PROGRAMA DE POSGRADO EPS-UAM, curso 2015-2016

Doctoral course: Bayesian Optimization

Lecturer: Dr. José Miguel Hernández Lobato
Harvard Intelligent Probabilistic Systems group
Harvard University
Cambridge, MA 02138, USA

Dates: 16-21 December 2015
Time: 11:00-13:00
Location: Escuela Politécnica Superior
Universidad Autónoma de Madrid
Ciudad Universitaria de Cantoblanco
Calle Francisco Tomás y Valiente, 11
28049 Madrid, SPAIN

Program

Lecture 1 [Wednesday, 2015/12/16, 11:00-13:0 @ LAB 16, 3rd fl., Bdg. A, EPS-UAM]

- An Introduction to Bayesian Optimization (1:15 hours).
- Practical: using spearmint for neural-network hyper-parameter tuning (45 minutes).

Lecture 2 [Thursday, 2015/12/17, 11:00-13:00 @ LAB 16, 3rd fl., Bdg. A, EPS-UAM]

- Bayesian Optimization with Constraints and with Multiple Objectives (1:15 hours).
- Practical: implement PAR-EGO in python (45 minutes).

Lecture 3 [Friday, 2015/12/18, 11:00-13:00 @ LAB 16, 3rd fl., Bdg. A, EPS-UAM]

- Information-theoretic Bayesian Optimization (1:15 hours).
- Practical: constrained hyper-parameter optimization with PESC (45 minutes).

The practical sections require installing Spearmint (https://github.com/HIPS/Spearmint) a Python package for Bayesian optimization and autograd (https://github.com/HIPS/autograd), an automatic differentiation package for python. Spearmint requires the database *mongodb*.

1. Install python, numpy, scipy, pymongo. For academic users, the anaconda distribution is great. Use numpy 1.8 or higher. We use python 2.7.
2. Download/clone the spearmint code
3. Install the spearmint package using pip: `pip install -e \</path/to/spearmint/root\>` (the -e means changes will be reflected automatically)
4. Download and install MongoDB: https://www.mongodb.org/

Install the pymongo package using e.g., pip `pip install pymongo` or anaconda `conda install pymongo`

Research seminar [Monday, 2015/12/21, 11:00-13:00, Sala de Grados A EPS-UAM]

Title: Probabilistic Backpropagation for Scalable Learning of Bayesian Neural Networks

Abstract: Large multilayer neural networks trained with backpropagation have recently achieved state-of-the-art results in a wide range of problems. However, using backprop for neural net learning still has some disadvantages, e.g., having to tune a large number of hyperparameters to the data, lack of calibrated probabilistic predictions, and a tendency to overfit the training data. In principle, the Bayesian approach to learning neural networks does not have these problems. However, existing Bayesian techniques lack scalability to large dataset and network sizes. In this work we present a novel scalable method for learning Bayesian neural networks, called probabilistic backpropagation (PBP). Similar to classical backpropagation, PBP works by computing a forward propagation of probabilities through the network and then doing a backward computation of gradients. A series of experiments on ten real-world datasets show that PBP is significantly faster than other techniques, while offering competitive predictive abilities. Our experiments also show that PBP provides accurate estimates of the posterior variance on the network weights.

Short bio: José Miguel Hernández Lobato is currently a postdoctoral fellow in the Harvard Intelligent Probabilistic Systems group at the School of Engineering and Applied Sciences of Harvard University, working with the group leader Prof. Ryan Adams. This position is funded through a post-doctoral fellowship awarded by the Rafael del Pino Foundation. Before that, he was a postdoctoral research associate in the Machine Learning Group at the Department of Engineering of Cambridge University (UK) from June 2011 to August 2014, working with Prof. Zoubin Ghahramani. During his first two years in Cambridge he worked in a collaboration project with the Indian multinational company Infosys Technologies. He also spent two weeks giving lectures on Bayesian Machine Learning at Charles University in Prague (Czech Republic). From December 2010 to May 2011, he was a teaching assistant at the Computer Science Department in Universidad Autónoma de Madrid (Spain), where he completed my Ph.D. and M.Phil. in Computer Science in December 2010 and June 2007, respectively. He also obtained a B.Sc. in Computer Science from this institution in June 2004, with a special prize to the best academic record on graduation. His research revolves around model based machine learning with a focus on probabilistic learning techniques and with a particular interest on Bayesian optimization, matrix factorization methods, copulas, Gaussian processes and sparse linear models. A general feature of his work is also an emphasis on fast methods for approximate Bayesian inference that scale to large datasets. The results of his research have been published in top machine learning journals (Journal of Machine Learning Research) and conferences (NIPS and ICML).

Organizado por el Grupo de aprendizaje Automático (GAA) del departamento de Ingeniería Informática de la Universidad Autónoma de Madrid (UAM) en el marco del programa de posgrado de la Escuela Politécnica Superior de la UAM (EPS-UAM). Esta actividad está co-financiada por EPS-UAM y por la Comunidad de Madrid dentro del proyecto "Conceptos y aplicaciones de los Sistemas Inteligentes" [CASI-CAM-CM S2013/ICE2845]

Escuela Politécnica Superior

UAM

Comunidad de Madrid
www.madrid.org

UNIÓN EUROPEA
Fondos Estructurales