Research Statement

Manuel Gomez Rodriguez

October 23, 2017

My main research interest lies in developing machine learning and large-scale data mining methods to model, understand and improve the functioning of large social and information systems. In this context, I pay particular attention to the creation, acquisition and dissemination of reliable knowledge and information in the Web and social media.

The above goal typically requires conceptual innovations and technical breakthroughs along three different dimensions:

1. Interpretable data-driven models of large social and information systems and dynamic processes that take place over them;
2. Scalable control algorithms that leverage these models to steer the functioning of large social and information systems; and,
3. Robust computational metrics to evaluate such models and algorithms in massive real-world datasets of gigabyte and terabyte scale.

Moreover, these conceptual innovations and technical breakthroughs often entails uncovering previously unexplored connections between fields, as exemplified by several of my current achievements and my vision for the future.

1 Current Achievements

Since November 2014, I have co-authored six papers in top conferences in machine learning—NIPS (4) and AISTATS (2)—and eleven papers in top conferences in data mining and the Web—WWW (3), WSDM (3), KDD (2), ICWSM (1), ICDM (1) and SDM (1). Additionally, I have also co-authored four articles in top journals in machine learning and data mining—JMLR (3), ACM TOIS (1). These publications spanned a wide variety of applications, which can be broadly categorized in the following broad research themes: information acquisition, information reliability, predicting and steering social processes, and fair machine learning.

Information acquisition. The advent of social media and social networking sites is changing dramatically the way in which people acquire information. In my work, I have focused on several important aspects of the information acquisition process in online settings, motivated by the following three questions:

(i) How efficient are people at curating information in social networking sites?

(ii) Can we understand the interplay between the structure of a social network and the information acquisition process?

(iii) Can we spot knowledge items (*e.g.*, questions and answers) in social media that systematically help people increase their expertise?

To answer the first question, together with researchers from the Networked Systems group lead by Krishna Gummadi, we introduced a computational framework to quantify the efficiency of a user as an information curator within a social networking site [1, 2]. The framework is general and applicable to any social networking site with an underlying information network, in which every user *follows* others to receive the information they produce. We find that social media users are sub-optimally efficient at acquiring information and this lack of efficiency is a consequence of the triadic closure mechanism by which users typically follow other users in social networking sites.

To answer the second question, in collaboration with researchers at Georgia Institute of Technology, we developed a temporal point process model for the joint dynamics of information diffusion and network evolution [13, 14]. This model has shed light on the interplay between the structure of a social network and the information acquisition process and it accurately predicts changes in the structure of a social network triggered by the information acquisition process. Moreover, the model has also served as a showcase of the rich set of possibilities offered by temporal point processes, which had been rarely explored before in large scale social networking modeling.

Finally, to answer the third question, we first focused on spotting specific knowledge items (*e.g.*, questions and answers) that systematically helped people increase their expertise. To this aim, we developed a probabilistic modeling framework that leverages the crowd to spot information with high (knowledge) value [32]. The key idea behind behind the framework is simple: if a knowledge item has high value, users who learn from the item will become more knowlegeable and thus their subsequent contributions be assessed more highly by others in terms of *e.g.*, upvotes, likes or shares. Thus, by jointly modeling learning events and contributing events, the framework can identify knowledge that leads to a measurable increase in expertise, as assessed by the crowd.

Then, we focused on spotting ordered sequences of knowledge items with high knowledge value. By uncovering these sequences, one can track users' interests and goals over time, and facilitate the design of automated curriculum building and, more broadly, personalized education. To this aim, we developed a novel modeling framework for efficiently clustering continuous-time grouped streaming data, the hierarchical Dirichlet Hawkes process (HDHP) [29], which is also one of the first example of models for complex social processes combining temporal point processes and Bayesian nonparametrics.

Information reliability. In the Web and social media, information is often not professionally curated. As a consequence, its high-quality, relevance and reliability is at stake and there is a need for smart algorithms to help curating information. In this context, I have worked towards that goal by developing a modeling framework to quantify information reliability and source trustworthiness [31] and a method to spot the source of a rumor [12].

— *Quantifying information reliability and source trustworthiness [31]:* online knowledge repositories, such as Wikipedia, Stack Overflow and Quora, put in place different evaluation mechanisms to increase the reliability of their content. For example, in Wikipedia, an editor can refute a questionable, false or incomplete statement in an article by removing it. In Stack Overflow, a user can accept or up-vote the answers provided by other users. However, these evaluation mechanisms only provide noise, often biased, measurements of the reliability of information and the trustworthiness of the information sources. In this context, in collaboration with researchers from the Max Planck Institute for Intelligent Systems, we developed a temporal point process modeling framework that leverage these noisy evaluations to distill robust, unbiased and interpretable measures of information reliability and source trustworthiness. The key idea is to disentangle to what extent a refutation (or a verification) is due to the intrinsic unreliability of the evaluated information or to the trustworthiness of the source providing the statement.

— *Spotting the source of a rumor [12]:* being able to spot the source that originally published a rumor or piece of malicious information in the Web and social media is critical for curtailing its spread and reducing the potential losses incurred. However, this is a very challenging problem since typically only incomplete information traces are observed and we need to unroll the incomplete traces into the past in order to pinpoint the source. To overcome this challenge, together with researchers at Georgia Institute of Technology, we developed a method that first learns a probabilistic diffusion network model based on historical diffusion traces and then identifies the source of an incomplete diffusion trace by maximizing the likelihood of the trace under the learned model.

Predicting and steering social processes. Since I joined MPI-SWS, I started to realize that my doctoral and postdoctoral work on network inference and influence maximization leveraged particular instances of a more general and powerful type of random processes, marked temporal point processes, which could be potentially used to design predictive models and control algorithms for a wide variety of social processes taking place on the Web and social media.

Since then, I have leveraged this realization to lead the design of:

(i) a new generation of data-driven predictive models based on marked temporal point process for a wide range of social processes over social and information networks, from product competition [33] and opinion dynamics [7] to spatiotemporal processes [8, 26]. In all cases, by exploiting the increasingly availability of fine grained user data, the models provide more accurate predictions than the state of the art.

(ii) a series of efficient off-line and online algorithms with provable guarantees to steer social processes both at a user and at a global level [11, 27, 37, 38]. These algorithms exploit an alternative representation of marked temporal point processes using stochastic differential equations (SDEs) with jumps and establish a previously unexplored connection between optimal control of SDEs with jumps and marked temporal point processes, which is of independent interest.

Fair machine learning. There is a growing concern that automated algorithmic decision-making, by now widely spread across a variety of online services, can lead to user discrimination, even if the absence of intent. In this context, the nascent field of fair machine learning aims to develop machine learning methods whose outcomes do not have a disproportionally large adverse impact on particular groups of people sharing certain sensitive traits such a race or sex.

In collaboration with researchers from the Networked Systems group lead by Krishna Gummadi, we have tackled the design of margin-based classifiers by introducing a variety of intuitive measures of decision boundary unfairness corresponding to different notions of fairness [34, 35, 36]. This work has received immediate international recognition by means of a best paper award honorable mention at the 26th International World Wide Web Conference (WWW), the flagship conference in Web research.

Other achievements. Together with a variety of collaborators from the Max Planck Institute for Intelligent Systems, Georgia Institute of Technology and Saarland University, I have extended my doctoral work on network inference and influence maximization in a collection of journals papers [9, 23, 24] and revisited this work in the context of privacy research [3].

Moreover, as an advocate for open knowledge and reproducibility, I have released companion websites for several of my publications, which typically include software packages and additional results.

Reference	Companion website	Content
[29]	https://learning.mpi-sws.org/hdhp/	Software package
[31]	http://learning.mpi-sws.org/reliability/	Software package, demo
[38]	http://learning.mpi-sws.org/redqueen/	Software package, demo
[34, 35]	http://learning.mpi-sws.org/fairness/	Software package

Table 1: Companion websites of several of my publications and their content

2 Previous Achievements

During my PhD, my research focused on understanding, predicting and controlling information diffusion over the Web and social media. This led to a series of papers [16, 17, 18, 20, 21, 22], which significantly advanced the state of the art in the network inference and influence maximization problems. One of these papers [17] received immediate international recognition by means of a best paper award honorable mention at KDD, the flagship conference in data mining. Since then, it has stimulated a large amount of follow-up work (>700 citations), now considered a subfield within the data mining community.

During my postdoctoral work, I continued working both on the network inference problem, establishing its theoretical foundations [6, 19], and the influence maximization problem, introducing a highly scalable algorithm [10]. The latter work received an outstanding paper award at NIPS, one of the premier conferences in machine learning.

3 Vision for the Future

In the upcoming years, my research agenda will be strongly influenced by complex social, technological and cognitive phenomena that emerge in an increasingly digital, always-on world. For example, the popularization of online tutoring systems and learning platforms has increased the availability of digital traces of human learning. Can we leverage these traces to design effective, personalized teaching algorithms? The amount of misinformation in social media and online social networking sites is rampant and it is often difficult for humans and machine learning algorithms alike to decide whether a piece of news is misinformation. Can we design machine learning methods with humans in the loop to effectively detect misinformation? Modern generative models such as variational auto encoders (VAEs) and generative adversarial networks (GANs) [25] have been proven successful at generating diverse collections of realistic images when parametrized by convolutional neural networks. Can we design VAEs and GANs that generate diverse collections of realistic graphs and temporal point processes when parametrized using graph convolutional networks [5] and recurrent marked temporal point processes [8], respectively?

My research agenda decomposes into many conceptual problems, which will often require technical innovations along three dimensions: (1) models, (2) algorithms and (3) experimental evaluation. Here, I briefly discuss four of them.

Optimizing human learning. Human learning has become an online activity, made popular by the emerging online tutoring systems and learning platforms. The promise of these online platforms is that automated fine-grained monitoring and greater degree of control and personalization will result in more effective learning. However, to fulfill this promise, it is necessary to empower these platforms with data-driven models of human learning and efficient teaching algorithms, which are still in its infancy. In the upcoming years, I plan to advance the state of the art on such models and algorithms, favoring principled, but practical, approaches with provable guarantees. More specifically, I will focus on two canonical problems at the heart of human learning: efficient memorization and curriculum learning. Moreover, while I will focus on human learning, the methodological innovations will be more broadly applicable towards developing teaching policies for other learning entities including machine learning algorithms and robots.

Detecting and preventing the spread of misinformation. In recent years, social media and online social networking sites has served as major disseminators of false facts, urban legends, fake news, or, more generally, misinformation. As a consequence, there are growing concerns that misinformation on these platforms has fueled the emergence of a *post-truth* society, where debate is perniciously framed by the repeated assertion of talking points to which factual rebuttals by the media or independent experts are ignored. In this context, I will develop machine learning methods to detect and prevent the spread of misinformation in social media and online social networking sites. These methods will leverage both *flags* provided by the crowd and fact-checks provided by *trusted third parties* and they will establish an unexplored connection between survival analysis, stochastic optimal control, and Bayesian inference, of independent interest. Initially, I will consider a optimistic scenario with not malicious users who unfaithfully flag genuine stories and then move on to a scenario with malicious users in which flags can be manipulated.

Temporal Point Processes and Graph Discovery. Despite their increasing predictive power, graph and point process models [4, 14, 28] typically make strong assumptions about the generative processes of the graphs and point process they model. This prevents from using these models to create realistic, while diverse, collections of graphs and point processes, which can be useful in a wide range of high-impact applications, *e.g.*, automatic chemical design [15] and automatic music generation [30]. Following up on the success of modern generative models such as VAEs and GANs at generating diverse collections of realistic images when parametrized by convolutional neural networks, I will design VAEs and GANs parametrized by graph convolutional networks and recurrent marked temporal point processes to generate diverse collections of realistic graphs and point processes.

Teaching Algorithms and Humans to be Fair. I will establish an unexplored connection between fair algorithmic and human decision making and the nascent field of machine teaching [39]. Machine teaching aims to design the optimal training data to *teach* a target model to a learning entity, which can be a machine learning algorithm or a human. In terms of algorithmic decision making, previous work has typically either argued for pre-processing the training data or modifying existing classifiers. By approaching the design fair classifiers from the perspective of machine teaching will allow me to combine and benefit from both strategies. In terms of human decision making, I will leverage machine teaching to, on the one hand, ensure a group of decisions makers is fair even if some of them are unfair and, on the other hand, teach individual human decision makers become fairer in their decisions.

References

[1] M. Babaei, P. Grabowicz, I. Valera, and M. Gomez-Rodriguez. Quantifying information overload in social media and its impact on social contagions. In *Proceedings of the 9th International AAAI Conference on Weblogs and Social Media*, 2015.

[2] M. Babaei, P. Grabowicz, I. Valera, K. Gummadi, and M. Gomez-Rodriguez. On the efficiency of the information networks in social media. In *Proceedings of the 9th ACM International Conference on Web Search and Data Mining*, 2016.

[3] M. Backes, M. Gomez-Rodriguez, P. Manoharan, and B. Surma. Reconciling privacy and utility in continuous-time diffusion networks. In *IEEE Computer Security Foundations Symposium*, 2017.

[4] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.

[5] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

[6] H. Daneshmand, M. Gomez-Rodriguez, L. Song, and B. Schölkopf. Estimating diffusion network structures: Recovery conditions, sample complexity & soft-thresholding algorithm. In *Proceedings of the 31th International Conference on Machine Learning*, 2014.

[7] A. De, I. Valera, N. Ganguly, S. Bhattacharya, and M. Gomez-Rodriguez. Learning and forecasting opinion dynamics in social networks. In *Advances in Neural Information Processing Systems*, 2016.

[8] N. Du, H. Dai, R. Trivedi, U. Upadhyay, M. Gomez-Rodriguez, and L. Song. Recurrent Marked Temporal Point Process: Embedding Event History to Vector. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, 2016.

[9] N. Du, Y. Liang, M.-F. Balcan, M. Gomez-Rodriguez, H. Zha, and L. Song. Scalable influence maximization for multiple products in continuous-time diffusion networks. *Journal of Machine Learning Research*, 2017.

[10] N. Du, L. Song, M. Gomez-Rodriguez, and H. Zha. Scalable influence estimation in continuous-time diffusion networks. In *Advances in Neural Information Processing Systems*, 2013.

[11] M. Farajtabar, N. Du, M. Gomez-Rodriguez, I. Valera, H. Zha, and L. Song. Shaping social activity by incentivizing users. In *Advances in Neural Information Processing Systems*, 2014.

[12] M. Farajtabar, M. Gomez-Rodriguez, N. Du, M. Zamani, H. Zha, and L. Song. Back to the past: Source identification in diffusion networks from partially observed cascades. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 2015.

[13] M. Farajtabar, Y. Wang, M. Gomez-Rodriguez, S. Li, H. Zha, and L. Song. Coevolve: A joint point process model for information diffusion and network co-evolution. In *Advances in Neural Information Processing Systems*, 2015.

[14] M. Farajtabar, Y. Wang, M. Gomez-Rodriguez, S. Li, H. Zha, and L. Song. Coevolve: A joint point process model for information diffusion and network co-evolution. *Journal of Machine Learning Research*, 2017.

[15] R. Gómez-Bombarelli, D. Duvenaud, J. M. Hernández-Lobato, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *arXiv preprint arXiv:1610.02415*, 2016.

[16] M. Gomez-Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the temporal dynamics of diffusion networks. In *Proceedings of the 28th International Conference on Machine Learning*, 2011.

[17] M. Gomez-Rodriguez, J. Leskovec, and A. Krause. Inferring Networks of Diffusion and Influence. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, 2010.

[18] M. Gomez-Rodriguez, J. Leskovec, and A. Krause. Inferring networks of diffusion and influence. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5(4):21, 2012.

[19] M. Gomez-Rodriguez, J. Leskovec, and B. Schölkopf. Modeling information propagation with survival theory. In *Proceedings of the 30th International Conference on Machine Learning*, 2013.

[20] M. Gomez-Rodriguez, J. Leskovec, and B. Schölkopf. Structure and Dynamics of Information Pathways in On-line Media. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, 2013.

[21] M. Gomez-Rodriguez and B. Schölkopf. Influence maximization in continuous time diffusion networks. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.

[22] M. Gomez-Rodriguez and B. Schölkopf. Submodular inference of diffusion networks from multiple trees. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.

[23] M. Gomez-Rodriguez, L. Song, H. Daneshmand, and B. Schoelkopf. Estimating diffusion networks: Recovery conditions, sample complexity & soft-thresholding algorithm. *Journal of Machine Learning Research*, 2016.

[24] M. Gomez-Rodriguez, L. Song, N. Du, H. Zha, and B. Schoelkopf. Influence estimation and maximization in continuous-time diffusion networks. *ACM Transactions on Information Systems*, 2016.

[25] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014.

[26] M. Jankowiak and M. Gomez-Rodriguez. Uncovering the spatiotemporal patterns of collective social activity. In *SIAM International Conference on Data Mining*, 2017.

[27] M. Karimi, E. Tavakoli, M. Farajtabar, L. Song, and M. Gomez-Rodriguez. Smart Broadcasting: Do you want to be seen? In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, 2016.

[28] J. Leskovec, D. Chakrabarti, J. Kleinberg, C. Faloutsos, and Z. Ghahramani. Kronecker graphs: An approach to modeling networks. *Journal of Machine Learning Research*, 11(Feb):985–1042, 2010.

[29] C. Mavroforakis, I. Valera, and M. Gomez-Rodriguez. Modeling the dynamics of online learning activity. In *Proceedings of the 26th International World Wide Web Conference*, 2017.

[30] I. Simon and S. Oore. Performance rnn: Generating music with expressive timing and dynamics. `https://magenta.tensorflow.org/performance-rnn`, 2017.

[31] B. Tabibian, I. Valera, M. Farajtabar, L. Song, B. Schoelkopf, and M. Gomez-Rodriguez. Distilling information reliability and source trustworthiness from digital traces. In *26th International World Wide Web Conference*, 2017.

[32] U. Upadhyay, I. Valera, and M. Gomez-Rodriguez. Uncovering the dynamics of crowdlearning and the value of knowledge. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining*, 2017.

[33] I. Valera and M. Gomez-Rodriguez. Modeling adoption and usage of competing products. In *IEEE International Conference on Data Mining*, 2015.

[34] B. Zafar, I. Valera, M. Gomez-Rodriguez, and K. Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *26th International World Wide Web Conference*, 2017.

[35] B. Zafar, I. Valera, M. Gomez-Rodriguez, and K. Gummadi. Training fair classifiers. In *20th International Conference on Artificial Intelligence and Statistics*, 2017.

[36] B. Zafar, I. Valera, M. Gomez-Rodriguez, K. Gummadi, and A. Weller. From parity to preference: Learning with cost-effective notions of fairness. In *Neural Information Processing Systems*, 2017.

[37] A. Zarezade, A. De, H. Rabiee, and M. Gomez-Rodriguez. Cheshire: An online algorithm for activity maximization in social networks. In *55th Annual Allerton Conference on Communication, Control, and Computing*, 2017.

[38] A. Zarezade, U. Upadhyay, H. Rabiee, and M. Gomez-Rodriguez. Redqueen: An online algorithm for smart broadcasting in social networks. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining*, 2017.

[39] X. Zhu. Machine teaching: An inverse problem to machine learning and an approach toward optimal education. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2015.