An ideal observer model for optimal inference in the presence of different types of uncertainty

Robert C. Wilson, Matthew R. Nassar and Joshua I. Gold

University of Pennsylvania

Whether getting a new job or a new president, life is full of "changepoints" that cause the rules of the game to change abruptly. Making predictions in such circumstances can be challenging because changepoints can render much of the past irrelevant. For example, the outcome of this year's presidential election suggests that policies from recent years will not necessarily be good predictors of policies in the coming years. Here we examine how to identify changepoints in a noisy environment and then use that information to adjust how to use past experiences to predict future events.

We present an ideal-observer model of a dynamic probabilistic estimation task designed to quantify the extent to which human subjects use new pieces of information in the presence of possible changepoints. At each time-step in this task, a bar appears on screen at a random location, and subjects are asked to predict the location of the bar on the next time-step. The subjects receive a cash reward based on the accuracy of their prediction. The location of the bar is determined by a changepoint process, which is sampled from a Gaussian distribution whose mean and variance change abruptly at changepoints but otherwise remain constant. Because the time of occurrence of changepoints are unknown to the subjects, they must contend with two different kinds of uncertainty: "expected" uncertainty corresponding to the variance of the Gaussian distribution and "unexpected" uncertainty corresponding to probability that a changepoint might occur [A.J.Yu and P.Dayan Neuron 46:681-692,2005].

When the frequency of changepoints (known as the hazard rate), but not their exact times of occurrence, is given, Adams and MacKay [Technical report, Cambridge University, 2007] and Fearnhead and Liu [J.R.Statist.Soc.B 69(4):589–605,2007] have developed online Bayesian models capable of making optimal predictions. However, in general, the hazard rate is unspecified and must be inferred from the data. We solve this problem by developing an online Bayesian algorithm capable of making optimal predictions in the face of expected and unexpected uncertainty and an unknown – and potentially variable – hazard rate. In particular, we model the occurrence of changepoints as a Bernoulli process and infer a constant hazard rate by keeping track of changepoints experienced over time. A hierarchical extension of the model in which the hazard rate is itself generated by a changepoint process allows us to make inferences about variable hazard rates.

We compare performance of the model and human subjects on the dynamic probabilistic estimation task. We focus on the weight given to each datum in determining the next estimate as the hazard rate changes. A key result is that this weighting differs substantially in a stable environment, in which changepoints are rare and the past is a good predictor of the future, versus a volatile environment, in which changepoints are common and past trends are typically not helpful. By quantifying these principles, the model establishes a benchmark both for comparing human performance on this kind of prediction task and for analyzing a host of changepoint data.