Journée Statistique / Apprentissage Paris-Saclay

8 Janvier 2016

9h00 Café d'accueil

10h00 **Pierre Alquier** *(ENSAE)*
PAC-Bayesian bounds for aggregation of estimators: introduction and computational issues.

PAC-Bayesian bounds were introduced by David McAllester (1998) in order to control the prevision risk associated to procedures of aggregation of estimators. When dealing with aggregation with exponential weights, these bounds lead in some settings to the proof of the optimality of the prevision in some sense. These aggregates are usually computed through Monte Carlo methods. However, in many practical applications with large data, the computational cost of Monte Carlo methods is prohibitive, and it is tempting to replace these by (faster) optimization algorithms that aim at approximating the exponentially weighted aggregate: we will refer to these methods as variational Bayes (VB) methods. In this talk I will introduce VB approximations and show, thanks to a PAC-Bayesian theorem, that they are well founded, in the sense that the loss incurred in terms of prevision risk is negligible. Joint work with James Ridgway (Bristol) and Nicolas Chopin (ENSAE)

10h50 Pause café

11h20 **Joseph Salmon** (*LTCI, CNRS, Télécom ParisTech)*
GAP safe screening rule for sparsity enforcing penalties

High dimensional regression benefits from sparsity promoting regularizations. In such a context, screening rules leverage the known sparsity of the solution by ignoring some variables during (or even before) the optimization process, hence speeding up solvers. Such rules are said to be "safe" when it cannot wrongly discard features. In this talk, new safe rules for generalized linear models with sparsity enforcing regularization will be proposed. Our proposed GAP Safe (screening) rules can cope with any iterative solver and we illustrate their performance on coordinate descent, demonstrating interesting speed ups for learning problems.

12h10 **Odalric-Ambrym Maillard** *(INRIA - Université Paris-Sud)*
Multi-armed bandits: Proof techniques and challenges

Over the last two decades, strong interest in Multi-armed Bandits from the Machine Learning community has renewed this old field of research, and lead to key advances in sequential decision making theory, reinforcement learning and applied recommender systems. In this talk, we will recap the state-of-the-art results the simplest stochastic setting, and focus on a few key successful proof techniques. We will also present results improving on the state of the art thanks to a modernized 30-year old proof technique.

13h00 Déjeuner

14h30 **Claire Lacour** *(Université Paris-Sud)*

Minimax adaptive estimation of nonparametric hidden Markov models

In this talk, we consider stationary hidden Markov models with finite state space and nonparametric modeling of the emission distributions. Only the number of hidden space is assumed to be known. Two methods are proposed for estimating the emission distributions. Both are based on projections of the emission distributions onto nested subspaces of increasing complexity. The first method uses spectral properties of the involved matrices. The second method consists in minimizing a penalized least-squares contrast. This new estimator achieves the optimal rate of convergence, up to a logarithmic term. Simulations show the improvement obtained when applying the least-squares minimization consecutively to the spectral estimation.

15h20 **Stéphane Gaiffas** *(Ecole Polytechnique)*

Concentration for matrix martingales and microscopic activity of social networks using Hawkes processes

We consider the problem of unveiling the implicit network structure of user interactions in a social network, based only on high-frequency timestamps. Our inference is based on the minimization of the least-squares loss associated with a multivariate Hawkes model, penalized by ℓ_1 and trace norms. We provide a first theoretical analysis of the generalization error for this problem, that includes sparsity and low-rank inducing priors. This result involves a new data-driven concentration inequality for matrix martingales in continuous time with observable variance, which is a result of independent interest. The analysis is based on a new supermartingale property of the trace exponential, based on tools from stochastic calculus. A consequence of our analysis is the construction of sharply tuned ℓ_1 and trace-norm penalizations, that leads to a data-driven scaling of the variability of information available for each users. Numerical experiments illustrate the strong improvements achieved by the use of such data-driven penalizations.

16h10 Pause café

16h40 **Yann Ollivier** *(CNRS - Université Paris-Sud)*

Learn as you go: Training recurrent neural networks and dynamical systems

We consider the problem of learning the best parameters of a dynamical system, to optimize some function over the resulting trajectories. Recurrent neural networks are an important example, used for instance in text or speech modelling. Currently known algorithms are deeply unsatisfactory: either they work "backwards in time" from the end of the trajectory in order to compute a gradient of the loss function, which prevents use in real-time systems, or they have an unacceptable, quadratic cost in the dimension (e.g., extended Kalman filter). The same problem occurs in all temporal, non-iid settings such as reinforcement learning, hidden Markov models... We present a new, stochastic algorithm which does not work backwards in time and produces, at each instant, a random but unbiased estimate of the gradient of the loss with respect to the parameters of the whole past trajectory. For recurrent neural networks, the cost of this algorithm is the same as running the network itself.

17h30 ***Discussion***