Bayesian Occam's Razor for structure selection in human motor learning

Tim Genewein
Max Planck Institute for Biological Cybernetics
Max Planck Institute for Intelligent Systems
Graduate Training Center for Neuroscience
Tübingen, Germany
Email: tim.genewein@tuebingen.mpg.de

Daniel A. Braun
Max Planck Institute for Biological Cybernetics
Max Planck Institute for Intelligent Systems
Tübingen, Germany
Email: daniel.braun@tuebingen.mpg.de

I. Abstract

Learning structure is a key-element for achieving flexible and adaptive control in real-world environments. However, what looks easy and natural in human motor control, remains one of the main challenges in today's robotics. Here we investigate in a quantitative manner how humans select between several learned structures when faced with novel adaptation problems.

One very successful framework for modeling learning of statistical structures are hierarchical Bayesian models, because of their capability to capture statistical relationships on different levels of abstraction. Another important advantage is the automatic trade-off between prediction error and model complexity that is embodied by Bayesian inference. This so called *Bayesian Occam's Razor* results from the marginalization over the model parameters when computing a model's evidence and has the effect of penalizing unnecessarily complex models—see Figure 1.

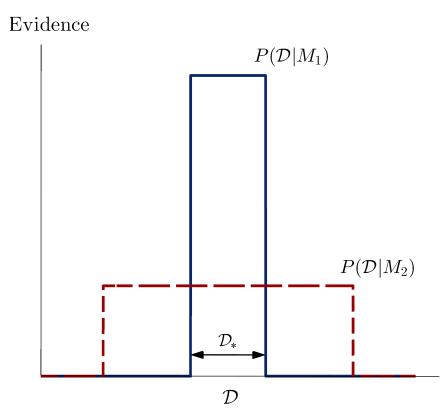

Fig. 1. Bayesian Occam's razor. Evidence $P(\mathcal{D}|M)$ for a simple model M_1 (blue, solid line) and a complex model M_2 (red, dashed line). Because both models have to spread unit probability mass over all compatible observations, the simpler model M_1 has a higher evidence in the overlapping region $\mathcal{D}_*$ and is thus the more probable model.

A standard paradigm to illustrate the trade-off between prediction error and model complexity is regression, where a curve has to be fitted to noisy observations with the aim of recovering an underlying functional relationship that defines a structure.

Here, we tested human behavior in a sensorimotor regression task, where participants had to draw a curve through noisy observations of an underlying trajectory generated by one of two possible Gaussian process (GP) models with different length-scales, a simple model with long length scale generating mostly smooth trajectories and a complex model with short length scale generating mostly wiggly trajectories. Participants were trained on both models, in order to be able to learn the two different structures. They then observed ambiguous stimuli that could be explained by both models and had to draw regression trajectories, which implied reporting their belief about the generating model.

In ambiguous trials where both models explained the observations equally well, we found that participants strongly preferred the simpler model. In all trials, Bayesian model selection provided a good explanation of subjects' choice and drawing behavior.

The approach presented in this work might also lend itself for application in robotic tasks, where sensory data has to be disambiguated or a goodness-of-fit versus complexity trade-off has to be performed.

Acknowledgments

This study was supported by the DFG, Emmy Noether grant BR4164/1-1.