Talks

Invited Talks

Nancy Reid, *Professor of statistical sciences, University of Toronto. Director of Canadian statistical sciences institute*

Title: Distributions for parameters

Abstract: There has been considerable recent controversy over the use of p-values and the phrase "statistically significant", both in subject matter settings and in the statistical literature. One approach to avoiding the dichotomization associated with hypothesis testing is to provide distributions for parameters. A familiar distribution is the posterior density of Bayesian inference, but there are renewed efforts to provide something like probability statements for parameter while avoiding specification of a prior probability. I will discuss the strengths and limitations of these procedures, with special attention to so-called objective Bayesian approaches.

Bio:

Nancy Reid, is a statistician recognized for her work on the theory of statistical inference, with an emphasis on likelihood-based methods and higher order asymptotics. She is known particularly for her development with D.R. Cox on adjustments to profile likelihood, and her work with D.A.S. Fraser on highly accurate approximations to significance functions. She is a Professor of Statistical Sciences at the University of Toronto, Canada Research Chair in Statistical Theory and Applications, and the director of the Canadian Statistical Sciences Institute. Her accolades include the COPSS Presidents' Award, the Krieger-Nelson Prize, and the Florence Nightingale David Award. She is a Fellow of the Royal Society of Canada. You can learn more about her work at http://utstat.toronto.edu/reid/

Doina Precup, *Professor of computer science, McGill University and Mila. Head of Deepmind Montreal.*

Title: Building Knowledge For AI Agents With Reinforcement Learning

Abstract: Reinforcement learning allows autonomous agents to learn how to act in a stochastic, unknown environment, with which they can interact. Deep reinforcement learning, in particular, has achieved great success in well-defined application domains, such as

Go or chess, in which an agent has to learn how to act and there is a clear success criterion. In this talk, I will focus on the potential role of reinforcement learning as a tool for building knowledge representations in AI agents whose goal is to perform continual learning. I will examine a key concept in reinforcement learning, the value function, and discuss its generalization to support various forms of predictive knowledge. I will also discuss the role of temporally extended actions, and their associated predictive models, in learning procedural knowledge. Finally, I will discuss the challenge of how to evaluate reinforcement learning agents whose goal is not just to control their environment, but also to build knowledge about their world.

Bio:

Doina Precup is associate dean of research at the faculty of science at McGill University and Mila, Canada research chair in machine learning and a senior fellow at the Canadian Institute for Advanced Research. She also heads the Montreal office of Deepmind. She conducts fundamental research on reinforcement learning at Deepmind, working in particular on AI applications in areas that have a social impact, such as health care, automated control and other fields. She's interested in machine decision-making in situations where uncertainty is high. She has extensive experience in organizing conferences and workshops over the last 20 years, including local chair of ICML, COLT, UAI and RLDM, as well as the program co-chair of ICML 2017.
You can read more about her at https://www.cs.mcgill.ca/~dprecup/

Naila Murray, *Scientific director at NAVER LABS Europe*

Title: Predicting aesthetic appreciation of images

Abstract: Image aesthetics has become an important criterion for visual content curation on social media sites and media content repositories. Previous work on aesthetic prediction models in the computer vision community has focused on aesthetic score prediction or binary image labeling. However, raw aesthetic annotations are in the form of score histograms and provide richer and more precise information than binary labels or mean scores. In this talk I will present recent work at NAVER LABS Europe on the rarely-studied problem of predicting aesthetic score distributions. The talk will cover the large-scale dataset we collected for this problem, called AVA, and will describe the novel deep architecture and training procedure for our score distribution model. Our model achieves state-of-the-art results on AVA for three tasks: (i) aesthetic quality classification; (ii) aesthetic score regression; and (iii) aesthetic score distribution prediction, all while using one model trained only for the distribution prediction task. I will also discuss our proposed method for modifying an image such that its predicted aesthetics changes, and describe how this modification can be used to gain insight into our model.

Bio:

Naila Murray is the Scientific Director at NAVER LABS Europe. She is working on topics including fine-grained visual categorization, image retrieval, and visual attention. From 2015 to 2019 she led the computer vision team at NLE. She currently serves as NLE's director of science. Her research interests include representation learning and multi-modal search. Naila has a B.Sc. in Electrical Engineering from Princeton University, and a PhD from the Unversitat Autonoma de Barcelona, in affiliation with the Computer Vision Center.
You can read more about her work at
https://europe.naverlabs.com/people_user/naila-murray/

Sara van de Geer, *Professor of statistics, ETH Zurich*

Title: Total variation regularization

Abstract: Consider the classical problem of learning a signal when observed with noise. One way to do this is to expand the signal in terms of basis functions and then try to learn the coefficients. The collection of basis functions is called a dictionary and the approach is sometimes called ``synthesis" because the signal is synthesised from the coefficients. Another learning approach, called ``analysis", is based on an l_1 regularization of a linear operator that describes the signal's structure. As an example one may think of a signal that lives on a graph, and the linear operator describes the change when going from one node to the next in the graph. The sum of the absolute values of the changes is called the total variation of the signal over the graph. A simple special case is the path graph, and a more complicated one is the two-dimensional grid. We will consider the regularized least squares estimator for such examples and also regularization using total variation of higher order discrete derivatives and Hardy Krause total variation. We will introduce the concept "effective sparsity" which is related to the dimensionality of the unknown signal. The regularized least squares estimator will be shown to mimic an oracle that trades off approximation error and ``estimation error", where the latter depends on the effective sparsity.

Bio:

Sara van de Geer has been Full Professor at ETH Zurich since September 2005. Her main field of research is mathematical statistics, with a special interest in high-dimensional problems. Her focus points are empirical processes, curve estimation, machine learning, model selection, and non-and semiparametric statistics. She is a member of the Research Council of The Swiss National Science Foundation, Leopoldina German National Academy of Sciences and the International Statistical Institute, and a fellow of the Institute of Mathematical Statistics. She is a correspondent of the Royal Dutch Academy of Sciences and the President of the Bernoulli Society.
You can read more about her work at
https://math.ethz.ch/research/seminar-for-statistics/sara-van-de-geer.html