Check for updates

Revealing trajectories of the mind via non-linear manifolds of brain activity

This work involved the design of a multi-view manifold learning algorithm that capitalizes on various types of structure in high-dimensional time-series data to model dynamic signals in low dimensions. The resulting embeddings of human functional brain imaging data unveil trajectories through brain states that predict cognitive processing during diverse experimental tasks.

The mission

Studying the brain in action is a powerful way to address age-old questions about the healthy human mind. Although cutting-edge technologies for measuring brain activity, like functional magnetic resonance imaging (fMRI), create computationally compelling data for understanding the dynamics of brain function, these data are riddled with spatiotemporal noise. This noise originates from both the method of measurement, as fMRI captures an indirect, vascular response that is delayed several seconds from neuronal activity, and the way that the brain represents information in spatially diffuse, redundant population codes (a way of representing information about a stimulus by the simultaneous, correlated activity of a large set of stimulus-responsive neurons)[1]. Together, these properties of brain imaging data indicate that brain activity encodes diverse states into a small number of activity patterns; that is, they could be captured as trajectories along a low-dimensional manifold. Furthermore, higher-order cognitive processes such as memory and narrative comprehension are represented in patterns of brain activation that unfurl over multiple timescales. The complex interactions between attributes of brain data and the ways that the brain encodes information highlight the need for an algorithm that jointly considers spatial and temporal signal structures to uncover the latent space of dynamic cognitive processes.

The solution

We designed a new manifold learning algorithm, called temporal potential of heat diffusion for affinity-based transition embedding (T-PHATE), for analyzing high-dimensional data that represent temporally dynamic processes. The T-PHATE algorithm learns two 'views' into the data: one view models the data's time-varying properties, and the other view learns the data's geometry via the PHATE manifold learning method[2]. These views are then combined to learn a manifold that represents both dynamic and geometric data properties (Fig. 1). We test this approach using data from three brain imaging (fMRI) experiments. Two experiments measured the brain during naturalistic tasks (as participants watched standard full-length movies)[3], and the other experiment had participants view static, random images[4]. These data capture the brain during complex cognitive tasks in which time plays an important role in understanding the movie, as well as a simpler task in which time does not interact with information processing.

We compared fMRI brain activity data embedded with T-PHATE with data embedded using several state-of-the-art dimensionality reduction algorithms, such as UMAP (uniform manifold approximation and projection for dimension reduction) and T-SNE (t-distributed stochastic neighbor embedding), as benchmarks against which to compare T-PHATE's performance. These embeddings were computed individually for each participant. Embeddings of fMRI activity with T-PHATE visually and quantitatively showed the clearest trajectories through the latent space[5]. The dynamic structure captured by T-PHATE generalized across study participants who were watching the same movie, even though T-PHATE embeddings were learned for each brain independently. Furthermore, the T-PHATE dynamics were related to how people consciously identify dynamics behaviorally during the tasks involving a continuous narrative stimulus, such as watching movies. In a separate task with no inherent dynamics, where participants viewed static images in no specific order, T-PHATE still revealed task-related structure by clustering the brain states related to certain object categories in proximal parts of the latent space. This shows the method's general-purpose ability to denoise high-dimensional time-series data.

Future directions

As big data become more ubiquitous, methods for denoising and uncovering latent structure in such data are increasingly valuable. T-PHATE has promising applications in other kinds of high-throughput time-series data, such as longitudinal health data, developmental trajectories, climate change data, economic trends, and language evolution.

At present, our results only cover fMRI data, which have known temporal properties that enable autocorrelation modeling (Fig. 1). We hope to extend the applications of T-PHATE to other forms of time-series data, which might require different approaches to modeling time in the temporal view of T-PHATE.

Given the noise inherent in brain imaging data at the single-person level, our approach holds clear promise for investigating the signatures of cognitive processes within individual brains. Furthermore, the individual-tailored nature of this approach will aid in the development of important applications for brain–computer interfaces. We are excited to expand along these avenues in the future.

Erica L. Busch & Smita Krishnaswamy
Yale University, New Haven, CT, USA

This is a summary of:
Busch, E. L. et al. Multi-view manifold learning of human brain-state trajectories. *Nat. Comput. Sci.* https://doi.org/10.1038/s43588-023-00419-0 (2023).

Published online: 27 March 2023

EXPERT OPINION

"This work presents an interesting approach for dimensional reduction in noisy spatiotemporal biological data. The researchers test their scheme on both synthetic data and labeled real fMRI data, which allows for an evaluation of the method performance. In relation to a series of benchmarks, the authors' method far outperforms other methods such as principal component analysis (PCA) and UMAP." **Vahid Shahrezaei, Imperial College London, London, UK.**

FIGURE

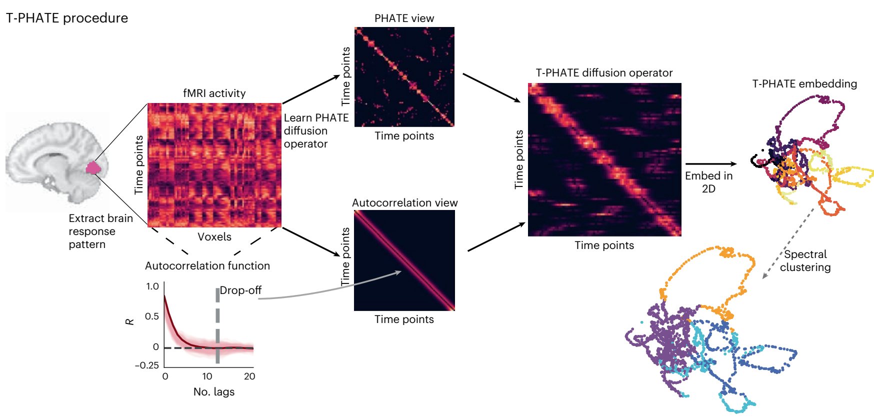

Fig. 1 | Multi-view manifold learning approach with T-PHATE. Voxel-wise fMRI activity data are extracted from a brain region of interest and used to learn two views of the data. The PHATE view captures the manifold geometry of the data, and the autocorrelation view learns the temporal dynamics of the data via an autocorrelation function, which computes the correlation coefficient of each voxel's activity against lagged versions of itself. These views are combined into a single T-PHATE diffusion operator, which is then embedded into low dimensions. This latent representation visually highlights trajectories through brain states as they emerge over time. © 2023, Busch, E. L. et al.

BEHIND THE PAPER

This project began with our interest in developing low-dimensional representations of brain activity for applications in non-invasive brain–computer interfaces. The approaches we initially tried were hampered either by the noise inherent in single-person fMRI data or by poorly modeled dynamic information. Our first few attempts at the dual-view manifold learning approach used a fixed-width temporal view, which essentially applied smoothing uniformly over a fixed number of time points. This approach outperformed a single-view algorithm but only for some brain regions. We finally settled on an approach that learned the width of the temporal view from the data itself. This approach is far more robust than our previous attempts and enables T-PHATE to adjust flexibly to the temporal integrative properties of different brain regions, different temporal sampling rates, and even different brain imaging modalities. **E.L.B.**

REFERENCES

1. Turk-Browne, N. B. Functional interactions as big data in the human brain. *Science* **342**, 580–584 (2013).
   A review article that presents the properties and complexity of functional neuroimaging data.
2. Moon, K. R. et al. Visualizing structure and transitions in high-dimensional biological data. *Nat. Biotechnol.* **37**, 1482–1492 (2019).
   This paper presents the PHATE algorithm, which our study extends.
3. Chen, J. et al. Shared memories reveal shared structure in neural activity across individuals. *Nat. Neurosci.* **20**, 115–125 (2017).
   This paper describes one of the data sets used in our study.
4. Hanke, M. et al. A *studyforrest* extension, simultaneous fMRI and eye gaze recording during prolonged natural stimulation. *Sci. Data* **3**, 160092 (2016).
   This paper describes one of the data sets used in our study.
5. Baldassano, C. et al. Discovering event structure in continuous narrative perception and memory. *Neuron* **95**, 709–721 (2017).
   This paper presents the event segmentation framework for modeling neural dynamics.

FROM THE EDITOR

"The authors propose an approach for reducing the dimensionality of fMRI data in naturalistic paradigms. The work stood out to me since the method accounts for temporal autocorrelation to improve data visualization, classification, and event segmentation of the data, which is an improvement over existing methods in the field." **Ananya Rastogi, Associate Editor, *Nature Computational Science.***