Research trajectory and main research lines

Vicenç Gómez i Cerdà

Last update: April 1, 2015

Abstract

My research focuses primarily on the development of machine learning methods for the analysis and/or control of complex systems. I take a multi-disciplinary approach from both theoretic and applied perspectives. From a theoretical point of view, my research lies at the interface between computer science, control theory and statistical physics. The application domains include social networks, robotics, brain computer interfaces and the smart grid.

1 Scientific Contributions

I describe here my scientific contributions organized by topic. First, I outline my main research lines, which include theoretical contributions on approximate inference and optimal control as well as applied work on robotics and knowledge discovery on big data sets. Finally, I will describe other important research lines.

1.1 Approximate Inference in Graphical Models

Pobabilistic inference in graphical models such as Bayesian networks is a central problem in machine learning and statistics. In general, exact inference is computationally hard for large networks, but this depends critically on the network connectivity. I have developed efficient algorithms for approximate inference based on loop calculus, a theoretical tool originated in the statistical physics community [5]. Loop Calculus is a reparametrization of the joint probability distribution associated with a graphical model that keeps its partition function invariant. It allows characterizing the fixed points of the popular belief propagation (BP) algorithm in terms of a loop series expansion, where each correction term to the BP solution corresponds to a loop in the graph. I introduced this tool into the machine learning community and developed principled ways to truncate this series, resulting in efficient and accurate approximate inference algorithms [26, 30].

1.2 Optimal Control and Reinforcement Learning

Stochastic optimal control problems deal with the problem of computing an optimal set of actions to attain some future goal. With each action and each state a cost is associated and the aim is to minimize some future goal. I have contributed to show the general equivalence between probabilistic inference in a graphical model and optimal control computation [12] for a class of non-linear stochastic optimal control problems introduced in [33]. This class of problems is known as Kullback-Leibler (KL) or Path-Integral (PI) control. Our contribution had important impact for diverse fields such as robotics [19, 32] and computational neuroscience [7, 18].

KL-control problems have inspired efficient reinforcement learning algorithms, such as Dynamic Policy Programming [2, 1], which enforces smooth policy updates during the optimization of the control policy. Such a feature has convenient theoretical properties in terms of asymptotic performance loss bounds. Smooth policy updates are also necessary for safe learning of motor skills in robotics, where state trajectories should change gradually.

Recently, in this application domain, I have (i) extended PI control to account for feedback policies [29], (ii) solved continuous, high-dimensional control problems by embedding KL-control in a hidden Markov model [17] and (iii) implemented hierarchical controllers for real-time control of a swarm of real quadrotors [31].

1.3 Learning communication patterns from online social interaction

Understanding how information spreads in a complex network is a very important problem. For example, characterizing the patterns of information flow in a social network plays a fundamental role in contexts like the spreading of innovations or diffusion of news and opinion. In this area, my research focuses on the analysis and modeling of the communication dynamics of special interest groups on the Internet.

We have been among the first researchers to analyze a large-scale social network from online communication data and to use simple mathematical models such as mixtures of log-normal distributions to explain the observed regular temporal activity patterns [24, 10, 8]. Our study comprised the analysis of billions of comments crawled from Slashdot, a popular technological news site, from the point of view of network macroscopic quantities, community detection and structure and evolution of discussions. This had implications, for example, in predicting the popularity of online content [21].

More recently, we have developed mathematical models and learning algorithms that are able to explain in detail the micro-structure and the evolution of such threaded discussions [27, 28, 3]. We have shown that a simple model of tree growth is able to explain the heterogeneity found in the discussion cascades of very diverse online websites, such as news forums or the Wikipedia talk pages. This model explains the evolution of the structure of a discussion by an interplay between popularity, novelty and certain trend bias to the original post.

1.4 Other contributions

In addition to the above results, I have also carried out research on other related topics to machine learning and complex networks. These are outlined below.

1.4.1 Brain-Computer Interfaces and Control

Brain Computer Interfaces (BCI) are a direct communication pathway between the brain and an external device, such as a speller or a wheel chair. Several problems limit the applicability of BCIs in real life. One of these problems is the presence of non-stationarities in the data, which causes patterns associated with each task during the training of the BCI to be different during testing, leading to a poor performance.

In this context, we have introduced a paradigm based on neural feedback for adaptive classification of the BCI. The analyzed feedback signal is the *interaction error potential*, a stereotyped evoked response that occurs in the brain when the BCI performs unexpectedly [6]. We have shown that such a signal can be decoded with high accuracy and used subsequently to adapt the BCI [16]. This paradigm has been implemented in several BCI applications, e.g. mental typewriters [20] (see also review [4]). Our original work considered a binary error signal for binary BCI tasks. We have extended the method for (i) probabilistic error signals, leading to a variety of learning algorithms that interpolate between supervised and unsupervised learning [13] and for (ii) multi-class problems [14]. More recently, we have evaluated the impact on the BCI performance of prior task selection, quantifying the potential increase in the number of users that can be expected to control a binary BCI [15].

1.4.2 Information Processing in Complex Networks

Finally, another research interest is to understand different dynamical emergent aspects of complex networks, such as synchronization, self-sustained activity and pattern formation and to find network adaptation mechanisms. In particular, during my PhD thesis, I have analyzed phase transition phenomena in spiking neural network models [9, 23] and derived local learning rules that lead to global self-organization towards critical states [25].

More recently, we have developed a novel efficient learning algorithm for sparse linear regression [11] and applied it in the context of power distribution networks. We showed how to reduce the load response uncertainty in open-loop control methods by learning the future price-elasticity of consumers based on their responses to previous pricing updates [22].

References

[1] M. G. Azar, **V. Gómez**, and H. J. Kappen. Dynamic Policy Programming with Function Approximation. *Journal of Machine Learning Research W&CP*, 15:119–127, 2011. AISTATS '11.

[2] M. G. Azar, **V. Gómez**, and H. J. Kappen. Dynamic policy programming. *Journal of Machine Learning Research*, 13(Nov):3207–3245, 2012.

[3] M. Cervallos. Online comments maybe not total waste of time. *Science news magazine*, November 12, 2010. `http://www.sciencenews.org/view/generic/id/65443/`.

[4] R. Chavarriaga, A. Sobolewski, and J. d. R. Millán. Errare machinale est: the use of error-related potentials in brain-machine interfaces. *Frontiers in neuroscience*, 8, 2014.

[5] M. Chertkov and V. Y. Chernyak. Loop series for discrete statistical models on graphs. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(06):P06009, 2006.

[6] P. W. Ferrez and J. d. R. Millán. Error-Related EEG Potentials Generated during Simulated Brain-Computer Interaction. *IEEE Transactions on Biomedical Engineerings*, 55(3):923–929, 2008.

[7] K. Friston. What is optimal about motor control? *Neuron*, 72(3):488–498, 2011.

[8] A. Kaltenbrunner, **V. Gómez**, and V. López. Description and prediction of slashdot activity. In *Proceedings of the 5th Latin American Web Congress (LA-WEB 2007)*, Santiago de Chile, 2007. IEEE Computer Society.

[9] A. Kaltenbrunner, **V. Gómez**, and V. López. Phase transition and hysteresis in an ensemble of stochastic spiking neurons. *Neural Computation*, 19(11):3011–3050, 2007.

[10] A. Kaltenbrunner, **V. Gómez**, A. Moghnieh, R. Meza, J. Blat, and V. López. Homogeneous temporal activity patterns in a large online communication space. 10th Int. Conf. on Business Information Systems. Workshop on Social Aspects of the Web, 2007.

[11] H. J. Kappen and **V. Gómez**. The variational garrote. *Machine Learning*, 96(3):269–294, 2014.

[12] H. J. Kappen, **V. Gómez**, and M. Opper. Optimal control as a graphical model inference problem. *Machine Learning*, 87(2):159–182, 2012.

[13] A. Llera, **V. Gómez**, and H. J. Kappen. Adaptive classification on brain computer interfaces using reinforcement signals. *Neural Computation*, 24(11):2900–2923, November 2012.

[14] A. Llera, **V. Gómez**, and H. J. Kappen. Adaptive multi class classification for brain computer interfaces. *Neural Computation*, 26(6):1108–1127, June 2014.

[15] A. Llera, **V. Gómez**, and H. J. Kappen. Quantitative analysis of task selection for Brain-Computer Interfaces. *Journal of Neural Engineering*, 11(5):056002, 2014.

[16] A. Llera, M. van Gerven, **V. Gómez**, O. Jensen, and H. Kappen. On the use of interaction error potentials for adaptive brain computer interfaces. *Neural Networks*, 24(10):1120 – 1127, 2011.

[17] T. Matsubara, **V. Gómez**, and H. J. Kappen. Latent Kullback Leibler control for continuous-state systems using probabilistic graphical models. *30th Conference on Uncertainty in Artificial Intelligence*, 2014.

[18] A. Pouget, J. M. Beck, W. J. Ma, and P. E. Latham. Probabilistic brains: knowns and unknowns. *Nature neuroscience*, 16(9):1170–1178, 2013.

[19] K. Rawlik, M. Toussaint, and S. Vijayakumar. On stochastic optimal control and reinforcement learning by approximate inference. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 3052–3056. AAAI Press, 2013.

[20] N. Schmidt, B. Blankertz, and M. Treder. Online detection of error-related potentials boosts the performance of mental typewriters. *BMC Neuroscience*, 13(1):19, 2012.

[21] G. Szabo and B. A. Huberman. Predicting the popularity of online content. *Communications of the ACM*, 53(8):80–88, Aug. 2010.

[22] **V. Gómez**, M. Chertkov, S. Backhaus, and H. J. Kappen. Learning price-elasticity of smart consumers in power distribution systems. In *Smart Grid Communications (SmartGridComm), 2012 IEEE Third International Conference on*, pages 647–652. IEEE, 2012.

[23] **V. Gómez**, A. Kaltenbrunner, and V. López. Event modeling of message interchange in stochastic neural ensembles. In *IEEE International Joint Conference on Neural Networks*, pages 81–88. IEEE World Congress of Computational Intelligence, Vancouver, BC, Canada, July 2006.

[24] **V. Gómez**, A. Kaltenbrunner, and V. López. Statistical analysis of the social network and discussion threads in slashdot. In *Proceeding of the 17th international conference on World Wide Web*, pages 645–654, New York, NY, USA, 2008. ACM.

[25] **V. Gómez**, A. Kaltenbrunner, V. López, and H. J. Kappen. Self-organization using synaptic plasticity. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 513–520. Curran Associates, Inc., 2008.

[26] **V. Gómez**, H. J. Kappen, and M. Chertkov. Approximate inference on planar graphs using loop calculus and belief propagation. *Journal of Machine Learning Research*, 11(Apr):1273–1296, 2010.

[27] **V. Gómez**, H. J. Kappen, and A. Kaltenbrunner. Modeling the structure and evolution of discussion cascades. In *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*, pages 181–190, New York, NY, USA, 2011. ACM.

[28] **V. Gómez**, H. J. Kappen, N. Litvak, and A. Kaltenbrunner. A likelihood-based framework for the analysis of discussion threads. *World Wide Web*, 16(5-6):645–675, 2013.

[29] **V. Gómez**, H. J. Kappen, J. Peters, and G. Neumann. Policy search for path integral control. *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 2014.

[30] **V. Gómez**, J. M. Mooij, and H. J. Kappen. Truncating the loop series expansion for belief propagation. *Journal of Machine Learning Research*, 8(Sept):1987–2016, 2007.

[31] **V. Gómez**, S. Thijssen, A. C. Symington, S. Hailes, and H. J. Kappen. Real-time stochastic optimal control for multi-agent quadrotor swarms. In *Robotics: Science and Systems, R4SIM workshop*, 2015.

[32] E. Theodorou, J. Buchli, and S. Schaal. A generalized path integral control approach to reinforcement learning. *Journal of Machine Learning Research*, 11:3137–3181, 2010.

[33] E. Todorov. Linearly-solvable Markov decision problems. *Advances in Neural Information Processing Systems 19*, pages 1369–1376, 2006.