

Multi-modal Variational Autoencoders for normative modelling across multiple imaging modalities

Ana Lawry Aguila, James Chapman, and Andre Altmann

University College London, London, WC1E 6BT ana.aguila.18@ucl.ac.uk

Abstract. One of the challenges of studying common neurological disorders is disease heterogeneity including differences in causes, neuroimaging characteristics, comorbidities, or genetic variation. Normative modelling has become a popular method for studying such cohorts where the ‘normal’ behaviour of a physiological system is modelled and can be used at subject level to detect deviations relating to disease pathology. For many heterogeneous diseases, we expect to observe abnormalities across a range of neuroimaging and biological variables. However, thus far, normative models have largely been developed for studying a single imaging modality. We aim to develop a multi-modal normative modelling framework where abnormality is aggregated across variables of multiple modalities and is better able to detect deviations than uni-modal baselines. We propose two multi-modal VAE normative models to detect subject level deviations across T1 and DTI data. Our proposed models were better able to detect diseased individuals, capture disease severity, and correlate with patient cognition than baseline approaches. We also propose a multivariate latent deviation metric, measuring deviations from the joint latent space, which outperformed feature-based metrics.

Keywords: Unsupervised learning · Normative modelling · Multimodal modelling · multi-view VAEs

1 Introduction

Normative modelling is a popular method to study heterogeneous brain disorders. Normative models assume disease cohorts sit at the tails of a healthy population distribution and quantify individual deviations from healthy brain patterns. Typically, a normative analysis constructs a normative model per variable, e.g., using Gaussian Process Regression (GPR)[9]. Recently, to model complex non-linear interactions between features, deep-learning approaches using adversarial (AAE) and variational autoencoder (VAE) models have been proposed[11, 8]. These models have a uni-modal structure with a single encoder and decoder network. So far, almost all deep-learning normative models have modelled only one modality. However, many brain disorders show deviations from the norm in features of multiple imaging modalities to a varying degree. Often it is unknown

which modality will be the most sensitive. Thus, it is advantageous to develop normative models suitable for multiple modalities.

Most previous deep-learning normative and anomaly detection models measure deviations in the feature space[11, 7, 4]. However, for multi-modal models built from modalities containing highly different, but complementary information (e.g., T1 and DTI features as used here), we may not expect to see significantly greater deviations in the feature space compared to uni-modal methods. Indeed previous work has shown that, when using VAEs, even for one modality, measuring deviation in the latent space outperforms metrics in the feature space[8] and provides a single measure of abnormality. As such, we develop a latent deviation metric suitable to measuring deviations in multi-modal data.

There are many approaches to extending VAEs to integrate information from multiple modalities and learn informative joint latent representations. Most multi-modal VAE frameworks learn separate encoder and decoder networks for each modality and aggregate the encoding distributions to learn a joint latent representation. Wu and Goodman[16] introduced a multi-modal VAE (mVAE) where each encoding distribution is treated as an ‘expert’ and the Product-of-Experts (PoE), which takes a product of the experts’ densities, is used to approximate a joint encoding distribution. The PoE approach treats all experts as equally credible taking a uniform contribution from every modality. In practice, however, different levels of noise, complexity and information are present in different modalities. Furthermore, if we have an overconfident miscalibrated expert, i.e. a sharp, shifted probability distribution, the joint distribution will have low density in the region observed by the other experts and a biased mean prediction. This can result in a suboptimal latent space and data reconstruction. Shi et al.[13] address this problem by combining latent representations across modalities using a Mixture-of-Experts (MoE) approach. For MoE, the joint distribution is given by a mixture of the experts’ densities so that the density is spread over all regions covered by the experts and overconfident experts do not monopolize the resulting prediction. However, MoE is less sensitive to consensus across modalities and will give lower probability to regions where experts are in agreement than PoE. Alternatively, we propose a mVAE modelling the joint encoding distribution as a generalised Product-of-Experts (gPoE)[2]. We optimise modality specific weightings to account for different information content between experts and enable the model to down-weight experts which cause erroneous predictions. Depending on the application, either MoE or gPoE will be most appropriate and so we consider both methods for normative modelling.

As far as we are aware, only one other multi-modal VAE normative modelling framework has been proposed in the literature which uses the PoE (PoE-normVAE)[7]. However, Kumar et al.[7] rely on measuring deviations in the feature space, which we argue does not leverage the benefits of multi-modal models. Here, we present an improved factorisation of the joint representation by modelling it as a weighted product or sum of each encoding distribution.

Our contributions are two-fold. Firstly, we present two novel multi-modal normative modelling frameworks, MoE-normVAE and gPoE-normVAE, which

capture the joint distribution between different imaging modalities. Our proposed models outperform baseline methods on two neuroimaging datasets. Secondly, we present a deviation metric, based on the latent space, suitable for detecting deviations in multi-modal normative distributions. We show that our metric better leverages the benefits of multi-modal normative models compared to feature space-based metrics.

2 Methods

Multi-modal Variational autoencoder (mVAE). Let $\mathbf{X} = \{\mathbf{x}_m\}_{m=1}^M$ be the observations of M modalities. We use a mVAE to learn a multi-modal generative model (Figure 1c), where modalities are conditionally independent given a common latent variable, of the form $p_\theta(\mathbf{X}, \mathbf{z}) = p(\mathbf{z}) \prod_{m=1}^M p_{\theta_m}(\mathbf{x}_m | \mathbf{z})$. The likelihood distributions $p_{\theta_m}(\mathbf{x}_m | \mathbf{z})$ are parameterised by decoder networks with parameters $\theta = \{\theta_1, \dots, \theta_M\}$. The goal of VAE training is to maximise the marginal likelihood of the data. However, as this is intractable, we instead optimise an evidence lower bound (ELBO):

$$\mathcal{L} = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{X})} \left[\sum_{m=1}^M \log p_\theta(\mathbf{x}_m | \mathbf{z}) \right] - D_{KL}(q_\phi(\mathbf{z} | \mathbf{X})|p(\mathbf{z})) \quad (1)$$

where the second term is the KL divergence between the approximate joint posterior $q_\phi(\mathbf{z} | \mathbf{X})$ and the prior $p(\mathbf{z})$. We model the posterior, likelihood, and prior distributions as isotropic gaussians.

Approximate joint posterior. To train the mVAE, we must specify the form of the joint approximate posterior $q_\phi(\mathbf{z} | \mathbf{X})$. Wu and Goodman[16] choose to factorise the joint posterior as a Product-of-Experts (PoE); $q_\phi(\mathbf{z} | \mathbf{X}) = \frac{1}{K} \prod_{m=1}^M q_{\phi_m}(\mathbf{z} | \mathbf{x}_m)$,

where the experts, i.e., individual posterior distributions $q_{\phi_m}(\mathbf{z} | \mathbf{x}_m)$, are parameterised by encoder networks with parameters $\phi = \{\phi_1, \dots, \phi_M\}$. K is a normalisation term. Assuming each encoder network follows a Gaussian distribution $q(\mathbf{z} | \mathbf{x}_m) = \mathcal{N}(\boldsymbol{\mu}_m, \boldsymbol{\sigma}_m^2 \mathbf{I})$, the parameters of joint posterior distribution can be computed[5]; $\boldsymbol{\mu} = \frac{\sum_{m=1}^M \boldsymbol{\mu}_m / \boldsymbol{\sigma}_m^2}{\sum_{m=1}^M 1 / \boldsymbol{\sigma}_m^2}$ and $\boldsymbol{\sigma}^2 = \frac{1}{\sum_{m=1}^M 1 / \boldsymbol{\sigma}_m^2}$ (see Supp. for proofs).

However, overconfident but miscalibrated experts may bias the joint posterior distribution (see Figure 1b) which is undesirable for learning informative latent representations between modalities[13].

Shi et al.[13] instead factorise the approximate joint posterior as a Mixture-of-Experts (MoE); $q_\phi(\mathbf{z} | \mathbf{X}) = \frac{1}{K} \sum_{m=1}^M \frac{1}{M} q_{\phi_m}(\mathbf{z} | \mathbf{x}_m)$.

In the MoE setting, each uni-modal posterior $q_\phi(\mathbf{z} | \mathbf{x}_m)$ is evaluated with the generative model $p_\theta(\mathbf{X}, \mathbf{z})$ such that the ELBO becomes:

$$\mathcal{L} = \sum_{m=1}^M \left[\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x}_m)} \left[\sum_{m=1}^M \log p_\theta(\mathbf{x}_m | \mathbf{z}) \right] - D_{KL}(q_\phi(\mathbf{z} | \mathbf{x}_m)|p(\mathbf{z})) \right]. \quad (2)$$

However, this approach only takes each uni-modal encoding distribution separately into account during training. Thus, there is no explicit aggregation of information from multiple modalities in the latent representation for reconstruction by the decoder networks. For modalities with a high degree of modality-specific variation, this enforces an undesirable upperbound on the ELBO potentially leading to a sub-optimal approximation of the joint distribution[3].

Generalised Product-of-Experts joint posterior. We propose an alternative approach to mitigate the problem of overconfident experts by factorising the joint posterior as a generalised Product-of-Experts (gPoE)[2]; $q_\phi(\mathbf{z} | \mathbf{X}) = \frac{1}{K} \prod_{m=1}^M q_{\phi_m}^{\alpha_m}(\mathbf{z} | \mathbf{x}_m)$

where α_m is a weighting for modality m such that $\sum_{m=1}^M \alpha_m = 1$ for each latent dimension and $0 < \alpha_m < 1$. We optimise α during training allowing the model to weight experts in such a way as to learn an approximate joint posterior $q_\phi(\mathbf{z} | \mathbf{X})$ where the likelihood distribution $p_\theta(\mathbf{X} | \mathbf{z})$ is maximised. This provides a means to down-weight overconfident experts. Furthermore, as α is learnt per latent dimension, different modality weightings can be learnt for different vectors, thus explicitly incorporating modality specific variation in addition to shared information in different dimensions of the joint latent space. Similarly to the PoE approach, we can compute the parameters of the joint posterior distribution; $\boldsymbol{\mu} = \frac{\sum_{m=1}^M \boldsymbol{\mu}_m \alpha_m / \sigma_m^2}{\sum_{m=1}^M \alpha_m / \sigma_m^2}$ and $\boldsymbol{\sigma}^2 = \sum_{m=1}^M \frac{1}{\alpha_m / \sigma_m^2}$.

Recently, a gPoE mVAE was proposed for learning joint representations of hand-poses and surgical videos[6]. However, we emphasize that our approach differs in application and offers a more lightweight implementation (Joshi et al.[6] require training of auxiliary networks to learn α per sample).

Multi-modal normative modelling. We propose two mVAE normative modelling frameworks shown in Figure 1a. MoE-normVAE, which uses a MoE joint posterior distribution, and gPoE-normVAE, which uses a gPoE joint posterior distribution. For both models, the encoder ϕ and decoder θ parameters are trained to characterise a healthy population cohort. normVAE models assume abnormality due to disease effects can be quantified by measuring deviations in the latent space[8] or the feature space[11]. At test time, the clinical cohort is passed through the encoder and decoder networks. Deviations of test subjects from the multi-modal latent space of the healthy controls and data reconstruction errors are measured. We compare our methods to the previously proposed PoE-normVAE[7] and three uni-modal models; two single modality and one multi-modality with a concatenated input.

To compare our normVAE models to a classical normative approach, we trained one GPR (using the PCNToolkit) per feature on a sub-set of 2000 healthy UK Biobank individuals and used extreme value statistics to calculate subject-level abnormality index[9]. We used a top 5% abnormality threshold (set using the healthy training cohort) to calculate a significance ratio (see Equation 6).

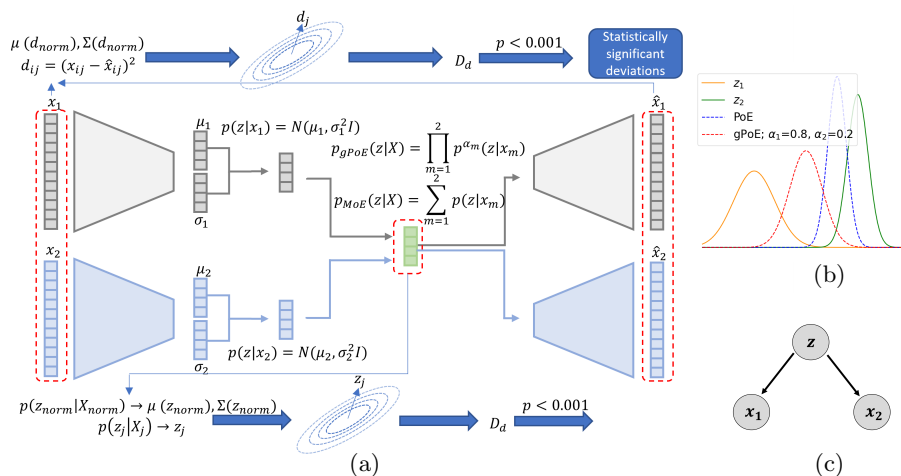


Fig. 1: (a) gPoE-normVAE and MoE-normVAE normative framework. All normVAE models were implemented using parameter settings; maximum epochs=2000, batch size=256, learning rate= 10^{-4} , early stopping=50 epochs, encoder layers=[20, 40], decoder layers=[20, 40]. A ReLU activation function was applied between layers. Models were trained with a range of latent space sizes (L_{dim}) from 5 to 20. Models with $L_{dim}=10$ were fine-tuned (maximum 100 epochs) using the ADNI healthy cohort. Learnt α values are given in Supp. Table 1. (b) Example PoE and gPoE joint distributions. (c) Graphical model.

Multi-modal latent deviation metric. Previous works using autoencoders as normative models mostly relied on feature-space based deviation methods[11, 7]. That is, they compare the input value for subject j for the i -th brain region x_{ij} to the value reconstructed by the autoencoder \hat{x}_{ij} : $d_{ij} = (x_{ij} - \hat{x}_{ij})^2$. Kumar et al.[7] propose the following normalised z-score metric on the data reconstruction (a univariate feature space metric):

$$D_{uf} = \frac{d_{ij} - \mu_{norm}(d_{ij}^{norm})}{\sigma_{norm}(d_{ij}^{norm})} \quad (3)$$

where $\mu_{norm}(d_{ij}^{norm})$ is the mean and $\sigma_{norm}(d_{ij}^{norm})$ the standard deviation of the deviations d_{ij}^{norm} of a holdout healthy control cohort.

However, in the multi-modal setting, feature space-based deviation metrics may not highlight the benefits of multi-modal models over their uni-modal counterparts. The goal of the joint latent representation is to capture information from all modalities. Thus, decoders for each modality must extract the information from the joint latent representation, which now carries information from all other modalities as well. Therefore, data reconstructions capture only information relevant to a particular modality and may also be poorer compared to uni-modal methods. As such, particularly when incorporating modalities with a high degree of modality-specific variation, we believe latent space deviation met-

rics would better capture deviations from normative behaviour across multiple modalities. Then, once an abnormal subject has been identified, feature space metrics can be used to identify deviating brain regions (e.g. Supp. Figure 3).

We propose a latent deviation metric to measure deviations from the joint normative distribution. To account for correlation between latent vectors and derive a single multivariate measure of deviation, we measure the Mahalanobis distance from the encoding distribution of the training cohort:

$$D_{\text{ml}} = \sqrt{(z_j - \mu(z^{\text{norm}}))^T \Sigma(z^{\text{norm}})^{-1} (z_j - \mu(z^{\text{norm}}))} \quad (4)$$

where $z_j \sim q(\mathbf{z}_j | \mathbf{X}_j)$ is a sample from the joint posterior distribution for subject j , $\mu(z^{\text{norm}})$ is the mean and $\Sigma(z^{\text{norm}})$ the covariance of the healthy cohort latent position. We use robust estimates of the mean and covariance to account for outliers within the healthy control cohort. For closer comparison with D_{ml} , we derive the following multivariate feature space metric:

$$D_{\text{mf}} = \sqrt{(d_j - \mu(d^{\text{norm}}))^T \Sigma(d^{\text{norm}})^{-1} (d_j - \mu(d^{\text{norm}}))} \quad (5)$$

where $d_j = \{d_{ij}, \dots, d_{Ij}\}$ is the reconstruction error for subject j for brain regions ($i = 1, \dots, I$), $\mu(d^{\text{norm}})$ is the mean and $\Sigma(d^{\text{norm}})$ the covariance of the healthy cohort reconstruction error.

Assessing deviation metric performance. For each model, we calculated D_{ml} and D_{mf} for a healthy holdout cohort and disease cohort. For each deviation metric, we identified individuals whose deviations were significantly different from the healthy training distribution ($p < 0.001$)[15]. Ideally, we want a model which correctly identifies disease individuals as outliers and healthy individuals as sitting within the normative distribution. As such, we use the following significance ratio (positive likelihood ratio) to assess model performance:

$$\text{significance ratio} = \frac{\text{True positive rate}}{\text{False positive rate}} = \frac{\text{TPR}}{\text{FPR}} = \frac{\frac{N_{\text{disease(outliers)}}}{N_{\text{disease}}}}{\frac{N_{\text{holdout(outliers)}}}{N_{\text{holdout}}}} \quad (6)$$

In order to calculate significance ratios, we calculated D_{uf} relative to the training cohort for the healthy holdout and disease cohorts (Bonferroni adjusted $p=0.05/N_{\text{features}}$)[7].

3 Experiments

Data processing. To train the normVAE models, we used 10,276 healthy subjects from the UK Biobank[14] (application number: 70047). We used pre-processed (provided by the UK Biobank[1]) grey-matter volumes for 66 cortical (Desikan-Killiany atlas) and 16 subcortical brain regions, and Fractional Anisotropy (FA) and Mean Diffusivity (MD) measurements for 35 white matter

tracts (John Hopkins University atlas). At test time, we used 2,568 healthy controls from a holdout cohort and 122 individuals with one of several neurodegenerative disorders; motor neuron disease, multiple sclerosis, Parkinson’s disease, dementia/Alzheimer/cognitive-impairment and other demyelinating disease.

We also tested the models using an external dataset. We extracted 213 subjects from the Alzheimer’s Disease Neuroimaging Initiative (ADNI)¹[10] dataset with significant memory concern (SMC; N=27), early mild cognitive impairment (EMCI; N=63), late mild cognitive impairment (LMCI; N=34), Alzheimer’s disease (AD; N=43) as well as healthy controls (HC; N=45). We used the healthy controls to fine-tune the models in a transfer learning approach. The same T1 and DTI features as for the UK Biobank were extracted for the ADNI dataset.

Rather than conditioning on covariates as done in some related work[7, 8], we adjusted for confounding effects prior to analysis. Non-linear age and linear ICV affects were removed from the DTI and T1 MRI features of both datasets[12]. Each brain ROI was normalised by removing the mean and dividing by the standard deviation of the healthy control cohort brain regions.

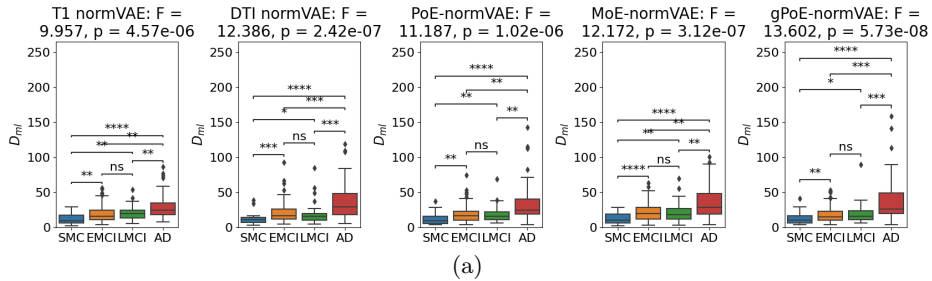
UK Biobank results. As expected, we see greater significance ratios for all models when using D_{ml} rather than D_{mf} (Table 1). When using D_{mf} or D_{uf} , all models perform similarly. Using D_{ml} over D_{mf} leads to a 4-fold increase in the significance ratio. Further, our proposed models give the best overall performance across different L_{dim} with the highest significance ratio for gPoE-normVAE with $L_{dim}=10$. Generally, all multi-modal normVAE showed better performance than the uni-modal models suggesting that by modelling the joint distribution between modalities, we can learn better normative models.

Table 1: Significance ratio calculated from D_{ml} , D_{mf} , and D_{uf} for the UK Biobank. See Supp. for results in figure form. Using GPR, we observed a significance ratio of 6.01, poorer performance than our models (using D_{ml}).

Latent dimension	Significance ratio, D_{ml}				Significance ratio, D_{mf}				Significance ratio, D_{uf}			
	5	10	15	20	5	10	15	20	5	10	15	20
gPoE-normVAE (ours)	7.89	9.24	7.02	7.6	1.62	1.62	1.61	1.7	1.37	1.41	1.37	1.56
MoE-normVAE (ours)	7.25	8.77	7.09	7.94	1.67	1.7	1.59	1.68	1.44	1.46	1.46	1.44
PoE-normVAE	7.4	8.06	5.71	6.96	1.6	1.62	1.66	1.71	1.45	1.42	1.39	1.41
concatenated normVAE	7.63	6.21	5.43	3.55	1.61	1.59	1.57	1.58	1.48	1.4	1.39	1.44
T1 normVAE	5.26	4.43	2.63	2.44								
DTI normVAE	6.82	7.35	3.96	2.61								
Average T1&DTI normVAE					1.63	1.7	1.64	1.64	1.47	1.53	1.45	1.45

¹ Data used in preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf

ADNI results. Previous work[8] explored the ability of a uni-modal T1 normVAE to detect deviations in the ADNI cohorts. Figure 2a shows the latent deviation D_{ml} for different diagnosis in the ADNI cohort for the T1 normVAE, DTI normVAE, PoE-normVAE, MoE-normVAE and gPoE-normVAE models. All models reflect the increasing disease severity with increasing disease stage. The gPoE-normVAE model showed greater sensitivity to disease stage as suggested by the higher F statistic and p-values from an ANOVA analysis. We measured the Pearson correlation with composite measures of memory and executive function (Figure 2b) and found that our proposed model exhibited greater correlation with both cognition scores than baseline approaches. Finally, we see that the sensitivity to disease severity for the gPoE-normVAE model extends to the feature space where we see a general increase in average D_{uf} from the LMCI to AD cohort (Supp. Figures 3a and 3b respectively).



model	Memory score (r)	Memory score (p)	Executive function score (r)	Executive function score (p)
T1 normVAE	-0.201754	8.9334e-03	-0.289296	1.495292e-04
DTI normVAE	-0.320889	2.3577e-05	-0.406776	4.898105e-08
PoE-normVAE	-0.255954	8.4200e-04	-0.404015	6.148406e-08
MoE-normVAE	-0.301793	7.3842e-05	-0.396141	1.1629e-07
gPoE-normVAE	-0.335123	9.5594e-06	-0.468410	1.7301e-10

(b)

Fig. 2: (a) D_{ml} by disease label ($L_{dim}=10$). Statistical annotations were generated using Welch’s t-tests between pairs of disease groups; ns : $0.05 < p \leq 1$, * : $0.01 < p \leq 0.05$, ** : $0.001 < p \leq 0.01$, *** : $0.0001 < p \leq 0.001$, **** : $p \leq 0.0001$. Robust estimates of the mean and covariance were not used to calculate D_{ml} due to the small healthy cohort size. (b) Pearson correlation between D_{ml} and patient cognition represented by age adjusted memory and executive function composite scores.

4 Discussion and further work

We have built on recent works[11, 7, 8] and introduced two novel mVAE normative models, which provide an alternative method of learning the joint normative distribution between modalities to address the limitations of current approaches. Our models provide a more informative joint representation compared to baseline methods as evidenced by the better significance ratio for the UK Biobank

dataset and greater sensitivity to disease staging and correlation with cognitive measures in the ADNI dataset. We also proposed a latent deviation metric suitable for detecting deviations in the multivariate latent space of multi-modal normative models which gave an approximately 4-fold performance increase over metrics based on the feature space.

Further work will involve extending our models to more data modalities, such as genetic variants, to better characterise the behaviour of a physiological system. We note that, for fair comparison across models, we remove the effects of confounding variables prior to analysis. However, confounding effects could be removed during analysis via condition variables[8]. Another limitation of normVAE models introduced here is the use of ROI level data. Data processing software, such as FreeSurfer, may fail to accurately capture abnormality in images, particularly if large lesions are present. Further work involves creating normative models designed for voxel level data to better capture disease effects.

Normative models have been successfully applied to the study of a range of heterogeneous diseases. Diseases often present abnormalities across a range of neuroimaging, biological and physiological features which provide different information about the underlying disease process. Normative systems that incorporate features from different data modalities offer a holistic picture of the disease and will be capable of detecting abnormalities across a broad range of different diseases. Furthermore, multi-modal normative modelling captures the relationship between different modalities in healthy individuals, with disruption to this relationship potentially leading to a disease signal. Code is publicly available at <https://github.com/alawryaguila/multimodal-normative-models>.

Acknowledgements. This work is supported by the EPSRC-funded UCL Centre for Doctoral Training in Intelligent, Integrated Imaging in Healthcare (i4health) and the Department of Health’s NIHR-funded Biomedical Research Centre at University College London Hospitals.

References

1. Alfaro-Almagro, F., Jenkinson, M., Bangerter, N.K., Andersson, J.L., Griffanti, L., Douaud, G., Sotiropoulos, S.N., Jbabdi, S., Hernandez-Fernandez, M., Vallee, E., Vidaurre, D., Webster, M., McCarthy, P., Rorden, C., Daducci, A., Alexander, D.C., Zhang, H., Dragonu, I., Matthews, P.M., Miller, K.L., Smith, S.M.: Image processing and quality control for the first 10,000 brain imaging datasets from uk biobank. *NeuroImage* **166**, 400–424 (2018). <https://doi.org/https://doi.org/10.1016/j.neuroimage.2017.10.034>, <https://www.sciencedirect.com/science/article/pii/S1053811917308613>
2. Cao, Y., Fleet, D.J.: Generalized product of experts for automatic and principled fusion of gaussian process predictions (2014). <https://doi.org/10.48550/ARXIV.1410.7827>, <https://arxiv.org/abs/1410.7827>
3. Daunhawer, I., Sutter, T.M., Chin-Cheong, K., Palumbo, E., Vogt, J.E.: On the limitations of multimodal vaes. *CoRR* **abs/2110.04121** (2021), <https://arxiv.org/abs/2110.04121>

4. Hendrycks, D., Mazeika, M., Dietterich, T.G.: Deep anomaly detection with outlier exposure. CoRR **abs/1812.04606** (2018), <http://arxiv.org/abs/1812.04606>
5. Hwang, H., Kim, G.H., Hong, S., Kim, K.E.: Multi-view representation learning via total correlation objective. In: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems*. vol. 34, pp. 12194–12207. Curran Associates, Inc. (2021), <https://proceedings.neurips.cc/paper/2021/file/65a99bb7a3115fdede20da98b08a370f-Paper.pdf>
6. Joshi, A., Gupta, N., Shah, J., Bhattarai, B., Modi, A., Stoyanov, D.: Generalized product-of-experts for learning multimodal representations in noisy environments. In: *Proceedings of the 2022 International Conference on Multimodal Interaction*. p. 83–93. ICMI '22, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3536221.3556596>, <https://doi.org/10.1145/3536221.3556596>
7. Kumar, S., Payne, P., Sotiras, A.: Normative modeling using multimodal variational autoencoders to identify abnormal brain structural patterns in alzheimer disease (2021). <https://doi.org/10.48550/ARXIV.2110.04903>, <https://arxiv.org/abs/2110.04903>
8. Lawry Aguila, A., Chapman, J., Janahi, M., Altmann, A.: Conditional vaes for confound removal and normative modelling of neurodegenerative diseases. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part I*. p. 430–440. Springer-Verlag (2022)
9. Marquand, A., Rezek, I., Buitelaar, J., Beckmann, C.: Understanding heterogeneity in clinical cohorts using normative models: Beyond case-control studies. *Biological Psychiatry* **80** (2016)
10. Petersen, R., Aisen, P., Beckett, L., Donohue, M., Gamst, A., Harvey, D., Jack, C., Jagust, W., Shaw, L., Toga, A., Trojanowski, J., Weiner, M.: Alzheimer’s disease neuroimaging initiative (adni): clinical characterization. *Neurology* **74**(3), 201–209 (2010). <https://doi.org/10.1212/wnl.0b013e3181cb3e25>, <https://europepmc.org/articles/PMC2809036>
11. Pinaya, W., Scarpazza, C., Garcia-Dias, R., Vieira, S., Baecker, L., Ferreira da Costa, P., Redolfi, A., Frisoni, G., Pievani, M., Calhoun, V., Sato, J., Mechelli, A.: Using normative modelling to detect disease progression in mild cognitive impairment and alzheimer’s disease in a cross-sectional multi-cohort study. *Scientific Reports* **11** (2021)
12. Pomponio, R., Erus, G., Habes, M., Doshi, J., Srinivasan, D., Mamourian, E., Bashyam, V., Nasrallah, I.M., Satterthwaite, T.D., Fan, Y., Launer, L.J., Masters, C.L., Maruff, P., Zhuo, C., Völzke, H., Johnson, S.C., Frupp, J., Koutsouleris, N., Wolf, D.H., Gur, R., Gur, R., Morris, J., Albert, M.S., Grabe, H.J., Resnick, S.M., Bryan, R.N., Wolk, D.A., Shinohara, R.T., Shou, H., Davatzikos, C.: Harmonization of large mri datasets for the analysis of brain imaging patterns throughout the lifespan. *NeuroImage* **208**, 116450 (2020). <https://doi.org/https://doi.org/10.1016/j.neuroimage.2019.116450>, <https://www.sciencedirect.com/science/article/pii/S1053811919310419>
13. Shi, Y., Siddharth, N., Paige, B., Torr, P.H.S.: Variational mixture-of-experts autoencoders for multi-modal deep generative models (2019). <https://doi.org/10.48550/ARXIV.1911.03393>, <https://arxiv.org/abs/1911.03393>
14. Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., Liu, B., Matthews, P., Ong, G., Pell, J., Silman,

- A., Young, A., Sprosen, T., Peakman, T., Collins, R.: Uk biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine* **12**, e1001779 (2015)
15. Tabachnick, B.G., Fidell, L.S.: Using multivariate statistics. Pearson, Upper Saddle River, NJ, 7 edn. (2018)
 16. Wu, M., Goodman, N.D.: Multimodal generative models for scalable weakly-supervised learning. *CoRR* **abs/1802.05335** (2018), <http://arxiv.org/abs/1802.05335>