Probabilistic Models of the Visual Cortex.

Prof. A.L. Yuille (Dept. Statistics)

Joint appointments: Computer Science, Psychiatry, Psychology.

Course Objectives and expanded course description:

The course gives an introduction to state of the art computational models of the mammalian visual cortex. It covers topics in low-, mid-, and high-level vision. It briefly discusses the relevant evidence from anatomy, electrophysiology, imaging (e.g., fMRI), and psychophysics. It concentrates on mathematical modelling of these phenomena taking into account recent progress in probabilistic models of computer vision and developments in machine learning. The course is aimed at students in the Phsyical Sciences and Engineering (e.g., Statistics,Computer Science, Physics, and Mathematics) who are interested in Neuroscience and at student in Life Science and the Medical School who are interested in computational modeling.

The course starts with an overview of the challenges of vision, the architecture of the visual system, and "big picture" theories about its function. It proceeds by describing basic properties of neurons and statistical methods for analyzing them. Next it reviews the basic mathematical concepts used throughout the course, namely probabilistic models defined over graphs. The course then proceeds through studies of low-, mid-, and high-level vision describing computational models and the behavioral and neural evidence for them. The course also describes the role of bottom-up and top-down processing and the coupling of visual cues. The course conjectures that properties of the visual system are driven by the statistics of the natural world and by the nature of the visual tasks which are performed by an agent acting in this world.

Week 1.

This week describes the basic challenges of vision, namely the complexity and ambiguity of natural images and the enormous range of visual tasks. We describe the basic architecture of the visual system and describes theories for its purpose and the types of neural computations it performs. We summarize the experimental findings and describe "big picture" theories of what it does.

Week 2.

This week describes models of neurons and groups of neurons. We discuss linear models of cells in the retina and LGN (e.g., center-surround cells). We introduce probabilistic models of neurons and show how they relate to thresholded linear models. We discuss models of populations of neurons and their interactions. We discuss GLMs and how they can be used to estimate properties of groups of neurons. We discuss models of neurons with computations on the dendritic tree (e.g., B. Mel).

Week 3.

This week gives background material on probability models on graphs. We introduce Markov Random Fields, Gibbs sampling, and mean field theory. We review the material about neurons (week 2) from this perspective and show, for example, how standard "neural models" can be derived asmean field approximations to MRF models.

Week 4.

This describes low-level vision starting with edge detection, sparse coding, divisive normalization (including simple illumination models to explain why normalization is good), associative fields and segmentation. We describe the psychophysical phenomena, the computational models, and the neural evidence.

Week 5.

This week introduces mid-level vision by introducing the ideas of surfaces, depth, and occlusion. This motivates visual representations such as the 2.1D sketch, intrinsic images, and the 2 1/2D sketch. We discuss visual cues and computational models for estimating these representations. We discuss how the computational models relate to the psychophysical phenomena and the neural evidence.

Week 6.

This week describes computational models of motion perception and their relations to psychophysics and neural evidence. This includes models for the first stages of motion measurement and how local motion estimates are integrated over space to yield the perception of motion flow. In addition we describe Bayes-Kalman models for integrating motion estimates over time.

Week 7.

This week describes how different visual can be coupled. We start by introducing weak coupling of visual cues and discuss the psychophysical evidence and possible neuronal mechanisms. Next we discuss strong coupling of visual cues and their relation to psychophysics.

Week 8.

This will describes perceptual grouping of low-level visual cues into larger structures and attentional mechanisms. This includes studies of classic Gestalt phenomena together with more recent studies of human visual abilities on natural scence. We describe computational models and how they relate to experimental phenomena.

Week 9.

This week introduces high level vision -- i.e. object detection and scene understanding. We discuss hierarchical models of vision and of the ventral stream (e.g., Hmax). We review the neural and psychophysical evidence for these models.

Week 10.

We describe hierarchical computational models of the visual system. This includes Deep Belief Learning and Compositional Models. We discuss the importance of hierarchies, for addressing the complexity problem, and the role of bottom-up and top-down processing.

Course Assignments:

There will be four homework assignments which include computational modeling and computer implementation of the models. There will be a course project.

Grading:

The Grades will be based equally on four homework assignments and the course project (i.e. twenty percent each).

Reading List:

The course will be based on the 100 page monograph on the visual cortex being prepared by A.L. Yuille and D. Kersten (U. Minnesota). This monograph will be required reading. In addition, a detailed reading listwill be prepared from references given in the monograph (which currently contains over 200 references) and other sources.

A.L. Yuille and D. Kersten. Computational Models of the Visual Cortex. Course Notes. In preparation 2013.

Some suggested readings:

I. Biederman. Recognition-by-components: a theory of human image understanding. Psychological Review, 94(2):115–147, Apr. 1987.

H. H. Bulthoff and H. A. Mallot. Integration of depth modules: stereo and shading. Journal of the Optical Society of America A, 5(10):1749–1758, Oct. 1988.

M. Carandini. Do We Know What the Early Visual System Does? Journal of Neuroscience, 25(46):10577–10597, Nov. 2005.

P. Dayan and L. Abbott. The MIT Press, Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems, 2001.

J. J. DiCarlo, D. Zoccolan, and N. C. Rust. How Does the Brain Solve Visual Object Recognition? Neuron, 73(3):415–434, Feb. 2012.

W. S. Geisler. Contributions of ideal observer theory to vision research. Vision Research, 51(7):771–781, 2011.

T.S. Lee and D. Mumford. Hierarchical Bayesian inference in the visual cortex. Journal of the Optical Society of America A, Optics, Image Science, and Vision, 20(7):1434–1448, July 2003.

T.S. Lee and A.L. Yuille. Efficient coding of visual scenes by grouping and segmentation: theoretical predictions and biological evidence. In Doya, K., Ishii, S., Pouget, A., and Rao, R. P., editors, Bayesian Brain: Probabilistic Approaches to Neural Coding, pages 1–29. 2006.

B. W. Mel. Information processing in dendritic trees. Neural Computation, 6(6):1031–1085, 1994.

B. A. Olshausen and D. Field. Sparse coding with an overcomplete basis set: A strategy employed by V1? Vision Research. 37(23):3311–3325, 1997.

T. Poggio. The Computational Magic of the Ventral Stream: Towards a Theory. Nature Precedings. 2011.

E. Simoncelli and B. A. Olshausen. Natural image statistics and neural representation. Annual Review of Neuroscience, 24:1193–1216, 2001.

S. Thorpe, D. Fize, and C. Marlot. Speed of processing in the human visual system. Nature, 381(6582):520–522. 1996.

S. Ullman. Sequence seeking and counter streams: a computational model for bidirectional information flow in the visual cortex. Cerebral Cortex, 5(1):1–11. 1995.