16-899A
The visual world as seen by neurons & machines

From a cognitive neuroscience perspective...

Quick intro again

Website: http://graphics.cs.cmu.edu/courses/16-899A/2014_spring/

People:
Abhinav Gupta: abhinavg@cs.cmu.edu
Elissa Aminoff: elissa@cnbc.cmu.edu

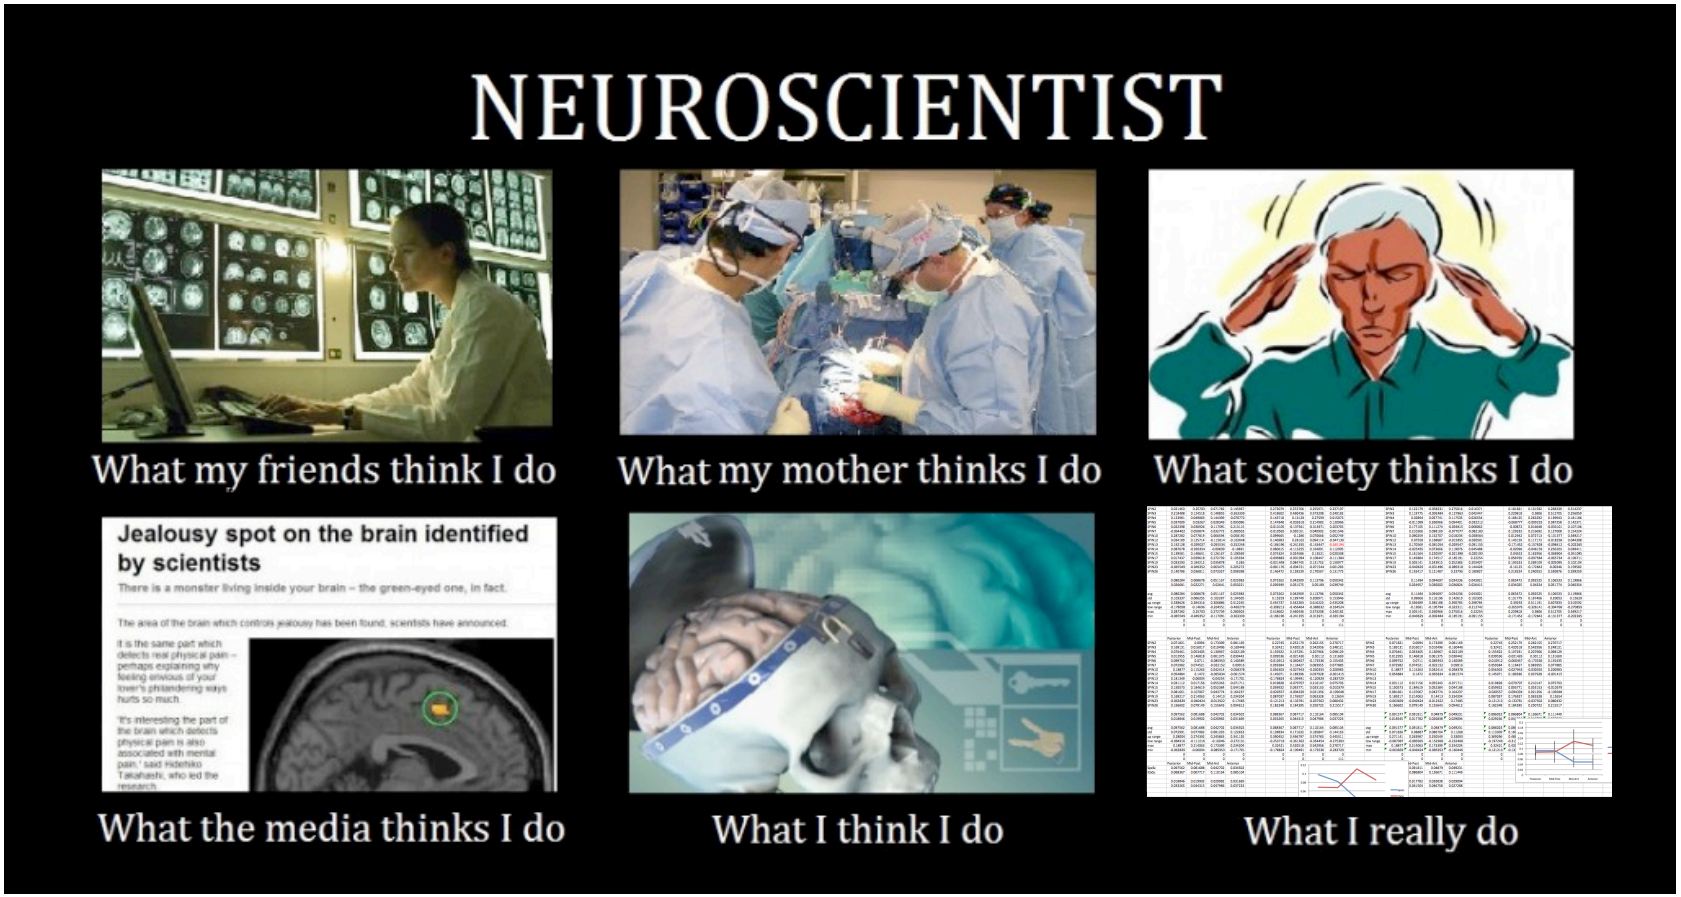

Visual cognition entails…

1. Receiving visual information
2. Filtering for meaningful information
3. Recognizing/interpreting/composing the scene
4. Integrating with context, memory, experience
5. Guiding motor behavior/thought

Black Box

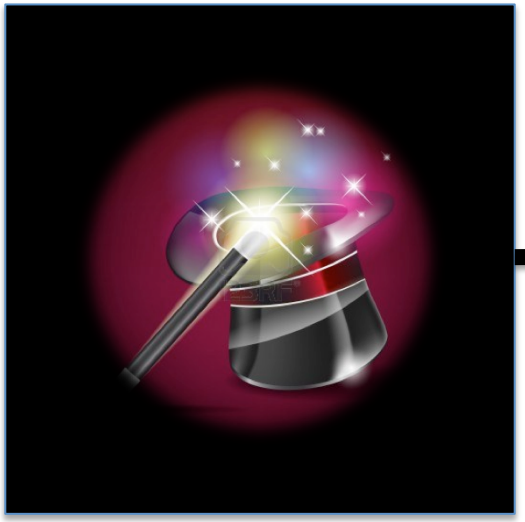

Input → Receiving visual information

→ **Output** Guiding motor behavior/thought

↑ Filtering for meaningful information

Recognizing/interpreting/composing the scene

Integrating with context, memory, experience

Black Box

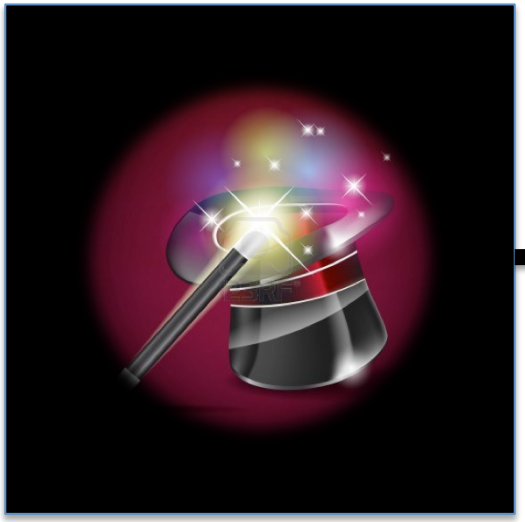

Input

Receiving visual information

Output

Guiding motor behavior/thought

Use rules/heuristics for fast processing

Marr's three levels

1. Computation
2. Representation and algorithm
3. Hardware implementation


STOP

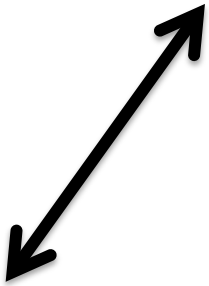

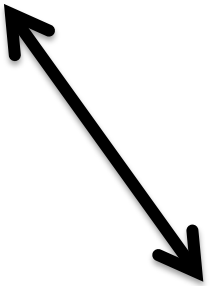

Why bring the brain into the discussion?

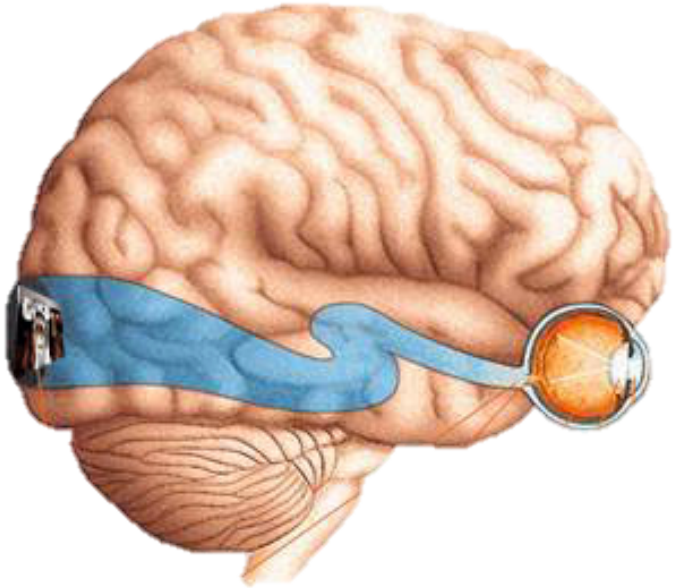

Jealousy spot on the brain identified by scientists

There is a monster living inside your brain – the green-eyed one, in fact.

The area of the brain which controls jealousy has been found, scientists have announced.

It is the same part which detects real physical pain – perhaps explaining why feeling envious of your lover's philandering ways hurts so much.

'It's interesting the part of the brain which detects physical pain is also associated with mental pain,' said Hidehiko Takahashi, who led the research.

Coke vs. Pepsi

How does your brain react to soft drinks?


587

Why bring the brain into the discussion?

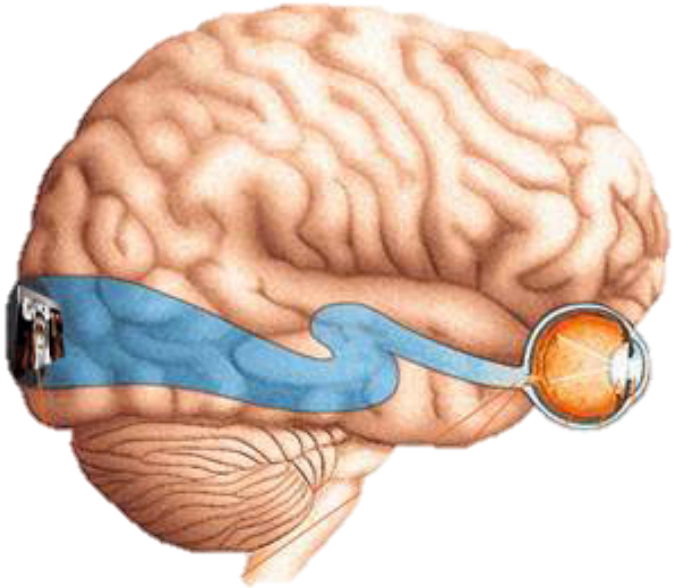

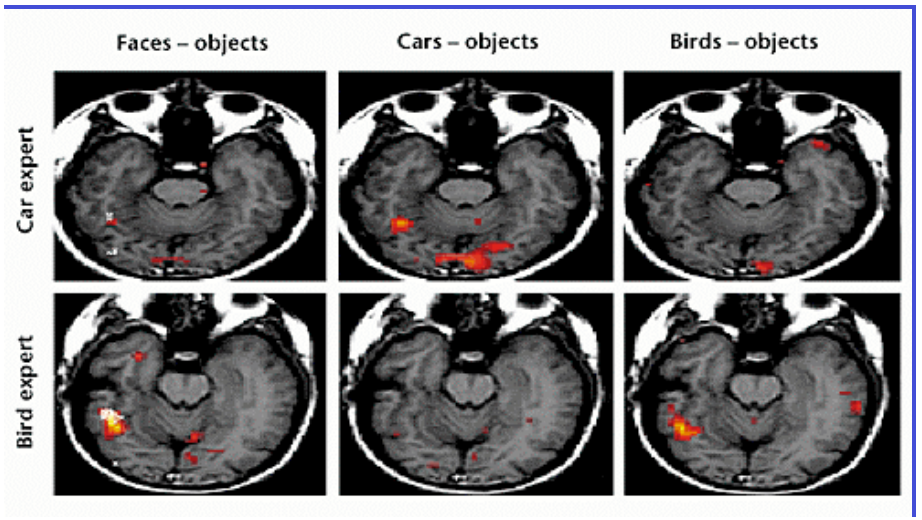

Gauthier et al., 2000

Why bring the brain into the discussion?

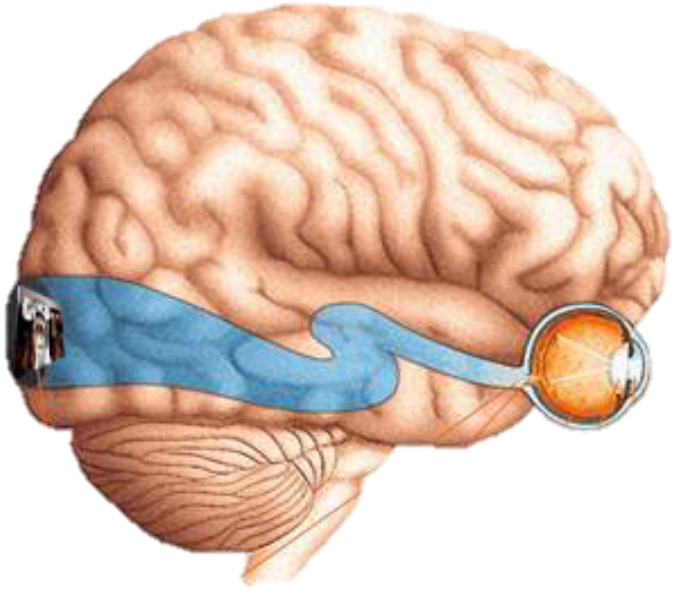

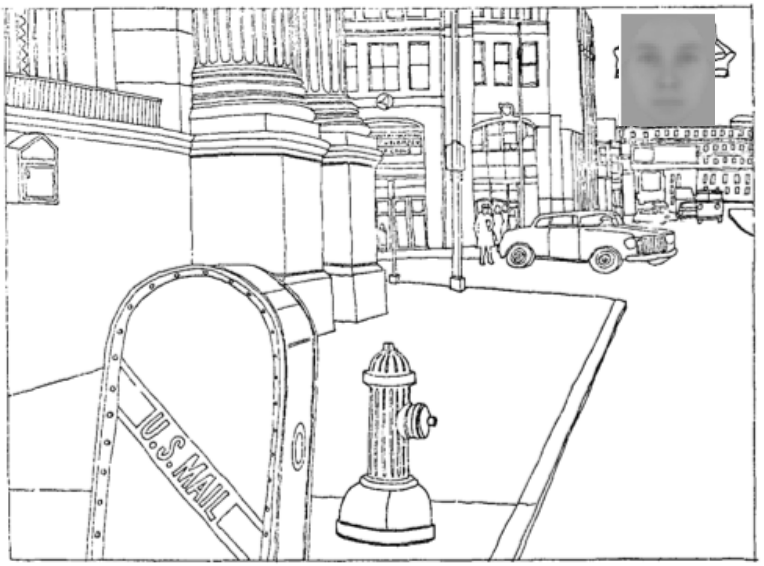

Why bring the brain into the discussion?

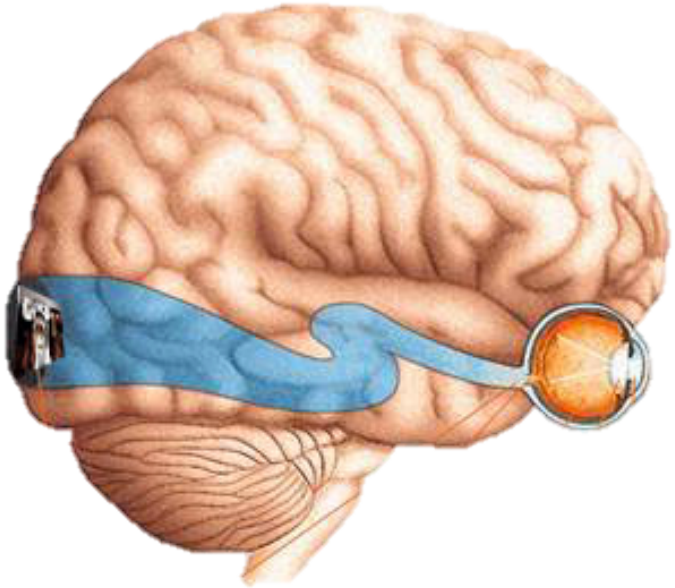

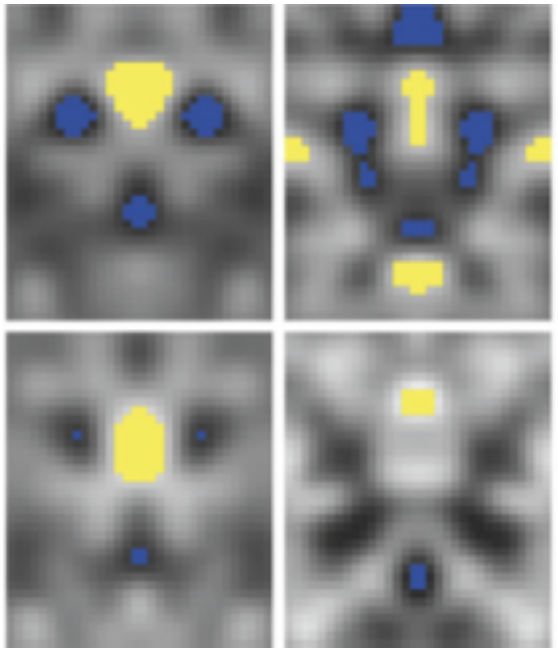

Nestor et al., 2012

Black Box

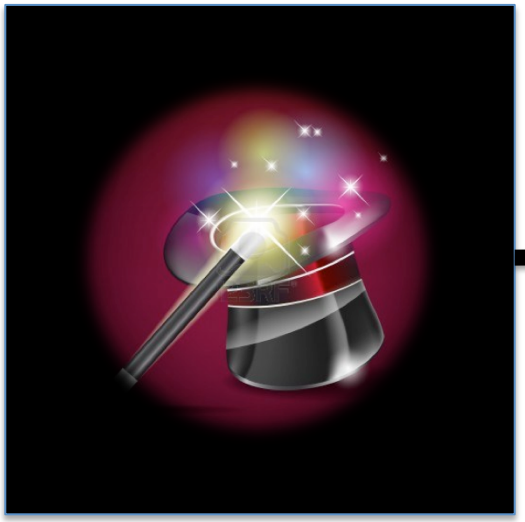

Input → → **Output**

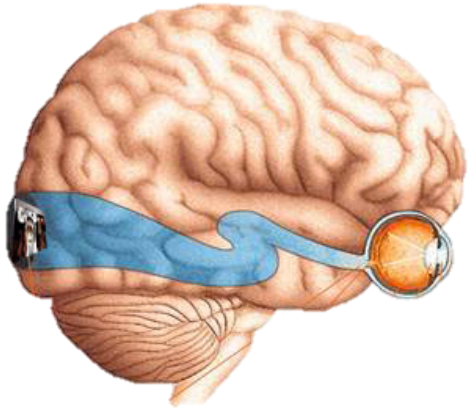

Black Box

Input →  → **Output**

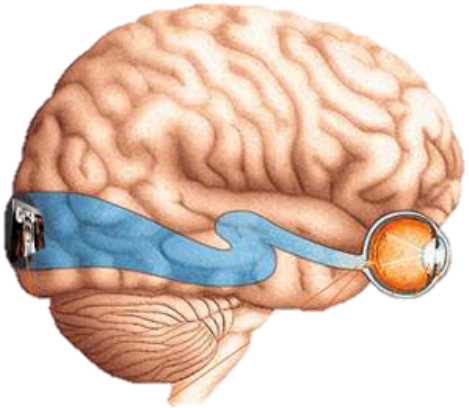

Road Blocks:

- Human bias
- Lack of models
- No means to:
  - Extract statistical regularities
  - Characterize prior experience

Example from our research

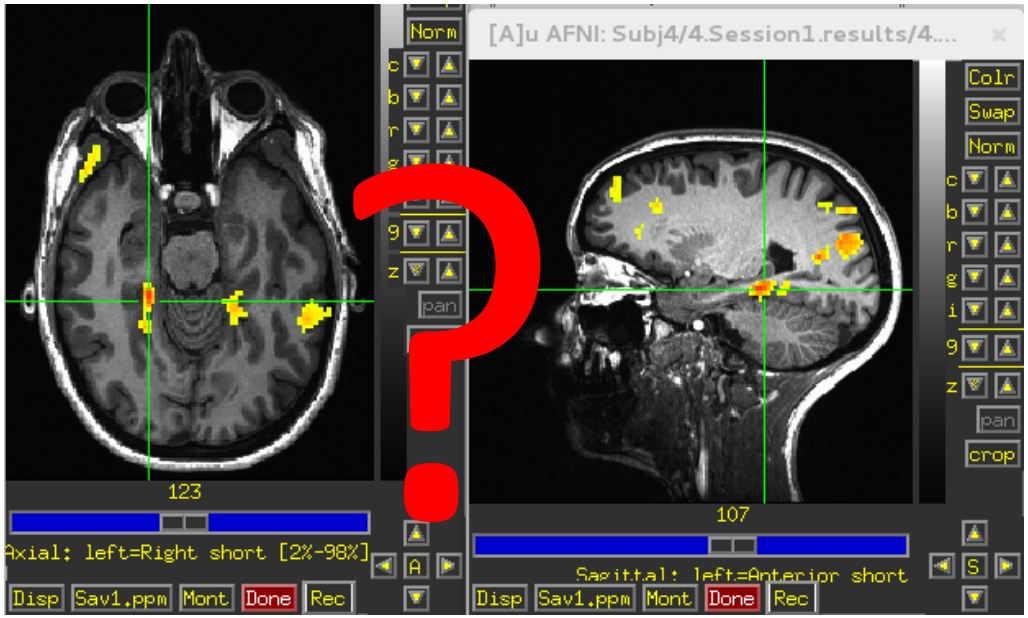

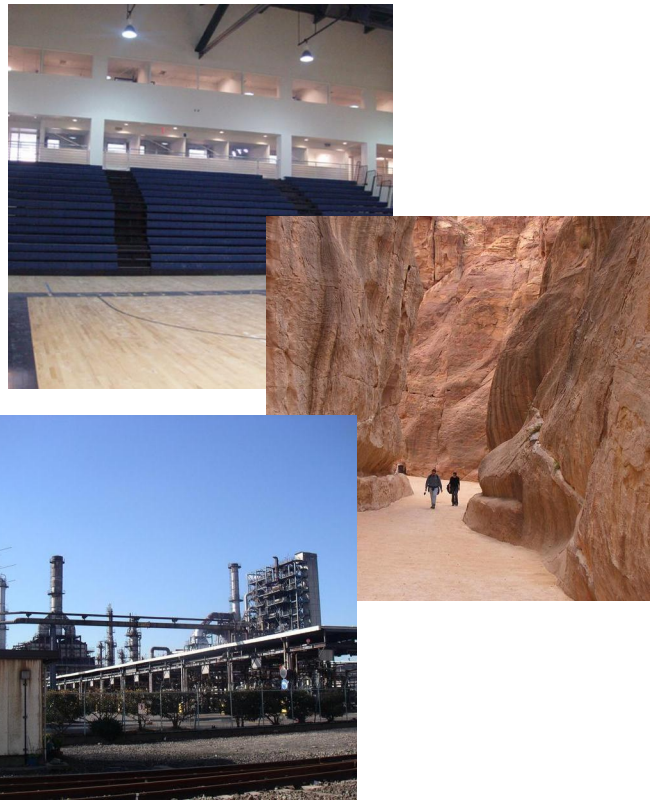

NEIL: Never Ending Image Learner

I Crawl, I See, I Learn.

Analyzed over **400,000** images

Running continuously for over 2.5 months, 350 cpu hours

- 1034 Scene Categories
- 1152 Object Categories
- 87 Attributes
- 1400 Commonsense Relationships

Numbers taken from 10/2013 -ish

Human labels for NEIL's visual attributes

amber
arch_shape
black
blue
brick_texture
brown
chain_texture
check_texture
chubby
clouds
cloudy_weather
cluttered
cold_scenery
cone_shape
crooked
crowded
cube_shape
cyan
cylinder_shape
diamond_shape
tartan
dots
dry
farming
feather_texture
fire
foggy_weather
foliage
furry_texture
gold
grass_texture
gray
green
horizontal_cylinder
horizontal_lines
ice_texture
indoor
magenta
modern
mountainous
mysterious
brass
sky
narrow
natural
open_area
orange
outdoor
pink
plain
purple
queue
railing
rainy_weather
rectangular_shape
red
round_shape
rugged_scene
running_water
rural
shingles
shiny_texture
shrubbery
silver
skin_texture
smoke
snowy_weather
speckled
square_shape
steep
still_water
stripe_texture
sunny_weather
symmetrical
turquoise
urban
vertical
vertical_cylinder
vertical_lines
violet
warm_scenery
wave
white
wide
wiry
yellow
plaid




Attribute Scores

.02 .84
.47 .5
.06 X .01 X = X .43 X .92
.7 .34 .11 .20
.32 .2 .00 .13
… … … …

?

Voxel Activation

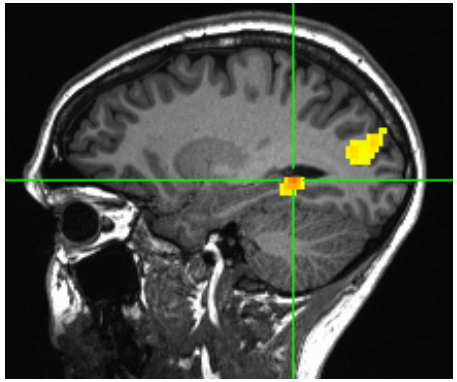

Can NEIL define features in cortical representation?

Compare similarity in attribute space with fMRI voxel space

NEIL: Attribute similarity

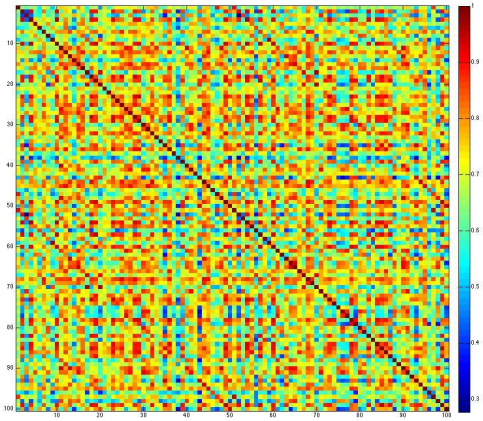

Scene Categories

fMRI: BOLD similarity

(area: PPA localized)

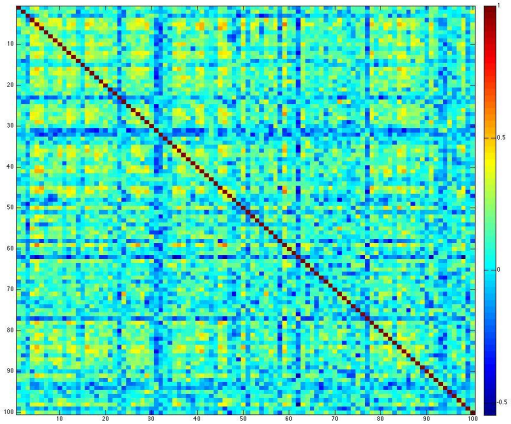

Scene Categories

Canonical Correlation Analysis: Find the linear combinations provide maximum correlation

Can NEIL define features in cortical representation?

Result: Categories that carry the strongest correlation with the linear combinations were similar in both attribute and fMRI datasets!

CCA Analysis

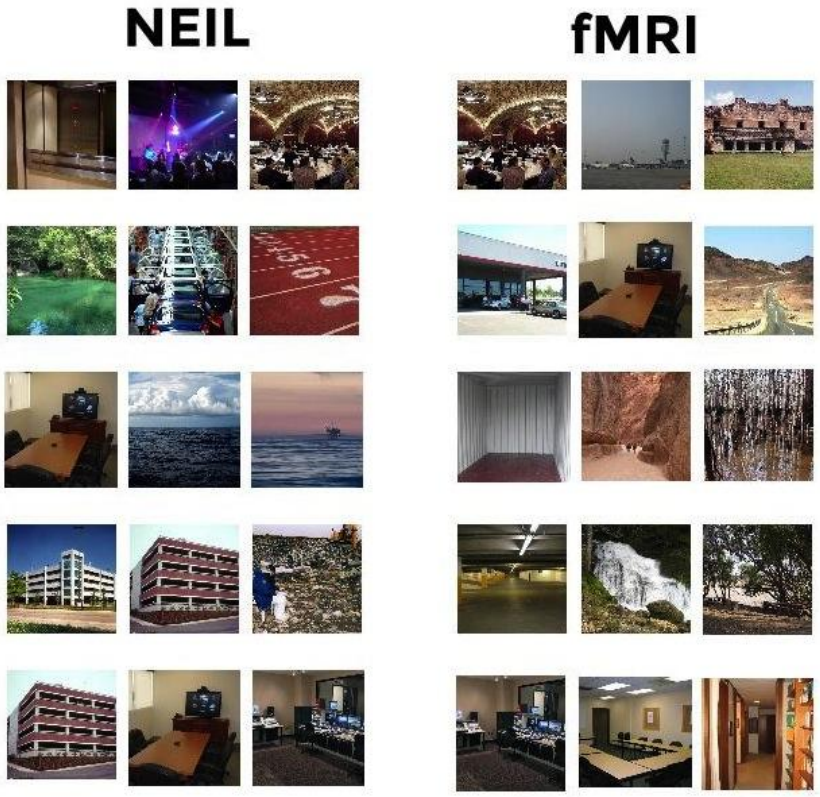


NEIL

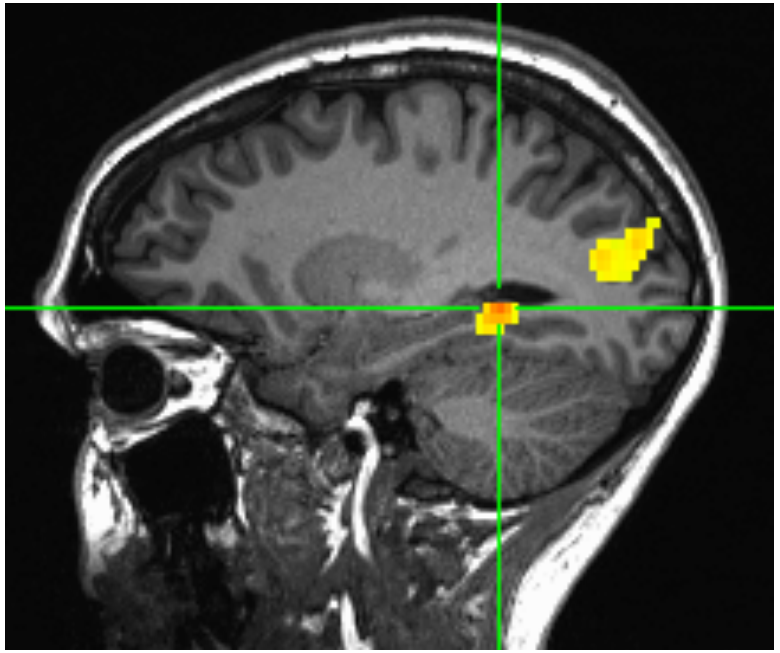

My goals for the course

- Learn CV perspective
- Learn the CV challenges
- Teach you some cog neuro of visual perception
- Peak your interest and confidence in cog neuro
- Get you comfortable with reading papers

Point of the course

- We all care about studying the visual world
- Some problems you are more equipped for solving; some we are.
- Stop reinventing the wheel
- It's a wide world out there – focus
- Collaborations

How I went about picking the papers

Questions? Comments?