International Society for Bayesian Analysis, 9^{th} World Meeting,
Hamilton Island, Australia, 2008.

Universal Bayesian Solution to the Induction Problem

Marcus Hutter

RSISE @ ANU and SML @ NICTA, Canberra, Australia
marcus@hutter1.net www.hutter1.net

Introduction. Bayesian reasoning is a well-studied and successful framework for inductive inference, which includes hypothesis testing and confirmation, parameter estimation, sequence prediction, classification, and regression. But standard statistical guidelines for choosing the model class and prior are not always available or can fail, in particular in complex situations. Finding tailor-made solutions to every particular (new) such problem might be possible but is cumbersome and prone to disagreement or contradiction. What is desirable is a formal general theory for inductive inference. Solomonoff completed the Bayesian framework by providing a rigorous, unique, formal, and universal choice for the model class and the prior. This "universal" Bayesian approach differs significantly from the classical objective as well as the subjective Bayesian philosophy. I show that *universal Bayes* (UB) essentially solves the long-standing induction problem, at least from a philosophical and statistical perspective. I know well that it is impossible to convince the reader in two pages about such a far-reaching claim.[1] But I expect that my talk will at least provoke a lively and critical discussion at ISBA'08, and hope it stimulates further investigations. All statements below have been mathematically formulated and rigorously proven. Strangely enough, it is forbidden to give a reference here.

The Model. In the Bayesian framework one has to specify the model class $\mathcal{M}$ and the prior $\mathcal{P}$. From this, all quantities of interest (evidence, posterior, predictive distribution) can in principle be computed. General guidelines are that $\mathcal{M}$ should be small but large enough to contain the "true" distribution, and $\mathcal{P}$ should reflect one's prior (subjective) belief or should be non-informative or neutral or objective if no prior knowledge is available. *Solomonoff* suggested to include in $\mathcal{M}_U := \mathcal{M}$ all constructive probability distributions in the sense that there exists an algorithm that computes the values. This class is small, since it is countable, but large since it includes all of today's valid physics theories. It does *not* contain many of the models beloved by statisticians, e.g. not even a Bernoulli(θ) process for incomputable θ.

[1] To the ISBA PC: Peter Grünwald is a respected Bayesian (he gave an invited talk at the Valencia 8 meeting) and could judge this work.

Nevertheless, UB is optimal in a strong sense even in those environments. *Ockham's razor* joined with *Epicurus principle* of multiple explanations, suggest to assign a high prior probability to simple models, and a small one to complex models. The *Kolmogorov complexity* of a function is the shortest computer program computing this function. It is an excellent complexity measure with various optimality properties, and can be used to define an essentially unique universal prior $\mathcal{P}_U$ over $\mathcal{M}_U$. $\mathcal{M}_U$ and $\mathcal{P}_U$ together define the universal Bayesian inference scheme UB.

What is shown. I consider all fundamental philosophical and statistical problems around induction I am aware of; more precisely, problems that the well-known approaches to inductive inference have. I then show that the aforementioned universal Bayesian theory has none of these problems. This is a significant progress if not solution to the induction problem.

The problems solved include:

- *The zero prior problem*: *Confirmation of (universal) hypotheses* in general, and the classical *Black ravens paradox* in particular (Maher's approach does not solve the problem).
- *Reparametrization invariance*: How to extend the symmetry principle from finite hypothesis classes (all hypotheses are equally likely) to infinite hypothesis classes (Jeffrey's prior does not always work).
- *Old-evidence problem / Ad-hoc hypotheses*: How can old evidence confirm a theory developed thereafter? How can we spot ad-hoc hypotheses, just tailored towards the past data?
- *Updating problem*: A Bayesian needs to choose the hypothesis/model class before seeing the data, which seldomly reflects scientific practice.
- Many other issues have been addressed: *consistency*, *efficiency*, *loss bounds*, magic numbers, *Carnap's confirmation theory*, Laplace rule, *continuous model classes*, how to incorporate *prior knowledge*, and others.

Critique. *Subjective Bayesians* will reject objective and hence also the universal prior, but there is a trick to incorporate prior knowledge into UB that provably works. *Objective Bayesians* will complain about the dependence on the underlying computational model. While this is indeed unfortunate, there are strong arguments why we can live with this. *Practitioners* will find the model useless, since it is incomputable, but UB provides a gold-standard that approximations and other practical approaches can, should and have aimed at. (And *Philosophers* find the math too tough and few *Computer Scientists* care about the issue).

Conclusions. In short, universal Bayesian induction solves or avoids or at least meliorates many if not all foundational and philosophical problems around induction, but has to be compromised in practice.