

$0 \leq C_{ij} \leq 3$. To transform these values into a probability distribution E , we use a simple Bayesian Thurstonian model presented in [1].

The core idea of this model is the following generative process for the values C_{ij} : we assume that each of the events e_i has a normal distribution associated with it, with mean $0 \leq \mu_i \leq 1$ and variance σ_i^2 . When a participant is asked to which of the events e_i and e_j is the most probable, they sample $s_i \sim N(\mu_i, \sigma_i^2)$ and $s_j \sim N(\mu_j, \sigma_j^2)$ and respond that event e_i is more probable if $s_i > s_j$, and that event e_j is more probable otherwise. Under this generative model, the values of C_{ij} are binomially distributed variables, with $C_{ij} \sim B(3, p_{ij})$, where the “success” probability p_{ij} is determined by μ_i, σ_i^2, μ_j and σ_j^2 according to the following equation:

$$p_{ij} = \Phi\left(\frac{(\mu_i - \mu_j)}{\sqrt{\sigma_i^2 + \sigma_j^2}}\right), \quad (1)$$

where Φ is the cumulative density function of the standard normal distribution $N(0, 1)$.

By placing a prior probability distribution over the set of all parameters, μ_i and σ_i^2 for $i = 1, 2, \dots$, we can use Bayes’ Law to compute a posterior probability over parameter values based on the values of C_{ij} yielded by our experiment. We place a component-wise prior on the parameters such that $P(\mu_i, \sigma_i^2) \propto \exp -\sigma_i^2$, with $P(\mu_i, \sigma_i^2) = 0$ if $\mu_i < 0$ or $\mu_i > 1$.

We perform the Bayesian inference numerically using a Metropolis Hastings algorithm to draw samples from the posterior distribution. Our proposal process for the MH algorithm selects a single parameter to change from a uniform distribution over all the parameters, and then proposes a new value for that parameter by sampling from a normal distribution centred on the parameter’s current value. The normal proposal distributions are not truncated, with the requirement that $0 \leq \mu_i \leq 1$ is enforced by the prior. To obtain an estimate of the values of μ_i , we take 10 samples from each of 10 randomly initialized chains, for a total of 100 samples, with a lag of 1000 iterations between samples. Each chain is allowed to “burn in” for 15,000 iterations. This value was chosen by examining plots of the log posterior probability versus iterations and observing a plateau after around 15,000 iterations.

At the end of this process we have recovered a value of μ_i for each of the events in our experimentally defined world. These values are transformed into an event distribution E via the straightforward process of setting each event’s probability to be directly proportional to its value of μ_i . This is the distribution used to produce Figure 4 and the mean deviation scores given in Section 4.

References

- [1] B. Miller, P. Hemmer, M. Steyvers, and M.D. Lee. The wisdom of crowds in rank ordering problems. In A. Howesa, D. Peebles, and R. Cooper, editors, *9th International Conference on Cognitive Modeling*, 2009.