

Recovering human category structure across development using sparse judgments

Pablo León-Villagr¹, Isaac Ehrlich², Christopher G. Lucas³, and Daphna Buchsbaum⁴

¹ University of Warwick, UK

² University of Toronto, Canada

³ University of Edinburgh, UK

⁴ Brown University, USA

Keywords: Categorization, experimental methods, child development.

Multidimensional scaling (MDS) has provided insight into the structure of human perception and conceptual knowledge. However, MDS usually requires participants to produce large numbers of similarity judgments, leading to prohibitively long experiments for most developmental research. Here we propose a method that combines simple grouping tasks with recent neural network models to uncover participants’ psychological spaces. We validate the method on simulated data and find that it can uncover the true structure even when given heterogeneous groupings. We then apply the method to data from the World Color Survey and find that it can uncover language-specific color organization. Finally, we apply the method to a novel developmental experiment and find age-dependent differences in conceptual organization.

1 Learning Similarities Implicitly

Our method is based on deep metric learning, a family of deep neural network architectures that learn similarities from groupings—items that have been sorted into groups taken to be the “same” type. These techniques produce a function that maps any stimuli in the domain to locations on a manifold. One such method is dimensionality reduction by learning an invariant mapping (DrLIM) [1]. In cognitive science, DrLIM has been used to elicit psychological spaces from same-different judgments. While DrLIM has been used before, the experiment in [2] required each participant to produce 90 same-different judgments, which makes the task too demanding for developmental settings.

Here we develop this method into a paradigm that can satisfy two critical requirements in developmental studies: recovering psychological spaces when participant data is heterogeneous and dealing with aggregations of participants that do not judge the same materials.

2 Validation of the Method

We constructed synthetic datasets from aggregations of heterogeneous categorizers to validate that our method can recover meaningful structure in these situations. We found that even for aggregates of agents using different categorization schemes, our method captured the high-level feature structure.

We then evaluated the applicability of DrLIM to heterogeneous datasets by training on data from the World Color Survey⁵. The WCS dataset contains color terms for 330 color chips from speakers of 110 languages in non-industrialized societies. We selected six languages to evaluate our model: two three-term languages (the smallest number of terms in the WCS), two seven-term languages (the most frequent), and two 11-term languages (the number of basic color terms in English). We then generated all pairs of chips for each speaker in each language, categorizing the pair as the same if the speaker assigned the same term for both colors. Similar to the simulation results and

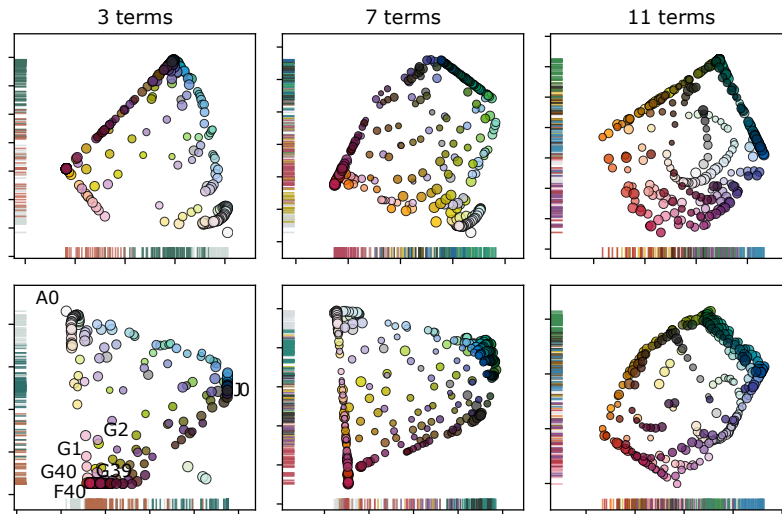


Fig. 1: Six languages from the WCS differing in the number of terms (columns). Rug plots display the distribution of terms within the 2D space. The size of the points represents agreement for the color term within the language.

consistent with the structure of perceptual color spaces, 2D solutions achieved acceptable loss, with subsequent dimensions offering only minor reductions. The resulting 2D spaces corresponded well to the number of terms within a language. Focal colors, which speakers deemed the most representative instances for each color term, were centered on clusters within the learned space. In contrast, colors for which speakers disagreed in their color terms often interpolated between the clusters, see Figure 1.

3 Developmental Experiment

Finally, we evaluate the item groupings of young children (4-5, 6-7 year-olds) and adults for a set of stylized fruits, a stimulus set previously used to uncover psychological spaces and latent category distributions [2]. By applying DrLIM to the groupings, we

⁵ <https://www1.icsi.berkeley.edu/wcs/>

