

Combining forecasts

Sebastian Funk

Note: the methodologies here have been discussed with DSTL and chosen following conversations with them.

Current approach

Forecasts are currently gathered as (1%, 5%, 25%, 75%, 95%, 99%) percentiles. These can be combined using mixture distributions (Wallis, 2005). At the moment, a normal distribution (which could be replaced by a more flexible skew-normal or non-parametric distribution) is fitted to each predictive probability distribution and these are combined as a mixture with equal weights.

Preferred approach

More sophisticated approaches create ensembles based on past predictive performance, i.e. the different predictive distributions are weighted differently based on how well they performed on past data. This can be achieved, for example, by Bayesian stacking (Yao et al., 2018) or ensemble post-processing (Wilks, 2017). It requires gathering a sufficient number of samples (of forecasts vs. data) and will replace the simple combination in the near future.

References

Wallis (2005). Combining Density and Interval Forecasts: A Modest Proposal. *Ox Bull Econ Stat* 67(s1) 983--994

Wilks (2017). Enforcing calibration in ensemble postprocessing. *Quart J Roy Met Soc* 114(710) 76-84

Yao, Vekhtari, Simpson, Gelman (2018). Using stacking to average Bayesian predictive distributions. [arxiv:1704.02030](https://arxiv.org/abs/1704.02030)