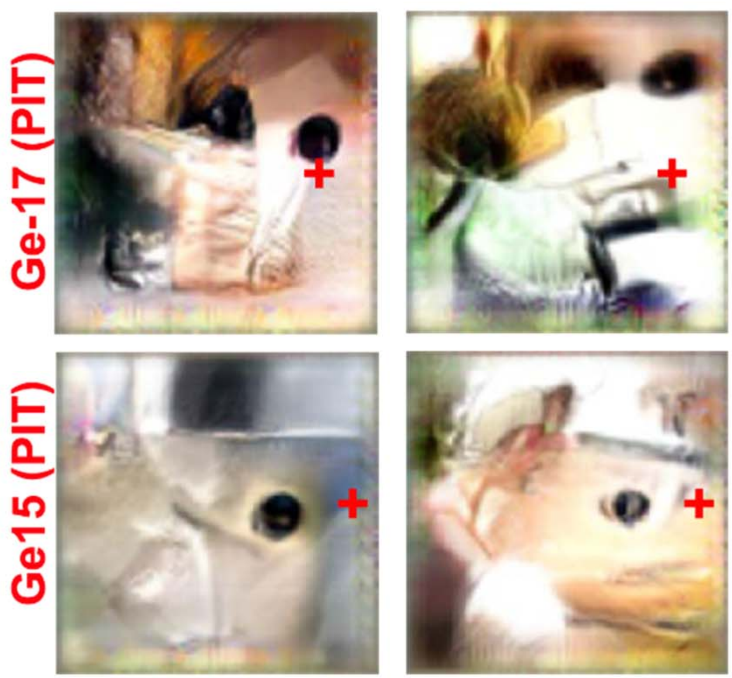

where did these representations come from?

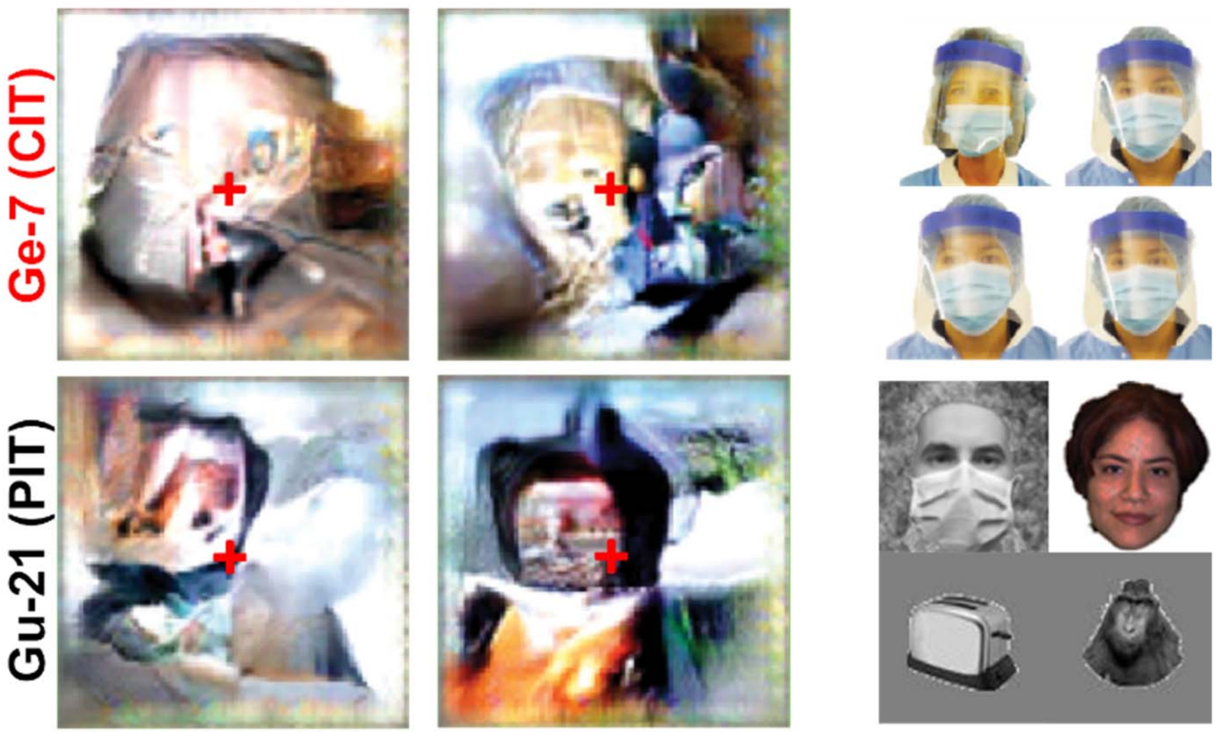

animal care staff that visit monkey Ge daily

evidence that early visual cortex abstracts shapes diagnostic of the immediate environment!

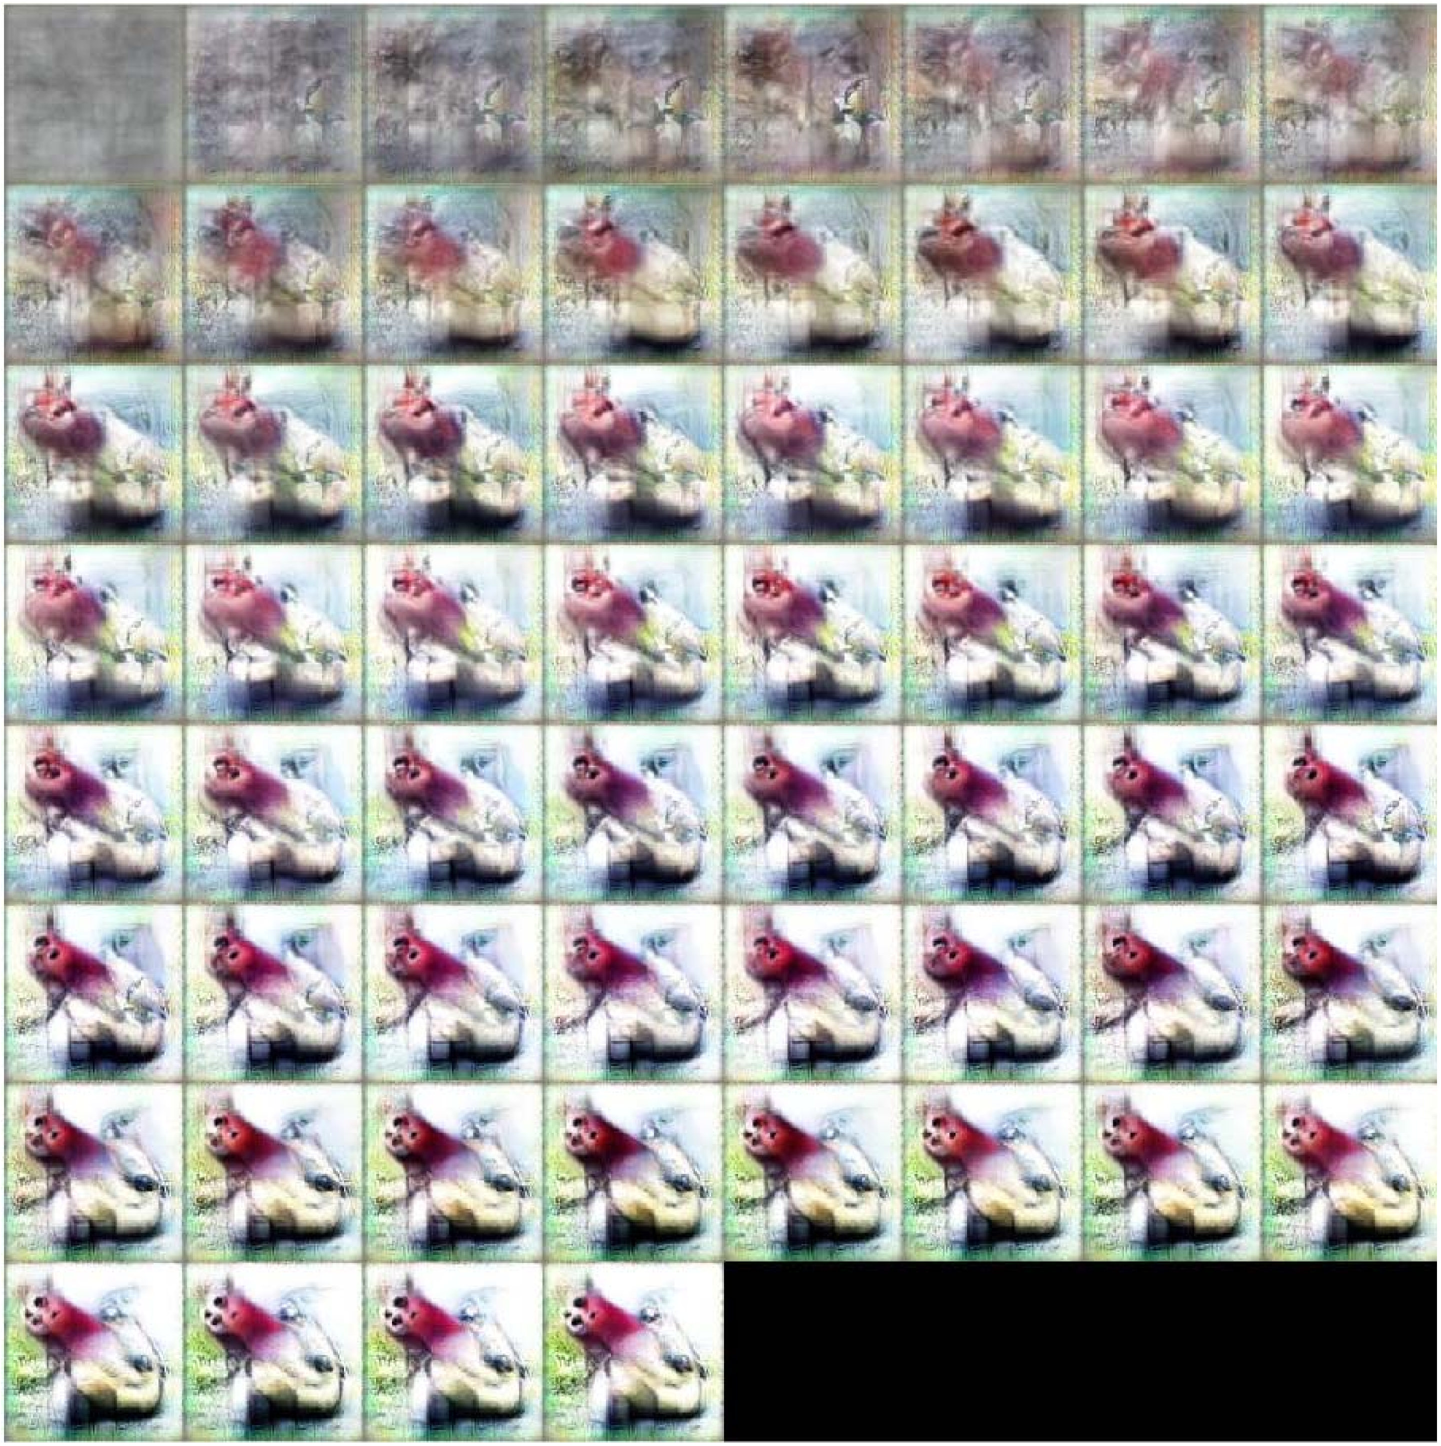

some shapes are undecipherable – this is likely a feature, not a bug

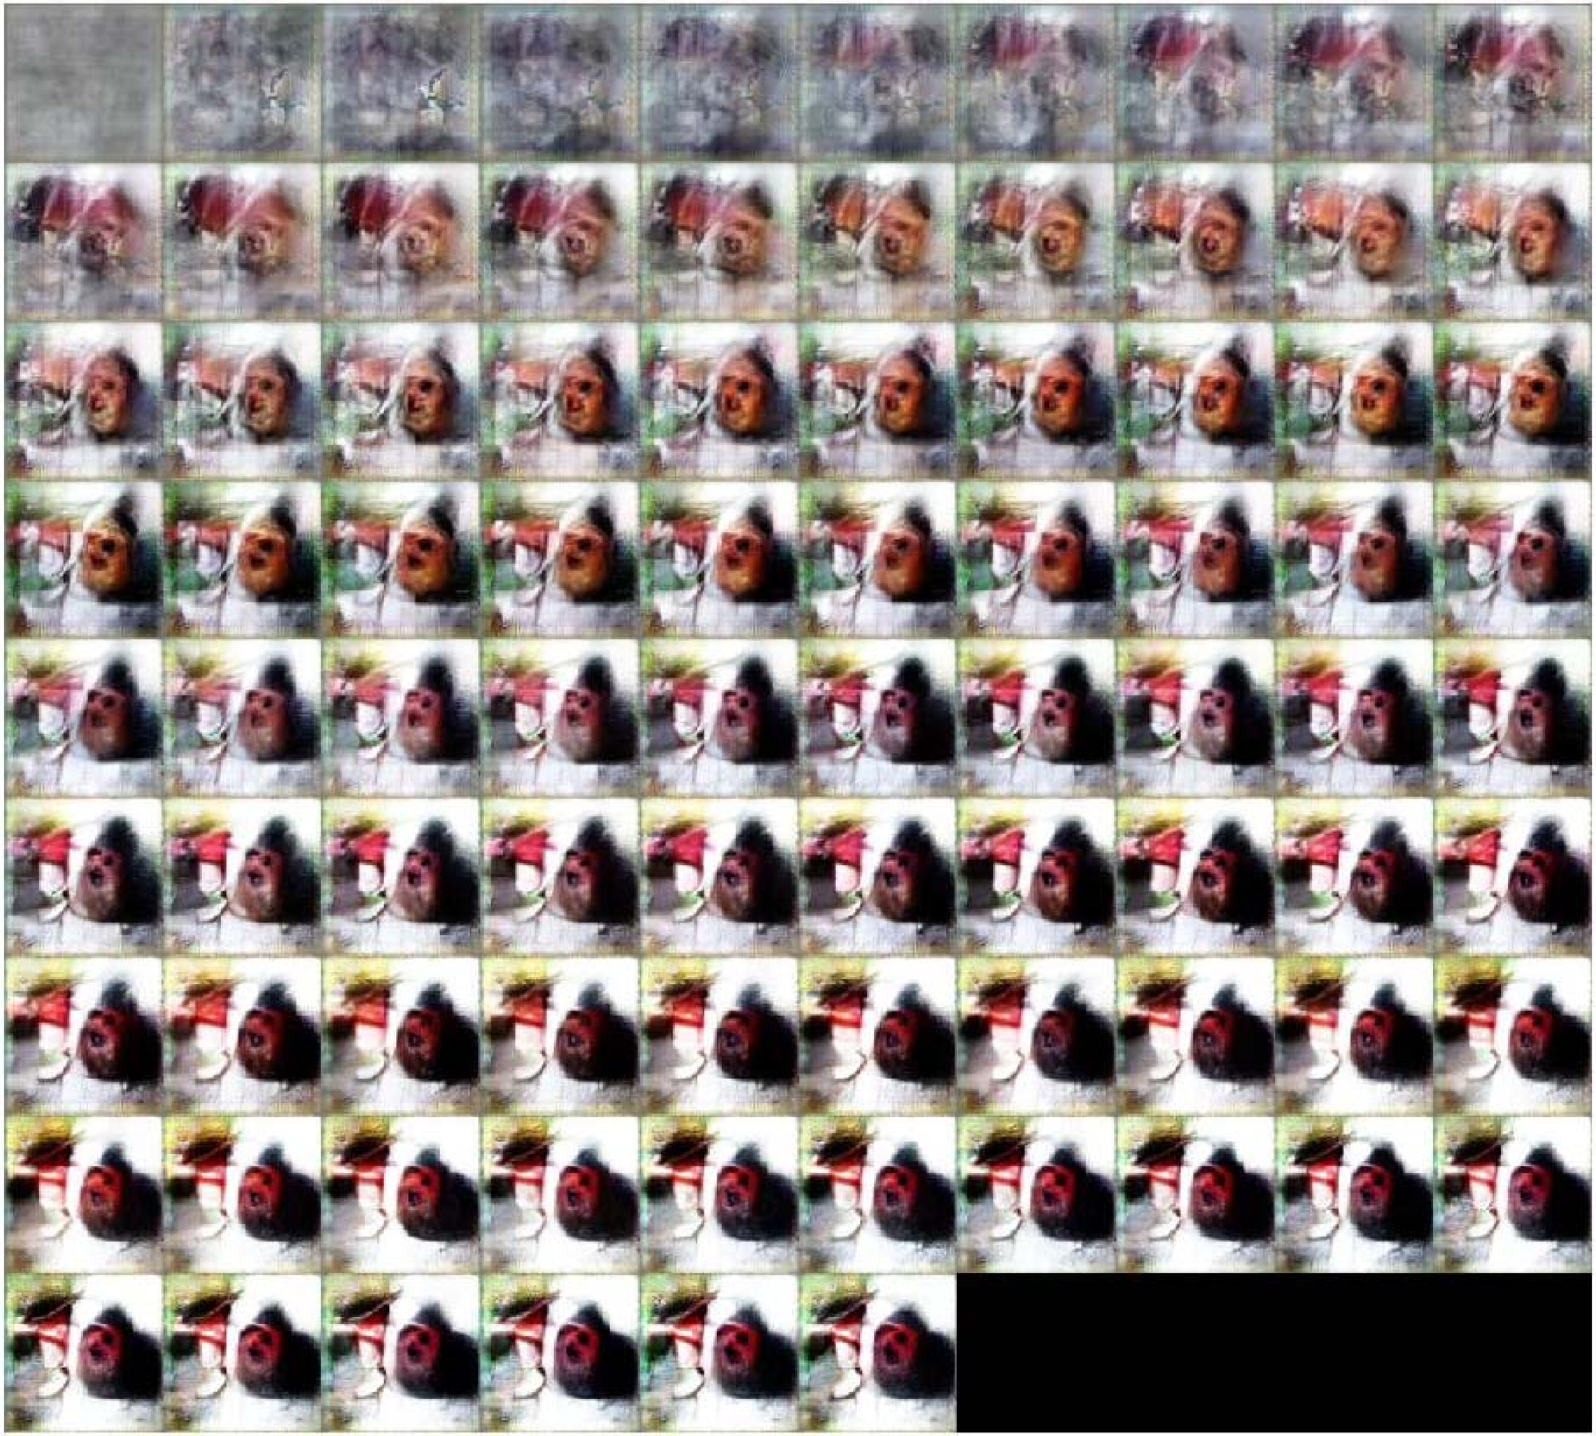

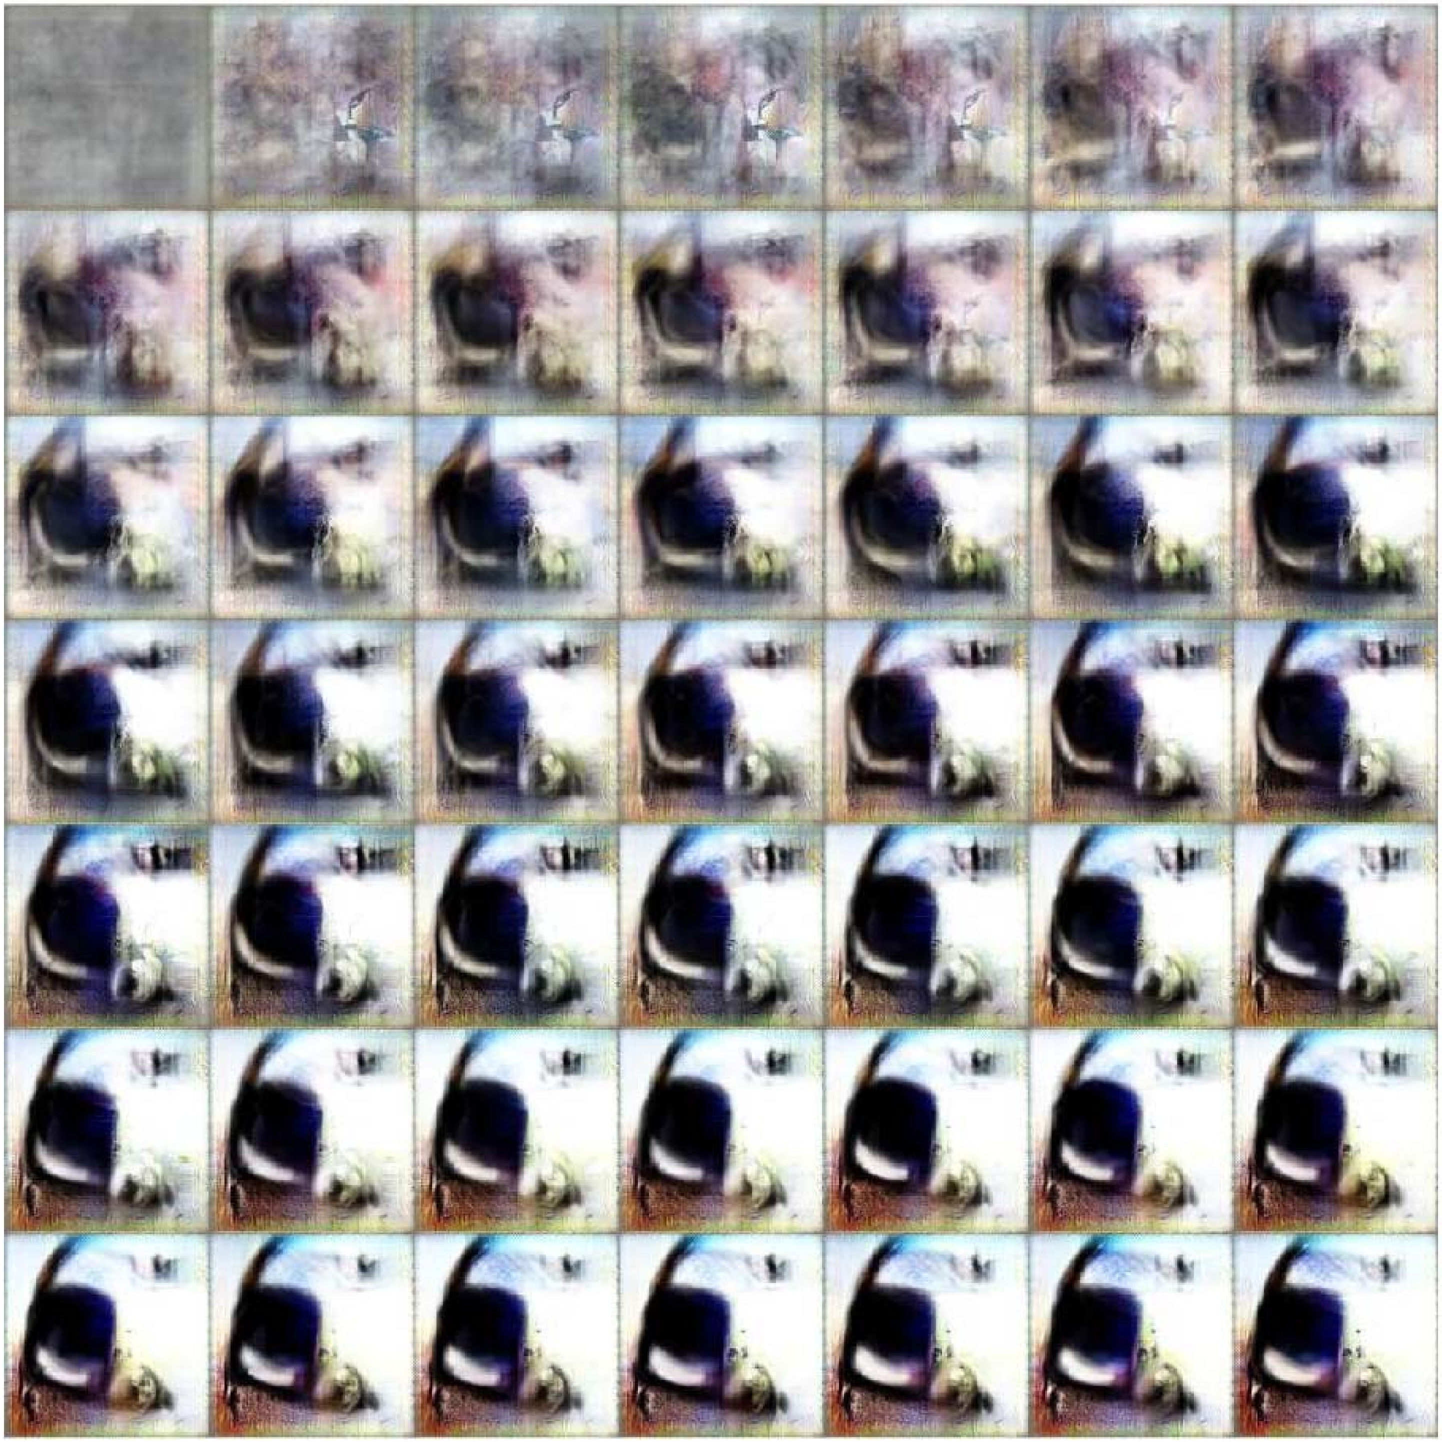

comparing with the Tanaka Alphabet

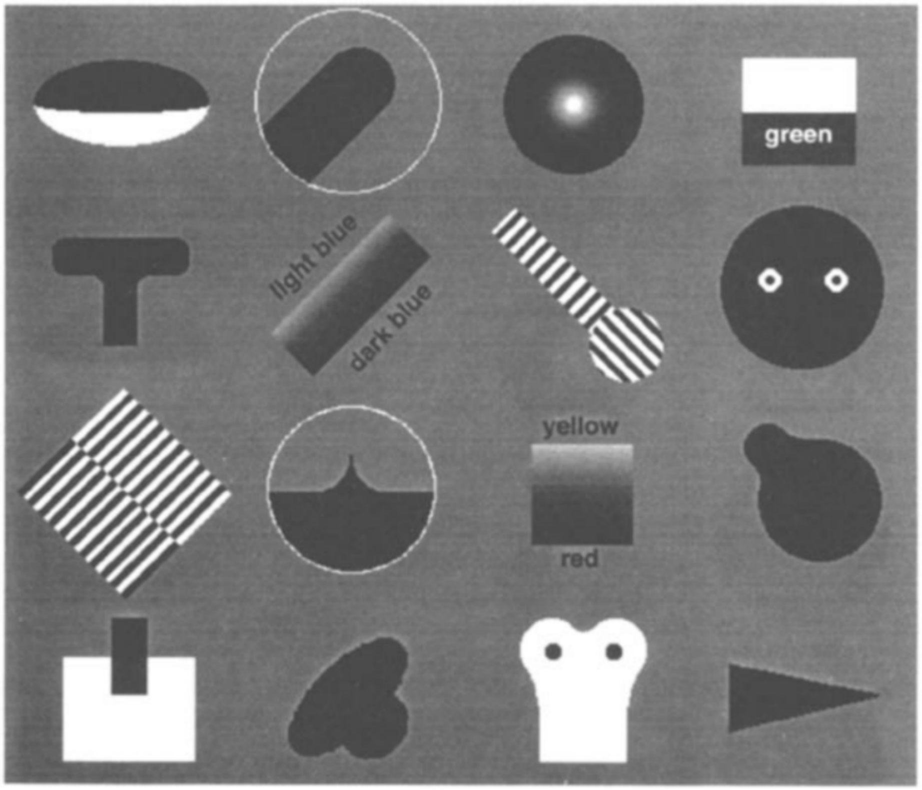

Figure 3 **Sixteen examples of the critical features of cells in TE. They are moderately complex.**

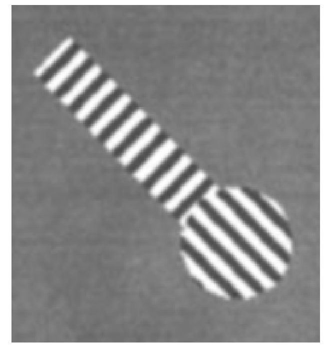

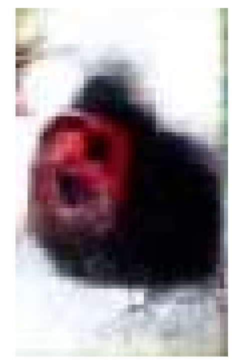

Uncovering the visual "alphabet"

Leslie G. Ungerleider *, Andrew H. Bell

Vision Research 51 (2011) 782–799

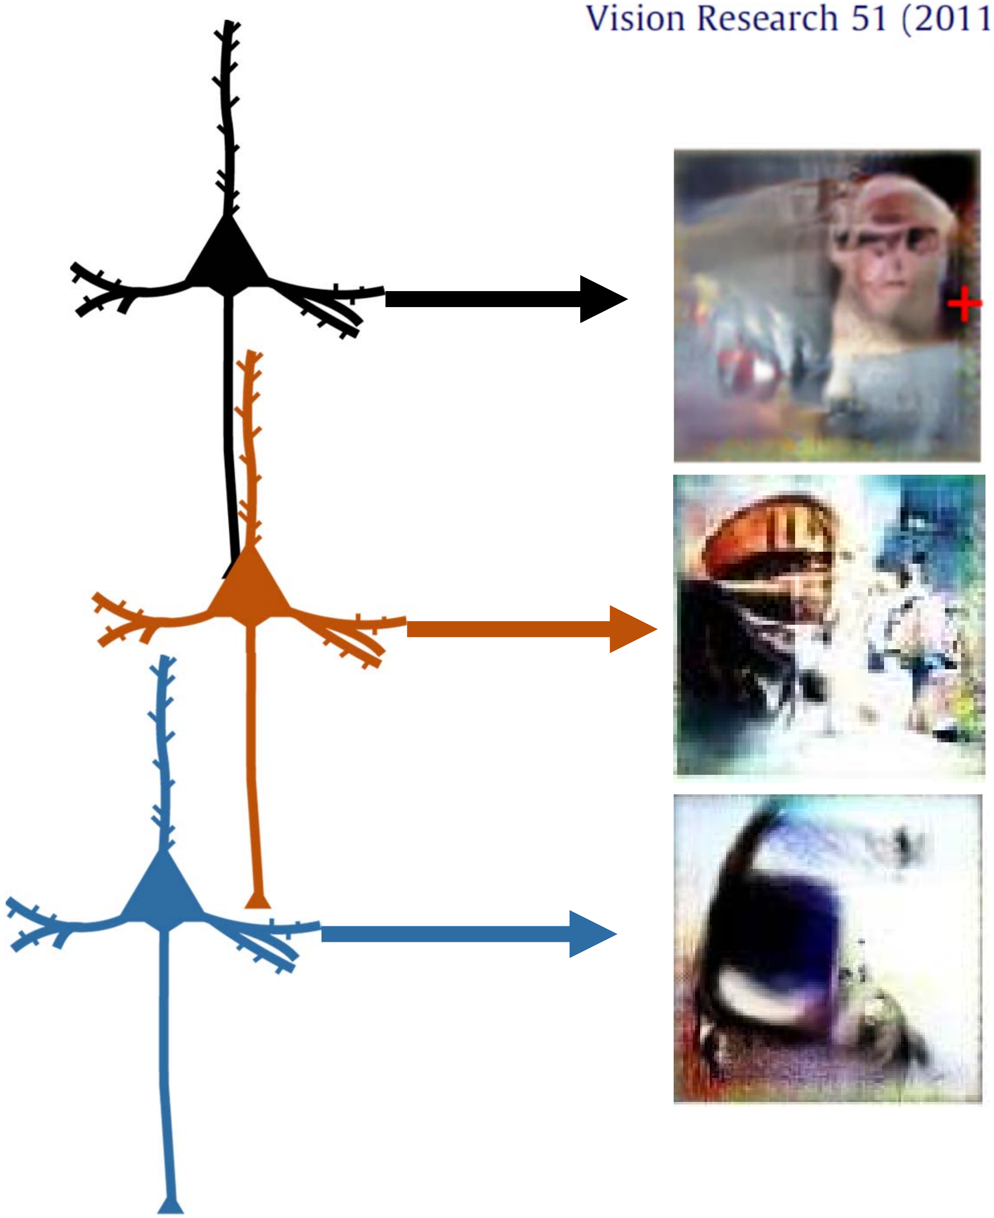

The path ahead:

1) Identify a critical density of representations

2) Explain how these representations relate to what monkeys care about (behavioral tasks)

3) Are these representations parameterized?

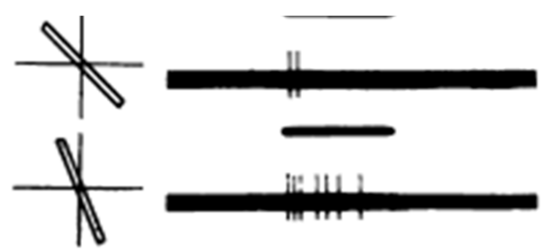

a review of the findings

We used generative networks and a genetic algorithm to decode neurons' representations in primate visual cortex

these resulting images are closer to the intrinsic optimal representation than most natural images we can find

Images are not photorealistic, suggesting they are highly abstracted features (but we have to try other generative networks)

Neurons appear to be particularly concerned with representing objects in the monkeys' immediate environment – monkey features and humans in protective attire

thank you!

Thanks to

Collaborators on phase 1 experiments

Will Xiao

Peter Schade

Till Hartmann

Gabriel Kreiman

Margaret Livingstone

Ponce lab members and collaborators

Mary Burkemper

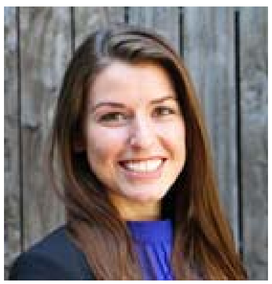

Olivia Bockler

Jeevun Kansupada

Katie Mueller

Binxu Wang

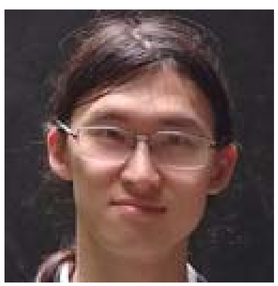

Heide Schoknecht

Center for Brains, Minds and Machines (you!)