Initial Study Plan — Concrete Problems

Ian wood

September 2, 2009

Abstract

This document forms a part of the initial study plan towards a PhD at ANU for Ian Wood. It describes two concrete problems intended for developing technical prowess in problems relating to algorithmic information theory and possible publication. It is hoped that working on these problems will help develop intuitions around the problem of (numerical) induction, as well as provide tools for materialising these intuitions in a rigorous way and an understanding of the formal constructs that have been developed for this task.

1 Setting : Algorithmic Information Theory and Universal Baesian Inference

1.1 Universal or Solomonoff induction

Baesian inference considers a given process, class of environments=hypotheses=models (assumed to contain a "true" model that governs the process) and a prior confidence assigned to each environment. Given some data obtained from the process, the Baesian framework tells us how to update our confidence in each model. Generally, the process and a class of environments are easy to identify, though the assumption that the true model is in the class can be harder to justify. How to determine a good prior is, however, not so clear.

In accordance with Occam's razor, we wish to give low probability to highly complex measures, and high probability to simple ones. Solomonoff defined a prior $M(x)$ to sequences x from some finite alphabet as the probability that a given universal Turing machine U produces x from uniformly random input, or equivalently that the most probable continuation of x is that given by the shortest program that produces x. Naturally, this definition is highly dependant on the choice of universal Turing machine, however asymptotically, it can be shown to perform prediction at least as well as any other prior for most computable environments.

2 Proposed problems

2.1 Reconstructing a sequence from it's probability under a universal prior

We wish to determine if it is possible to recover a sequence x from it's probability under the universal prior. It is known that a related quantity, Kolmogorov complexity, K is **not** injective: there are generally many sequences with a given Kolmogorov complexity. However, $M(x)$ encodes much more information than $K(x)$. It can be constructed from the complexities of **all** computable sequence measures ν for which $\nu(x) \neq 0$.

More formally, we wish to determine if M is injective:

$$\forall x, y \in \mathcal{X}^* : M(x) = M(y) \Rightarrow x = y$$

To disprove this, it would be sufficient to find sequences x, y and a universal Turing machine U for which $M(x) = M(y)$. For example, it may be interesting to investigate binary sequences x and $\bar{x}$, the complement of x and try to constuct a siutable universal Turing machine.

Alternatively, we could approach the problem directly. For example attemtpt to show that if x and y do not coincide, then the Turing machine used to define M is not universal.

2.2 Find calculable approximations to the universal prior for Bernoulli sequences

Universal induction provides a powerfull theoretical framework for solving inductive problems, however it is completely incomputable. Despite this, it can be used to inspire computable (indeed calculable in a practical sense) approaches to induction.

This is interesting from two perspectives: firstly, it may result in effective and efficient implementations. Secondly, it may lead to a deeper understanding of the nature of universal induction and ways in which it can be utilised to help solve practical real world problems. A step toward this goal is to consider the relatively simple (and yet profound) example of predicting Bernoulli sequences.

A Bernoulli sequence is a sequence of zeros and ones that was generated by flipping a biased coin. The coin produces a 1 with probability $\theta \in [0, 1]$. We wish to investigate calculable approximations to the universal prior on θ which assign a non-zero probability to all potentially interesting theta (such as 0 and 1, 1/2, 1/6, ...). This is in contrast to continuous priors, for which the prior probability for specific values of θ is identically zero.

I will be investigating discrete priors that give zero probability to θ outside a countable set, and non-zero probability to elements of that set. Ideally, the chosen set should be dense in $[0, 1]$. The rational numbers or $\{k/2^n : k, n \in \mathbb{N}, k \leq 2^n\}$ are possible candidates.

The quality of each prior will be assessed in terms of its closeness to the universal prior and efficient computability of the quantities of interest.