Modular priors for partially identified models

Ioanna Manolopoulou

This work is motivated by the challenges of drawing inferences from presence-only data. For example, when trying to determine what habitat sea-turtles "prefer" we only have data on where turtles were observed, not data about where the turtles actually are. Therefore, if we find that our sample contains very few turtles living in regions with tall sea grass, we cannot conclude that these areas are unpopular with the turtles, merely that we are unlikely to observe them there. Similar issues arise in forensic accounting: attempts to determine which companies are apt to misreport their official earnings based on a history of which firms were censured by the SEC are confounded by the fact that we only observe which firms got caught cheating, not which firms cheat (many of whom do not get caught). This sort of confounding is insurmountable from a point-estimation perspective, but the data are not entirely uninformative either. Our present work is devoted to parametrizing observation models in a way that isolates which aspects of the model are informed by the data and which aspects are not. This approach allows us to construct priors which are informative with respect to the unidentified parts of the model without simultaneously (and unintentionally) biasing posterior estimates of the identified parameters; these priors do not "fight against" the data. In addition, their modularity allows for convenient sensitivity analysis in order to examine the extent to which our ultimate conclusions are driven by prior assumptions as opposed to our data.

Joint work with Richard Hahn