Multidimensional local scoring rules

Matthew Parry*

Dept of Mathematics & Statistics, University of Otago, New Zealand
mparry@maths.otago.ac.nz

A scoring rule is a principled way of assessing a probabilistic forecast. The key requirement of a scoring rule is that it rewards honest statements of ones beliefs. A scoring rule is said to be local if it assigns a score based on the observed outcome and on outcomes that are in some sense "close" to the observed outcome. In practice, almost all scoring rules can be derived from a concave entropy functional. The property of locality then follows when the entropy is 1-homogeneous (up to an additive constant). Consequently, except for the log score, a local scoring rule has the remarkable property that it is 0-homogeneous; in other words, it assigns a score that is independent of the normalization of the quoted probability distribution. In many statistical applications, it is not plausible to treat observed outcomes as independent, e.g. time series data or multicomponent measurements. By using an appropriate entropy, we show that local scoring rules can be easily extended to multidimensional outcome spaces. Furthermore, we are able construct local scoring rules that are extensive, i.e. the score of independent outcomes is a sum of independent scores. Previously, only the log score was known to have this property. We end with an application of multidimensional local scoring rules to sequential data.

Key Words: entropy, normalization, log score, extensive