Brains, Minds and Machines Summer School

MBL, Woods Hole
May 29 – June 12, 2014

Teaching Fellows

Full time

Yasmine Meroz (Mahadevan)
Andrei Barbu (Ullman/Katz)
Leyla Isik (Poggio/Kreiman)
Emily Mackevicius (Fee)

Partial:

Danny Harari (Katz) June 3, 4, 6, 7, 9, 10, 11, 12
Georgios Evangelopoulos (Poggio) May 28, 29, 30, 31, June 9, 10, 11, 12
Owen Lewis (Poggio) June 2-12

Intermittent:

Ben Deen (Kanwisher) May 29-31, June 5 and June 11-12
Norman-Haignere (Kanwisher) May 29-31, June 5 and June 11-12
Alex Kell (Kanwisher) May 29-31, June 5 and June 11-12
Ethen Meyers (Poggio) June 3-12
Maximilian Nickel (Poggio) not May 28
Hector Penagos (Wilson) June 2-12
Stephen Voinea (Poggio) No Doodle response
Tomer Ullman (Tenenbaum)

THRUST 1 – Josh Tenenbaum (Tomer Ullman)

General requirements

All projects will be based partly on probmods.org, and so will need a few hours to go over the basics of probabilistic programming, conditioned queries and Church syntax.

General goals

I) Students will be able to design and test toy-world models of developmental data, which can then be scaled as necessary.
II) Students will understand the logic behind generative models as a tool for modeling core knowledge and development.

[1] "Planning, inference and Theory of Mind"

Theory of Mind as a psychological construct can be implemented computationally using tools developed in economics, robotics and decisions making. We will construct a generative planning algorithm and then invert it to reason about hidden variables such as goals and beliefs, from observable actions.

a) Implement a planning-as-inference algorithm in Church, moving through a few states and using several actions to get to a goal.
b) Use different transition functions and goals.
c) Embed planning algorithm in a query to reason about agent goals.
d) Introduce beliefs into planning algorithm and reason about beliefs, false belief.
e) Suggest 'lesions' of model to account for developmental trajectory. Is a notion of false belief necessary? Implement alternative "increasing inhibition" account.

[2] "Recursive reasoning and social development"

Building on the previous example, goals and desires become increasingly more interesting and complicated when other agents are involved, but this requires agents to reason about agents reasoning over agents…

a) Consider simple recursive reasoning modifications to the examples in probmods.org – Is there a way to make them more 'psychological'?
b) Implement multi-agent planning and use this for 'helping/hindering' examples. Consider effects of false belief in the case of either the social agent or the protagonist.
c) Again, suggest lesions of model to account for developmental trajectory. For example, why do kids seem more sensitive to hinderers first?

[3] "Social cliques, in-groups and out-groups"

Recent data from Spelke, Baillargeon and Hamlin (independently) shows infants have an early/core understanding of social groups that is different from agency. For example, they expect agents that are similar (in their appearance, behavior or preference) to be nicer to one another, and “dissimilar” agents to punish and hinder one another.

a) Define different group graphs and hierarchies, what fits with infant data
b) Implement a simple generative model for creating groups of agents
c) Specify how group membership affects planning and decision making, implement this in Church using previous ‘agent model’
d) Query simple observed data from past experiments, suggest new experiments.

[4] “Learning intuitive physics with probabilistic physics”

Two decades of research shows infants understand *something* about physical objects (they shouldn’t interpenetrate, they should follow smooth paths, etc.). We will use a recent suggestion for modeling adults’ understanding of intuitive physics and expand it to infants.

a) Run a simple generative-physics problem in Church - balls rolling and colliding, objects falling off a table, objects under different forces [For more advanced students who know python or are comfortable around scheme, a fuller example is suggested]
b) Embed the generative physics in a query and reason about unobserved physical qualities – forces, masses, etc.
c) Suggest learning trajectory of physics to account for various infant data (Baillargeon in particular, focusing on momentum and stability).
d) **Advanced** topics include integrating the intuitive physics section with the intuitive psychology section – create a simple model to tell apart an object from an agent based on their motion.

THRUST 2 – Gabriel Kreiman

(2.1 Ethan Meyers, 2.2 Hanlin Tang and Gabriel Kreiman, 2.3a Hector Penagos, 2.3b Isik, Meyers, Desimone, Kreiman)

---

[1] “What is there?” The neuronal, behavioral and computational viewpoints

Consider a set of single-object images {p1,…,pN} for which we have behavioral data, neurophysiological data (in monkeys and humans), and computational model responses.

a. Use classifiers to evaluate how well we can discriminate among them in single trials based on physiological data, based on computational models
b. Compare different classifiers

c. Compare different computational models
d. Compare different neural codes
e. Compare errors in computer models versus physiological data versus behavioral data
f. Tolerance to image transformations (scale, position, viewpoint)
g. Extension to objects embedded in background

[2] Object completion

In the real world, objects often appear partially occluded by other objects. This often poses a significant challenge for recognition and requires inference (extrapolation) to perform “object completion”.

a. Evaluate how well feed-forward architectures fare in identifying partially occluded objects

b. Analytically investigate the problem of recognition from partial information in feed-forward architectures

c. Show how Hopfield networks can perform object completion in toys examples

d. Combine Hopfield networks and feed-forward architectures

[3] “What will happen next?”

Given a state or an image, humans can make educated guesses about what will happen next. Here we examine this in simplified contexts:

a. Consider behavioral (position) and neurophysiological data (hippocampus ensemble recordings) from rats navigating the environment. How well can we predict the rat’s future behavior?
b. Consider behavioral (eye movements) and neurophysiological data (recordings in V4, IT) from humans and monkeys in visual search tasks. How well can we predict the next saccade?

THRUST 3 – Shimon Ullman and Boris Katz

(Andrei Barbu, Danny Harary)

---

[1] Why object detectors work and fail

Object detectors don't just underperform humans, they also fail in surprising ways. They can work on one image but fail on other, very similar, images.

a. Implement several object detectors (HOG, kNN, deep network, etc)

b. Determine how and why each object detector works

c. Can humans use the representations of these detectors to recognize objects?

d. Determine when each detector fails

e. Do humans fail in related but more extreme cases?

f. While these detectors seem very different from each other boosting doesn't help much, why?

g. What is missing? Do we need better representations, different features, better training procedures, or something else? Implement one of these extensions.

[2] Object detection in time

You have a lot of knowledge about how objects usually move. They rarely jump around the scene, disappearing and reappearing at random. But most object detectors can't use this information and operate on a frame by frame basis, while most trackers can't exploit the particulars of any object detector.

a. Implement an object detector and a tracker

b. Determine why each fails

c. Jointly detect the objects and track them

d. Tune the parameters of the object detector while tracking. For example: tune HOG features on the fly to deal with object rotation, or employ a kNN detector that looks at volumes in time.

[3] Features and detectors

Most approaches to computer vision tasks choose a feature set and an inference mechanism/classifier. Even for existing feature sets and classifiers it's likely the best combinations have yet to be found.

a. Implement several approaches to one problem such as object detection, gaze detection, pose detection, etc. This will require several different feature detectors.

b. Try different approaches to detection with different feature sets. For example one could take pose detectors that use edge features and modify them to use deep networks.

c. Why do some combinations work better than others?

d. Can we exploit the particulars of some combinations to perfor even better? By for example combining some steps of feature detection and inference, or by capturing the feature detector's uncertainty in the inference algorithm.

THRUST 4 – Nancy Kanwisher

(Alex Kell, Sam Norman-Haignere and Ben Deen)

Data-driven fMRI Analyses

Kanwisher (with Alex Kell, Sam Norman-Haignere, and Ben Deen)

FMRI research over the last 15 years has successfully identified several dozen robust functional divisions of the brain through the use of traditional hypothesis-driven methods, which test specific mental functions hypothesized in advance. However, there is no guarantee that all of the important ways the brain carves up the problem of cognition correspond to subdivisions scientists will think to test. To address this problem, a newer set of data-driven data fMRI analysis methods have been devised. These analysis methods entail the collection of fMRI responses to large, relatively unconstrained stimulus sets designed to broadly sample the space of stimuli/mental processes, rather than to test any particular hypothesis. Methods to discover structure in the resulting neural activations include clustering, in which sets of voxels are identified that have similar profiles of response across stimuli (or vice versa) and a variety of linear analysis methods including PCA and ICA, which respectively identify primary dimensions of variation and the most statistically independent dimensions in the fMRI response across stimuli.

In this project students will analyze either or both of two stimulus sets, each containing a pattern of response across voxels to a large number of movie stimuli. One data set consists of a single subject from the study described in Huth, Nishimoto, & Gallant (Neuron, 2012); another consists of a single subject scanned in the Kanwisher lab on a subset of these stimuli.

Numerous analyses are possible with these data, and students participating in this project are welcome to innovate and try new analyses. Among the analyses we envision are the following. We can ask i) how variants of the published analyses from Huth et al affect the results obtained (for example the use of alternate semantic models, or the inclusion versus exclusion of hierarchical regressors from the Word Net model), ii) to what extent the results described in Huth et al derive from the contribution of known functionally specific cortical regions, and iii) whether the resulting PCs map smoothly across the cortex, using a stronger test than that in the published study.

We can also compare the results derived by analyzing these data with PCA/ICA to the results derived from analysis in which voxels are clustered according to their response across timepoints/stimuli. A central question here is whether these fMRI data are better fit by clustering (in which each voxel is assigned to a unique functional profile) or by linear analyses (in which the functional response of each voxel is fit by a linear weighted sum of components).

The broader question asked in this project is whether, when, and how data-driven analyses of fMRI data can discover new functional organization of the human brain in the not already discovered by hypothesis-driven methods.

THRUST 5 – Tomaso Poggio

(5.1 Fabio Anselmi, Georgios Evangelopoulos, 5.2 Owen Lewis, Stephen Voinea, 5.3 Yasmine Meroz, Orit Peleg)

---

[1] "Magic theory" Simulations (Fabio Anselmi, Georgios Evangelopoulos)

Show basic properties of the theory. Simulations should be try simulations but lead to nice slides to explain theory. You will use and modify appropriately software from CBMM

a. Show empirically that invariance (eg to scale) decrease number of labeled examples required by a classifier
b. Show that recognition of parts of an image requires networks with more than 1 layer. Compare 2 vs 1 layers
c. Show that networks with more than one layer are globally invariant to more than affine transformations
d. Characterize clutter-tolerance in 2 vs 1 layer

[2] Classifiers as associative memory-based OS (O. Lewis + S. Voinea)

Invariant signatures of relevant objects and their parts must be stored (these are image patches that were segmented by motion). There are interesting issues about storage capacity and interference between memories; there are also interesting possibilities in how retrieval of parts could be used to support recognition of a full object and viceversa possibly pointing out the use of backprojections managed by a smart memory which learns and stores not only labels and pointers to other labels bust also procedures/routines.

[3] Collective intelligence for decision-making (Yasmine Meroz + Orit Peleg)

1. Decision-making processes abound in different contexts; other than neural networks, it is found also in the collective decision-making of insect colonies such as bee swarms,

and on the other hand also on the molecular level such as the polarization of cells in chemotaxis.

(a) All these models include cross inhibition of some sort.

a.1 What are the similarities and differences of the decision-making process amongst these three systems?

a.2 What is a minimal dynamical abstraction that lends itself to these different systems ?

a.3 At the population level, how might we write down stochastic DE for the process-similar to the TAFC problem and its variants (two alternative forced choice dynamical system).

(b) The role of noise in decision making.

b.1 Incorporate noise and simulate the system - what happens? How does it affect what happens when the two possibilities are similar?

b.2 Try to write a Fokker-Planck equation for the time-evolution of the probability distribution function (PDF) of such a process, and see whether it fits the PDF yielded by the numerics.