



My name is Tim Januschowski. I manage a team of machine learning scientists at Amazon Web Services (AWS), in particular the AWS Artificial Intelligence (AI) Labs. My work at Amazon is focussed on forecasting and lately also on more general time series analysis tasks such as anomaly detection. Over the course of my tenure at Amazon, I have helped launch a number of practical forecasting solutions, first internally for the Amazon retail supply chain, then for internal Amazon teams and finally to anyone via DeepAR, an algorithm available on Amazon SageMaker and the algorithms for Amazon Forecast. Both Amazon SageMaker and Amazon Forecast are AWS AI services.

Before joining Amazon, I worked for SAP also on forecasting. It was during my work at SAP when Stephan Kolassa introduced me to the community. I attended my first ISF (2012 Boston) and wrote my first Foresight article. My academic background is in mathematical programming and AI and I obtained my PhD in the intersection of the two disciplines.

My personal interests in forecasting are currently in bringing modern deep learning-based methods together with classic methods. Over the years, my interest in the engineering and technological aspects of building real-world forecasting software and identifying good software frameworks has remained steady, starting from in-memory databases, to distributed computing via map-reduce and most recently deep learning frameworks.

Over the course of the last years, I have given talks at ISF and written further articles with my team for Foresight and the IJF. Together with my team, I publish also in leading machine learning venues such as NeurIPS and AISTATS on forecasting and have at the major systems conferences such as VLDB and SIGMOD. For this year's ISF, I'm organizing a workshop on deep learning for forecasting and will give another practitioner talk. On my initiative, Amazon became a sponsor of the ISF, sponsored the M4 competition and the corresponding conference in December in New York (where I also gave a talk). I've actively promoted forecasting in the systems community, where I've been involved in tutorials for forecasting at VLDB last year and this year's SIGMOD conference.

My aim, if elected to the Board of Directors, would be to (i) help bridge the gap between forecasting and the machine learning community and (ii) help bridge the gap between forecasting conducted in academia and in industry. For (i), I will continue to deliver tutorials and talks in non-forecasting venues on forecasting and promote best practices, break throughs and software from the forecasting community in other disciplines. Conversely, I will continue to invite machine learning practitioners and academics to the ISF and personally help in bringing best practices from machine learning to forecasting such as deep learning software frameworks. For (ii), identifying meaningful benchmark data sets (either synthetic or ideally real-world) for researchers in academia is one avenue, inviting academics to industrial venues/visits and conversely inviting more practitioners to the ISF.

Thank you very much, Tim Januschowski