



A multimodal variational autoencoder for estimating progression scores from imaging and microRNA data in rare neurodegenerative diseases

Virgilio Kmetzsch, Emmanuelle Becker, Dario Saracino, Vincent Anquetil, Daisy Rinaldi, Agnès Camuzat, Thomas Gareau, Isabelle Le Ber, Olivier Colliot

► To cite this version:

Virgilio Kmetzsch, Emmanuelle Becker, Dario Saracino, Vincent Anquetil, Daisy Rinaldi, et al.. A multimodal variational autoencoder for estimating progression scores from imaging and microRNA data in rare neurodegenerative diseases. SPIE Medical Imaging 2022: Image Processing, Feb 2022, San Diego, California, United States. pp.376-382, 10.1117/12.2607250 . hal-03576117

HAL Id: hal-03576117

<https://hal.science/hal-03576117>

Submitted on 16 Feb 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A multimodal variational autoencoder for estimating progression scores from imaging and microRNA data in rare neurodegenerative diseases

Virgilio Kmetzsch^{a, b}, Emmanuelle Becker^c, Dario Saracino^{a, b, d, e}, Vincent Anquetil^b, Daisy Rinaldi^{b, d, e}, Agnès Camuzat^{b, f}, Thomas Gareau^b, Isabelle Le Ber^{b, d, e, g}, Olivier Colliot^{b, a}, and The PREV-DEMALS study group

^aInria, Aramis project-team, F-75013, Paris, France

^bSorbonne Université, Paris Brain Institute – Institut du Cerveau – ICM, Inserm U1127, CNRS UMR 7225, AP-HP - Hôpital Pitié-Salpêtrière, Paris, France

^cUniv Rennes, Inria, CNRS, IRISA, F-35000 Rennes, France

^dCentre de référence des démences rares ou précoces, IM2A, AP-HP, Paris, France

^eDMU Neurosciences, AP-HP - Hôpital Pitié-Salpêtrière, Paris, France

^fEPHE, PSL Research University, Paris, France

^gParis Brain Institute – Institut du Cerveau – ICM, FrontLab, Paris, France

ABSTRACT

Frontotemporal dementia (FTD) is a rare neurodegenerative disease, often of genetic origin, with no effective treatment. There is a substantial pathophysiological overlap with amyotrophic lateral sclerosis (ALS), mutations in the *C9orf72* gene being their most common genetic cause. In these disorders, no single biomarker can accurately measure progression, thus it is crucial to combine complementary information from multiple modalities to evaluate new therapeutic interventions. In particular, neuroimaging and transcriptomic (microRNA) data have been shown to have value to track FTD and ALS progression. As these conditions are rare, large samples are not available, hence the need for methods to fuse multimodal data from small samples. In this paper, we propose a method for computing a disease progression score (DPS) from cross-sectional multimodal data, based on variational autoencoders (VAE). We show that unsupervised training leads to the estimation of meaningful latent spaces, where subjects with similar disease states are clustered together and from which a DPS may be inferred. Models were evaluated on 14 patients, 40 presymptomatic mutation carriers and 37 healthy controls from the PREV-DEMALS study. Since there is no ground truth for the DPS, we used the inferred scores to perform pairwise classification as a proxy metric. Presymptomatic subjects and patients were classified with an average area under the ROC curve of 0.83 and 0.94, respectively without and with feature selection. The proposed approach has the potential to leverage cross-sectional multimodal datasets with small sample sizes in order to objectively measure disease progression.

Keywords: Multimodal, neuroimaging, transcriptomics, microRNA, variational autoencoder, deep learning, disease progression score, neurodegenerative disease

1. INTRODUCTION

Frontotemporal dementia (FTD) is a rare heterogeneous neurodegenerative disease characterized by progressive behavioral changes, executive dysfunction and language impairments.¹ A large proportion of FTD cases are due to genetic mutations, the most frequent being an expansion in the *C9orf72* gene.^{2,3} *C9orf72* expansions are also an important genetic cause of amyotrophic lateral sclerosis (ALS), a motor neuron disease leading to muscle atrophy, progressive weakness and eventual paralysis.⁴ These fatal disorders, which may occasionally co-occur in *C9orf72*-mutated individuals, have no effective treatment to date.

Further author information: (Send correspondence to Olivier Colliot.)

Olivier Colliot: E-mail: olivier.colliot@sorbonne-universite.fr, Telephone: +33 1 57 27 43 65

Presymptomatic carriers of the *C9orf72* mutation, with no current clinical symptoms, are an ideal population for the evaluation of new disease-modifying treatments, before any irreversible brain damage has occurred. Previous work demonstrated the importance of neuroimaging⁵ and transcriptomics (microRNA)⁶ biomarkers to better understand the *C9orf72* disease progression. However, when analysed independently, neuroimaging and microRNA data provide incomplete views of FTD and ALS. Therefore, in order to monitor the effect of experimental therapies, it is critical to leverage the complementary information provided by these modalities.

Since different biomarkers characterize a disease in different stages, several biomarkers could be combined to represent the entire progression with a single disease progression score (DPS). Many approaches have been developed for data-driven disease progression modeling, including event-based models (EBM),^{7,8} a vertex-wise model of brain pathology fitted with expectation-maximisation,⁹ non-linear mixed-effects models,^{10,11} alternating least squares to fit sigmoid functions,¹² Gaussian processes,¹³ Recurrent Neural Networks¹⁴ and M-estimation.¹⁵ Most of these approaches require a large amount of longitudinal data, which is not available for FTD/ALS. The only published methods that infer a disease progression score from cross-sectional data are event-based models,^{7,8} but these approaches do not scale well for hundreds of biomarkers, such as microRNA data.

In the present work, we proposed a method for inferring a disease progression score (a latent trait) based on a multimodal variational autoencoder (VAE).¹⁶ VAEs are powerful generative models that project data in a regularized latent low dimensional space and have been shown to be effective in high dimensional low sample size settings.¹⁷ These models have already been used with multimodal data,¹⁸ although not with the goal of inferring a DPS. We hypothesized that the inferred score, based on cross-sectional neuroimaging and microRNA data, could represent the distance traveled along the underlying FTD/ALS pathophysiological pathway, and thus be used to monitor disease progression and evaluate novel treatments.

2. MATERIALS AND METHODS

2.1 Studied population

Participants were recruited through the PREV-DEMALS study (<https://clinicaltrials.gov>, ID NCT02590276), a cohort focused on *C9orf72* expansion carriers, comprising neuroimaging and microRNA sequencing data. MicroRNAs (miRNAs) are a class of noncoding RNAs that negatively regulate gene expression,¹⁹ being detected in blood plasma and correlating with the progression of many neurodegenerative diseases,²⁰ including FTD and ALS.

Our study comprised 110 individuals, divided into three groups: 22 symptomatic carriers of a pathogenic expansion (patient group), 45 asymptomatic carriers (presymptomatic group) and 43 asymptomatic non-carriers (control group). Written informed consents were obtained from all participants and the study was approved by the ethics committee (Comité de Protection des Personnes CPP Ile-De-France VI, CPP 68-15 and ID RCB 2015-A00856-43).

2.2 Data acquisition and preprocessing

All individuals had transcriptomic data available, consisting of the expression levels of 589 miRNAs. However, only 91 (14 patients, 40 presymptomatic carriers and 37 controls) had also neuroimaging data available, consisting in grey matter volumes extracted from anatomical MRI (T1) including 68 cortical regions of interest (ROIs) (Desikan atlas) and 18 subcortical ROIs (Aseg atlas) as well as the estimated total intracranial volume, thus resulting in 87 imaging features. Details regarding features and population can be found in Ref. 6 and Ref. 5. Subjects were divided into two datasets: 19 subjects with only microRNA data, used as a discovery set for feature selection, and 91 subjects with multimodal neuroimaging and microRNA data, used as input to our models. Features were rescaled from 0 to 1 and ordered via principal component analysis in the transposed data matrix: we projected features into the first principal component and used the coordinate values to sort them.

We also conducted experiments with two simulated datasets, based on the real one. To build the simulated data matrices, we simply increased (or decreased) each feature value by 5% or 15% for all patients and healthy controls, to accentuate their means' difference. The presymptomatic participants remained unchanged.

2.3 Multimodal variational autoencoder

In order to build disease progression scores, we propose a multimodal variational autoencoder for estimating a latent space representation. Let $x \in \mathcal{X}$ represent a set of multimodal data, where each point is a vector with concatenated neuroimaging and microRNA data. A variational autoencoder (VAE)¹⁶ is a generative model which aims to learn the training data distribution using a latent representation model:

$$p(x) = \int p(x|z)p(z)dz, \quad (1)$$

where $z \in \mathcal{Z}$ is a lower dimensional latent variable and $p(z)$ is its prior distribution (commonly a multivariate unit Gaussian). VAEs learn two mappings in the form of neural networks: an encoder $q_\phi(z|x)$ which maps data x to its latent representation z , and a decoder $p_\theta(x|z)$ which maps from the latent representation z back to the input space. Since the marginal log-likelihood of the data is intractable, VAEs are trained to maximize the variational lower bound of the marginal log-likelihood, known as ELBO (Evidence Lower Bound):

$$\log p(x) \geq \mathbb{E}_{z \sim q_\phi}[\log p_\theta(x|z)] - D_{KL}[q_\phi(z|x)||p(z)], \quad (2)$$

where $D_{KL}[q_\phi(z|x)||p(z)]$ is the Kullback-Leibler divergence between the approximated posterior $q_\phi(z|x)$ and the prior distribution $p(z)$ and acts as a regularization term.

Our encoder consisted of a 1-dimensional convolution layer, followed by two fully-connected layers, while the decoder was implemented with two fully-connected layers followed by a 1-dimensional transposed convolutional layer. After each layer, batch normalization²¹ was applied for its regularization properties and to avoid vanishing or exploding gradients. The nonlinear activation function was the rectified linear unit (ReLU) $f(x) = \max(0, x)$ in all layers except the decoder's last one, which used a sigmoid function $f(x) = \frac{1}{1+e^{-x}}$ in order to have the output normalized between 0 and 1. The loss function was optimized using Adam.²²

Two slightly different networks were used as our final models. For the experiments with no feature selection (589 miRNAs + 87 neuroimaging features), we identified that 64 channels and a kernel of dimension 80 with a stride of 10 was a good parametrization, along with a hidden layer of 400 units and a latent space of dimension 5. For the experiments with the discovery set and feature selection (68 miRNAs + 87 neuroimaging features), we chose 32 channels, a kernel of dimension 20 with a stride of 5, along with a hidden layer of 50 units and 2-dimensional latent space. The VAEs were implemented in Python 3.8.5 using PyTorch 1.8.1, and trained with batches of 32 subjects for 250 epochs using a learning rate of 10^{-3} .

2.4 Computing disease progression scores in the latent space

We used a stratified 5-fold cross-validation strategy, training the VAE with four folds and testing with the remaining fold in each iteration. Training was unsupervised: no clinical labels were used. Our hypothesis was that the VAE would identify a meaningful latent space, placing subjects with the same clinical status (and similar disease stage among presymptomatic individuals and patients) closer together.

Once each model was trained, we projected the training data in the latent space and used the clinical labels (patient, presymptomatic subject or control) to compute the centroid of each group. We then defined the trajectory to traverse the latent space as the line passing through the centroids of the presymptomatic and the patient groups. Finally, we encoded the test fold in the latent space and computed the DPS for each subject as the coordinate of their projection in this line.

Since there is no ground truth for the DPS, we applied a proxy metric to assess model performance: the inferred scores were used to classify subjects according to their clinical status. Therefore, labels were used during test time to compute the area under the receiver operating characteristic curve (ROC AUC) averaged over the five folds.

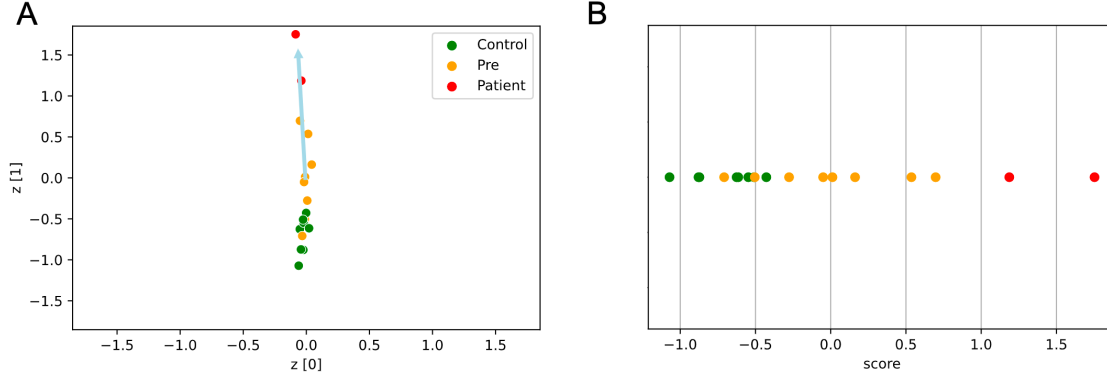


Figure 1. Example of (A) a 2-dimensional latent space encoding a test fold comprising 2 patients, 8 presymptomatic individuals and 8 controls and the trajectory to traverse this latent space (blue arrow); (B) the corresponding disease progression scores.

3. RESULTS

Figure 1 depicts an example of a two-dimensional latent space obtained after training the VAE with four folds and using the trained model to encode the remaining test fold. In this particular fold, we observe a perfect separation between patients and the other two groups, and a clear (although not perfect) distinction between presymptomatic individuals and controls.

Table 1 displays the mean and standard deviation of the area under the ROC curve obtained after a 5-fold cross-validation, for each pairwise comparison between clinical groups. The models were initially trained without any feature selection, with the expression levels of 589 miRNAs and the grey matter volumes of 87 ROIs. Then, we used 19 subjects as a discovery set to identify the most differentially expressed miRNAs between clinical groups, reducing the dimension of the microRNA data to 68. Classification performance improved when feature selection was applied. Table 1 also shows the results with the two simulated datasets. As expected, performance increases when a dataset with more discriminating features is used as input.

Table 1. Area under the ROC curve (mean \pm SD) for each pairwise classification, obtained using the inferred disease progression scores respectively without feature selection, with feature selection and with two simulated datasets.

Comparison Features	589 miRNAs, 87 T1 ROIs	68 miRNAs, 87 T1 ROIs	Simulated dataset (5%)	Simulated dataset (15%)
Control vs Presymptomatic	0.57 ± 0.15	0.62 ± 0.20	0.76 ± 0.19	0.89 ± 0.12
Control vs Patient	0.88 ± 0.11	0.95 ± 0.07	0.99 ± 0.02	1.00 ± 0.00
Presymptomatic vs Patient	0.83 ± 0.19	0.94 ± 0.12	0.95 ± 0.10	0.98 ± 0.03

Finally, Fig. 2 presents the visualization of the inferred scores for all 91 subjects after a 5-fold cross validation. The scores were computed for each subject when included in a test fold, without or with feature selection in the microRNA data. There is a superior performance (better separation between groups) when miRNAs are selected using the discovery set.

4. DISCUSSION

We proposed a multimodal variational autoencoder for combining imaging and transcriptomic (microRNA) data. It allowed inferring a single score to represent disease progression, using only cross-sectional neuroimaging and microRNA data from less than a hundred subjects. We showed that variational autoencoders built with shallow 1-dimensional convolutional neural networks were able to infer meaningful latent spaces, putting closer together subjects from the same clinical groups (patients, presymptomatic individuals and controls) without using any

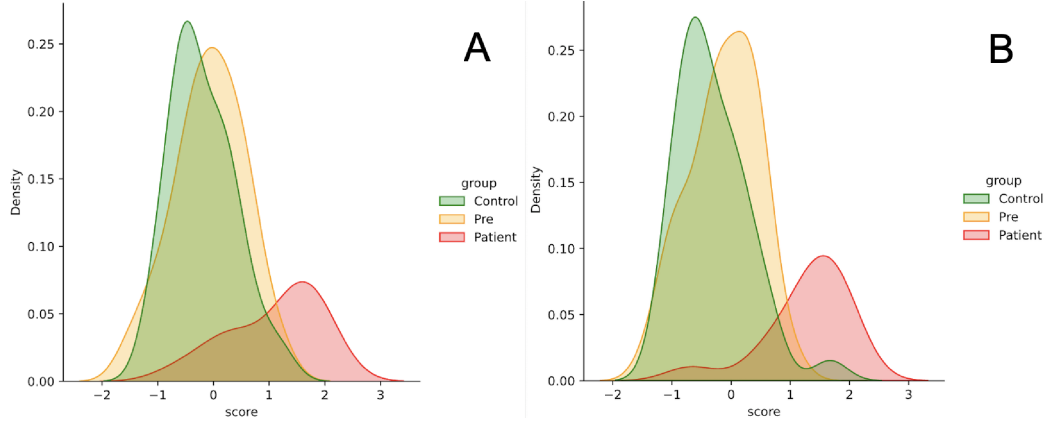


Figure 2. Disease progression score densities conditioned on clinical status, obtained from 5-fold cross-validation with 91 individuals, when using (A) all 589 miRNAs and all 87 T1 MRI ROIs and (B) selected 68 miRNAs and all 87 T1 MRI ROIs.

labels during training. We were able to encode individuals from the test sets into the latent spaces and compute their corresponding DPS. Then using only the computed scores, presymptomatic subjects and patients were distinguished with an average ROC AUC of 0.83 and 0.94, respectively without and with feature selection.

Our experiments with the simulated datasets showed that more informative features will lead to better results. In addition, the presented approach is generic enough to be used with datasets from other neurodegenerative diseases, even though our experiments focused only on *C9orf72*-associated FTD and ALS. So our results motivate further experiments with other neurodegenerative diseases with well established biomarkers.

The current study has a limitation: the absence of ground truth for the progression scores, which led us to use classification performance as a proxy metric. Long-term longitudinal data would be needed to confirm the accuracy of the inferred DPS. For instance, we hypothesize that, for presymptomatic subjects, a higher DPS implies an earlier disease onset, and we would need long-term follow-up data to confirm this hypothesis. Future work could explore different network architectures (e.g. 2-dimensional inputs, different number of layers, feature maps and kernel sizes), investigate the integration of different prior information to order the input data and analyze different methods to traverse the latent space.

In summary, our results encourage the use of the proposed approach as a tool to measure disease progression in rare neurodegenerative diseases and evaluate potential treatments.

ACKNOWLEDGMENTS

The PREV-DEMALS study group includes: Eve Benchetrit (Hôpital de la Salpêtrière, Paris), Anne Bertrand (Hôpital de la Salpêtrière, Paris), Anne Bissery (Hôpital de la Salpêtrière, Paris), Marie-Paule Boncoeur (CHU Dypuytren, Limoges), Stéphanie Bombois (CHU Roger Salengro, Lille), Agnès Camuzat (ICM, Paris), Mathieu Chastan (CHU Charles Nicolle, Rouen), Yaohua Chen (CHU Roger Salengro, Lille), Marie Chupin (ICM, Paris), Olivier Colliot (ICM, Paris), Philippe Couratier (CHU Dypuytren, Limoges), Xavier Delbeuck (CHU Roger Salengro, Lille), Vincent Deramecourt (CHU Roger Salengro, Lille), Christine Delmaire (CHU Roger Salengro, Lille), Emmanuel Gerardin (CHU Charles Nicolle, Rouen), Claude Hossein-Foucher (CHU Roger Salengro, Lille), Bruno Dubois (Hôpital de la Salpêtrière, Paris), Marie-Odile Habert (Hôpital de la Salpêtrière, Paris), Didier Hannequin (CHU Charles Nicolle, Rouen), Géraldine Lautrete (CHU Dypuytren, Limoges), Thibaud Lebouvier (CHU Roger Salengro, Lille), Isabelle Le Ber (Hôpital de la Salpêtrière, Paris), Benjamin Le Toullec (ICM, Paris), Richard Levy (Hôpital de la Salpêtrière, Paris), Olivier Martinaud (CHU Charles Nicolle, Rouen), Kelly Martineau (ICM, Paris), Marie-Anne Mackowiak (CHU Roger Salengro, Lille), Jacques Monteil (CHU Dypuytren, Limoges), Florence Pasquier (CHU Roger Salengro, Lille), Gregory Petyt (CHU Roger Salengro, Lille), Pierre-François Pradat (Hôpital de la Salpêtrière, Paris), Assi-Hervé Oya (Hôpital de la Salpêtrière, Paris), Armelle Rametti-Lacroux (Hôpital de la Salpêtrière, Paris), Daisy Rinaldi (Hôpital de la Salpêtrière, Paris), Adeline

Rollin-Sillaire (CHU Roger Salengro, Lille), François Salachas (Hôpital de la Salpêtrière, Paris), Sabrina Sayah (Hôpital de la Salpêtrière, Paris), David Wallon (CHU Charles Nicolle, Rouen).

The research leading to these results has received funding from the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute) and reference ANR-10-IAIHU-06 (Agence Nationale de la Recherche-10-IA Institut Hospitalo-Universitaire-6), from Agence Nationale de la Recherche (project PREV-DEMALS, grant number ANR-14-CE15-0016-07), and from the Inria Project Lab Program (project Neuromarkers).

We thank Justine Guégan, from iCONICS (ICM bioinformatics facility) for microRNA raw reads to counts pipeline handling. We also thank Yannick Marie and Delphine Bouteiller, from iGenSeq (ICM sequencing facility), for library preparation and sequencing. Part of this work was carried out on the DNA and cell bank of ICM. We gratefully acknowledge Philippe Martin-Hardy and Ludmila Jornea for technical assistance. The study was conducted with the support of the Centre d’Investigation Clinique Neurosciences (CIC 1422), GH Pitié Salpêtrière, Paris, and the Centre pour l’Acquisition et le Traitement des Images platform.

REFERENCES

- [1] Rascovsky, K., Hodges, J. R., Knopman, D., Mendez, M. F., Kramer, J. H., Neuhaus, J., van Swieten, J. C., Seelaar, H., Dopper, E. G. P., Onyike, C. U., Hillis, A. E., Josephs, K. A., Boeve, B. F., Kertesz, A., Seeley, W. W., Rankin, K. P., Johnson, J. K., Gorno-Tempini, M.-L., Rosen, H., Prioleau-Latham, C. E., Lee, A., Kipps, C. M., Lillo, P., Piguet, O., Rohrer, J. D., Rossor, M. N., Warren, J. D., Fox, N. C., Galasko, D., Salmon, D. P., Black, S. E., Mesulam, M., Weintraub, S., Dickerson, B. C., Diehl-Schmid, J., Pasquier, F., Deramecourt, V., Lebert, F., Pijnenburg, Y., Chow, T. W., Manes, F., Grafman, J., Cappa, S. F., Freedman, M., Grossman, M., and Miller, B. L., “Sensitivity of revised diagnostic criteria for the behavioural variant of frontotemporal dementia,” *Brain* **134**, 2456–2477 (Sept. 2011).
- [2] DeJesus-Hernandez, M., Mackenzie, I. R., Boeve, B. F., Boxer, A. L., Baker, M., Rutherford, N. J., Nicholson, A. M., Finch, N. A., Flynn, H., Adamson, J., Kouri, N., Wojtas, A., Sengdy, P., Hsiung, G.-Y. R., Karydas, A., Seeley, W. W., Josephs, K. A., Coppola, G., Geschwind, D. H., Wszolek, Z. K., Feldman, H., Knopman, D. S., Petersen, R. C., Miller, B. L., Dickson, D. W., Boylan, K. B., Graff-Radford, N. R., and Rademakers, R., “Expanded GGGGCC hexanucleotide repeat in noncoding region of C9ORF72 causes chromosome 9p-linked FTD and ALS,” *Neuron* **72**, 245–256 (Oct. 2011).
- [3] Renton, A. E., Majounie, E., Waite, A., Simón-Sánchez, J., Rollinson, S., Gibbs, J. R., Schymick, J. C., Laaksovirta, H., van Swieten, J. C., Myllykangas, L., Kalimo, H., Paetau, A., Abramzon, Y., Remes, A. M., Kaganovich, A., Scholz, S. W., Duckworth, J., Ding, J., Harmer, D. W., Hernandez, D. G., Johnson, J. O., Mok, K., Rytén, M., Trabzuni, D., Guerreiro, R. J., Orrell, R. W., Neal, J., Murray, A., Pearson, J., Jansen, I. E., Sondervan, D., Seelaar, H., Blake, D., Young, K., Halliwell, N., Callister, J. B., Toulson, G., Richardson, A., Gerhard, A., Snowden, J., Mann, D., Neary, D., Nalls, M. A., Peuralinna, T., Jansson, L., Isoviiita, V.-M., Kaivorinne, A.-L., Hölttä-Vuori, M., Ikonen, E., Sulkava, R., Benatar, M., Wu, J., Chiò, A., Restagno, G., Borghero, G., Sabatelli, M., ITALSGEN Consortium, Heckerman, D., Rogaeva, E., Zinman, L., Rothstein, J. D., Sendtner, M., Drepper, C., Eichler, E. E., Alkan, C., Abdullaev, Z., Pack, S. D., Dutra, A., Pak, E., Hardy, J., Singleton, A., Williams, N. M., Heutink, P., Pickering-Brown, S., Morris, H. R., Tienari, P. J., and Traynor, B. J., “A hexanucleotide repeat expansion in C9ORF72 is the cause of chromosome 9p21-linked ALS-FTD,” *Neuron* **72**, 257–268 (Oct. 2011).
- [4] Pasinelli, P. and Brown, R. H., “Molecular biology of amyotrophic lateral sclerosis: insights from genetics,” *Nature Reviews Neuroscience* **7**, 710–723 (Sept. 2006). Number: 9 Publisher: Nature Publishing Group.
- [5] Bertrand, A., Wen, J., Rinaldi, D., Houot, M., Sayah, S., Camuzat, A., Fournier, C., Fontanella, S., Routier, A., Couratier, P., Pasquier, F., Habert, M.-O., Hannequin, D., Martinaud, O., Caroppo, P., Levy, R., Dubois, B., Brice, A., Durrleman, S., Colliot, O., Le Ber, I., and Predict to Prevent Frontotemporal Lobar Degeneration and Amyotrophic Lateral Sclerosis (PREV-DEMALS) Study Group, “Early Cognitive, Structural, and Microstructural Changes in Presymptomatic C9orf72 Carriers Younger Than 40 Years,” *JAMA neurology* **75**(2), 236–245 (2018).

- [6] Kmetzsch, V., Anquetil, V., Saracino, D., Rinaldi, D., Camuzat, A., Gareau, T., Jornea, L., Forlani, S., Couratier, P., Wallon, D., Pasquier, F., Robil, N., Grange, P. d. l., Moszer, I., Ber, I. L., Colliot, O., and Becker, E., “Plasma microRNA signature in presymptomatic and symptomatic subjects with C9orf72-associated frontotemporal dementia and amyotrophic lateral sclerosis,” *Journal of Neurology, Neurosurgery & Psychiatry* (Nov. 2020). Publisher: BMJ Publishing Group Ltd Section: Neurodegeneration.
- [7] Fonteijn, H. M., Modat, M., Clarkson, M. J., Barnes, J., Lehmann, M., Hobbs, N. Z., Scahill, R. I., Tabrizi, S. J., Ourselin, S., Fox, N. C., and Alexander, D. C., “An event-based model for disease progression and its application in familial Alzheimer’s disease and Huntington’s disease,” *NeuroImage* **60**, 1880–1889 (Apr. 2012).
- [8] Venkatraghavan, V., Bron, E. E., Niessen, W. J., and Klein, S., “Disease progression timeline estimation for Alzheimer’s disease using discriminative event based modeling,” *NeuroImage* **186**, 518–532 (Feb. 2019).
- [9] Marinescu, R. V., Eshaghi, A., Lorenzi, M., Young, A. L., Oxtoby, N. P., Garbarino, S., Crutch, S. J., Alexander, D. C., and Alzheimer’s Disease Neuroimaging Initiative, “DIVE: A spatiotemporal progression model of brain pathology in neurodegenerative disorders,” *NeuroImage* **192**, 166–177 (May 2019).
- [10] Schiratti, J.-B., Allasonnière, S., Colliot, O., and Durrleman, S., “A Bayesian Mixed-Effects Model to Learn Trajectories of Changes from Repeated Manifold-Valued Observations,” *Journal of Machine Learning Research* **18**(133), 1–33 (2017).
- [11] Koval, I., Bône, A., Louis, M., Lartigue, T., Bottani, S., Marcoux, A., Samper-González, J., Burgos, N., Charlier, B., Bertrand, A., Epelbaum, S., Colliot, O., Allasonnière, S., and Durrleman, S., “AD Course Map charts Alzheimer’s disease progression,” *Scientific Reports* **11**, 8020 (Apr. 2021). Bandiera_abtest: a Cc_license.type: cc_by Cg_type: Nature Research Journals Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject.term: Computational science;Computer science Subject.term.id: computational-science;computer-science.
- [12] Jedynak, B. M., Lang, A., Liu, B., Katz, E., Zhang, Y., Wyman, B. T., Raunig, D., Jedynak, C. P., Caffo, B., and Prince, J. L., “A computational neurodegenerative disease progression score: Method and results with the Alzheimer’s disease neuroimaging initiative cohort,” *NeuroImage* **63**, 1478–1486 (Nov. 2012).
- [13] Lorenzi, M., Filippone, M., Frisoni, G. B., Alexander, D. C., and Ourselin, S., “Probabilistic disease progression modeling to characterize diagnostic uncertainty: Application to staging and prediction in Alzheimer’s disease,” *NeuroImage* **190**, 56–68 (Apr. 2019).
- [14] Mehdipour Ghazi, M., Nielsen, M., Pai, A., Cardoso, M. J., Modat, M., Ourselin, S., and Sørensen, L., “Training recurrent neural networks robust to incomplete data: Application to Alzheimer’s disease progression modeling,” *Medical Image Analysis* **53**, 39–46 (Apr. 2019).
- [15] Mehdipour Ghazi, M., Nielsen, M., Pai, A., Modat, M., Jorge Cardoso, M., Ourselin, S., and Sørensen, L., “Robust parametric modeling of Alzheimer’s disease progression,” *NeuroImage* **225**, 117460 (Jan. 2021).
- [16] Kingma, D. P. and Welling, M., “Auto-Encoding Variational Bayes,” *arXiv:1312.6114 [cs, stat]* (May 2014). arXiv: 1312.6114.
- [17] Chadebec, C., Thibeu-Sutre, E., Burgos, N., and Allasonnière, S., “Data Augmentation in High Dimensional Low Sample Size Setting Using a Geometry-Based Variational Autoencoder,” *arXiv:2105.00026 [cs, stat]* (Apr. 2021). arXiv: 2105.00026.
- [18] Antelmi, L., Ayache, N., Robert, P., and Lorenzi, M., “Sparse Multi-Channel Variational Autoencoder for the Joint Analysis of Heterogeneous Data,” in *[International Conference on Machine Learning]*, 302–311, PMLR (May 2019). ISSN: 2640-3498.
- [19] Huntzinger, E. and Izaurralde, E., “Gene silencing by microRNAs: contributions of translational repression and mRNA decay,” *Nature Reviews Genetics* **12**, 99–110 (Feb. 2011). Number: 2 Publisher: Nature Publishing Group.
- [20] Grasso, M., Piscopo, P., Confaloni, A., and Denti, M. A., “Circulating miRNAs as biomarkers for neurodegenerative disorders,” *Molecules (Basel, Switzerland)* **19**, 6891–6910 (May 2014).
- [21] Ioffe, S. and Szegedy, C., “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift,” *arXiv:1502.03167 [cs]* (Mar. 2015). arXiv: 1502.03167.
- [22] Kingma, D. P. and Ba, J., “Adam: A Method for Stochastic Optimization,” *arXiv:1412.6980 [cs]* (Jan. 2017). arXiv: 1412.6980.