



Alzheimer's Disease Modelling and Staging Through Independent Gaussian Process Analysis of Spatio-Temporal Brain Changes

Clement Abi Nader^{1(✉)}, Nicholas Ayache¹, Philippe Robert^{2,3}, Marco Lorenzi¹, and for the Alzheimer's Disease Neuroimaging Initiative

¹ UCA, Inria Sophia Antipolis, Epione Research Project, Sophia Antipolis, France
{clement.abi-nader,nicholas.ayache,marco.lorenzi}@inria.fr

² UCA, CoBTeK, Nice, France

³ Centre Memoire, CHU de Nice, Nice, France
probert@unice.fr

Abstract. Alzheimer's disease (AD) is characterized by complex and largely unknown progression dynamics affecting the brain's morphology. Although the disease evolution spans decades, to date we cannot rely on long-term data to model the pathological progression, since most of the available measures are on a short-term scale. It is therefore difficult to understand and quantify the temporal progression patterns affecting the brain regions across the AD evolution. In this work, we present a generative model based on probabilistic matrix factorization across temporal and spatial sources. The proposed method addresses the problem of disease progression modelling by introducing clinically-inspired statistical priors. To promote smoothness in time and model plausible pathological evolutions, the temporal sources are defined as monotonic and independent Gaussian Processes. We also estimate an individual time-shift parameter for each patient to automatically position him/her along the sources time-axis. To encode the spatial continuity of the brain substructures, the spatial sources are modeled as Gaussian random fields. We test our algorithm on grey matter maps extracted from brain structural images. The experiments highlight differential temporal progression patterns mapping brain regions key to the AD pathology, and reveal a disease-specific time scale associated with the decline of volumetric biomarkers across clinical stages.

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.

© Springer Nature Switzerland AG 2018

D. Stoyanov et al. (Eds.): MLCN 2018/DLF 2018/iMIMIC 2018, LNCS 11038, pp. 3–14, 2018.

https://doi.org/10.1007/978-3-030-02628-8_1

1 Introduction

Neurodegenerative disorders such as Alzheimer’s disease (AD) are characterized by morphological and molecular changes of the brain, and ultimately lead to cognitive and behavioral decline [8]. To date there is no clear understanding of the dynamics regulating the disease progression. Consequently several attempts have been made to model the disease evolution in a data-driven way, using sets of biomarkers extracted from different imaging acquisition techniques, such as Magnetic Resonance Imaging (MRI) [12]. However available data are mostly represented by cross-sectional measures or time-series acquired on a short-term time span, while the ultimate goal is to unveil the “long-term” disease evolution spreading over decades. Therefore there is a critical need to define the AD evolution in a data-driven manner with respect to an absolute time scale associated to the natural history of the pathology.

To this end, in [9] the authors introduce a disease progression score for each patient in order to identify a data-driven disease scale. This score is based on a set of biomarkers and was shown to correlate with the decline of brain cognitive abilities. A similar approach was proposed by [12] and [6] with scalar biomarkers. In [3], a disease progression score was estimated using higher-dimensional biomarkers from molecular imaging. However these methods don’t provide information about the brain structures involved in AD, and how the disease affects them along time. To overcome these limitations, [13] proposes a spatio-temporal model of disease progression explicitly accounting for different temporal dynamics across the brain. This is done by decomposing cortical thickness measurements as a mixture of spatio-temporal processes, by associating each vertex to a temporal progression modeled by a sigmoid function. They also estimate a disease progression score for each subject as a linear transformation of time. However since the proposed formulation does not account for spatial correlation between vertices, it may be potentially sensitive to spatial variation and noise, thus leading to poor interpretability.

The challenge of spatio-temporal modelling in brain images is a classical problem widely addressed via Independent Component Analysis (ICA [7]), especially on functional MRI (fMRI) data [4]. ICA aims at decomposing the data via matrix factorization, looking for a reduced number of spatio-temporal latent sources. Although successful in fMRI analysis, ICA cannot find straightforward applications to the modelling of AD progression. First, ICA retrieves maximally independent latent sources best explaining the data. However, although brain regions can exhibit different atrophy rates, this doesn’t necessarily imply statistical independence between them. Second, differently from fMRI data, the absolute time axis of AD spatio-temporal observations is unknown. Thus estimating the pathology timing is a key step in order to model the disease progression, and cannot be performed with standard dimensionality reduction methods such as ICA. Finally, fMRI time series are defined over hundreds of time points, while we work essentially in a cross-sectional setting with one or a few images per-subject.

In this work we present a novel spatio-temporal generative model of disease progression aimed at quantifying the independent dynamics of changes in the

brain. We model the observed data through matrix factorization across temporal and spatial sources, with a plausibility constraint introduced by clinically-inspired statistical priors. To promote smoothness in time and model steady evolution from normal to pathological stages, the temporal sources are defined as monotonic independent Gaussian Processes (GPs). We also estimate an individual time-shift parameter for each patient to automatically position him along the sources time-axis. To encode the spatial continuity of the brain sub-structures, the spatial sources are modeled as Gaussian random fields. The framework is efficiently optimized through stochastic variational inference. In the next sections we detail the method formulation and show its application on synthetic and real data composed by a large dataset of MRIs from the Alzheimer’s Disease Neuroimaging Initiative (ADNI). Further information can be found in the Appendix.¹

2 Method

We assume that the spatio-temporal data $Y(x, t) = [Y_1(x, t_1), Y_2(x, t_2), \dots, Y_P(x, t_p)]$ is stored in a matrix with dimensions $P \times F$, where P is the number of patients, F the number of image features, and $Y_i(x, t_i)$ is the image of an individual i observed at position x and at time t_i . We postulate a generative model in order to decompose the data in N_s spatio-temporal sources such that:

$$Y_p(x, t_p) = S(\theta, t + t_p)A(\psi, x) + \mathcal{E} \quad (1)$$

where S is a $P \times N_s$ matrix where each column represents a temporal trajectory, t_p the individual time-shift parameter, and θ the set of parameters related to the temporal sources. A is a $N_s \times F$ matrix where each row represents a spatial map, and ψ is a set of spatial parameters. \mathcal{E} is a $\mathcal{N}(0, \sigma^2 I)$ Gaussian noise. According to the generative model the likelihood is:

$$p(Y|A, S, \sigma) = \prod_{p=1}^P \frac{1}{(2\pi\sigma^2)^{\frac{F}{2}}} \exp\left(-\frac{1}{2\sigma^2} \|Y_p - S(\theta, t + t_p)A(\psi, x)\|^2\right) \quad (2)$$

For each row A_n of A we specify a $\mathcal{N}(0, I)$ prior, while each column S_n of S is a GP modeled as in [5]. This setting leverages on kernel approximation through sampling of basis functions in the spectral domain [14]. For specific choices of the covariance, such as the Radial Basis Function used in our work, the GPs can be approximated as a Bayesian neural network with form: $S(t) = \phi(\Omega t)W$. Where Ω is the projection in the spectral domain, ϕ the non-linear basis function activation, and W the regression parameter. The GPs inference problem thus amounts at estimating approximated distributions for Ω and W .

To account for the steady increase of the sources from normal to pathological stages we introduce a monotonicity prior over the GPs. To do so, we constrain the space of the temporal sources to the set $\mathcal{C} = \{S(t) \mid S'(t) \leq 0 \quad \forall t\}$, following

¹ Appendix: <https://hal.archives-ouvertes.fr/hal-01849180/document>.

[11]. This leads to a second likelihood term constraining the dynamics of the temporal sources:

$$p(\mathcal{C}|S', \lambda) = (1 + \exp(-\lambda S'(t)))^{-1} \quad (3)$$

We jointly optimize (2) according to priors and constraints, by maximizing the data evidence:

$$\log(p(Y, \mathcal{C}|\sigma, \lambda)) = \log\left[\int_A \int_S \int_{S'} p(Y|A, S, \sigma)p(\mathcal{C}|S', \lambda)p(A)p(S, S'|\lambda)dAdSdS'\right] \quad (4)$$

Since this integral is intractable, we tackle the optimization of (4) via stochastic variational inference. Following [10] and [5] we introduce approximations $q_1(A)$ and $q_2(\Omega, W)$ to derive the lower bound:

$$\begin{aligned} \log(p(Y, \mathcal{C}|\sigma, \lambda)) &\geq \mathbb{E}_{A \sim q_1, (\Omega, W) \sim q_2} [\log(p(Y|A, \Omega, W, \sigma))] + \mathbb{E}_{(\Omega, W) \sim q_2} [\log(p(\mathcal{C}|\Omega, W, \lambda))] \\ &\quad - \mathcal{D}[q_1(A)||p(A)] - \mathcal{D}[q_2(\Omega, W)||p(\Omega, W)] \end{aligned} \quad (5)$$

where \mathcal{D} refers to the Kullback-Leibler divergence.

We specify the approximated distribution of the spatial activation maps q_1 such that $q_1(A) = \prod_{n=1}^{N^s} \mathcal{N}(\mu_n, \Sigma(\alpha, \beta))$. To introduce spatial correlations in the maps we choose $\Sigma_{i,j}(\alpha, \beta) = \alpha \exp(-\|u_i - u_j\|^2/2\beta)$ to model a smooth decay across voxels with coordinates (u_i, u_j) . We follow [5] and [11] to also define a variational lower bound on the constrained GPs parameterizing the temporal processes. Thanks to the proposed framework, (4) can be efficiently optimized by stochastic variational inference through backpropagation. We chose to alternate the optimization between the spatio-temporal parameters and the time-shift. We set λ to the minimum value that gives monotonic sources, while σ was arbitrarily determined from the data. A detailed derivation of the model and lower-bound can be found in the Appendix.

3 Results

3.1 Benchmark on Synthetic Data

We tested the algorithm on synthetic data to assess its ability to separate spatio-temporal sources from mixed data, and to provide a model selection via the variational lower bound. We generated three monotonically increasing functions $S_i(t)$ such that $S_i(t) = 1/(1 + \exp(-t + \alpha_i))$, and three synthetic Gaussian activation maps A_1, A_2, A_3 with a 30×30 resolution, to mimic grey matter brain areas (Figs. 1a and b). The data was generated as $Y_{p,j} = S(t_p)A + \mathcal{E}_j$ over 40 time points t_p , where t_p is uniformly distributed in $[0,1]$. We sampled 50 images at instants t_p and applied our method. To simulate a pure cross-sectional setting the time associated to each input image was set to zero. Figures 1c and d show the estimated spatio-temporal processes when fitting the model with three latent sources. In Fig. 2, we see that the individual time-shift parameter estimated for each subject correlates with the original time used to generate the data. This means that the algorithm correctly positions each subject on the temporal trajectories.

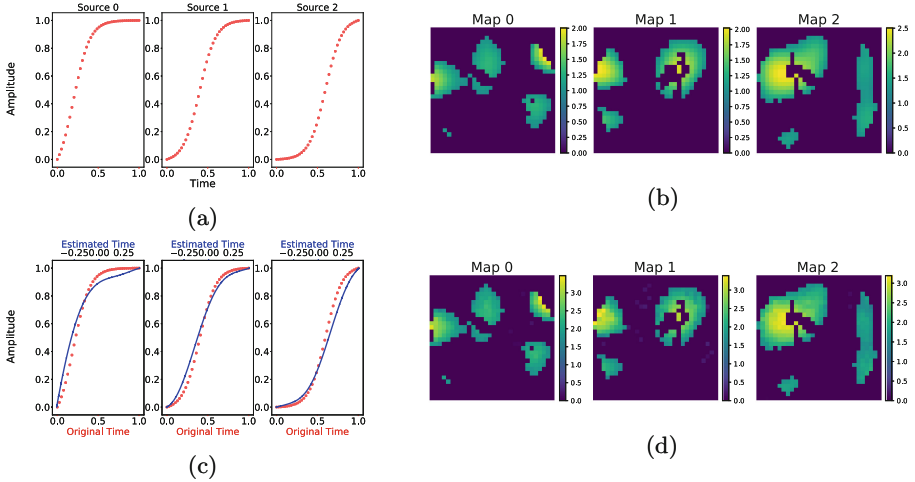


Fig. 1. (a)–(b) Ground truth temporal and spatial sources. (c) Red: raw temporal sources against the original time axis. Blue: recovered temporal sources against the estimated time scale. (d) Estimated spatial maps. (Color figure online)

To test the model selection, we generated the data as described above using respectively one, two, or three sources over ten folds. For each fold we ran the algorithm looking for one to four sources. Figure 3 shows mean and standard deviation of the lower bound. We observe that when the number of sources is under-estimated the lower bound is higher. When the number of sources is over-estimated, although the lower bound for model selection is more uncertain, by looking at the extracted spatial maps we observe that the additional sources are mainly set to zero or have low weights (see the map of Fig. 3). These experimental results indicate that the optimal number of sources should be selected by inspection of both the lower bound and the extracted spatial sources.

The method was also compared to ICA in a simplified setting by assigning the ground truth parameter t_p beforehand. This simplification is necessary since standard ICA can't be applied when the time associated to each image is unknown. We observed that ICA recovered the spatio-temporal sources, by providing however more noisy estimations than the ones we obtained. This result highlights the importance of the priors and constraints introduced in our method (see Appendix).




Fig. 2. The red points represent the values of the estimated subjects' time-shift against their associated ground truth value. (Color figure online)




Fig. 3. (a)–(b)–(c): Distribution of the lower bound against the number of fitted sources. (d): 4th extracted spatial map with data generated by 3 latent sources.

3.2 Application on Real Data

Data used in the preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. For up-to-date information, see www.adni-info.org.

In this section we present an application of the algorithm on real data, using grey matter maps extracted from structural MRI. We selected a cohort of 555 subjects from ADNI composed by 94 healthy controls, 343 MCI, and 118 AD patients. We processed the baseline MRI of each subject to obtain high-dimensional grey matter density maps in a standard space [1]. We extracted the 90×100 middle coronal slice for each patient, to obtain a data matrix Y with dimensions 555×9000 , and applied our algorithm looking for three spatio-temporal sources (see Fig. 4). The middle spatial map shows a strong activation of the hippocampus, while the left and right plots show an activation on the temporal lobes, with two similar temporal behaviours, characterized by a less pronounced grey matter loss compared to the hippocampus. More specifically, we observe that the hippocampal trajectory has a strong acceleration in opposition to the other brain areas. This pattern quantified by our model in a pure data-driven manner is compatible with empirical evidence from clinical studies [2]. In Fig. 5 we observe the estimated time of each patient against standard

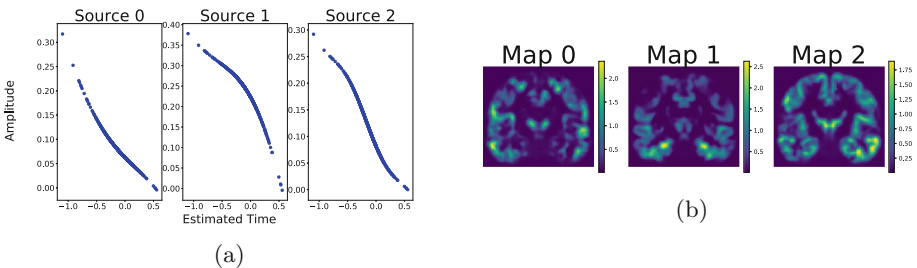


Fig. 4. (a)–(b) Temporal and spatial sources extracted from the data.

volumetric and clinical biomarkers. We see a strong correlation between brain volumetric measures and the estimated time, as well as a non-linear relation in the evolution of ADAS11. The latter result indicates an acceleration of clinical symptoms along the estimated time course.




Fig. 5. Evolution of volumetric and clinical biomarkers along the estimated time.

4 Conclusion

We presented a method for analyzing spatio-temporal data, which provides both independent spatio-temporal processes at stake in AD, and a disease progression scale. Applied on grey matter maps, the model highlights different brain regions affected by the disease, such as the hippocampus and the temporal lobes, along with their differential temporal trajectory. We also show a strong correlation between the estimated disease progression scale and different clinical and volumetric biomarkers. We are currently extending the approach to scale to 3D volumetric images by parallelization on multiple GPUs. The lower bound properties will be also further investigated to better assess its reliability, in order to improve the model comparison. Moreover the method will be extended beyond the cross-sectional application of Sect. 3.2, to account for time-series of brain images, as well as for multimodal imaging biomarkers. Finally we will investigate the use of the approach for prognosis purposes, to provide a data-driven assessment of disease severity in testing patients.

Acknowledgements. This work has been supported by the French government, through the UCA^{JEDI} Investments in the Future project managed by the National Research Agency (ref.n ANR-15-IDEX-01), the grant AAP Santé 06 2017-260 DGA-DSH, and by the Inria Sophia Antipolis - Méditerranée, “NEF” computation cluster.

Appendix

A. Lower bound derivation

In this section we detail the derivation of the lower bound:

$$\begin{aligned}
 \log(p(Y, C | \sigma, \lambda)) &= \log\left[\int_A \int_S \int_{S'} p(Y|A, S, \sigma) p(C|S', \lambda) p(A) p(S, S' | \lambda) dA dS dS'\right] \\
 &= \log\left[\int_A \int_S \int_{S'} p(Y|A, S, \sigma) p(C|S', \lambda) p(A) p(S'|S, \lambda) p(S) dA dS dS'\right]
 \end{aligned} \tag{1}$$

If we know S this completely determines S' , thus we have $\int p(S'|S, \lambda) dS' = 1$ which gives us:

$$\begin{aligned}
\log(p(Y, \mathcal{C}|\sigma, \lambda)) &= \log\left[\int_A \int_S p(Y|A, S, \sigma)p(\mathcal{C}|S', \lambda)p(A)p(S)dAdS\right] \\
&= \log\left[\int_A \int_S p(Y|A, S, \sigma)p(\mathcal{C}|S', \lambda)p(A)p(S)\frac{q_1(A)q_2(S)}{q_1(A)q_2(S)}dAdS\right] \\
&= \log\left[\mathbb{E}_{A\sim q_1, S\sim q_2}\left[\frac{p(Y|A, S, \sigma)p(\mathcal{C}|S', \lambda)p(A)p(S)}{q_1(A)q_2(S)}\right]\right] \\
&\geq \mathbb{E}_{A\sim q_1, S\sim q_2}\left[\log\left[\frac{p(Y|A, S, \sigma)p(\mathcal{C}|S', \lambda)p(A)p(S)}{q_1(A)q_2(S)}\right]\right]
\end{aligned} \tag{2}$$

This is obtained thanks to Jensen's inequality. Finally this leads us to:

$$\begin{aligned}
\mathbb{E}_{A\sim q_1, S\sim q_2}\left[\log\left[\frac{p(Y|A, S, \sigma)p(\mathcal{C}|S', \lambda)p(A)p(S)}{q_1(A)q_2(S)}\right]\right] &= \mathbb{E}_{A\sim q_1, S\sim q_2}[\log[p(Y|A, S, \sigma)]] \\
&\quad + \mathbb{E}_{S\sim q_2}[\log(P(\mathcal{C}|S', \lambda))] \\
&\quad - \mathcal{D}[q_1(A|Y)||p(A)] \\
&\quad - \mathcal{D}[q_2(S|Y)||p(S)]
\end{aligned} \tag{3}$$

In the Method section we introduced the approximation $q_1(A) = \prod_{n=1}^{Ns} \mathcal{N}(\mu_n, \Sigma(\alpha, \beta))$. The covariance matrix is shared by all the spatial processes which gives us the set of spatial parameters:

$$\psi = \{\mu_n, n \in [1, Ns], \alpha, \beta\} \tag{4}$$

Following [5] we introduce for each GP two vectors, Ω_n with a prior $p(\Omega_n) = \mathcal{N}(0, \frac{1}{l_n}I)$ for each element and W_n with a prior $p(W_n) = \mathcal{N}(0, I)$, such that $S_n(t) = \Phi(t\Omega_n)W_n$. Where Φ is chosen to obtain a RBF kernel as explained in [5]. We define the approximated distributions $q_3(W_n) = \prod_j \mathcal{N}(m_{n,j}, s_{n,j}^2)$ and $q_4(\Omega_n) = \prod_j \mathcal{N}(\alpha_{n,j}, \beta_{n,j}^2)$ of $p(W_n)$ and $p(\Omega_n)$. Using these approximations and following [5], we can derive a lower bound for S with the same technique than above. We have the set of temporal parameters:

$$\theta = \{m_n, s_n, \alpha_n, \beta_n, l_n, n \in [1, Ns]\} \tag{5}$$

Now we can obtain every term of (3). The Kullback-Leibler of a multivariate Gaussian has a closed-form:

$$\mathcal{D}[q_1(A|X)||p(A)] = \frac{1}{2} \sum_{n=1}^{Ns} Tr(\Sigma) + \mu_n^T \mu_n - F - \log[\det(\Sigma)] \tag{6}$$

Using the factorized form of q_2 and the fact that the different Gaussian processes are independent from each other we can write:

$$\mathcal{D}[q_2(S|X)||p(S)] = \sum_{n=1}^{N_s} \mathcal{D}[q_3(W_n)|p(W_n)] + \mathcal{D}[q_4(\Omega_n)|p(\Omega_n)] \quad (7)$$

Since the approximations q_3 and q_4 and their respective priors are normally distributed we have an analytic formula for both Kullback-Leibler divergences.

$$\mathcal{D}[q_3(W_n)|p(W_n)] = \frac{1}{2} \sum_j s_{n,j}^2 + \mu_{n,j}^2 - 1 - \log(s_{n,j}^2) \quad (8)$$

$$\mathcal{D}[q_4(\Omega_n)|p(\Omega_n)] = \frac{1}{2} \sum_j \beta_{n,j}^2 l_n + \alpha_{n,j}^2 l_n - 1 - \log(\beta_{n,j}^2 l_n) \quad (9)$$

As in [10] we employ the reparameterization trick to have an efficient way of sampling the expectations of (3). Thus we have:

$$\begin{aligned} - W_{n,j} &= m_{n,j} + s_{n,j} * \epsilon_{n,j} \\ - \Omega_{n,j} &= \alpha_{n,j} + \beta_{n,j} * \zeta_{n,j} \\ - A_n &= \mu_n + \Sigma_n^{\frac{1}{2}} * \kappa_n \end{aligned}$$

Which gives us:

$$\mathbb{E}_{A \sim q_1, S \sim q_2}[\log(p(Y|A, S, \sigma))] = \mathbb{E}_{\epsilon, \zeta, \kappa}[\log(p(Y|m, s, \alpha, \beta, \mu, \Sigma, \sigma))] \quad (10)$$

$$\mathbb{E}_{S \sim q_2}[\log(p(\mathcal{C}|S', \lambda))] = \mathbb{E}_{\epsilon, \zeta}[\log(p(\mathcal{C}|m, s, \alpha, \beta, \lambda))] \quad (11)$$

Where $\epsilon_{n,j} \sim \mathcal{N}(0, 1)$, $\zeta_{n,j} \sim \mathcal{N}(0, 1)$ and $\kappa_n \sim \mathcal{N}(0, I)$.

B. Kronecker factorization

Here we detail how to split the covariance matrix in a Kronecker product of three matrices along each spatial dimensions. We have:

$$\Sigma_{i,j}(\alpha, \beta) = \alpha \exp\left(-\frac{\|u_i - u_j\|^2}{2\beta}\right) \quad (12)$$

We can use the separability properties of the exponential to decompose the covariance between two locations $u_i = (x_i, y_i, z_i)$ and $u_j = (x_j, y_j, z_j)$:

$$\Sigma_{i,j}(\alpha, \beta) = \alpha \exp\left(-\frac{(x_i - x_j)^2}{2\beta}\right) \exp\left(-\frac{(y_i - y_j)^2}{2\beta}\right) \exp\left(-\frac{(z_i - z_j)^2}{2\beta}\right) \quad (13)$$

So Σ can be decomposed into the Kronecker product of 1D processes:

$$\Sigma = \Sigma_x \otimes \Sigma_y \otimes \Sigma_z \quad (14)$$

Allowing us to deal with large-size matrices.

C. Comparison with ICA

We performed a comparison of our algorithm with ICA on a similar example than in Sect. 3.1. However the data was generated in a simplified setting since ICA can't be applied when the time associated to each image is unknown. To do so we assigned the ground truth parameter t_p beforehand. The goal was to compare the separation performances of both our algorithm and ICA, on data generated with three latent spatio-temporal processes. In Fig. 6 we observe that the sources estimated by ICA are more noisy and uncertain than the ones estimated by our method, highlighting the performances of our algorithm in terms of sources separation.




Fig. 6. First row: raw sources. Second row: sources estimated by our method. Third row: sources estimated by ICA.

D. ADNI

Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) and DOD ADNI. ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and

Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer’s Association; Alzheimer’s Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd. and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer’s Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

References

1. Ashburner, J.: A fast diffeomorphic image registration algorithm. *NeuroImage* **38**(1), 95–113 (2007)
2. Bateman, R.J., et al.: Clinical and biomarker changes in dominantly inherited Alzheimer’s disease. *New Engl. J. Med.* **367**(9), 795–804 (2012). PMID: 22784036
3. Bilgel, M., et al.: Temporal trajectory and progression score estimation from voxelwise longitudinal imaging measures: application to amyloid imaging. *Inf.Process. Med. Imaging* **24**, 424–436 (2015)
4. Calhoun, V.D., et al.: A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and ERP data. *Neuroimage* **45**(1 Suppl), S163–S172 (2009)
5. Cutajar, K., et al.: Random feature expansions for deep Gaussian processes. In: Precup, D., Teh, Y.W. (eds.) *Proceedings of the 34th International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 70, pp. 884–893. PMLR, International Convention Centre, Sydney, Australia, 06–11 August 2017
6. Donohue, M.C., et al.: Estimating long-term multivariate progression from short-term data. *Alzheimer’s Dementia* **10**(Suppl. 5), S400–S410 (2014)
7. Hyvärinen, A., Oja, E.: Independent component analysis: algorithms and applications. *Neural Netw.* **13**, 411–430 (2000)
8. Jack, C.R.: Hypothetical model of dynamic biomarkers of the Alzheimer’s pathological cascade. *Lancet Neurol* **9**(1), 119–128 (2010)
9. Jernigan, B.M.: A computational neurodegenerative disease progression score: method and results with the Alzheimer’s disease neuroimaging initiative cohort. *Neuroimage* **63**(3), 1478–1486 (2012)
10. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *CoRR* abs/1312.6114 (2013)

11. Lorenzi, M., Filippone, M.: Constraining the dynamics of deep probabilistic models. In: Dy, J., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning, Proceedings of Machine Learning Research, vol. 80, pp. 3233–3242. PMLR, Stockholmsmässan, Stockholm, Sweden, 10–15 July 2018
12. Lorenzi, M., et al.: Probabilistic disease progression modeling to characterize diagnostic uncertainty: application to staging and prediction in Alzheimer’s disease. *NeuroImage* (2017)
13. Marinescu, R.V., et al.: A vertex clustering model for disease progression: application to cortical thickness images. In: Niethammer, M., et al. (eds.) IPMI 2017. LNCS, vol. 10265, pp. 134–145. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59050-9_11
14. Rahimi, A., Recht, B.: Random features for large-scale kernel machines. In: Platt, J.C. (ed.) *Advances in Neural Information Processing Systems*, vol. 20, pp. 1177–1184. Curran Associates Inc., New York (2008)