Bayesian oddities: homeopathic priors and the mirage of admissibility

James Theiler
Space Remote Sensing and Data Science Group
Los Alamos National Laboratory

ABSTRACT:

While non-Bayesian approaches (notably, the Generalized Likelihood Ratio Test or GLRT) have been widely used for target detection problems in spectral imagery, there are theoretical advantages provided by Bayesian detectors. In trying to exploit these advantages in the development of new target detection algorithms, some oddities were observed.

For the class of target detection problems considered here, the target signature is well known (eg, the absorption spectrum of a greenhouse gas) -- but its magnitude (eg, the gas concentration) is not known. The problem is often formulated in the language of hypothesis testing, where the null hypothesis is that the target is absent, and the alternative that the target is present. When the target's magnitude is unknown, however, the alternative hypothesis is composite: it is a union of simple hypotheses, each associated with a specific target strength. Composite hypothesis tests do not, except in rare cases, have unambiguously optimal solutions. A detector that works really well for finding targets at one target strength might not be as effective (compared to other detectors) at another target strength.

In searching for operationally useful detectors, it makes sense to restrict attention to "admissible" detectors. A detector is admissible if no other detector is "uniformly more powerful" than it is. [A detector is more powerful than another if it achieves higher detection rates at a given false alarm rate, and a detector is uniformly more powerful if it is more powerful at all target strengths.]

Theorems tell us that every Bayesian detector is admissible, and that every admissible detector either is Bayesian or is the limit of Bayesian detectors. Thus, searching for an admissible detector is equivalent to searching for a Bayesian prior. Indeed, this work was motivated by the possibility of "sculpting" priors to optimize application-specific performance criteria. In this work, however, no such prior will be exhibited. Instead, two obstacles in the way of achieving such priors will be discussed.

One of these is the homeopathic prior. This prior can be expressed as the sum of two components, one of which has infinitesimal magnitude. But the inclusion of this infinitesimal component leads to a tangible effect on the Bayesian detector that is built from this prior.

The second oddity is the apparent failure of theory when applied to certain performance criteria based on the receiver operating characteristic (ROC) curve. An example is constructed for which there is no prior that allows the associated Bayesian detector to outperform or even equal the non-Bayesian GLRT detector. The failure is only apparent, of course: a theorem is still a theorem. But in explaining the source of

this discrepancy, a case is made for inadmissible detectors in operational scenarios.