

Mark Steyvers
Cognitive Sciences, UCI

The Wisdom of Crowds in the Aggregation of Rankings

We analyze the collective performance of individuals in a number of ranking tasks, such as reconstructing the order of historic events from memory (e.g. the order of US presidents according to terms in office) as well as predicting the rankings of sports teams. We also study iterated learning environments where individuals observe the ranking from the previous person in a chain.

We propose a Bayesian approach to the aggregation of rank order judgments based on Thurstonian models. We assume that individuals' reported rankings are based on random permutations of a latent truth that is shared among all individuals and that each individual is associated with a latent level of expertise that determines the similarity of the reported ranking to the latent truth. We demonstrate a wisdom of crowds effect – averaged across multiple problems, the aggregate ranking is better than the best individual in the group and also better than standard aggregation methods such as Borda count. We also show that our model-based measure of expertise outperforms self-report measures, taken both before and after completing the ordering of items.

Finally, we introduce a Bayesian non-parametric extension of the model, based on Dirichlet process mixture models, that can identify subgroups of individuals associated with different latent truths. This approach simultaneously estimates the number of latent truths as well as the latent truths themselves. We show that this model is useful to separate individuals who use qualitatively different sources of information to produce their answers.