

Disease Progression Modelling and Stratification for detecting sub-trajectories in the natural history of pathologies: application to Parkinson’s Disease trajectory modelling

Alessandro Viani¹, Boris A. Gutman², Emile d’Angremont³, and Marco Lorenzi¹

¹ Epione Research Team, Inria Center of Université Côte d’Azur, Sophia Antipolis, France,

² Illinois Institute of Technology, Department of Biomedical Engineering, USA

³ Amsterdam University Medical Center, Department of Anatomy and Neurosciences, The Netherlands
`alessandro.viani@inria.fr`

Abstract. Modelling the progression of Degenerative Diseases (DD) is essential for detection, prevention, and treatment, yet it remains challenging due to the heterogeneity in disease trajectories among individuals. Factors such as demographics, genetic conditions, and lifestyle contribute to diverse phenotypical manifestations, necessitating patient stratification based on these variations. Recent methods like Subtype and Stage Inference (SuStaIn) have advanced unsupervised stratification of disease trajectories, but they face potential limitations in robustness, interpretability, and temporal granularity. To address these challenges, we introduce Disease Progression Modelling and Stratification (DP-MoSt), a novel probabilistic method that optimises clusters of continuous trajectories over a long-term disease time-axis while estimating the confidence of trajectory sub-types for each biomarker.

We validate DP-MoSt using both synthetic and real-world data from the Parkinson’s Progression Markers Initiative (PPMI). Our results demonstrate that DP-MoSt effectively identifies both sub-trajectories and sub-populations, and is a promising alternative to current state-of-the-art models.

Keywords: Disease Progression Modelling · Expectation Maximisation · Parkinson Disease

1 Introduction

Modelling the progression of Degenerative Diseases (DD) is crucial for detection, prevention, and treatment purposes [9]. This task is challenging due the generally large heterogeneity of disease trajectories observed across subjects. Despite a common degenerative process, the manifestation of symptoms and the configuration of biomarkers may vary widely among individuals, due for example to

demographics, genetics conditions and life-style [17]. For this reason researchers are steering their attention towards the problem of *patients stratification* based on phenotypical manifestations of the disorder [4].

Typical disease progression modelling approaches generally focus on estimating long-term biomarkers evolution from short-term patients observation. For instance, the Gaussian Process Progression Model (GPPM) [8] and the Personalized Input-Output Hidden Markov Model [14] characterise the transition of biomarkers over time from normal to pathological stages, based on the assumption of an underlying disease trajectory defined by an absolute time axis. On the other hand, other approaches focus on the detection of sub-populations, for example biomarkers values [21] or genetic observations [16].

In the last decade, innovative methodologies have attempted to automatically stratify sub-types of disease trajectories. SuStaIn is a popular unsupervised method detecting sub-populations and their respective trajectories within a given dataset of patients and control population [19,20,18]. SuStaIn has been demonstrated in a variety of applications, showing its ability to identify sub-populations exhibiting common patterns of biomarkers changes [1,22]. Nevertheless, from an analytical perspective, SuStaIn presents some relevant limitations in terms of robustness and interpretability. First, disease progression is described as a discrete series of events, which limits the interpretability and temporal granularity of the estimated trajectories. Second, SuStaIn assumes the existence of cutoff values optimising separation between disease stages across biomarkers; the assumptions behind this statistical construct (e.g. Gaussian) are non-necessarily realistic about the biomarkers distribution across stages, and may negatively affect the robustness of the clustering task. Third, SuStaIn does not directly quantify the uncertainty in whether each biomarker exhibits a distinct trajectory between subtypes. For example, although certain biomarkers may not be discriminating between sub-types, SuStaIn will attempt at estimating group-specific cutoff values which may lead to poor interpretability and generalisation of the results. Finally, SuStain is designed to model cross-sectional information, without accounting for the temporally correlated nature of patients’ time-series. Recent extensions of SuStain (t-SuStaIn) [18], attempt to overcome this latter limitation by accounting for multiple measures per subjects. However, this approach still presents the above mentioned limitations, as it relies on the estimation of sequences of events occurring in a discrete space.

Disease Course Mapping (DCM) [12] is an orthogonal approach to SuStain, in which continuous parametric disease trajectories are optimized by accounting for patient’s random effects represented by time-warp functions. DCM was recently extended to account for mixture of trajectories (MM-DCM) [11]. However, the mixture model there proposed assumes that all biomarkers’ trajectories should be split into sub-progressions. This approach thus does not allow uncertainty quantification of the split across biomarkers, and ultimately does not account for the specificity of certain biomarkers in characterising disease sub-types.

To address these limitations, in this paper we present Disease Progression Modelling and Stratification (DP-MoSt), a novel probabilistic method to identify

differential disease progression trajectories in heterogeneous cohorts. DP-MoSt relies on the optimisation of clusters of continuous trajectories across a long-term disease time-axis, while also estimating the confidence for the existence of trajectories sub-types for each biomarker.

We validate our model on both synthetic and real-world data from the Parkinson’s Progression Markers Initiative (PPMI) [2]. Our results show that our model is a promising alternative to the state-of-the-art, effectively identifying sub-trajectories and sub-populations, while providing interpretable and clinically meaningful solutions.

The manuscript is structured as follows: in Section 2.1 we introduce DP-MoSt, and in Section 2.2 we describe the optimisation procedure and the statistical assumptions. In Section 3.1 we present results on a panel of synthetic benchmarks, and in Section 3.2 we provide a comparison between DP-MoSt and SuStaIn on clinical data from the PPMI database.

2 Disease Progression Modelling and Stratification model

In this section, we provide the mathematical details of DP-MoSt, by describing the underlying statistical assumptions along with the associated optimization procedure.

2.1 Model definition

DP-MoSt is based on the optimization of two complementary problems: (i) estimating an absolute long-term disease time axis from short-term observations, and (ii) identifying along this disease time axis the existence of sub-populations with respective sub-trajectories.

Considering problem (i), for each individual j we define the observations across all biomarkers as $\mathbf{x}^j = (\mathbf{x}_b^j)_{b=1}^B$; where $\mathbf{x}_b^j = (\mathbf{x}_b^j(\tilde{t}_1), \dots, \mathbf{x}_b^j(\tilde{t}_{k_j}))$ and B is the number of biomarkers. Without loss of generality, for notational convenience, we assume that the sampling times are common among all subjects and biomarkers. To map the individual observations to a common disease time scale, we parameterize the individual time axis via a translation by a time-shift $\delta\tilde{t}_j$, i.e. $\mathbf{t}_j = \tilde{t}_{1:k_j} + \delta\tilde{t}_j$. In this work, we evaluate the time shifts relying on the Gaussian process theory of GPPM [8], which is based on the monotonic description of biomarkers trajectories from normal to pathological stages.

Considering problem (ii), given the measured observations $\mathbf{x} = \mathbf{x}^{1:J}$, where J is the number of subjects, and the estimated absolute time $\mathbf{t} = \mathbf{t}_{1:J}$, we define a trajectory mixture model to identify the existence of sub-populations.

To achieve our goal, we assume that the evolution of each biomarker b can be split into multiple sub-trajectories with probability ξ_b ($b = 1, \dots, B$). Once a sub-trajectory is considered, we assume that each subject j is issued from this trajectory with probability π_j ($j = 1, \dots, J$). We observe that both $\xi = (\xi_b)_1^B$ and $\pi = (\pi_j)_1^J$ are independent with respect to time. This allows our model to link information deriving from longitudinal data; we also note that

the probability for a subject belonging to one sub-trajectory must be consistent across all the biomarkers.

To simplify the inference process, compatibly with the monotonic assumption of GPPM, we adopt a parametric approach for the disease trajectories assuming that biomarkers follow increasing sigmoidal functions over time. We furthermore assume that the given measures are perturbed by additive Gaussian noise with standard deviation $\sigma = (\sigma_b)_{b=1}^B$. Given the assumptions above, the posterior distribution for our model can be written as:

$$p(\theta, \sigma, \xi, \pi | \mathbf{x}) \propto p(\theta, \sigma, \xi, \pi) \prod_{j,b} p(\mathbf{x}_b^j | \theta_b, \sigma_b, \xi_b, \pi_j) \quad (1)$$

where for simplicity we omitted the conditioning on the time points. We observe that Equation (1) implicitly assumes independence between the unknown parameters as well as independence between different subjects and biomarkers.

We can rewrite the equation by expanding the likelihood function in order to highlight the two-level mixture model formulation. In this setting, a first level deals with the sub-trajectory discovery task, while a second one determines the probability of a subject to belong to the sub-trajectory:

$$p(\mathbf{x} | \theta, \sigma, \xi, \pi) = \prod_{j,b} \left[p(\mathbf{x}_b^j | \theta_b^0, \sigma_b) \xi_b + \left(\pi_j p(\mathbf{x}_b^j | \theta_b^1, \sigma_b) + (1 - \pi_j) p(\mathbf{x}_b^j | \theta_b^2, \sigma_b) \right) (1 - \xi_b) \right], \quad (2)$$

where $p(\mathbf{x}_b^j | \theta_b^i, \sigma_b) = \prod_{\ell=1}^{k_j} \text{NormPDF} \left(x_b^j(t_\ell), f(t_\ell | \theta_b^i), \sigma_b \right)$ due to the assumption of additive Gaussian noise, and $f(t_\ell | \theta_b^i)$ is a Sigmoid function with parameters θ_b^i evaluated at t_ℓ .

2.2 Two-levels Expectation-Maximization

Model (1) accounts for a substantial number of parameters. For each biomarker it includes: three parameters for each of the three Sigmoid functions $\theta_b^{0:2}$; the noise parameters σ_b ; the probability for the existence of sub-trajectories ξ_b ; and the membership probabilities for each subject π . To limit the computational cost, we focus on Maximum a Posteriori (MAP) estimation through Expectation Maximisation (EM), exploiting the two-levels mixture nature of the model:

$$\begin{aligned} \hat{\theta}, \hat{\xi}, \hat{\pi}, \hat{\sigma} &= \arg \max \ln(p(\theta, \sigma, \xi, \pi | \mathbf{x})) \\ &= \arg \max \ln(p(\mathbf{x} | \theta, \sigma, \xi, \pi)) + \beta \sum_b \xi_b - \beta_N \sum_b \left(\ln(\sigma_b) + \frac{1}{\sigma_b} \right). \end{aligned} \quad (3)$$

We rely on the following prior assumptions:

- $p(\theta) \propto 1$ and $p(\pi) \propto 1$ as improper uniform prior for the Sigmoid parameters and sub-population subdivision in order to encode lack of information;

- $p(\sigma) = \prod_b p(\sigma_b)$, where $p(\sigma_b)$ is an Inv-Gamma distribution of shape parameter $\beta_N - 1$ and scale parameter $\beta_N > 1$ as classical prior distribution for the noise standard deviation [3] in order to penalise small values;
- $p(\xi) = \prod_b p(\xi_b)$, where $p(\xi_b)$ is a Laplace distribution [7] of location 1 and scale parameter $1/\beta$ restricted to the interval $[0, 1]$, as regularisation term for penalising the introduction of a sub-trajectory to prevent overfitting.

DP-MoSt is designed to provide interpretable progression dynamics: first, we estimate the parameters for the continuous biomarkers trajectories; second, we estimate the probability for each biomarker to present sub-trajectories; finally, we estimate for each subject its probability to belong to each sub-population.

3 Results

In this section, we validate the model on two different experimental scenarios⁴: first, we demonstrate the effectiveness of the method in identifying sub-trajectories and sub-populations on an extensive synthetic benchmark; second, we apply the model to the real data from the PPMI dataset, comparing the solution to state-of-the-art solutions as provided by SuStaIn. The decision to validate only the clustering component of the model stems from the fact that the time-shift is assessed using GPPM, whose effectiveness has already been proven.

3.1 Experiment on synthetic data

Data generation. We analyze the performance of DP-MoSt on the task of sub-trajectories and sub-populations identification by evaluating its performances across increasing levels of data complexity. This involves altering both the number of biomarkers and the Signal to Noise Ratio (SNR) [5] between sub-trajectories. Specifically, we consider three different sets of biomarkers with $B = 2, 5, 10$ and assess performance under three distinct SNR conditions: low, normal, and high. This systematic variation allows us to test the robustness and accuracy of our model across a range of realistic scenarios of increasing difficulty.

For each configuration, we generate 100 datasets with the following parameters:

- $J = 100$ subjects
- $k_{1:J} = 1$, i.e. one time point for each subject, randomly sampled in $[0, 20]$.
- $\sigma_b = 0.5$ to ensure a reasonable amount of noise in the data.
- Three fixed thresholds (low = 0.1, normal = 0.5, and high = 1) for the Mean Squared Error (MSE) between sub-trajectories to achieve a controlled SNR.
- Half of the biomarkers exhibit sub-trajectories (in the case of 5 biomarkers, 3 of them show the split).

⁴ The code for DP-MoSt can be found at <https://github.com/alessandro-viani/DP-MoSt.git>.

- Equal partitioning of subjects between sub-populations.
- Parameters for the Sigmoid function describing the biomarker trajectories are randomly sampled from Gaussian distributions with parameters ensuring a positive supremum and rate of growth.

Model setup. For each simulated dataset, we evaluate the parameters of the model described by Equation (1) using the EM method described in Section 2.2.

We initialize the model as follows to avoid local maxima for the posterior distribution:

- set $\xi_b = 0.5$ and $\pi_j = 0.5$, assuming a complete lack of information on the sub-trajectories and sub-population probabilities;
- Initialize the noise standard deviation as the standard deviation of the data;
- randomly select the parameters for the Sigmoid functions from Gaussian distributions ensuring to provide a positive rate of growth and supremum;
- Set the prior parameters β and β_N as the 15% of the number of subjects; this choice ensures effective regularisation while maintaining the values within a reasonable range.

Performance metrics. To validate the proposed method, we use different metrics to evaluate the error in parameter approximations. For evaluating the error on the biomarker trajectories, we employ the Optimal Subpattern Assignment (OSPA) metric [13,15]; this metric is particularly suitable because it accounts for potential differences in the number of true and estimated configurations. The OSPA metric measures the minimum MSE between the actual configuration of Sigmoids and the estimated one:

$$\text{OSPA}(\theta, \hat{\theta}) = \min_{\phi} \sum_{i=1}^{\min\{\hat{d}, d\}} \frac{1}{N} \sum_j \left\| f(t_j | \hat{\theta}_i) - f(t_j | \theta_{\phi(i)}) \right\|_2^2 \quad (4)$$

where $\hat{\theta}$ and \hat{d} represent the estimated parameters and number of sub-trajectories, respectively; θ and d denote the true values of the parameters; and the symbol ϕ represents all possible permutations.

Results. The obtained results are summarised in Figure 1, where we show:

- First Panel, First Row: the logarithm of the OSPA error subdivided by the number of features considered and SNR levels. We observe that variations in data complexity do not significantly impact the trajectory approximation; indeed the error remains roughly constant over different data configurations.
- Second Panel, First Row: the logarithm of the relative error on the noise standard deviation subdivided by the number of features considered and SNR levels. We can observe that the approximation error for the noise std increases with the data complexity.

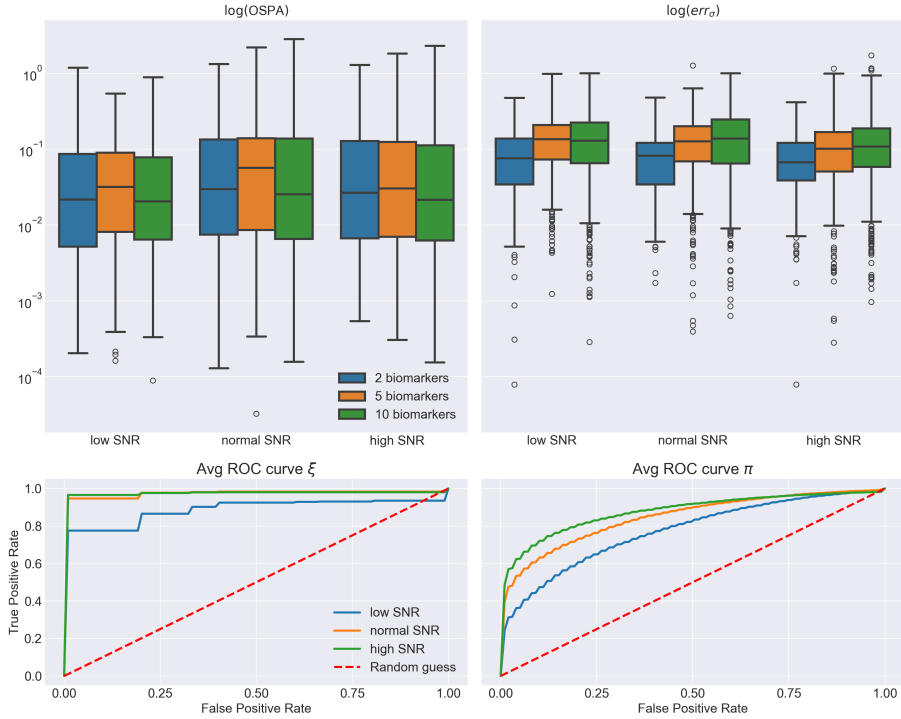


Fig. 1. The figure shows performance metrics obtained on synthetic data. In the first row we show the OSPA error for the trajectories approximation and the error on the noise standard deviation; in the second row we show the ROC curves respectively for trajectory and individual clustering, as identified by parameters ξ and π .

- First/Second Panel, Second Row: the ROC curve for the parameters ξ and π . We observe that the method performs better in estimating the number of sub-trajectories compared to the estimation of sub-populations, probably due to the fewer number of parameters to be estimated.

3.2 PPMI data

In this section we provide the results obtained considering the PPMI dataset, a comprehensive, multi-center longitudinal study for Parkinson’s research.

The PPMI dataset includes extensive clinical, imaging, and biological data collected from PD patients, individuals with PD risk factors, and healthy controls. The overall data considered in this work is composed of a total of 3559 patients and 93 measured biomarkers. Each patient is labelled according to the following clinical scores: tremor-dominant (TD), postural instability gait disorder predominant (PIGD), and intermediate [2,10].

We focus our analysis on a specific set of measures:

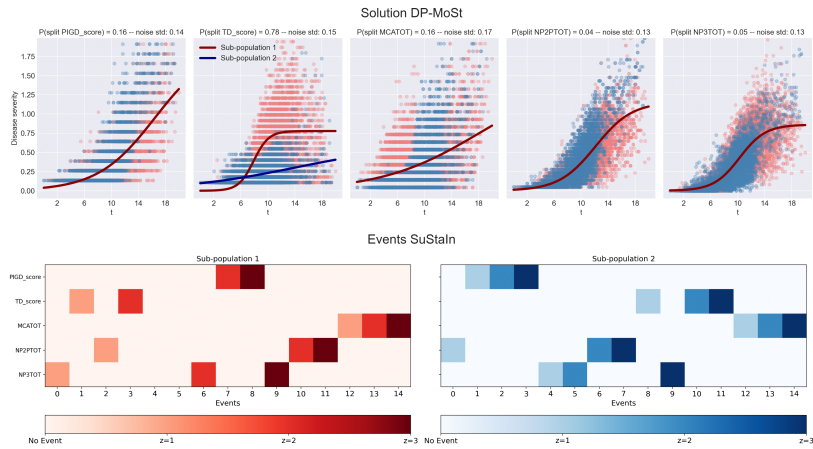


Fig. 2. The figure shows the results obtained on the PPMI dataset. The first row shows the results obtained with DP-MoSt: (sub-)trajectories are represented as solid coloured lines, and subjects are colour-coded based on their estimated subgroups. The second row shows the event progression estimated by SuStain as two coloured matrices, one for each sub-population.

- *PIGD_score*: an indicator that asserts the gravity of the PIGD classification;
- *TD_score*: an indicator asserts the gravity of the TD classification;
- *MCAOT*: the total amount of Montreal Cognitive Assessment scores [6];
- *NP2PTOT* & *NP3PTOT*: are the total score of MDS-UPDRS part 2 and 3.

After preprocessing the data composed by the 5 variables of interest, we obtain a dataset including 10 198 longitudinal data from 1954 patients, with an average of 5 time points for each subject. Data is analysed by applying DP-MoSt and SuStain. When applying DP-MoSt, we did not account for the temporal correlation between observations across subjects. This allows for a more fair comparison with SuStain, where data are treated disregarding temporal dependency.

SuStain setup. SuStain estimates the maximum likelihood (ML) solution for the number of sub-types in the dataset as well as the associated sequence of events (i.e., stages of severity increase). Based on previous works, we decided to implement the z-score SuStain with three different progression stages associated to the quartiles (z_1, z_2, z_3) and 100 000 MCMC steps.

DP-MoSt setup. We assess DP-MoSt by considering the same parameter values as for the experiment on synthetic data (Section 3.1). However, differently from the synthetic experiment in which we evaluated solely the clustering step of our method, we also optimise individual time-shift parameters, establishing an absolute time axis for the disease progression.




Fig. 3. The figure shows the distribution for the biomarker values across the two different sub-populations detected by SuStAin and DP-MoSt, where we highlight the quartiles of the distributions as dotted lines.

Results. In Figure 2 (first row) we show the solution for DP-MoSt, representing biomarker trajectories as solid lines with respect to the absolute time axis for disease progression (in years), and subjects as dots colour-coded based on their estimated sub-group. We observe that there is one clinical score that primarily governs the split into two sub-groups ($\xi_b > 0.5$): the TD_score . The two sub-groups that result from this split show a clear subdivision with respect to TD_score values: sub-population 1 is associated with higher values for this biomarker compared to sub-population 2, with a clear differentiation of trajectories after $t=6$ years on the reparameterized time axis. We also observe that, in spite of the split identified for the TD_score trajectory, the progressions of the other clinical scores do not exhibit clear partitioning in sub-trajectories ($\xi_b < 0.5$). This aspect is illustrated in Figure 3 where, besides the TD_score , the clinical scores’ distributions between sub-populations show a substantial overlap. We finally note that, due to the continuous nature of DP-MoSt trajectories, the resulting clinical scores’ distributions can be multi-modal, as they describe the entire disease history across pathological stages (e.g. TD_score in Sub-population 1). This is less evident for SuStAin (Figure 3, left panels), which enforces a stronger separation between distributions across sub-groups, probably due to the discrete assumption for the events ordering.

In Table 1 (first column), we show the distribution of subjects into different sub-populations based on their label available at the last visit (Intermediate, PIGD, TD). We can observe that the sub-trajectories identified by DP-MoSt are respectively associated to subjects labeled as TD and PIGD, providing a clinically meaningful partitioning of the subjects. More specifically, DP-MoSt associates to sub-population 1 the majority of TD subjects (76%) and to sub-population 2 the majority of PIGD subjects (62%). The subjects labelled as Intermediate are not clearly assigned into any specific sub-population, probably because of their mixed composition.

In Figure 2 (second row) we show the solution identified by SuStAin as two coloured matrices, one for each sub-population, indicating the event ordering for each biomarker. SuStAin identifies 2 sub-populations, with sub-population 1 characterised by a faster increase of the TD_score and a slower change of the $PIGD_score$ compared to sub-population 2. Notably, we can observe that

Condition	DP-MoSt		SuStaIn	
	Sub-pop 1	Sub-pop 2	Sub-pop 1	Sub-pop 2
Intermediate	0.56	0.44	0.84	0.16
PIGD	0.38	0.62	0.78	0.22
TD	0.76	0.24	0.48	0.52
N°% data	56%	44%	72%	28%

Table 1. The Table shows the subdivision between different sub-populations considering solution provided by DP-MoSt and SuStaIn.

the *MCATOT* score does not show any differences between sub-populations, suggesting that this biomarker is not relevant to disease sub-typing.

In Table 1 (second column) we observe that the clustering obtained with SuStaIn identifies two sub-populations that are unbalanced, with the majority of subjects (72%) in the first sub-population. Therefore, differently from DP-MoSt, the stratification provided by SuStaIn does not provide a clear subdivision between clinical groups, substantially including the majority of subjects into the first sub-population.

4 Conclusions

In this paper, we introduced DP-MoSt, an innovative probabilistic method for identifying continuous biomarker trajectories and stratifying sub-populations in the context of DPMs. DP-MoSt addresses key limitations of current state-of-the-art methods by ensuring interpretability and robustness through a two-level mixture model that captures sub-population clusters while incorporating temporal information from longitudinal data. Additionally, DP-MoSt provides uncertainty quantification for both sub-trajectories and sub-population composition, enabling the characterization of pathological trajectory patterns.

We validated the model’s performance with synthetic data, demonstrating its effectiveness. Furthermore, its application to the PPMI dataset yielded interpretable and clinically relevant results; this feature can be due to the continuous formulation of DP-MoSt, improving the model’s reliability and interpretability.

It is important to observe that SuStaIn provides an estimation of the number of sub-populations, a feature currently not included in DP-MoSt. However, we note that the introduction of multiple sub-populations with their respective sub-trajectories may also be obtained with DP-MoSt: this primarily involves more complex notation and higher computational costs, and is one of the improvements that will be implemented in future work.

Overall, DP-MoSt’s ability to capture detailed sub-population characteristics makes it a promising tool for analysing heterogeneous disease progression patterns in longitudinal studies.

Acknowledgments. This work has been supported by the Michael J. Fox Foundation for Parkinson’s Research (MJFF), and to the French government, through the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA- 0002, by the TRAIN project ANR-22-FA11-0003-02, and by the ANR JCJC project Fed-BioMed 19-CE45-0006-01. Data used in the preparation of this article were obtained on 05/01/2024 from the Parkinson’s Progression Markers Initiative (PPMI) database.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Aksman, L.M., Wijeratne, P.A., Oxtoby, N.P., Eshaghi, A., Shand, C., Altmann, A., Alexander, D.C., Young, A.L.: pysustain: a python implementation of the subtype and stage inference algorithm. *SoftwareX* **16**, 100811 (2021)
2. Aleksovski, D., Miljkovic, D., Bravi, D., Antonini, A.: Disease progression in parkinson subtypes: the ppmi dataset. *Neurological Sciences* **39**, 1971–1976 (2018), <https://link.springer.com/article/10.1007/s10072-018-3522-z>
3. Calvetti, D., Hakula, H., Pursiainen, S., Somersalo, E.: Conditionally gaussian hypermodels for cerebral source localization. *SIAM Journal on Imaging Sciences* **2**(3), 879–909 (2009)
4. Greenland, J.C., Williams-Gray, C.H., Barker, R.A.: The clinical heterogeneity of parkinson’s disease and its therapeutic implications. *European Journal of Neuroscience* **49**(3), 328–338 (2019)
5. Johnson, D.H.: Signal-to-noise ratio. *Scholarpedia* **1**(12), 2088 (2006)
6. Julayanont, P., Nasreddine, Z.S.: Montreal cognitive assessment (moca): concept and clinical review. *Cognitive screening instruments: A practical approach* pp. 139–195 (2017)
7. Kaban, A.: On bayesian classification with laplace priors. *Pattern Recognition Letters* **28**(10), 1271–1282 (2007)
8. Lorenzi, M., Filippone, M., Frisoni, G.B., Alexander, D.C., Ourselin, S., Initiative, A.D.N., et al.: Probabilistic disease progression modeling to characterize diagnostic uncertainty: application to staging and prediction in alzheimer’s disease. *NeuroImage* **190**, 56–68 (2019)
9. Mould, D.: Models for disease progression: new approaches and uses. *Clinical Pharmacology & Therapeutics* **92**(1), 125–131 (2012)
10. Pötter-Nerger, M., Dutke, J., Lezius, S., Buhmann, C., Schulz, R., Gerloff, C., Kuhle, J., Choe, C.u.: Serum neurofilament light chain and postural instability/gait difficulty (pigd) subtypes of parkinson’s disease in the mark-pd study. *Journal of Neural Transmission* **129**(3), 295–300 (2022)
11. Poulet, P.E., Durrleman, S.: Mixture modeling for identifying subtypes in disease course mapping. In: *International Conference on Information Processing in Medical Imaging*, pp. 571–582. Springer (2021)
12. Schiratti, J.B., Allasonnière, S., Colliot, O., 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)
13. Schuhmacher, D., Vo, B.T., Vo, B.N.: A consistent metric for performance evaluation of multi-object filters. *IEEE Transactions on Signal Processing* **56**(8), 3447–3457 (2008). <https://doi.org/10.1109/TSP.2008.920469>

14. Severson, K.A., Chahine, L.M., Smolensky, L., Ng, K., Hu, J., Ghosh, S.: Personalized input-output hidden markov models for disease progression modeling pp. 309–330 (2020)
15. Viani, A., Luria, G., Bornfleth, H., Sorrentino, A.: Where bayes tweaks gauss: Conditionally gaussian priors for stable multi-dipole estimation. arXiv preprint arXiv:2006.04141 (2020)
16. Whitwell, J.L., Weigand, S.D., Boeve, B.F., Senjem, M.L., Gunter, J.L., DeJesus-Hernandez, M., Rutherford, N.J., Baker, M., Knopman, D.S., Wszolek, Z.K., et al.: Neuroimaging signatures of frontotemporal dementia genetics: C9orf72, tau, progranulin and sporadics. *Brain* **135**(3), 794–806 (2012)
17. Yang, E., Farnum, M., Lobanov, V., Schultz, T., Raghavan, N., Samtani, M.N., Novak, G., Narayan, V., DiBernardo, A., Initiative, A.D.N., et al.: Quantifying the pathophysiological timeline of alzheimer’s disease. *Journal of Alzheimer’s Disease* **26**(4), 745–753 (2011)
18. Young, A.L., Aksman, L.M., Alexander, D.C., Wijeratne, P.A., Initiative, A.D.N.: Subtype and stage inference with timescales. In: *International Conference on Information Processing in Medical Imaging*. pp. 15–26. Springer (2023)
19. Young, A.L., Marinescu, R.V., Oxtoby, N.P., Bocchetta, M., Yong, K., Firth, N.C., Cash, D.M., Thomas, D.L., Dick, K.M., Cardoso, J., et al.: Uncovering the heterogeneity and temporal complexity of neurodegenerative diseases with subtype and stage inference. *Nature communications* **9**(1), 4273 (2018)
20. Young, A.L., Vogel, J.W., Aksman, L.M., Wijeratne, P.A., Eshaghi, A., Oxtoby, N.P., Williams, S.C., Alexander, D.C., Initiative, A.D.N.: Ordinal sustain: subtype and stage inference for clinical scores, visual ratings, and other ordinal data. *Frontiers in artificial intelligence* **4**, 613261 (2021)
21. Zhang, X., Mormino, E.C., Sun, N., Sperling, R.A., Sabuncu, M.R., Yeo, B.T., Initiative, A.D.N.: Bayesian model reveals latent atrophy factors with dissociable cognitive trajectories in alzheimer’s disease. *Proceedings of the National Academy of Sciences* **113**(42), E6535–E6544 (2016)
22. Zhou, C., Wang, L., Cheng, W., Lv, J., Guan, X., Guo, T., Wu, J., Zhang, W., Gao, T., Liu, X., et al.: Two distinct trajectories of clinical and neurodegeneration events in parkinson’s disease. *npj Parkinson’s Disease* **9**(1), 111 (2023)