Submission for Conference on Novel Predictions in Düsseldorf

Jan-Willem Romeijn
Faculty of Philosophy
University of Groningen

Specificity, accommodation, and the sub-family problem

A pressing problem for probabilistic models of confirmation is that the preference for theories that have high probability seems at variance with the preference for theories that are specific. If high probability were the only objective, then our best theory would simply be the logical tautology. But of course, a tautology is of no use when it comes to making predictions.

In this paper I show how the tension between probability and specificity can be resolved. I begin by observing that theories can be framed as sets of specific hypotheses, each associated with a probabilistic relation to the data. I argue that the proper probabilistic framework for modelling the confirmation of theories is model selection, in which sets of hypotheses rather than individual ones are confronted with the data. The attractive feature of model selection is that it enables us to turn specificity into a probabilistic virtue. A more specific theory may turn out to have a higher marginal likelihood, or be on average closer to the truth, hence it will have a better predictive performance on past data. By carefully defining the probability space in which the competing theories are represented as sets of hypotheses, we can even maintain that the more specific theory has a higher posterior probability.

This presents a solution to the problem that specificity is difficult to capture in a probabilistic model for theory appraisal. Model selection presents a wide range of tools in which specificity is a factor. Notably, specificity is not merely factored in as the number of free parameters in a theory. A number of model selection criteria (AIC, BIC, and MDL) have recent reformulations that allow for comparisons of model size, if the models themselves do not differ in dimensionality.

Unfortunately, the use of model selection for theory appraisal also invites critical reflection. If specificity indeed runs parallel to probability, then what stops scientists from striving for maximal specificity, as long as they retain the predictive performance on past data? Or in model selection terms, what stops scientists from overfitting? In part, the answer lies in the very same model selection criteria. As I argue, all model selection criteria guard against overfitting because the expected predictive performance of overly specific theories is lower than that of their less specific rivals.

This is an unsurprising answer, because model selection tools were designed to deal with the risk of overfitting. Moreover, the answer fails to address one important way in which scientists might overshoot in their drive towards specificity. It is known as the sub-family problem in model selection, and as the problem of accommodation in confirmation theory: how can we rule out the post-hoc construction of a predictively accurate yet parsimonious theory? Rather than focusing on the theory being post-hoc, I will employ robustness, as spelled out in the context of model selection, to rule out theories that have been constructed with hindsight. I argue that both model selection and confirmation theory can benefit from this notion of robustness.