Locally scale invariant proper scoring rules

Jonas Wallin[1] **and David Bolin**[2]

[1] *Lund University, Sweden, jonas.wallinstat.lu.se*
[2] *King Abdullah University of Science and Technology, Saudi Arabia, david.bolinkaust.edu.sa*

Averages of proper scoring rules are often used to rank probabilistic forecasts. In many cases, the variance of the individual observations and their predictive distributions vary in these averages. We show that some of the most popular proper scoring rules, such as the continuous ranked probability score (CRPS) which is the go-to score for continuous observation ensemble forecasts, up-weight observations with large uncertainty which can lead to unintuitive rankings. To describe this issue, we define the concept of local scale invariance for scoring rules. A new class of generalized proper kernel scoring rules is derived, and as a member of this class, we propose the scaled CRPS (SCRPS). This new proper scoring rule is locally scale-invariant and therefore works in the case of varying uncertainty. Like CRPS it is computationally available for output from ensemble forecasts and does not require the ability to evaluate the density of the forecast. The theoretical findings are illustrated in a few different applications, where we in particular focus on models in spatial statistics.