Active Learning in Changing Environments

Nicolas Collignon & Christopher G. Lucas

As humans, we have a remarkable ability to adapt our beliefs and strategies when we are confronted with new problems. One reason for our successes is that we are not passive observers. We direct our attention and choose actions in ways that allow us to update our beliefs more efficiently than if we had to learn through observation alone (Gureckis & Markant, 2009; Oaksford & Chater, 1994; L. E. Schulz & Bonawitz, 2007). This is true even when our goals are pragmatic, rather than being about learning per se. We strike a balance between exploring options and gaining information to serve future goals, and exploiting knowledge to obtain immediate rewards (Mehlhorn et al., 2015; E. Schulz, Konstantinidis, & Speekenbrink, 2015).

However, active learning only goes so far by itself: efficient exploration and exploitation depend on having a high-level model of the environment or the task to be solved. Without a sense of how actions relate to one another or what variables might be relevant to actions' outcomes, learning will be aimless and rewards will be comparatively scant. The distinctively human talent is not active learning itself – which has been studied extensively and yielded valuable algorithms in many domains – but active learning in a world where the problems we must solve are constantly changing.
These changes may be gradual, requiring minor adaptation to the environment, but there may also be structural changes that require drastic adaptation. Sensitivity to these structural changes allows an agent to transfer knowledge across domains, select informative actions, and bootstrap learning without requiring extensive domain-specific knowledge.

Our project seeks to understand active learning in settings where the structure of the task is subject to abrupt change, using a combination of computational models and psychological experiments. Our approach builds on recent developments in active learning and sequential decision making algorithms, and allows us to understand group-level and individual human behaviour. We use Bayesian optimization (BO), a family of techniques that has proven useful in solving practical active learning problems (Snoek, Larochelle, & Adams, 2012). BO uses Gaussian processes to build rich probabilistic models of the environment, and both have been successful in predicting aspects of human behaviour (Borji & Itti, 2013; Lucas, Griffiths, Williams, & Kalish, 2015). To explore how participants learn in a dynamical setting, we combine these tools with sequential Monte Carlo algorithms that have shed light on human judgments in causal reasoning and categorisation (Abbott, Griffiths, et al., 2011; Sanborn, Griffiths, & Navarro, 2006), and show how a resource-contrained agent might make approximately optimal inferences in practice. Using these methods, we can explicitly compare different theories of human active learning, as expressed in priors over problem structures and environmental change as well as mechanisms underlying belief revision.

To assess the predictions of our models and gather data about human active learning more generally, we introduce a flexible new experimental framework. Participants are presented a sequence of grids made of tiles with different observable features, and are told to select tiles in order to maximize their score. Each click reveals a reward, providing indirect information about the problem's structure, such as whether brightness is relevant, or whether there is a hidden spatial pattern that can be exploited.

Our results show that participants are able to learn efficiently and adaptively in changing environments, but can also suffer from garden-path effects and sometimes fail to adapt. Our models successfully predict human decisions, and characterise the learning behaviours and inductive biases of individual participants. In so doing, they provide a means to mimic the best learners' adaptive behaviour, and simultaneously provide insights into how and why human learners occasionally fail.