Multivariate Prediction of Human Behavior in Task fMRI

Alessandra Camassa
Salk Institute for Biological Studies
acamassa@salk.edu

Joseph Park
Okinawa Institute of Science and Technology
josephpark@ieee.org

Margot Wagner
Salk Institute for Biological Studies, UCSD
mwagner@ucsd.edu

Stanislav Smirnov
Université de Genève
stanislav.smirnov@unige.ch

Terrence J. Sejnowski
Salk Institute for Biological Studies, UCSD
terry@salk.edu

Gerald M. Pao *
Okinawa Institute of Science and Technology
gerald.pao@oist.jp

Abstract

A central problem of research at the intersection between neuroscience, machine learning and artificial intelligence is to understand and model the relationship between brain activity and behavior in humans. Here we propose an algorithm able to identify the brain regions that are causally relevant to a specific behavior and use them to build a predictive model for time series data. Our goal is to predict the behavioral output of human subjects executing a task within the fMRI scan using the hemodynamic response of different brain regions as an indirect measure of neural activity. We use the Simplex prediction algorithm and a new iterative feature evaluation method: this approach is based on state space reconstruction, and the variables chosen to build a low dimensional representation of the behavior are selected according to their relationship to the target time series (i.e. the behavioral time series). To discover these relationships, we use convergent cross mapping (CCM), a Takens theorem based causal inference method that can detect nonlinear causation between time series. We iteratively select the variables that enhance forecasting skill and with convergent causal coupling to the target to construct a multivariate local manifold able to predict the behavior through multivariate Simplex projection. This method facilitates the integration of empirical observations of brain activity into forecasting of realistic behaviors. Our approach allows to keep a direct correspondence between observations and model variables while optimizing the prediction accuracy.

1 Introduction

To understand and model human behavior is one of the central problems of research in neuroscience, machine learning and artificial intelligence. In this context, there is a rising need for explainable algorithms able not only to provide a good forecast of behavioral time series, but also to identify the key variables that are contributing to the behavior dynamics. The brain may be seen as a nonlinear dynamical system, where brain functions depend on interactions among a large number of neural populations, regulated by causal relationships and distributed structural and functional connectivity.

*corresponding author

NeurIPS 2024 Workshop on Behavioral Machine Learning.

Organized spontaneous activity is prevalent in many brain regions, where encoding for a particular stimulus or behavior is typically distributed across a large population of neurons. Recent evidence suggests that the activity of these numerous neuronal populations can actually be described by low-dimensional dynamics. Dimensionality reduction methods, such as PCA, reveal that a small number of components derived from the original dynamical system's activity contain the majority of the information [8, 20, 21, 5, 7, 33]. Thus, low-dimensional embeddings can accurately reconstruct the system dynamics [23, 26]. Data-driven approaches to model and generate biologically plausible dynamical data require extensive, standardized, and high-quality observations to ensure accuracy and generalizability of results. Magnetic resonance imaging (MRI and fMRI) data play a crucial role in the context of understanding human behavior dynamics, offering non-invasive techniques for structural reconstruction and functional studies of neural ensembles and their connectivity during behavior (resting state or task). Here we propose an empirical method for the prediction of human behavior in task fMRI data based on the Takens theorem [35] and it's extension, the generalized Takens theorem[13]. This approach leverages brain activity and manifold learning techniques providing an observable low dimensional representation of the behavior which allows direct identification of variables that contribute to the dynamics of the target time series. In this way, it is possible to obtain a better understanding of brain behavior dynamics and integrate this knowledge in a predictive model.

2 Methods

2.1 Data

We use neuroimaging data from the ABCD Study at baseline [18, 6], our sample includes $N = 10$ human subjects. The minimally preprocessed fMRI data from the ABCD Study undergoes several preprocessing steps, including image registration, motion correction, gradient distortion correction, and alignment [15, 12, 1, 28, 17]. Here, we use the Stop Signal Task (SST) [19] fMRI data, designed to test different cognitive processes including decision making and response time (Figure S1), and the Emotional N-back task (EN-back), which involves both memory and emotion regulation processes (Figure S4). The tasks are programmed to accept input from the dominant hand to a detection device. The motor output in this fMRI tasks is not directly available in the form of a time series. Consequently, we process the data in order to be able to extract this signal directly form the brain scans of the participants. From each task fMRI session, using independent component analysis (ICA) we extract a number of time series representing the brain activity in different cortical and subcortical regions (Figure S2A and B). We tested this approach using a variable number of independent components (20, 50, and 100 components). Alternatively, seed-based methods can be applied to segment the brain in a number of Regions of Interest (ROIs) defined by anatomical atlases. Here we used SynthSeg+ deep learning algorithm for brain segmentation and cortical parcellation [3, 4]. The independent component or ROI corresponding to the brain stem and spinal cord is manually identified by comparison with anatomical atlases and the corresponding time series is used as a surrogate for actual movement, i.e. a fictive behavioral time series. This component is used as the target variable and the independent components or ROIs extracted from the rest of the brain are used as candidate variables to be embedded into a multivariate manifold able to predict the target through attractor reconstruction [27, 22]. The time series were z scored but not filtered in order to avoid altering the system dynamics.

2.2 Manifold Dimensional Expansion (MDE) algorithm

We propose an algorithm that iteratively selects brain regions with causal relationships to behavioral time series and reconstructs a low-dimensional manifold for accurate behavior prediction within task fMRI data. The Manifold Dimensional Expansion (MDE) algorithm iteratively applies the Simplex projection method [30] to time series data to identify and select variables that improve prediction accuracy (Figure1). The algorithm runs independently on each subject's data, taking as input the ICs time series extracted from the fMRI data with shape ($n_{\text{timesamples}}$, n_{ICs}). The same algorithm can be run using time series from different anatomical ROIs extracted from the imaging data through segmentation. The time series data are divided into a training set and test set, where half of the time samples are used for the training and the other half as test set as shown in Figure 2D-F. Initially, a univariate Simplex model is created building a delay embedding of the target variable (e.g. x(t), x(t-1), x(t-2)) using the training set (Figure2E). The accuracy of the univariate prediction is computed as the correlation coefficient between the model's prediction and the actual observations in the test

set. Then, a multivariate embedding is created substituting the delayed versions of the target variable (e.g. x(t-1), x(t-2) in Figure 2E) with different variables (e.g. y(t) and z(t) in Figure 2F) having a causal relationship with the target, using principles from the generalized Takens theorem [13]. The algorithm then tests additional variables to be potentially included in the predictive model by incrementally adding them to the model, recalculating the correlation between observed and predicted values after each addition. The first inclusion criterion is the improvement of the prediction accuracy, as measured by an increase in the correlation coefficient between observed and predicted. Secondly, a convergence test is performed using Convergent Cross Mapping (CCM) [32] to determine whether the added variable is causally linked to the original target variable. The variable is retained and added to the model if the cross-mapping prediction skill of the target improves with the library size ensuring causality between the target and the variable; otherwise, it is discarded. This process continues until no further variables can improve the prediction, resulting in a set of optimally selected variables that enhance the predictive model's performance. In the multivariate state space reconstruction, predictions are made using again the Simplex projection, a nearest neighbor algorithm. Briefly, a "simplex" is constructed around the query point k nearest neighbors projected n points into the future. The prediction is computed as the weighted sum of the neighbors, and the number of neighbors can be optimized to obtain better performance. Alternatively, the MDE algorithm can be implemented using S-Map forecasting algorithm [31]. Similar to Simplex, S-Map prediction is based on locally weighted nearest neighbors. However, S-Map fits a local linear model to the data, assigning different weights to each neighbor according to its distance from the current state. This adaptive approach makes S-Map more effective in handling noisy data and capturing dynamic variations in the system.

3 Results

Following this method, a local manifold for multivariate prediction is constructed as an E dimensional attractor built around the target variable, having optimal embedding dimension equal to E. E-1 is the number of variables that are embedded together with the target variable to obtain the attractor. Different combinations of time series data from interacting components can form a multivariate embedding. Here, the optimal set of variables is selected according to the strength of their causal interaction with the target and their ability to improve the prediction accuracy (FigureS2C-E). Noisy components are automatically eliminated from consideration because they will not contain any information and so will show no causal link with the rest of the components. This local manifold is used for forecasting (Figure2F).

We ran the algorithm on SST fMRI data transformed into 20 Independent Components using ICA. We compared our results with predictions made using a GLM model, the state-of-the-art technique for analyzing task fMRI data. An example of the improvement from univariate prediction to multivariate, with sequentially added variables selected by the algorithm for one subject, is shown in Figure 2B, more details are provided in FigureS1. Additionally, Figure 2 presents, for an example subject, the prediction obtained with the GLM (Figure 2D), the delay embedding and univariate prediction (Figure 2E), and the manifold constructed by our method along with the corresponding target prediction (Figure 2F). At the population level, the prediction accuracy achieved with the GLM was 31.05% ± 22.4% (mean ± std), univariate prediction accuracy was 60.55% ± 12.83%, while MDE resulted in an accuracy of 81.02% ± 7.15% (Figure 2A). Hence, MDE predictions were on average 25.5% more accurate than univariate predictions and 61.68% more accurate than those made by the GLM. On average, 4.5 ± 1 variables were selected for multivariate prediction, providing an estimate of the optimal embedding dimension of the target variables, which suggests that behavior is low dimensional (Figure 2C). We tested MDE with S-Map prediction algorithm on the same data (Figure S3). The S-Map version of MED achieved an accuracy of 81.33% ± 8.76%, showing a slight improvement over the Simplex method. However, this difference was not statistically significant (Table 3).

Additionally, we applied the MDE algorithm to the EN-back task fMRI data from the ABCD study, which was segmented into 98 cortical and subcortical regions of interest (ROIs) based on the Desikan-Killiany atlas. Our analysis focused on predicting the time series of two target ROIs: the brainstem (BSM), serving as a proxy for motor output, and the amygdala (Amy), which plays a key role in emotional processing during cognitive tasks. MDE achieved an accuracy of 56.42% ± 20.85% for predicting the BSM time series and 68.86% ± 16.71% for predicting the Amy time series across N = 10 subjects (Figure S5A, C). Importantly, different sets of ROIs were selected for predicting activity in the BSM and Amy. For BSM, the selected regions were linked to motor control, cognitive functions,

and visual processing, particularly in face and object recognition. In contrast, the ROIs associated with Amy prediction included areas involved in memory, emotional regulation, sensory integration, and attention, as well as visual processing (Figure S5B). These findings are consistent with the cognitive demands of the EN-back task, which engages visual processing, emotional regulation, memory, and motor control.

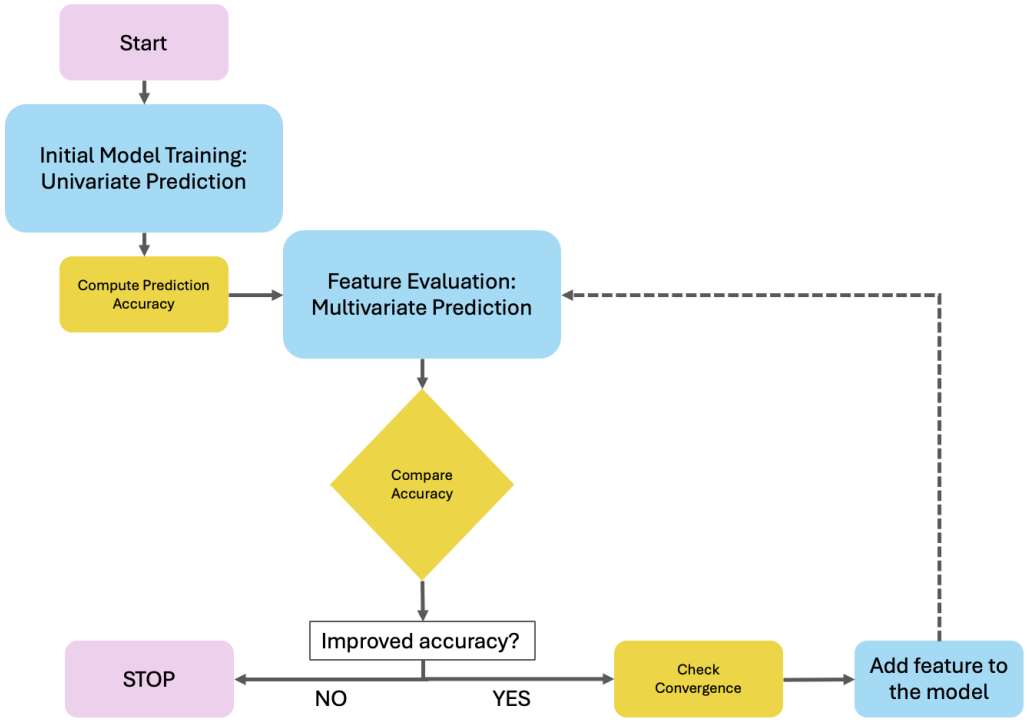

Figure 1: Manifold Dimensional Expansion (MDE) algorithm starts with an initial model training, performing a univariate prediction of the target variable using a delay embedding and the Simplex projection algorithm. Calculate the initial prediction accuracy as the correlation coefficient between the model's prediction and the actual observations. Iterate through each potential feature in the dataset and add it to the model along with the target variable to perform a multivariate prediction. Train the model with the new feature and calculate the prediction accuracy for each combination of the feature and the target variable. If the prediction accuracy is improved and the variable passes the convergence check, the new feature is added to the model. The process continues until no additional feature improves the model's accuracy or all possible features have been evaluated.

4 Discussion

We presente here a method to model and predict human behavior with task fMRI data. Our method enables the detection of a set of features for predicting a target variable. When applied to human behavior prediction, this approach not only produces accurate forecasts of future behavior but, more importantly, helps identify the brain regions involved in generating that behavior, i.e. the variables causally linked to the target. This offers a powerful tool for disentangling brain-behavior dynamics. At the same time, its applications are not confined to behavior prediction but can be extended to any brain feature. Specifically, this approach can be used to predict other Independent Component time series from the brain or can be applied to neuroimaging data segmented into anatomical regions of interest (ROIs) as well as raw voxelwise fMRI data. A high-performance implementation of Empirical Dynamic Modeling (EDM), kEDM [34], efficiently handles high-dimensional datasets and can run on both CPUs and GPUs. It is particularly well-suited for high-dimensional data, such as voxelwise fMRI, and it has been successfully applied to datasets containing over 100,000 time series, each with

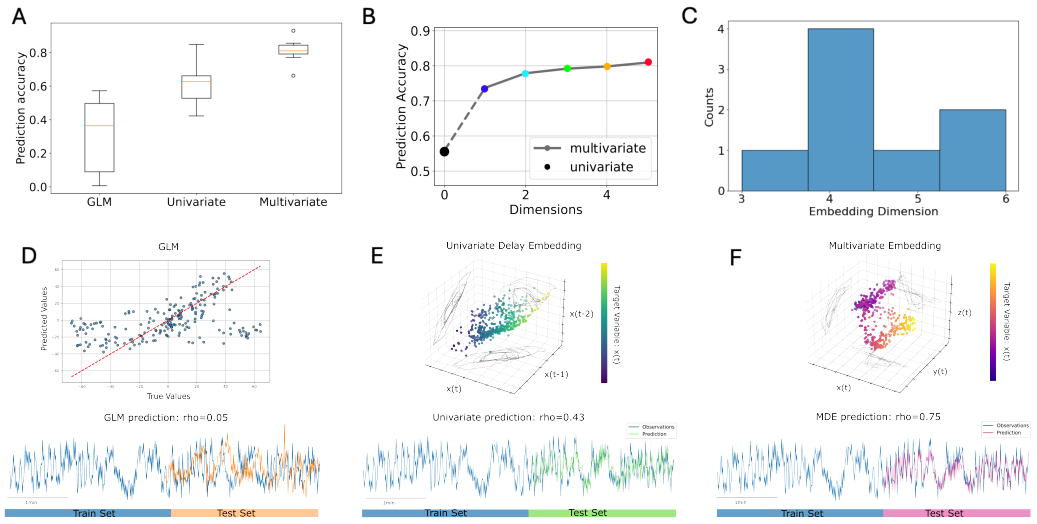

Figure 2: (A) Boxplot showing the increase in prediction accuracy from GLM, univariate, to multivariate methods, illustrating the enhanced reconstruction of behavioral time series. Differences are statistically significant (p-value<0.05, see appendix Table 1). (B) Prediction accuracy for the univariate model (black dot) and its improvement as each selected variable is added (see Fig. S2 for details). (C) Embedding dimension of the target variable for all subjects in this study (N=10), estimated using our method. (D) Example of GLM prediction, demonstrating an accuracy of 5% on the test set. (E) Example of univariate delay embedding with a prediction accuracy of 43% on the test set. (F) Example of multivariate embedding achieving a prediction accuracy of 75% on the test set. Panels (C)-(E) are derived from the same subject and target variable.

more than 8,000 samples. A major advantage of this method is that it does not rely on latent variables, ensuring a direct correspondence between observed variables and model features. Compared to other machine learning and deep learning models, it allows for clearer result interpretation and facilitates experimental validation. Additionally, the CCM convergence check conducted in this analysis ensures that the variables selected by the algorithm are causally linked to the behavior, offering a mechanistic insight into the multivariate manifold surrounding the target, something that cannot be achieved with other methods. We demonstrated that our prediction method outperforms GLM, which is commonly used to model the relationship between brain activity and external stimuli or tasks in fMRI data acquisition. The GLM assumes a linear relationship between the behavior and the blood-oxygen-level-dependent (BOLD) signal, limiting its ability to capture complex, dynamic brain interactions. In contrast, our approach, based on the non-linear Simplex forecasting algorithm, allows for the modeling of more complex brain dynamics, thereby increasing the prediction accuracy compared to linear models. One limitation of our approach is its sensitivity to noise, as noisy data can distort the nearest neighbor search, reducing prediction accuracy. To handle noisy data, an equivalent version of this algorithm can be implemented using the S-Map forecasting algorithm [31]. The S-Map forecasting algorithm uses locally weighted linear approximations to make predictions that can down weight distant or less relevant data points, allowing it to focus on nearby points that are more representative of the underlying system. This flexibility helps S-Map handle noise better than the Simplex algorithm [25]. Our model performs best on low-dimensional, stationary data or long time series where the full range of the system's dynamics is represented in the training set. However, as data size and dimensionality increase, the algorithm becomes computationally expensive but sill scales linearly.

References

[1] Jesper LR Andersson, Stefan Skare, and John Ashburner. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *Neuroimage*, 20 (2):870–888, 2003.

[2] Deanna M Barch et al. Function in the human connectome: Task-fmri and individual differences in behavior. *NeuroImage*, 80:169–189, 2013.

[3] Benjamin Billot, Colin Magdamo, Steven E Arnold, Sudeshna Das, and Juan Eugenio Iglesias. Robust segmentation of brain mri in the wild with hierarchical cnns and no retraining. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 538–548. Springer, 2022.

[4] Benjamin Billot, Douglas N Greve, Oula Puonti, Axel Thielscher, Koen Van Leemput, Bruce Fischl, Adrian V Dalca, Juan Eugenio Iglesias, et al. Synthseg: Segmentation of brain mri scans of any contrast and resolution without retraining. *Medical image analysis*, 86:102789, 2023.

[5] Connor Brennan and Alexander Proekt. A quantitative model of conserved macroscopic dynamics predicts future motor commands. *Elife*, 8:e46814, 2019.

[6] Betty Jo Casey, Tariq Cannonier, May I Conley, Alexandra O Cohen, Deanna M Barch, Mary M Heitzeg, Mary E Soules, Theresa Teslovich, Danielle V Dellarco, Hugh Garavan, et al. The adolescent brain cognitive development (abcd) study: imaging acquisition across 21 sites. *Developmental cognitive neuroscience*, 32:43–54, 2018.

[7] Rishidev Chaudhuri, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature neuroscience*, 22(9):1512–1520, 2019.

[8] Mark M Churchland, John P Cunningham, Matthew T Kaufman, Justin D Foster, Paul Nuyujukian, Stephen I Ryu, and Krishna V Shenoy. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, 2012.

[9] Nancy Cohen et al. Processing facial emotion in children and adolescents with anxiety and depression. *Developmental Neuroscience*, 2016.

[10] Nancy Cohen et al. Emotion regulation and attention control in children: Relationships to anxiety and depression. *Journal of Child Psychology and Psychiatry*, 2016.

[11] Mark I Conley et al. Emotion regulation in adolescents: Neural mechanisms and behavioral outcomes. *Psychological Science*, 28:1584–1593, 2017.

[12] Robert W Cox. Motion and functional mri. In *Informal notes for Boston'96 Workshop on Functional MRI*, 1996.

[13] Ethan R Deyle and George Sugihara. Generalized theorems for nonlinear state space reconstruction. *Plos one*, 6(3):e18295, 2011.

[14] Dylan G Gee et al. Developmental shifts in the balance between prefrontal and subcortical control of emotion. *The Journal of Neuroscience*, 33(10):4584–4593, 2013.

[15] Donald J Hagler Jr, SeanN Hatton, M Daniela Cornejo, Carolina Makowski, Damien A Fair, Anthony Steven Dick, Matthew T Sutherland, BJ Casey, Deanna M Barch, Michael P Harms, et al. Image processing and analysis methods for the adolescent brain cognitive development study. *Neuroimage*, 202:116091, 2019.

[16] Todd A Hare et al. Biological substrates of emotional reactivity and regulation in adolescents during an emotional go-nogo task. *Biological Psychiatry*, 63(10):927–934, 2008.

[17] Jorge Jovicich, Silvester Czanner, Douglas Greve, Elizabeth Haley, Andre van Der Kouwe, Randy Gollub, David Kennedy, Franz Schmitt, Gregory Brown, James MacFall, et al. Reliability in multi-site structural mri studies: effects of gradient non-linearity correction on phantom and human data. *Neuroimage*, 30(2):436–443, 2006.

[18] Nicole R Karcher and Deanna M Barch. The abcd study: understanding the development of risk for mental and physical health outcomes. *Neuropsychopharmacology*, 46(1):131–142, 2021. doi: https://doi.org/10.1038/s41386-020-0736-6.

[19] Gordon D Logan. On the ability to inhibit thought and action: A users' guide to the stop signal paradigm. *American Psychological Association*, 1994.

[20] Ryan J Low, Sam Lewallen, Dmitriy Aronov, Rhino Nevers, and David W Tank. Probing variability in a cognitive map using manifold inference from neural dynamics. *BioRxiv*, page 418939, 2018.

[21] Niru Maheswaranathan, Alex Williams, Matthew Golub, Surya Ganguli, and David Sussillo. Universality and individuality in neural dynamics across large populations of recurrent networks. *Advances in neural information processing systems*, 32, 2019.

[22] N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw. Geometry from a time series. *Phys. Rev. Lett.*, 45:712–716, Sep 1980. doi: 10.1103/PhysRevLett.45.712. URL `https://link.aps.org/doi/10.1103/PhysRevLett.45.712`.

[23] Chethan Pandarinath, Daniel J O'Shea, Jasmine Collins, Rafal Jozefowicz, Sergey D Stavisky, Jonathan C Kao, Eric M Trautmann, Matthew T Kaufman, Stephen I Ryu, Leigh R Hochberg, et al. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature methods*, 15(10):805–815, 2018.

[24] Marius V Peelen and Paul E Downing. The neural basis of visual object recognition: Evidence for domain specificity. *Trends in Cognitive Sciences*, 9(8):349–355, 2005.

[25] Charles T Perretti, George Sugihara, and Stephan B Munch. Nonparametric forecasting outperforms parametric methods for a simulated multispecies system. *Ecology*, 94(4):794–800, 2013.

[26] Patrick T Sadtler, Kristin M Quick, Matthew D Golub, Steven M Chase, Stephen I Ryu, Elizabeth C Tyler-Kabara, Byron M Yu, and Aaron P Batista. Neural constraints on learning. *Nature*, 512(7515):423–426, 2014.

[27] T. D. Sauer. Attractor reconstruction. *Scholarpedia*, 1(10):1727, 2006. doi: 10.4249/scholarpedia.1727. revision #91017.

[28] Stephen M Smith, Mark Jenkinson, Mark W Woolrich, Christian F Beckmann, Timothy EJ Behrens, Heidi Johansen-Berg, Peter R Bannister, Marilena De Luca, Ivana Drobnjak, David E Flitney, et al. Advances in functional and structural mr image analysis and implementation as fsl. *Neuroimage*, 23:S208–S219, 2004.

[29] Leah H Somerville et al. The neural basis of emotional reactivity and regulation: From adolescence through adulthood. *Psychological Science*, 22(10):1326–1335, 2011.

[30] George Sugihara and Robert M May. Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series. *Nature*, 344(6268):734–741, 1990.

[31] George Sugihara and Robert M May. Nonlinear forecasting for the classification of natural time series. *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences*, 348(1688):477–495, 1994.

[32] George Sugihara, Robert May, Hao Ye, Chih-hao Hsieh, Ethan Deyle, Michael Fogarty, and Stephan Munch. Detecting causality in complex ecosystems. *science*, 338(6106):496–500, 2012.

[33] Satohiro Tajima, Toru Yanagawa, Naotaka Fujii, and Taro Toyoizumi. Untangling brain-wide dynamics in consciousness by cross-embedding. *PLoS computational biology*, 11(11):e1004537, 2015.

[34] Keichi Takahashi, Wassapon Watanakeesuntorn, Kohei Ichikawa, Joseph Park, Ryousei Takano, Jason Haga, George Sugihara, and Gerald M Pao. kedm: a performance-portable implementation of empirical dynamic modeling using kokkos. In *Practice and Experience in Advanced Research Computing*, pages 1–8. Practice and Experience in Advanced Research Computing, 2021.

[35] F. Takens. Detecting strange attractors in turbulence. In D. A. Rand and L. S. Young, editors, *Dynamical Systems and Turbulence, Lecture Notes in Mathematics*, volume 8, pages 366–381. Springer-Verlag, 1981.

[36] Nim Tottenham et al. The nimstim set of facial expressions: Judgments from untrained research participants. *Psychiatry Research*, 168:242–249, 2009.

A Appendix / supplemental material

A.1 Statistical analysis

In our statistical analysis, we utilized the Friedman test to assess the differences in accuracy between the three models (GMN, Univariate and Multivariate prediction model). Following this, we conducted pairwise comparisons using the Wilcoxon signed-rank test to determine the significance of differences between each pair of models. The p-values obtained with Wilcoxon signed-rank test for the SST task fMRI data are shown in Table 1, and for EN-back task fMRI data in Table 2. Table 3 shows the statistical comparison of the results obtained with Simplex MDE versus SMap MDE shown in Figure S3.

Table 1: Statistical comparison of models prediction accuracy on SST task fMRI data

Comparison of Models	
Comparison	p-value
GLM vs Univariate	0.015
GLM vs Multivariate	0.007
Univariate vs Multivariate	0.007

Table 2: Statistical comparison of models prediction accuracy on EN-back task fMRI data

Comparison of Models	
Comparison	p-value
Multivariate:BSM vs Univariate:BSM	0.041
Multivariate:Amy vs Univariate:Amy	0.001
Univariate:BSM vs Univariate:Amy	0.160
Multivariate:Amy vs Multivariate:BSM	0.064

Table 3: Statistical comparison of Simplex MDE vs SMap MDE prediction

Comparison of Models	
Comparison	p-value
Multivariate:Simplex vs Univariate:Simplex	0.007
Multivariate:SMap vs Univariate:SMap	0.003
Univariate:Simplex vs Univariate:SMap	0.148
Multivariate:Simplex vs Multivariate:SMap	0.460

A.2 Data Compliance

In this study, we utilized minimally processed task-based fMRI recordings from the ABCD Study [18]. In conducting this research, all responsibilities related to human subjects ethics and compliance with relevant regulations are delegated to the ABCD Research Consortium. We rely on the Consortium's established protocols and ethical guidelines for managing and overseeing research involving human subjects, ensuring that all aspects of ethical conduct and data handling are appropriately addressed. As a research institution having access to this dataset, we are responsible for the responsible use of human subjects' data, including safeguarding their privacy and ensuring the confidentiality of their information, as outlined in the terms and conditions of our NIMH Data Archive Data Use Certification (DUC).

A.3 Computational Resources

Data preprocessing was performed in Python using the Nilearn package, specifically leveraging the CanICA function for Independent Component Analysis (ICA) and normalization.

The prediction algorithm was implemented in Python using functions from the pyEDM package, including Simplex, SMap and Convergent Cross Mapping (CCM).

The prediction of a behavioral variable, based on a dataset comprising 20 Independent Components, completed in 175.68 seconds (approximately 2.93 minutes) on an Intel® Xeon® CPU E5-2687W @ 3.10 GHz.

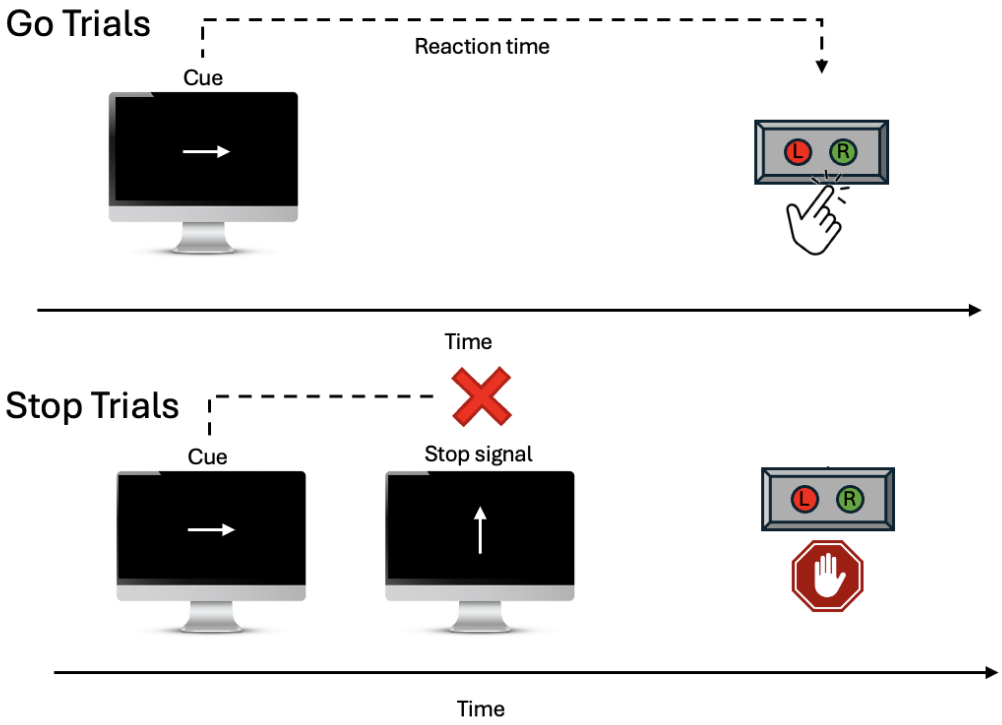

Figure S1: Schematic illustration of the Stop Signal Task (SST) design. Briefly, participants are instructed to quickly respond to directional cues (left or right arrow) via a two button panel. Participants should not respond when a stop signal is showed (upward arrow). This paradigm assess cognitive processes such as decision-making, response time and inhibitory control. (Adapted from [6])

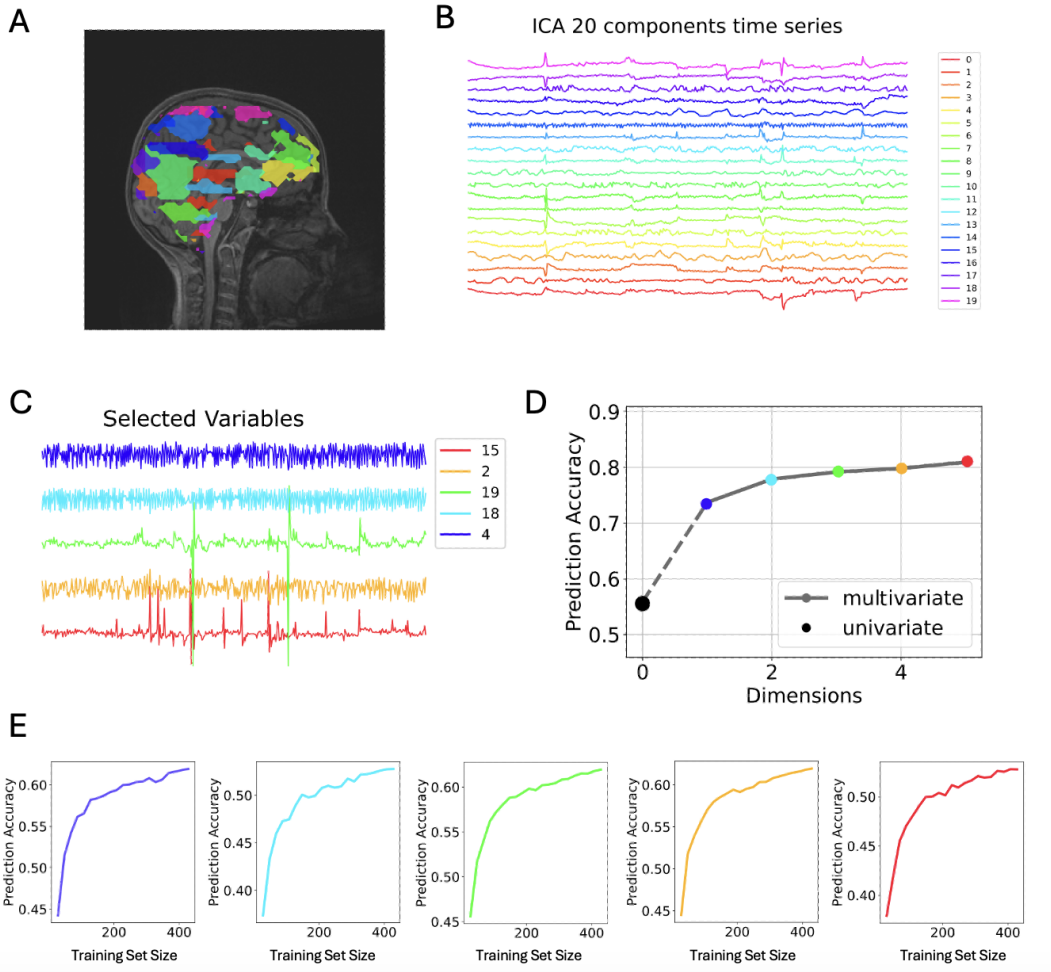

Figure S2: Task-based fMRI data were processed to derive 20 spatially distributed Independent Components, as shown in (A). For each Independent Component, the corresponding time series was extracted (B). (C) displays the variables selected by the algorithm for multivariate prediction. (D) shows the prediction accuracy for the univariate model (black dot) and the incremental improvement as each selected variable is added, color-coded to match (C). (E) presents a convergence check for each selected variable, also color-coded as in (C).

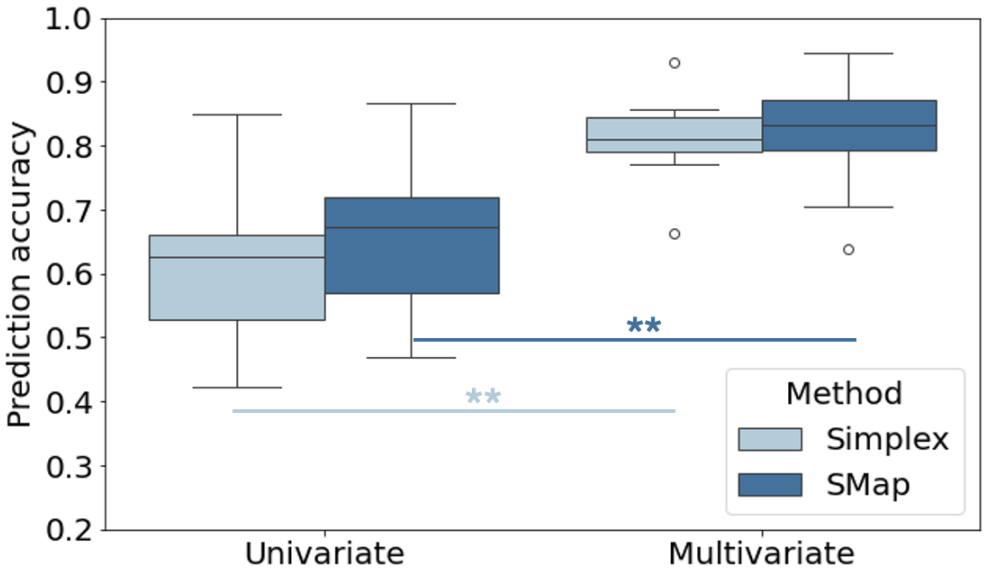

Figure S3: Boxplot of prediction accuracy in Univariate and Multivariate models with Simplex prediction algorithm (light blue) and SMap prediction algorithm (dark blue). The multivariate prediction accuracy is significantly higher in Multivariate than Univariate models, while no significant differences were found between Simplex and SMap prediction algorithms (see statistic in Table 3).

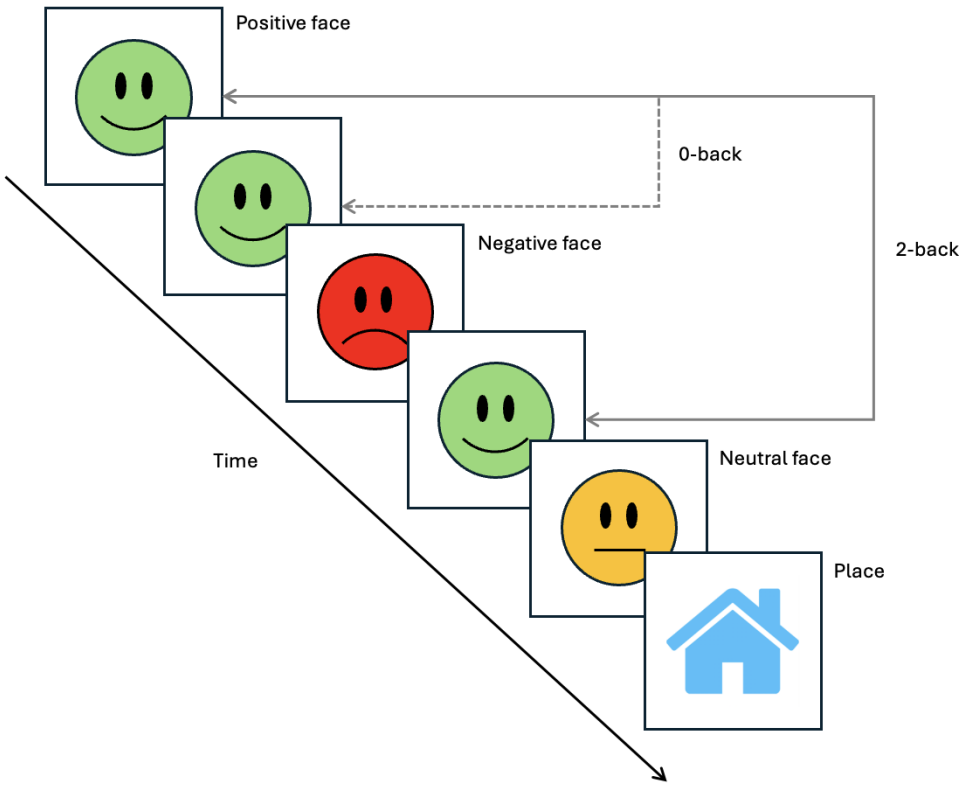

Figure S4: Schematic illustration of the Emotional N-back (EN-back) task design. The EN-back task [9, 10] is a variant of the HCP n-back task [2] designed to engage both memory and emotion regulation processes. Similar to the traditional n-back task, it includes high and low memory load conditions (2-back solid gray line and 0-back dashed gray line), which allow for the assessment of activation related specifically to working memory as opposed to general cognitive function. The stimuli used in the EN-back task include happy, fearful, and neutral facial expressions [11, 36], unlike the neutral stimuli used in earlier versions of the n-back task. In the schematic illustration of the task, these emotional faces are represented as smile icons that serve as placeholders for the real facial expressions used during the actual task that are drawn from the NimStim emotional stimulus set [36] and the Racially Diverse Affective Expressions (RADIATE) set [11]. These emotionally evocative stimuli activate fronto-amygdala circuitry involved in emotion reactivity and regulation [16, 14], as well as ventral fronto-striatal pathways linked to reward processing [29]. The task's design allows for contrasting neural responses to emotional faces versus neutral faces, assessing activation specificity to emotionally salient stimuli. Non-emotional and non-social place stimuli are also included, providing a reliable baseline for comparison [24].(Adapted from [6])

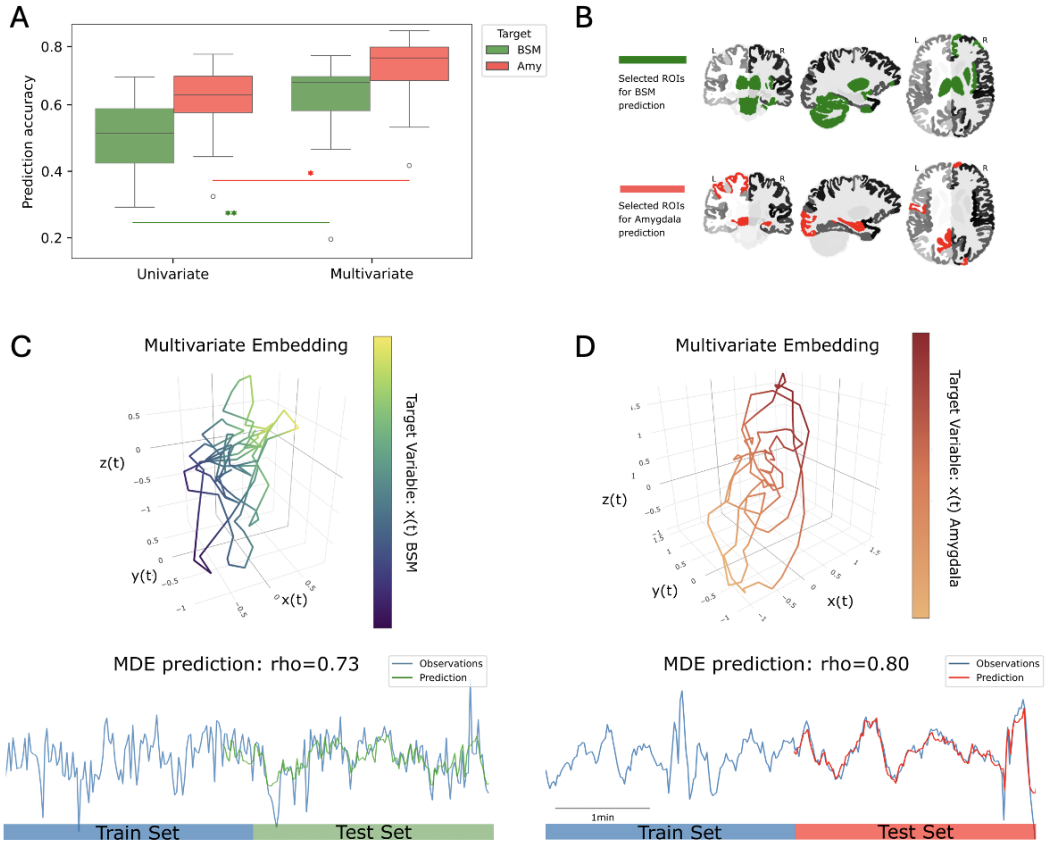

Figure S5: (A) Boxplot showing the increase in prediction accuracy from univariate, to multivariate methods, for two different target variables, brainstem (BSM, in green) and amygdala (Amy, in red). Differences between methods are statistically significant (p-value<0.05, see appendix Table 2). (B) Map of the brain segmented into 98 Regions of Interest (ROIs), highlighting the areas selected for the prediction of the BSM in green (Left Globus Pallidus, Right Orbitofrontal Cortex, Right Putamen, Left Thalamus, Right Insula, Right Cerebellar Cortex, Right Fusiform Gyrus, Right Frontopolar Cortex, Right Medial Orbitofrontal Cortex, Right Thalamus), and Amy in red (Superior Frontal Cortex, Precuneus, Amygdala, Frontopolar Cortex, Hippocampus, Lingual Gyrus, Precentral Gyrus, Isthmus Cingulate Gyrus, Ventral Diencephalon, Lateral Occipital Cortex). (C) Example of multivariate embedding predicting the BSM activity with accuracy of 73% on the test set. (D) Example of multivariate embedding predicting the Amy activity with accuracy of 80% on the test set.