Unfolding mental spaces onto a plane

Frédéric Gosselin (frederic.gosselin@umontreal.ca)
Département de psychologie, Université de Montréal, Canada

Aya Azbane (aya.azbane@umontreal.ca)
Département de psychologie, Université de Montréal, Canada

Nikolaus Kriegeskorte (nk2765@columbia.edu)
Zuckerman Mind Brain Behavior Institute, Columbia University, New York, USA

Ian Charest (ian.charest@umontreal.ca)
Département de psychologie, Université de Montréal, Canada

Abstract

Methods for collecting similarity judgments face significant challenges, including task efficiency and changes in strategy. Here we introduce the *unfolding task*, which enables efficient acquisition of similarity judgments protected against strategy change. On each trial, the participant is presented with a planar graph of the items whose similarities are to be judged. The task is to adjust the graph with mouse drag-and-drop operations such that the lengths of the edges accurately reflect the similarities among the items. Critically, only a subset of the item pairs are connected by edges on each trial, and the similarities among connected items can be conveyed without distortion. All pairwise judgments are acquired across multiple trials. Like the multiarrangement task (MAT), the *unfolding task* uses 2D arrangement of items. However, unfolding the graph avoids the distortion of distances in the 2D arrangement. We demonstrate the accuracy of our method with a simple experiment with known ground-truth similarities. We compared the ability of the *unfolding task* to recover the ground-truth similarities with the current state-of-the-art MAT. The next step will be to test the task with larger, more complex naturalistic stimulus spaces.

Keywords: similarity judgements, psychophysics

Introduction

Pairwise similarity judgments have been widely used in psychology to estimate mental spaces. However, accurately measuring pairwise similarity judgments can be challenging (Attneave 1950; Tsogo, Masson, and Bardot 2000; Kriegeskorte and Mur 2012). Two notable issues with pairwise similarity judgments are: 1) susceptibility to changes in strategy (e.g., a change in the scale of the judgments) and 2) feasibility limited to a relatively small number of items, as a minimum of $\frac{n(n-1)}{2}$ responses are required for *n* items. The multiarrangement task (MAT), which is gaining in popularity, was designed to address these two limitations (Goldstone 1994; Kriegeskorte and Mur 2012). In a nutshell, participants are invited to arrange items inside a circular arena using drag and drop operations in such a way that the distances between them reflect their similarities as accurately as possible. This method is applicable to relatively large sets of items because each of them needs to be placed as little as once. In addition, participants can judge each pair in the context of other pairs and can revisit and adjust earlier placements until they are satisfied with each arrangement. That said, the MAT has limitations of its own. Specifically: a) concerns regarding participants' ability to fulfil the tasks as instructed (e.g., considering all items when placing another); b) when there are more than three items in a trial, the distances in the 2D arrangement provide a distorted (due to dimensionality reduction) reflection of the underlying mental spaces and these effects must be corrected by acquiring multiple arrangements and inferring the underlying similarities.

Here, we introduce and test the *unfolding task*, which combines some of the advantages of pairwise judgements and MAT. As in pairwise judgements, responses reflect similarities without distortion. As in MAT, items are judged in context and a single item placement conveys more than one similarity. In fact, the minimum number of required actions in the *unfolding task* is $\frac{n^2(n-1)}{4n-6}$, roughly half of the required actions in pairwise judgements. In the *unfolding task*, participants are asked to arrange items on a plane exactly as in MAT. However, only the distances between items connected by a line need adjustment to ensure they are proportional to their similarities in the mental space. Notably, connected items form triangles (which implies that two items at most must be considered when placing another), and these triangles are linked together at their edges with the conditions of avoiding loops in the chain of triangles. These constraints define collections of similarities/distances between items in mental spaces that can always be exactly unfolded onto a plane, irrespective of the dimensionality of the mental space.

Methods

We compared the distance measurements obtained using the *unfolding task* and the MAT on a simple experiment with known distances. Fifteen participants took part in this experiment. The stimuli consisted of blue spheres positioned at every vertex of a cube. Participants were instructed to arrange these stimuli (only those connected in the case of the *unfolding task*; see Figure 1) to best match the distances between their blue spheres relative to the cube (the length of either an edge, the diagonal on a face of the cube, or the long diagonal across faces). The tasks were administered on a laptop computer connected to a 28” monitor in a dimly lit room. The order of the tasks was counterbalanced across participants. For the MAT, we used the Meadows implementation (https://mead.ac). This task ended when sufficient evidence had been gathered for all pairs. For the *unfolding task*, we used a custom Python program developed using PyGame. Subjects judged all 28 possible pairwise distances in three *unfolding task* trials, the first two containing ten connections, and the last one, eight.

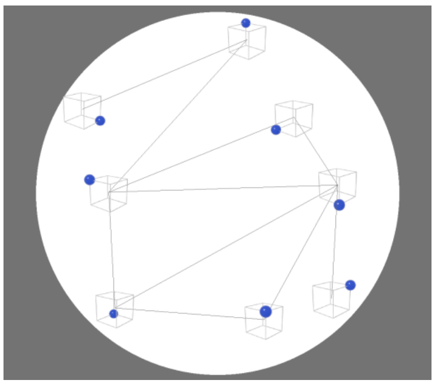

Figure 1. Example trial from the unfolding task with our stimulus set.

Results

We computed ground-truth pairwise distances based on the known coordinates of the blue spheres relative to the cube. Comparing these distances with those estimated using both the MAT and the *unfolding task* showed strong correlations (MAT: mean=.61, std=.03; unfolding task: mean=.68, std=.04; see Figure 2). The top 4 correlations were observed for the *unfolding task* (p=.0495, Monte Carlo method). However, the difference observed in mean correlations between the two tasks did not reach statistical significance ($t(14)$=1.51 for Fisher-transformed rs, p=.15). It's worth noting that this comparison is somewhat biased since, in the *unfolding task*, each distance was measured once, while in the MAT, distances were measured once in the first trial and some were measured again in subsequent trials. The correlations between the distance estimations after the first MAT trial and the ground-truth distances (MAT trial 1: mean=.48, std=.02) were lower than those reported for the unfolding task ($t(14)$=4.20 for Fisher-transformed rs, p<.001). However, this comparison is also biased because the distances measured in the first MAT trial contain distortions that are corrected in subsequent trials (rather well given the performance difference), unlike those measured in the *unfolding task*.

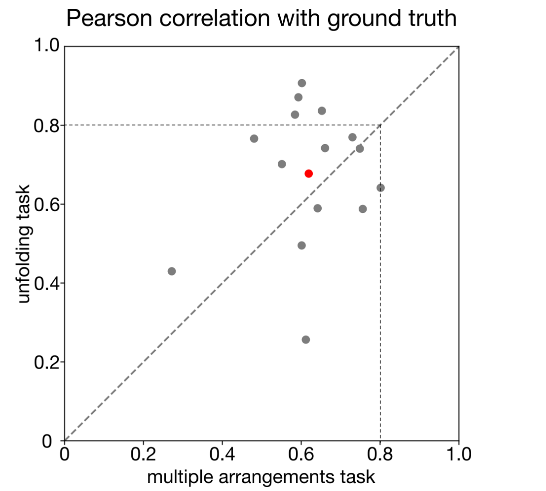

Figure 2: Dots show individual correlation to ground truth for the unfolding and MAT tasks (grey dots). The red dot indicates the mean for both tasks.

Discussion

We introduced the *unfolding task* for making similarity judgments. We compared it with the multiarrangement task for recovering known similarities. Both tasks performed well, with a slight advantage for the *unfolding task*.

In the *unfolding task*, our participants evaluated all 28 pairwise distances/similarities. However, sometimes a subset of all these distances is sufficient to recover the full set. In the *unfolding task*, we measure the stability of the full set of pairwise distances inferred from the subset of measured distances after each trial. Specifically, we use gradient descent to position n points in a high-dimensional space, minimising the error between the points' distances and the distances measured experimentally thus far. The optimised points are used to recover one full set of distances. The task concludes when this set achieves sufficient stability across multiple optimisations. Furthermore, in addition to the criteria enumerated above, the connections between items in the subsequent trial prioritise those items for which the corresponding points in these optimisations exhibit the steepest gradient—those requiring the most pronounced "rigidification".

Future work will include 1) assessing the efficiency of the *unfolding task* for larger stimulus sets, and 2) assessing the *unfolding task* with more complex stimuli (e.g., images of natural scenes; Allen et al. 2022). The multidimensional coordinate system thus recovered could be used to predict brain responses to the same visual scenes.

Acknowledgments

This project was supported by a Canadian Foundation for Innovation John Evans Leaders fund (Charest), NSERC Discovery grants (Gosselin and Charest), and a Courtois Chair in computational neuroscience (Charest).

References

Allen, E. J., St-Yves, G., Wu, Y., Breedlove, J. L., Prince, J. S., Dowdle, L. T., Nau, M. et al. 2022. "A Massive 7T fMRI Dataset to Bridge Cognitive Neuroscience and Artificial Intelligence." *Nature Neuroscience* 25 (1): 116–26.

Attneave, F. 1950. "Dimensions of Similarity." *The American Journal of Psychology* 63 (4): 516–56.

Goldstone, R. 1994. "An Efficient Method for Obtaining Similarity Data." *Behavior Research Methods, Instruments, & Computers: A Journal of the Psychonomic Society, Inc* 26 (4): 381–86.

Kriegeskorte, N., and Mur. M. 2012. "Inverse MDS: Inferring Dissimilarity Structure from Multiple Item Arrangements." *Frontiers in Psychology* 3.

Tsogo, L., Masson, M. H., and Bardot, A. 2000. "Multidimensional Scaling Methods for Many-Object Sets: A Review." *Multivariate Behavioral Research* 35 (3): 307–19.