

KDD 2016 Speakers

Learning Sparse Models at Scale

Ralf Herbrich: Director of Machine Learning / Amazon

Tuesday, August 16 - 10:00 am to 12:00pm (Yosemite)

Applied Data Science Invited Talks

Recently, learning deep models from dense data has received a lot of attention in tasks such as object recognition and signal processing. However, when dealing with non-sensory data about real-world entities, data is often sparse; for example people interaction with products in e-Commerce, people interacting with each other in social networks or word sequences in natural language. In this talk, I will share lessons learned over the past 10 years when learning predictive models based on sparse data: 1) how to scale the inference algorithms to distributed data setting, 2) how to automate the learning process by reducing the amount of hyper-parameters to zero, 3) how to deal with Zipf distributions when learning resource-constrained models, and 4) how to combine dense and sparse-learning algorithms. The talk will be drawing from many real-world experiences I gathered over the past decade in applications of the techniques in gaming, search, advertising and recommendations of systems developed at Microsoft, Facebook and Amazon.

Speaker Bio

Ralf is Director of Machine Learning at Amazon and Managing Director of the Amazon Development Center Germany. He works on problems of demand forecasting, scalable machine learning, computer vision and linking structured content. In 2011, he worked at Facebook leading the Unified Ranking and Allocation team. From 2000 - 2011, he worked at Microsoft Research and was co-leading the Applied Games and Online Services and Advertising group which engaged in research at the intersection of machine learning and computer games and in the areas of online services, search and online advertising combining insights from machine learning, information retrieval, game theory, artificial intelligence and social network analysis. Ralf was Research Fellow of the Darwin College Cambridge from 2000 - 2003. He has a diploma degree in Computer Science (1997) and a Ph.D. in Statistics (2000). Ralf's research interests include Bayesian inference and decision making, computer games, kernel methods and statistical learning theory. He is one of the inventors of the Drivatars system in the Forza Motorsport series as well as the TrueSkill ranking and matchmaking system in Xbox 360 Live. He also co-invented the adPredictor click-prediction technology launched in 2009 in Bing's online advertising system.



ACM Code of Conduct