Mike West
The Arts & Sciences Professor of Statistics & Decision Sciences,
Department of Statistical Science, Duke University

Bayesian Forecasting and Predictive Decision Synthesis

Goal-focused perspectives on model uncertainty expand traditional statistical thinking about calibrating, comparing, and combining statistical model outputs for forecasting and resulting decisions. One theme in the recent literature links back to 1980s/90s developments of foundational concepts on combining forecasts from multiple sources. In recent years, this has evolved to define an encompassing, foundational Bayesian framework for evaluating, calibrating, comparing, and combining probabilistic forecasts from varied sources: multiple agents, agencies or models. This links into core questions of probabilistic model uncertainty analysis. Bayesian predictive synthesis (BPS) builds on the earlier developments and defines new approaches to forecast calibration and pooling. BPS includes traditional Bayesian model averaging and other methods for "pooling" multiple predictive distributions as special cases.

Then, the more recent development of Bayesian predictive decision synthesis (BPDS) takes this further to complete the inference-decision circle, reflecting the reality that forecasting models are often built with end-use decisions as a primary goal. The theoretical foundation of BPDS allows for explicit integration of decision outcomes as well as raw predictive performance in the model and forecast pooling enterprise, and a broader view of decision-focused forecasting. BPDS is generating applications in areas including optimal design for prediction and control linked to reinforcement learning, forecasting for macroeconomic policy decision analysis, and financial portfolio forecasting analysis, among others.