I^2 WORKSHOP - ABSTRACTS

THE VISUAL WORLD

Pawan Sinha
Title: Visual Object Discovery
Abstract: A key learning task the infant brain has to accomplish is to discover objects in complex sensory inputs. Within the first few months of life, well before they acquire motor or linguistic proficiency, children begin to exhibit significant object recognition skills. This developmental process transforms a child's sensorium from an amorphous collection of primitive attributes, to one where these attributes are integrated into cliques corresponding to distinct objects. Understanding how this is accomplished is a grand challenge, with implications that extend beyond developmental neuroscience, into the study of adult recognition strategies, as well as neurological disorders. I shall describe some of our experimentally motivated work on this question.

Bill Freeman
Title: Vision for Birds (joint with Russ Tedrake (EECS) and Emilio Frazzoli (Aero-Astro))
Abstract: A bird exhibits an impressive competence at visually guided flying--avoiding obstacles, exploiting varying aerial conditions, and landing on a perch. Russ Tedrake (EECS), Emilio Frazzoli (Aero-Astro) and I, with collaborators outside MIT, plan to build visually guided flying robots to explore various aspects of "bird-brained intelligence". The demands of the visual task, combined with the strong resource limitations, make this a good test-bed for making algorithms and systems that exhibit some aspects of what we consider to be intelligence.

Edward Adelson
Title: Intelligent photography
Abstract: Traditional photography is about pixels. Future photography will be about surfaces, objects, people, and scenes. There is already a lively interchange of ideas between researchers in human vision, computer graphics, and machine vision. The notion of a camera taking a picture at a particular position and time is replaced by the notion of gathing data about scenes, and combining it with knowledge about the world. As our representations get closer to those used by the human brain, we will be able to search, edit, and combine images using mid-level or high-level visual representations.

Earl Miller
Title: Multiple category representations in the prefrontal cortex
Abstract: Items are categorized differently depending on the behavioral context. For instance, a lion can be an African animal or a type of cat. We recorded lateral prefrontal cortex (PFC) neural activity while monkeys switched between categorizing the same image set along two different category schemes with orthogonal boundaries. We found that each category scheme was largely represented non-overlapping PFC neurons populations and that activity reflecting a category distinction was weaker, but not absent, when that category was irrelevant. By contrast, when monkeys were trained on independent, non-competing, category distinctions for two different image sets, many PFC neurons represented both category schemes. This suggests that the PFC separates competing category representations to reduce interference between them but that PFC neurons can represent multiple categories when the categories are not in competition. This also indicates that category representations in the PFC are based exclusively on top-down, not bottom-up, information.

Aude Oliva
Title: Visual Scene Understanding
Abstract: Visual scene understanding is central to our interactions with the world. We need to be able to recognize our current environment in order to act meaningfully within it: to choose a path for walking, to know where objects are likely to appear, or to know which behaviors are socially acceptable within a particular space. Converging evidence from behavioral, computational and cognitive neuroscience studies suggests that the human brain may employ a strategy for representing the gist or meaning of scenes that is independent of the visual complexity of the image, and that, under specific learning conditions, human observers have a phenomenal long-term memory capacity for the visual details of images. Together, these remarkable feats of human perception and memory suggest new avenues of research for understanding and modeling the mechanisms that underlie visual intelligence.

ACTION AND SPACE

Matt Wilson
Title: Thinking from experience - mechanisms and principles of sequential event memory
Abstract: Building models of the world through experience requires mechanisms for forming and evaluating memories of past time sequenced events. By recording patterned neural activity in the hippocampus while animals explore their environment, we have found that memories of spatial sequences are replayed while animals stop and "think" about past and future events. Such replayed event memories take the form of compressed chains of shorter memory chunks, suggesting a connection with models of cognition that involve linked state representations that can be used to carry out statistical inference on knowledge gained from experience. By combining this work with computational approaches to learning and cognition, we hope to understand the principles by which the brain develops knowledge through experience and uses this knowledge to guide behavior.

Nicholas Roy
Title: Bridging the gap between robot and human intelligence
Abstract: In the last few years, how robots understand the world around them has advanced considerably. Examples include the autonomous vehicles in the DARPA Grand Challenges and Urban Challenge, the considerable work in robot mapping, and the growing interest in home and service robots.
However, these example technologies and systems are still mostly restricted to research prototypes. One obstacle to getting more widely useful robots is that the way robots reason about their world is still pretty different to how people reason. Robots think in terms of point features, dense occupancy grids and cost maps. People think in terms of landmarks, segmented objects and tasks (among other representations). There are good reasons why these are different, and robots are unlikely to ever reason about the world in the same way that people do. But, there has been recent work in bridging the gap between the low-level geometric representations and higher-level semantic representations. I will talk about this work, what the open challenges are and why addressing these challenges will be important in general for addressing the problem of intelligence.

Leslie Kaelbling
Title: Household Intelligence
Abstract: What would it take for a robot to be a useful assistant and participant in a household? Perception, manipulation, reasoning, planning, learning. We know something about how to do all of this, but are nowhere near putting together a robust, flexible, intelligent household robot. I

will sketch an approach to the underlying representational issues that is based on a combination of logic and probability, and talk about how it can support these aspects of intelligence.

SOCIAL AND COLLECTIVE INTELLIGENE

Whitman Richards
Title: The fine structure of networks
Abstract: Networks are ubiquitous: in social interactions, in brain and nervous systems, in advanced computational systems. It is difficult to conceive of an intelligent system that is not built upon networks. What variations in network structures are there? Can they be classified? In the bio- and social worlds, what are the seeds and motifs for networks and how do they evolve? To begin to answer such questions, we decompose networks into "atomic" subgraphs that compose the larger "molecules" that in turn form still larger aggregates. A three parameter Simplex representation shows how these scale-space differences can reveal network characteristics not currently captured by popular measures.

Rebecca Saxe
Title: Neural circuits and computations underlying distinctively human cognition
Abstract: Human brains are composed of the same kinds of brain cells as other animals' brains, from monkeys to sea slugs. And yet, what we can do with our human brains seems qualitatively different. We not only pass on our knowledge to our young; we build whole institutions devoted to teach perfect strangers. We not only make moral and social judgments about the people we encounter; we produce and consume novels and plays, comedies and tragedies. Using modern non-invasive neuroimaging, these aspects of human social behavior can be traced back to the structure and function of our human brains. How the computations underlying these behaviors are implemented in neural circuits remains an open question.

Thomas Malone
Title: Collective Intelligence
Abstract: Intelligence is not just something that happens inside individual brains (or computers). It also arises in groups of people (or computers). And it occurs in some especially interesting new ways when groups of people are connected by computers. This talk will describe some of the basic design patterns embodied in these new types of collective intelligence (such as Google and Wikipedia), and it will identify open research questions such as how to measure collective intelligence and how to design new examples of it.

Sandy Pentland
Title: Honest Signals
Abstract: Social species, including humans, coordinate by signaling behavior. A particularly important class are reciprocal signals called `honest signals', and I have found that in humans these signals are important determinants of decision making performance of dyads, groups, and organizations. These signals thus serve to bind groups of humans together into a sort of collective intelligence.

Andrew W. Lo
Title: Neuroscience, Evolution, and Financial Crises
Abstract: The current financial crisis has exposed some significant gaps in the traditional economic paradigm of efficient markets and rational expectations. These gaps can be explained by recent research in the neuroscience of decision making, and when coupled with evolutionary models of behavior, a more complete understanding of financial market dynamics

and economic interactions emerges. Markets are neither always efficient nor always irrational---they are adaptive. This new paradigm is capable of explaining booms, busts, crises, and also suggests directions for regulatory reform.

LANGUAGE

Patrick Winston
Title: Language, the Great Differentiator
Abstract: If we are to understand human intelligence, we must understand how language differentiates human from nonhuman primates.

Going down, language directs and marshals visual and motor perceptions and actions. Shout "watch out for the car on the right" and eyes go to the right, consequent to the listener's language faculty instructing the visual faculty. Demand "pick up the hammer" and the listener's hand goes out to follow the instruction.

Going sideways, language cooperates with perception, making it possible to recognize objects by function. A cup, for example, is something we can hold, drink from, and set on a table, so we can recognize as a cup an animal horn affixed to some sort of base.

Going up, language enables description and description enables story telling and story telling lies at the center of human education, from our early exposure to fairy tales that keep us from wandering into the woods to our later reading of literature and history, and our still later personal experiences in life and surrogate experiences in law, medicine, business, and military schools.

Deb Roy
Title: New Horizons in the Study of Language Development
Abstract: Emerging forms of naturalistic longitudinal recordings of human behavior and social interaction promise fresh perspectives on age-old questions of child development. In a pilot effort, 240,000 hours of audio and video recordings of one child's life at home are being analyzed with a focus on language development. To study a corpus of this scale and richness, current methods of developmental sciences are insufficient. New data analysis algorithms and methods for interpretation and computational modeling are under development. Preliminary speech analysis reveals surprising levels of linguistic "finetuning" by caregivers that may provide crucial support for word learning. Ongoing analysis of various other aspects of the corpus aim to model detailed aspects of the child's language development as a function of learning mechanisms combined with everyday experience. Plans to collect similar corpora from more children based on a streamlined recording system are underway.

Ted Gibson
Title: Cognition of language above the word
Abstract: Research in my lab investigates the nature of linguistic representations and processes, mostly above the word level. I will here outline three recent lines of work. In one line of work, motivated and informed by information theory, we examine whether the presence of ambiguity at different levels of structure in language may be a desirable feature of an efficient communication system. In a second line of work, we are investigating the extent to which the human sentence processing mechanism fits different possible cue combination models that are based on different Bayesian inference methods. In this project, we include information about word frequency, syntactic structure frequency and world knowledge, and we attempt to model

reading times from the processing of different linguistic structures. In a third line of work, we propose a novel method to quantitatively evaluate similarity and differences between different kinds of linguistic representations and processes: Inter-Subject Analysis of Covariation. The method is a quantitative version of an approach that has been traditionally used in the syntactic literature, and has the potential to uncover important properties of the functional architecture of the language system.

Regina Barzilay

Title: Reinforcement Learning for Automatic Language Interpretation

Abstract: An essential part of language understanding is the ability to ground the meaning of words and sentences in the world they describe. In our recent work, we have demonstrated that automatic language processing systems can build on this connection to improve linguistic analysis. Specifically, we developed a reinforcement learning framework that maps imperative language to actions by executing candidate interpretations in a virtual environment. Ultimately, the goal is to enable machines to take advantage of human knowledge expressed in natural language documents, without the need for human intervention.