Proper scoring rules for point processes

Claudio Heinrich[1], Thordis Thorarinsdottir[2], Peter Guttorp[3] and Max Schneider[4]

[1] *Norwegian Computing Center, Norway, claudio@nr.no*
[2] *Norwegian Computing Center, Norway, thordis@nr.no*
[3] *University of Washington, USA, peter@stat.washington.edu*
[4] *University of Washington, USA, maxs15@uw.edu*

Probabilistic predictions of the occurrence of events in space or space-time take the form of point processes. Examples include the prediction of earthquakes, crimes or distribution of species. While there is a wide variety of diagnostic tools available to assess the fit of point process models to observed data, it can be very challenging to rank competing point process models according to their predictive performance. The main reason for this is that many point process models are defined iteratively and their density is only known up to a constant, obstructing the use of the log-likelihood score. We present a class of proper scoring rules that overcomes this issue and can be evaluated from simulated draws of a point process model. The class of scores we consider is flexible and can target different properties of the prediction, such as homogeneity or point interaction. The principle we use for constructing these scores is not restricted to point processes and is useful for constructing scoring rules whenever the observation space is involved.

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.