Seminar on Deep Learning at Campus Nord (UPC)

Monday 23rd May 2016 Aula Màster - Campus Nord (UPC)

**11:00 – 12:00: Convolutional Neural Networks against the Curse of Dimensionality
                by Joan Bruna (UC Berkeley)**

Abstract: Convolutional Neural Networks are a powerful class of non-linear representations that have shown through numerous supervised learning tasks their ability to extract rich information from images, speech and text, with excellent statistical generalization. These are examples of truly high-dimensional signals, in which classical statistical models suffer from the curse of dimensionality, referring to their inability to generalize well unless provided with exponentially large amounts of training data.

In this talk we will start by studying statistical models defined from wavelet scattering networks, a class of CNNs where the convolutional filter banks are given by complex, multi-resolution wavelet families. The reasons for such success lie on their ability to preserve discriminative information while being stable with respect to high-dimensional deformations, providing a framework that partially extends to trained CNNs. We will give conditions under which signals can be recovered from their scattering coefficients, and will discuss a family of Gibbs processes defined by CNN sufficient statistics, from which one can sample image and auditory textures. Although the scattering recovery is non-convex and corresponds to a generalized phase recovery problem, gradient descent algorithms show good empirical performance and enjoy weak convergence properties.

We will discuss connections with non-linear compressed sensing, applications to texture synthesis, inverse problems such as super-resolution, as well as an application to sentence modeling, where convolutions are generalized using associative trees to generate robust sentence representations.

Bio: Joan Bruna graduated cum-laude from Universitat Politècnica de Catalunya in both Mathematics and Telecommunications Engineering, before graduating in Applied Mathematics from ENS Cachan (France). He then became a Sr. Research Engineer in an Image Processing startup, developing real-time video processing algorithms. In 2013 he obtained his PhD in Applied Mathematics at École Polytechnique (France). After a postdoctoral stay at the Computer Science department of Courant Institute, NYU, he became a Postdoctoral fellow at Facebook AI Research. Since Jan 2015 he is an Assistant Professor at UC Berkeley, Statistics Department. His research interests include invariant signal representations, deep learning, stochastic processes, and its applications to computer vision.

12:00 – 13:00: Recurrent Neural Nets and Differentiable Memory Mechanisms
by Oriol Vinyals (Google DeepMind)

Abstract: This past year, RNNs have seen a lot of attention as powerful models that are able to decode sequences from signals. The key component of such methods are the use of a recurrent neural network architecture that is trained end-to-end to optimize the probability of the output sequence given those signals. In this talk, I'll define the architecture and review some recent successes in my group on machine translation, image understanding, and beyond.

On the second part of the talk, I will introduce a new paradigm "differentiable memory" that has enabled learning programs (e.g., planar Traveling Salesman Problem) using training instances via a powerful extension of RNNs with memory. This effectively turns a machine learning model into a "differentiable computer". I will conclude the talk giving a few examples (e.g., AlphaGo) on how these recent Machine Learning advances have been the main catalyst in Artificial Intelligence in the past years.

Bio: Oriol Vinyals is a Research Scientist at Google DeepMind, working on Deep Learning. Oriol holds a Ph.D. in EECS from University of California, Berkeley, a Master degree from University of California, San Diego, and a double degree in Mathematics and Telecommunication Engineering from UPC, Barcelona. He is a recipient of the 2011 Microsoft Research PhD Fellowship. He was an early adopter of the new deep learning wave at Berkeley, and in his thesis he focused on non-convex optimization and recurrent neural networks. At Google Brain and Google DeepMind he continues working on his areas of interest, which include artificial intelligence, with particular emphasis on machine learning, language, and vision.