A More Visualizations for Autism Dataset

We have visualized real and reconstructed brain images from the nine subsets of autism dataset (Caltech, Leuven, MaxMun, NYU, SBL, Stanford, UM, USM, and Yale sites) along with z_0^w after training DMFA on the full autism dataset, in Fig. S1. DMFA clustered together temporal latent variables associated with each imaging site without supervision in z_0^w . As depicted, the variation among different imaging sites dominates the cognitive differences between ASD group and control, hence, a downstream connectivity matrix classification (using the learned temporal weights, W) helps better in differentiating ASD group from control in multi-site analysis.

B Hyperparameters for Real Data

We chose K = 100 in real experiments to make it roughly at par with the number of regions in the AAL atlas. In practice, K should be selected by a practitioner depending on the desired granularity. We set the size of temporal latents z_t^w to be 2-dimensional as we found it helpful for visualization purposes and because larger sizes had marginal effect on the results. We selected the size of the hidden layer for the temporal transition network D_t from {2, 3, 5}, and picked the model with highest prediction accuracy on a small validation set. We set the size of the hidden layer for the temporal emission network D_e to the geometric mean between the size of z_t^w and K (i.e., size of w_t), which is $D_e = \sqrt{2 \times 100} \sim 15$.

C Background on TFA Models

TFA approximates each fMRI sequence Y_n as a product between temporal weights $W_n = \lfloor w_{n,k,t} \rfloor$ and spatial factors $F_n = \lfloor f_{n,k,v} \rfloor$, and defines a hierarchical Gaussian prior over each of these latent variables:

$$\begin{split} Y_n &\sim \operatorname{Norm} \left(W_n^\top F_n, \sigma^Y \mathbf{I} \right), \\ w_{n,k,t} &\sim \operatorname{Norm} (\mu_{n,k}^w, \sigma_{n,k}^w), \qquad \qquad \mu_{n,k}^w \sim p(\mu^w), \qquad \qquad \sigma_{n,k}^w \sim p(\sigma^w) \\ f_{n,k,v} &= \operatorname{RBF}(v; \, \rho_{n,k}, \, \gamma_{n,k}), \qquad \qquad \rho_{n,k} \sim p(\rho), \qquad \qquad \gamma_{n,k} \sim p(\gamma). \end{split}$$

TFA treats each fMRI sequence as independent. HTFA works similarly to TFA, but places an additional constraint over the factors to bias all of the sequences to exhibit similar factors:

$$\begin{split} \mu_{n,k}^w &\sim p(\mu_{n,k}^w \mid \bar{\mu}_k^w), \qquad \bar{\mu}_k^w \sim p(\bar{\mu}^w), \qquad \sigma_{n,k}^w \sim p(\sigma_{n,k}^w \mid \bar{\sigma}_k^w), \qquad \bar{\sigma}_k^w \sim p(\bar{\sigma}^w), \\ \rho_{n,k} &\sim p(\rho_{n,k} \mid \bar{\rho}_k), \qquad \bar{\rho}_k \sim p(\bar{\rho}), \qquad \gamma_{n,k} \sim p(\gamma_{n,k} \mid \bar{\gamma}_k), \qquad \bar{\gamma}_k \sim p(\bar{\gamma}). \end{split}$$

In this way, whereas TFA attempts to find the factors that best explain an individual sequence, HTFA assumes that the factors across sequences vary around a shared Gaussian prior. NTFA extends HTFA for task fMRI by incorporating neural networks onto its framework. NTFA assumes separate latent embeddings for participants and stimuli and from there maps into the temporal and spatial latents with neural networks. However, it is evident from the priors in TFA methods that the temporal weights are independent as a function of time and therefore these models do not encode temporal dynamics.

D Likelihood Approximation

The model's likelihood is intractable, however, we computed an importance sampling based lowerbound approximation of the likelihood which is also adopted in Krishnan et al. [2017], Kingma and Welling [2014]. Let's denote Y as the train data and \overline{Y} as the test data. We approximated the test set posterior-predictive likelihood from L = 100 samples as follows:

$$\log p(\bar{Y} \mid Y) \ge \mathbb{E}[\bar{\mathcal{L}}] \approx \frac{1}{\bar{L}} \sum_{l} \sum_{t} \log p_{\theta}(\bar{Y}_{t} \mid \bar{w}_{t}^{(l)}, \bar{\rho}^{(l)}, \bar{\gamma}^{(l)}, \bar{z}_{t}^{w(l)}, z_{t-1}^{w(l)}, \mathbf{C}^{(l)}), \quad (1)$$

where superscript (l) indexes a sample. We sampled z_{t-1}^{w} , z^{F} , and C from their variational distributions and the remaining latent variables from their priors, i.e., the generative model:

$$\begin{split} z_{t-1}^{\mathbf{w}(l)} &\sim q(z_{t-1}^{\mathbf{w}}), \qquad z^{\mathbf{F}(l)} \sim q(z^{\mathbf{F}}), \qquad \mathbf{C}^{(l)} \sim q(\mathbf{C}), \\ \bar{z}_{t}^{\mathbf{w}(l)} &\sim p(\bar{z}_{t}^{\mathbf{w}} \mid z_{t-1}^{\mathbf{w}(l)}), \qquad \bar{w}_{t}^{(l)} \sim p(\bar{w}_{t} \mid \bar{z}_{t}^{\mathbf{w}(l)}), \qquad \bar{\rho}^{(l)}, \, \bar{\gamma}^{(l)} \sim p(\bar{\rho}, \, \bar{\gamma} \mid z^{\mathbf{F}(l)}) \end{split}$$

The likelihood is estimated per subject and the average is reported in Table 2 of the paper.

Algorithm S1 DMFA Generative Model

1: for n in 1, ..., N do $z_n^{\rm F} \sim \operatorname{Norm}(0, \mathbf{I})$ 2: $\rho_{n,1}, \ldots, \rho_{n,K}, \gamma_{n,1}, \ldots, \gamma_{n,K} \sim \operatorname{Norm}(\mu_{\theta}^{\mathsf{F}}(z_n^{\mathsf{F}}), \sigma_{\theta}^{\mathsf{F}}(z_n^{\mathsf{F}}))$ for k in $1, \ldots, K$ do 3: 4: for v in $1, \ldots, V$ do 5: $f_{n,k,v} \leftarrow \operatorname{RBF}(v; \rho_{n,k}, \gamma_{n,k})$ 6: $F_n \leftarrow \begin{bmatrix} f_{n,k,v} \end{bmatrix}$ $C_n \sim \operatorname{Cat}(\pi), \quad \pi = \{\pi_1, \dots, \pi_S\}$ $z_{n,0}^{\mathsf{W}} \sim \operatorname{Norm}(\mu_{\mathsf{c}_n}, \Sigma_{\mathsf{c}_n})$ for $t \text{ in } 1, \dots, T \text{ do}$ 7: 8: 9: 10: $z_{n,t}^{\mathsf{w}} \sim \operatorname{Norm}\left(\mu_{\theta}^{\mathsf{z}}(z_{n,t-1}^{\mathsf{w}}), \, \sigma_{\theta}^{\mathsf{z}}(z_{n,t-1}^{\mathsf{w}})\right)$ 11: $w_{n,t} \sim \operatorname{Norm}(\mu_{\theta}^{W}(z_{n,t}^{W}), \sigma_{\theta}^{W}(z_{n,t}^{W}))$ 12: $W_n \leftarrow [w_{n,t}]$ 13: $Y_n \sim \operatorname{Norm}(W_n^\top F_n, \sigma^{\mathrm{Y}} \mathrm{I})$ 14:

E DMFA Generative Algorithm

We have summarized the generative model of DMFA in Algorithm S1.

F Autism Separation Accuracy

Table S1 includes the autism separation accuracy for all the subsets of autism dataset and all the methods.

G Sensitivity Analysis

Table S2 provides test set prediction errors for $D_t \in \{2, 3, 5\}$ and $\sigma^Y \in \{0.001, 0.005, 0.01\}$ on different datasets. DMFA proves stable w.r.t. to these hyperparameters. Note that D_t is the dimension of hidden layer for the temporal transition network and σ^Y is the observation noise.

H Societal Impact

Analysing brain imaging scans are intended for research aimed at understanding the pathophysiology of neurodegenerative disease and the development of treatments for use in the presymptomatic phase. However, an important neuroethical issue is the predictive and diagnostic imaging for progressive diseases that lack effective treatments, such as Alzheimer's disease. These imaging scans could be used for other reasons by the worried employers or insurers. In such cases, the benefits of foreknowledge, for example the greater opportunity to plan, must be weighed against the psychological burden of this knowledge and its potential impact on employability or insurability.

Dataset	Caltech	Leuven	MaxMun	NYU	SBL	Stanford	UM_1	USM	Yale	Full
DMFA	70	50	57	52	62	71	59	51	56	51
DMFA+SVM	48±6	$ 55\pm 5$	60±6	59± 3	$ 63\pm7$	77±6	66 ± 4	$ 65\pm 6$	$ 52\pm 6$	$ 64\pm 2$
ROI+SVM	42± 8	54±6	54±7	66±3	46± 9	63±8	67±4	$ 67\pm 6$	47± 7	$ 60\pm 2$
ICA+SVM	4 3± 8	$ 54\pm 6$	55±8	$ 66\pm 3$	45±9	64±7	66±4	$ 67\pm 6$	$ 48\pm 8$	59± 3

Table S1: Comparison of autism separation accuracy (%).

The best results are highlighted in bold fonts.

Hyperparameter Dataset	$\left D_t = 2 \right D_t = 3 \left D_t = 5 \right \sigma^{\mathrm{Y}} =$	$= 0.001 \left \sigma^{\mathrm{Y}} = 0.005 \right \sigma^{\mathrm{Y}} = 0.01$
Autism (Caltech)	5.75 5.39 5.44 5	.32 5.41 5.44
Depression	6.50 6.48 6.35 6	.35 6.43 6.78
Synthetic	2.07 2.12 2.28 2	.06 2.09 2.07

Table S2: Prediction error (NRMSE%) of DMFA versus D_t and σ^{Y} on the test sets.

In other cases, brain imaging analysis raises new ethical, legal, and social issues that stem directly from the special relationship between brain and mind. The ability of brain imaging to deliver information about our psyches—about who we are and what we might be thinking or feeling while in the scanner—opens up a range of ethical challenges with few, if any, direct precedents. In other words, to the extent that brain imaging can actually deliver useful information about a person's mental states or traits, the issue of privacy gets important. To the extent that it cannot, but people believe that it can, the issue of public misunderstanding gets important. On the face of things, brain imaging poses a novel challenge to privacy in that it can in principle deliver information about thoughts, attitudes, beliefs, and traits even when someone offers no behavioral responses. On the other hand, studies suggest that laypersons may attribute greater objectivity and certainty to brain images than to other types of information about the human mind. This may contribute to the premature commercialization of brain imaging for various real-world applications [Farah, 2012].

References

- Martha J Farah. Neuroethics: the ethical, legal, and societal impact of neuroscience. *Annual review* of psychology, 63:571–591, 2012.
- Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun, editors, 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings, 2014. URL http://arxiv.org/abs/1312.6114.
- Rahul Krishnan, Uri Shalit, and David Sontag. Structured inference networks for nonlinear state space models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.

Checklist

- 1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] See section 5.3. In general, the mean-field approximation of posterior distributions imposes limitation on capturing temporal dynamics. Also, the variation among different imaging sites dominates the cognitive differences between clinical groups and makes the multi-site analysis hard (see figure S1 in supplementary).
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] Kindly see supplementary.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
- 2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
 - (b) Did you include complete proofs of all theoretical results? [N/A]
- 3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code and experiments are submitted as supplementary.

- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Kindly see section 4 and section 5 throughout.
- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See figure 3.
- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See section 4.
- 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
- 5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]



Figure S1: **Real and reconstructed** brain images from the nine subsets of Autism dataset (Caltech, Leuven, MaxMun, NYU, SBL, Stanford, UM, USM, and Yale sites) showing the smoothing given by sparse factorization. **Visualizing** z_0^{w} after training DMFA on the full autism dataset. DMFA clustered together temporal latent variables associated with each acquisition site without supervision. As depicted, the variation among different imaging sites dominates the variation in cognitive state of the brain (ASD group vs. control), hence, a downstream connectivity matrix classification helps better in differentiating ASD group from control in multi-site analysis.