What Mental States? Exploring How Dimensionality Reduction Might Contribute to the Refinement of Cognitive Models

Kenneth C. Whang
Department of Computer Science
Bryn Mawr College
Bryn Mawr, PA 19010
kcwhang@brynmawr.edu

Abstract

Questions in cognitive neuroscience are often framed in terms of correspondences between known types: How is brain state X related to cognitive state Y? What are the correlations or mappings between particular structures and functions? Such framings are well suited for confirmatory testing of coarse-grained hypotheses. They are not necessarily informative, however, for the purpose of exploring finer physical and functional structure. To the contrary, physical states are typically aggregated over anatomical regions of interest, while tasks are designed to optimize one or a few functional contrasts of interest rather than to cover a fuller behavioral or cognitive range.

Recent advances in analysis and modeling of neuroimaging, multiunit, and spike timing data suggest significant opportunities for the space of physical brain states to be more finely parsed [1]-[3]. I am interested in the complementary problem of how finer models of behavioral and cognitive states can be developed, and how neural data can drive such refinements. I view this as a problem of how best to use what we know (from empirical measurement) of the structure of neural activity to help fill in what we don't know (or are less sure of) of the structure of behavioral or cognitive representations.

Multidimensional scaling (MDS) is a dimensionality reduction method that has traditionally been used for an analogous class of problems, where pairwise data on

perceptual similarities are the basis for deriving the structure of perceptual representations. MDS fits a matrix of pairwise proximities to a configuration of points in space, such that interpoint distances in the configuration best fit the corresponding proximities [4]-[5]. The arrangement of points in the configuration space reflects the dimensions that are used in the perceptual judgment. In that way, variables of interest are derived from data rather than being assumed in a prior model. MDS has been used to characterize representations of color, phonemes, and many other perceptual domains [6]. Methods have also been developed to handle non-metric data and individual differences among observers [7]-[8].

I will present examples that suggest that MDS can be productively applied to analyses of neural and behavioral data from perceptual and motor tasks, where differences in neural activity across task conditions are treated as similarities for neural "observers," and individual difference models account for variations in neural tuning. In a simple reaching task, MDS yields results that agree with known neurophysiology [9]. In slightly more complex cases, MDS pulls out relationships among dozens of task conditions, indicating task variables that are related to neural activity differences, and discovering features relevant to specific task contexts and individual differences.

These examples are intended to present a plausibility case and to invite further discussion. MDS is one of many approaches to dimensionality reduction that could potentially be applied to neural data [10]-[11]: can its good features, such as proximity mapping and individual difference modeling, be incorporated into newer methods or vice versa? What are the best ways to model individual differences and time? Can individual difference models be used as a tool to facilitate comparisons of data across experiments or across data modalities? Can proximities be measured directly (e.g., adaptation)? How might methods aimed at refining physical and cognitive models interact? What latitude might these methods present for the design of new experiments? I will present some preliminary thoughts on these issues.

Acknowledgments

I thank Apostolos Georgopoulos for introducing me to multidimensional scaling and for his data on 3-dimensional reach, and the Emergent Systems discussion group of Bryn Mawr and Swarthmore Colleges for helpful discussions of these ideas.

References

[1] Kamitano, Y. & Tong, F., 2005. Decoding the visual and subjective contents of the human brain. Nature Neuroscience 8(5): 679-685.

[2] Haynes, J.-D. & Rees, G., 2006. Decoding mental states from brain activity in humans. Nature Reviews Neuroscience 7: 523-534.

[3] Deneve, S., 2006. Submitted manuscript.

[4] Cox, T. F. & Cox, M. A. A., 2001. Multidimensional scaling. Chapman & Hall.

[5] Borg, I. & Groenen, P., 1997. Modern multidimensional scaling. Springer.

[6] Shepard, R. N., 1980. Multidimensional scaling, tree-fitting, and clustering. Science 210(4468): 390-398.

[7] Kruskal, J. B., 1964. Multidimensional scaling by optimizing goodness-of-fit to a nonmetric hypothesis. Psychometrika 29: 1-27.

[8] Caroll, J. D. & Chang, J. J., 1970. Analysis of individual differences in multidimensional scaling via an *n*-way generalization of "Eckart-Young" decomposition. Psychometrika 35: 283-319.

[9] Schwartz, A. B., Kettner, R. E., & Georgopoulos, A. P., 1988. Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. relations between single cell discharge and direction of movement. Journal of Neuroscience 8(8): 2913-2927.

[10] Duda, R. O., Hart, P. E., & Stork, D. G., 2000. Pattern classification. Wiley-Interscience.

[11] Saul, L. K., Weinberger, K. Q., Sha, F., Ham, J., & Lee, D. D., 2006. Spectral methods for dimensionality reduction. To appear in Chapelle, O., Schoelkopf, B., and Zien, A. (eds.), Semisupervised Learning. MIT Press.