

Multiple Orderings of Events in Disease Progression

Alexandra L Young¹, Neil P Oxtoby¹, Jonathan Huang², Razvan V Marinescu¹, Pankaj Daga¹, David M Cash^{1,3}, Nick C Fox³, Sebastien Ourselin^{1,3}, Jonathan M Schott³, and Daniel C Alexander¹, for the Alzheimer’s Disease Neuroimaging Initiative

¹Centre for Medical Image Computing, University College London; ²Google Inc., California; ³Dementia Research Centre, University College London

Abstract. The event-based model constructs a discrete picture of disease progression from cross-sectional data sets, with each event corresponding to a new biomarker becoming abnormal. However, it relies on the assumption that all subjects follow a single event sequence. This is a major simplification for sporadic disease data sets, which are highly heterogeneous, include distinct subgroups, and contain significant proportions of outliers. In this work we relax this assumption by considering two extensions to the event-based model: a generalised Mallows model, which allows subjects to deviate from the main event sequence, and a Dirichlet process mixture of generalised Mallows models, which models clusters of subjects that follow different event sequences, each of which has a corresponding variance. We develop a Gibbs sampling technique to infer the parameters of the two models from multi-modal biomarker data sets. We apply our technique to data from the Alzheimer’s Disease Neuroimaging Initiative to determine the sequence in which brain regions become abnormal in sporadic Alzheimer’s disease, as well as the heterogeneity of that sequence in the cohort. We find that the generalised Mallows model estimates a larger variation in the event sequence across subjects than the original event-based model. Fitting a Dirichlet process model detects three subgroups of the population with different event sequences. The Gibbs sampler additionally provides an estimate of the uncertainty in each of the model parameters, for example an individual’s latent disease stage and cluster assignment. The distributions and mixtures of sequences that this new family of models introduces offer better characterisation of disease progression of heterogeneous populations, new insight into disease mechanisms, and have the potential for enhanced disease stratification and differential diagnosis.

1 Introduction

The sequence in which biomarkers become abnormal provides a simple, intuitive description of disease progression, giving insights into the underlying disease biology and a potential mechanism for disease staging. The sequence of biomarker abnormality in sporadic neurodegenerative diseases, e.g. Alzheimer’s disease, has

been a topic of intense debate amongst neurologists [1]. Reconstructing this sequence for sporadic neurodegenerative diseases is difficult because the position of subjects with respect to the full disease time course is unknown. Typically clinical diagnoses are used as a time proxy, but this limits the temporal resolution of the sequence, e.g. in Alzheimer’s disease there are only three stages: cognitively normal, mild cognitive impairment and Alzheimer’s disease [2]. Additional complications arise due to the long disease time course [3] and inherent heterogeneity of sporadic disease datasets. Many different factors contribute to this heterogeneity [4, 5], for example genetic disease subtypes, mixed pathology, environmental factors, and misdiagnosed subjects.

The event-based model [6] considers disease progression as a series of events, where each event corresponds to a new biomarker becoming abnormal. By considering cross-sectional patient data as snapshots of a single common event sequence, the event-based model is able to probabilistically reconstruct the ordering of events across subjects, without relying on a-priori disease staging. Taking samples of the posterior probability of this sequence provides insight into the uncertainty in this single event ordering. The application of this model has been demonstrated in familial Alzheimer’s disease and Huntington’s disease [6] to determine the sequence in which regional brain volumes become abnormal, and in sporadic Alzheimer’s disease to determine the sequence in which cerebrospinal fluid (CSF) markers, cognitive test scores, and a limited set of regional atrophy and brain volume biomarkers become abnormal [7]. Young et al. [7] found that this sequence is different in APOE4 positive individuals, with increased genetic risk of Alzheimer’s disease, compared to the whole population, suggesting that the whole population contains a proportion of subjects who do not follow the single ordering of events encoded by the event-based model.

The assumption made by the event-based model in [6] and [7] of a single ordering of events in all subjects is a major simplification for heterogeneous sporadic disease datasets. In this work we relax this assumption by considering a family of models that allow for multiple and distributed orderings of events. The first is a generalised Mallows model [8], which parameterises the variance in the single ordering, allowing subjects to deviate from the central event sequence. The second is a Dirichlet process mixture model [9], which allows for subgroups of subjects that follow different event sequences. Previous work [10] on generalised Mallows event-based models relied on a well-defined control population and a complete set of biomarkers for each subject. Here we re-formulate this model to remove the reliance on a well-defined control population, allowing the model to be fitted to heterogeneous sporadic disease datasets, and to handle missing data, providing a multi-modal picture of disease progression. We formulate a Gibbs sampling technique that further provides samples of the uncertainty in the model parameters. We additionally introduce a new model: Dirichlet process mixtures of generalised Mallows event-based models, and develop a Gibbs sampler to estimate its parameters [11]. We apply these models to determine the sequence in which FDG-PET, CSF markers, cognitive test scores, and a large set of regional brain volumes become abnormal in sporadic Alzheimer’s disease.

2 Models

2.1 The Event-Based Model

The event-based model of disease progression consists of a set of events $\{e_1, \dots, e_N\}$ and an ordering $\sigma = (\sigma(1), \dots, \sigma(N))$, where $\sigma(k) = i$ means that event e_i occurs in position k . If a subject is at stage k in the sequence σ the events $e_{\sigma(1)} \dots e_{\sigma(k)}$ have occurred and events $e_{\sigma(k+1)} \dots e_{\sigma(N)}$ have yet to occur. The occurrence of event e_i in subject j is informed by biomarker measurement x_{ij} . In practise we only observe a snapshot of the event sequence for each subject. This adduces a partition of the event set, or partial ranking, $\gamma_k = e_{\sigma(1)}, \dots, e_{\sigma(k)} | e_{\sigma(k+1)}, \dots, e_{\sigma(N)}$, where the vertical bar indicates that the first set of events precedes the second. The generative model of the biomarker data is

$$k_j \sim P(k),$$

$$x_{\sigma(i),j} \sim p(x_{\sigma(i),j} | e_{\sigma(i)}) \text{ if } i \leq k_j,$$

$$x_{\sigma(i),j} \sim p(x_{\sigma(i),j} | \neg e_{\sigma(i)}) \text{ otherwise.}$$

$p(x|e)$ and $p(x|\neg e)$ are probability density functions on observing biomarker measurement x given that event e has or has not occurred respectively. $P(k)$ is a prior on the disease stage k .

2.2 The Generalised Mallows Event-Based Model

We formulate the generalised Mallows event-based model by using a generalised Mallows model to parameterise the variance in a central event sequence π through the spread parameter $\theta = (\theta_1, \dots, \theta_{N-1})$. Each subject then has their own latent ordering σ_j , which is assumed to be a sample from a generalised Mallows model. The generative model of the biomarker data in the event-based model is therefore preceded by

$$\pi, \theta \sim P(\pi, \theta | \nu, \mathbf{r}),$$

$$\sigma_j \sim GM(\pi, \theta).$$

$GM(\pi, \theta) = \frac{1}{\psi(\theta)} \exp[-d_\theta(\pi, \sigma)]$ is a generalised Mallows distribution with $\psi(\theta) = \prod_{j=1}^{n-1} \psi_{n-j}(\theta_j) = \prod_{j=1}^{n-1} \frac{1 - e^{-(n-j+1)\theta_j}}{1 - e^{-\theta_j}}$. $d_\theta(\pi, \sigma)$ is the generalised Kendall's tau distance [8], which penalises the number of pairwise disagreements between sequences. $P(\pi, \theta | \nu, \mathbf{r})$ is a conjugate prior over the generalised Mallows distribution parameters of the form $P(\pi, \theta | \nu, \mathbf{r}) \propto \exp\left(-\nu \sum_j [\theta_j r_j + \ln \psi_{n-j}(\theta_j)]\right)$ [12].

2.3 Dirichlet Process Mixtures of Generalised Mallows Event-Based Models

Dirichlet process mixtures of generalised Mallows models assume that each subject has their own central ordering π_j and spread parameters θ_j , which are sampled from a discrete distribution G that is drawn from a Dirichlet process [9]. A Dirichlet process mixture is a generative clustering model where the number of clusters is a random variable, meaning that the number of clusters is detected automatically depending on the concentration parameter α . The generative model of the biomarker data in the event-based model is now preceded by the process

$$G \sim DP(\alpha, P(\pi, \theta | \nu, \mathbf{r})),$$

$$\pi_j, \theta_j \sim G,$$

$$\sigma_j \sim GM(\pi_j, \theta_j),$$

where $DP(\alpha, P(\pi, \theta | \nu, \mathbf{r}))$ is a Dirichlet process [9]. Each data point π_j can be characterised by an association with a cluster label $c_j \in 1, \dots, C$ and each cluster c with a set of generalised Mallows parameters σ_c and θ_c .

3 Inference

3.1 The Event-Based Model

Inference in the event-based model can be performed by taking Markov Chain Monte Carlo (MCMC) samples of $P(\sigma | X) = \frac{P(X|\sigma)P(\sigma)}{P(X)}$ where

$$P(X|\sigma) = \prod_{j=1}^J \left[\sum_{k=0}^K P(k) \left(\prod_{i=1}^k p(x_{\sigma(i),j} | e_{\sigma(i)}) \prod_{i=k+1}^N p(x_{\sigma(i),j} | \neg e_{\sigma(i)}) \right) \right]. \quad (1)$$

3.2 The Generalised Mallows Event-Based Model

We use Gibbs sampling to infer the parameters of the generalised Mallows event-based model. This consists of two stages. First, generating a set of sample event sequences $\sigma_{1:J}$. We sample from an augmented model [10], by alternating between sampling a subject's ordering σ_j and disease stage k_j , which are used to deterministically reconstruct their partial ranking γ_j . The Gibbs sampling updates are therefore

$$\sigma^{(j)} \sim P(\sigma | \gamma = \gamma_j, \pi, \theta),$$

$$k^{(j)} \sim P(k | \sigma = \sigma_j, X_j).$$

Second, sampling the model parameters given the set of sample orderings $\sigma_{1:J}$ using the updates

$$\pi \sim P(\pi | \theta, \nu, \mathbf{r}, \sigma_{1:J}),$$

$$\theta_k \sim P(\theta_k | \pi, \nu, \mathbf{r}, \sigma_{1:J}).$$

3.3 Dirichlet Process Mixtures of Generalised Mallows Event-Based Models

We formulate another Gibbs sampler to infer the parameters of Dirichlet process mixtures of generalised Mallows event-based models. We generate a set of candidate sample orderings $\sigma_{1:J,1:C}$, disease stages $k_{1:J,1:C}$, and partial rankings $\gamma_{1:J,1:C}$, which are conditioned on the parameters for each cluster via the updates

$$\sigma^{(j,c)} \sim P(\sigma | \gamma = \gamma_{jc}, \pi_c, \theta_c),$$

$$k^{(j,c)} \sim P(k | \sigma = \sigma_{jc}, X_j).$$

From these samples we sample the cluster assignment c_j of each subject conditioned on the cluster assignments of the other subjects c_{-j} , where c_{-j} is the set of cluster assignments for all subjects except subject j , the subject's sample ordering for each cluster $\sigma_{j,1:C}$, disease stage $k_{j,1:C}$ and their biomarker data X_j . We then update the generalised Mallows model parameters for each cluster, π_c and θ_c , from the set of subject orderings assigned to each cluster, σ_c . So we have the updates

$$c^{(j)} \sim P(c | c_{-j}, \sigma_{j,1:C}, \theta, \alpha, \nu, \mathbf{r}, X_j, k_{j,1:C}),$$

$$\pi^{(c)} \sim P(\pi | \theta = \theta_c, \nu, \mathbf{r}, \sigma_c),$$

$$\theta_k^{(c)} \sim P(\theta_k | \pi = \pi_c, \nu, \mathbf{r}, \sigma_c).$$

4 Implementation

4.1 ADNI Dataset

We considered 382 subjects (135 cognitively normal subjects, 149 mild cognitive impairment, 98 Alzheimer's disease) who had a 1.5T structural MRI (T1) scan at baseline. We calculated the total volume (left plus right hemisphere) of 42 regions in the Neuromorphometrics parcellation (<http://neuromorphometrics.org:8080/>) for each subject, correcting for head size variance by regressing against total intracranial volume. Segmentation was performed using the Geodesic Information Flow framework [13]. We retained the 35 regions having significant differences between cognitively normal and Alzheimer's disease subjects using the Wilcoxon rank sum test with $p < 0.01$. We downloaded biomarker values from the ADNI database (adni.loni.usc.edu) for CSF markers ($A\beta_{1-42}$, tau, p-tau), cognitive test scores (MMSE, RAVLT, ADAS-Cog), and global FDG-PET metabolism.

4.2 Model Fitting

We compare the result of fitting the event-based model, generalised Mallows event-based model and Dirichlet process mixtures of generalised Mallows event-based models to the ADNI data set. Following previous work [6] we model the

probability that a biomarker is normal, $p(x|\neg e)$, as a Gaussian distribution, and the probability that a biomarker is abnormal, $p(x|e)$, as a uniform distribution covering the full range of observed values to reflect the range of severity that corresponds to an abnormal biomarker. We use a mixture model to fit these distributions to the data to account for a proportion of outliers in the control population. In subjects that had missing data points we imputed the biomarker values such that $p(x|e) = p(x|\neg e)$, i.e. it is equally probable that the event e has or has not occurred. The prior probability that a subject is at a particular disease stage $P(k)$ is assumed to be uniform. To fit the generalised Mallows model we need to sample σ from $P(\sigma|\gamma, \pi, \theta)$. We approximate this by sampling from a generalised Mallows model for each of the event sets in the partial ranking γ separately; the set of events γ_e that have occurred and the set of events $\gamma_{\neg e}$ that have yet to occur. We sample

$$\sigma_e \sim GM(\pi_{\gamma_e}, \theta_{\gamma_e}), \text{ and}$$

$$\sigma_{\neg e} \sim GM(\pi_{\gamma_{\neg e}}, \theta_{\gamma_{\neg e}}).$$

This means that the precedence of events specified by the partial ranking is preserved, and that the central ordering of the generalised Mallows model for each event set, π_{γ_e} and $\pi_{\gamma_{\neg e}}$, has the minimal Kendalls tau distance [8] from the central ordering π of the full generalised Mallows model. We sample k from $P(k|\sigma, X_j)$ using equation 1, i.e.

$$P(k|\sigma, X_j) \propto \prod_{i=1}^k p(x_{\sigma(i),j}|e_{\sigma(i)}) \prod_{i=k+1}^N p(x_{\sigma(i),j}|\neg e_{\sigma(i)}).$$

The remaining sampling updates follow the algorithm in [11]. We sample π exactly using a stage-wise algorithm, and θ using a beta function approximation. We used the Beta-Gibbs algorithm [11] to update the Dirichlet process mixture model cluster assignments c_j , weighting the probability each subject belongs to each cluster by $P(X_j|\sigma_{j,c}, k_{j,c})$, and the generalised Mallows model parameters π_c, θ_c for each cluster. We fix the priors to be $\nu = 1, \mathbf{r} = \mathbf{1}, \alpha = 1$. We initialise π randomly, γ_e as the set of events with $p(x|e) > p(x|\neg e)$, $\gamma_{\neg e}$ as the set of events with $p(x|e) \leq p(x|\neg e)$, and the Dirichlet process mixture to have 25 clusters.

5 Results and Discussion

5.1 The Event-Based Model

Figure 1 shows a positional variance diagram of the MCMC samples of the single ordering of events returned by the event-based model. We visualise a few key stages of this sequence in the top row of figure 3 to show the spatial correspondence of the sequence of regional volume loss estimated by the model. We find that CSF markers are the first to become abnormal, followed by cognitive test scores, then memory-related brain regions, then FDG-PET, and then

other Alzheimer’s disease-related brain regions. This sequence complements the findings of other studies, but provides a much more detailed picture of the regional progression of volume changes than has been seen previously in sporadic Alzheimer’s disease, and a direct comparison of the sequence of regional changes relative to a multi-modal set of biomarkers. Fonteijn et al. [6] looked at the regional progression of volume loss but in familial Alzheimer’s disease and using atrophy rates. The results in Young et al. [7] show a multi-modal sequence of biomarker abnormality in sporadic disease but for a small set of regional volumes, and hippocampal and whole brain atrophy rates from short-term longitudinal MRI. Here we show the first multi-modal sequence of biomarker abnormality in sporadic Alzheimer’s disease, including a large set of regional volumes. We are able to construct this picture from entirely cross-sectional data, and incorporate biomarkers with missing values.




Fig. 1: Central ordering estimated by the event-based model: Positional variance diagram of the MCMC samples of the maximum likelihood event sequence σ . The events on the y-axis are ordered by the maximum likelihood sequence estimated by the model. Each entry of the positional variance diagram represents the proportion of samples in which a particular event appears in a particular position in the central ordering, ranging from 0 in white to 1 in black. A black diagonal corresponds to high certainty in the ordering of events, whereas grey blocks in the diagram mean that the events permute.

5.2 The Generalised Mallows Event-Based Model

The generalised Mallows event-based model estimates both the central ordering of the events and the variance in this single event ordering across the population (Figure 3). Figure 3 compares the central ordering π and variance θ estimated

by the generalised Mallows event-based model, i.e. the range of event sequences across the population, with the central ordering estimated by the event-based model. The central event sequence has a similar ordering to the event-based model, but the variance in this central ordering of events increases, as shown by the increase in the number of orange regions in Figure 3. By using our Gibbs sampling technique we further obtain estimates of the uncertainty in each of the model parameters, as well as the latent variables included in the model, for example a subject’s disease stage (Figure 4). Fitting the generalised Mallows event-based model means that the uncertainty in this stage accounts for the variance in the ordering of the events across the population.




Fig. 2: Key for figures 3 and 5, generated using the BrainColorMap software [14].

5.3 Dirichlet Process Mixtures of Generalised Mallows Event-Based Models

We fitted a Dirichlet process mixture of generalised Mallows event-based models to allow for clusters of subjects that follow different sequences of events, of which each cluster has its own central ordering π_c and variance θ_c . The Dirichlet process mixture model identifies three main clusters in the data, with an average proportion of $0.48 (\pm 0.02)$, $0.24 (\pm 0.10)$, and $0.29 (\pm 0.10)$ subjects being assigned to each cluster respectively over the Gibbs samples. Figure 5 compares the estimated central ordering and variance for each of the clusters. The first two clusters look more Alzheimer’s disease-like than the third cluster, producing a similar event sequence to the event-based model and generalised Mallows model (figure 3), with CSF biomarkers and memory-related brain regions becoming abnormal early in the sequence. The third cluster likely captures outliers that do not fit the Alzheimer’s disease sequence of events. The ordering of events for the third cluster consists of only mild cognitive deficits and no CSF abnormalities, perhaps representing a normal aging event sequence, or simply reflecting that




Fig. 3: Comparison of the central ordering estimated by the event-based model (top) with the generalised Mallows model (bottom) (see key in figure 2). We display the results for six stages: stage 6, 12, 18, 24 and 36, where each stage number corresponds to the number of biomarkers that have become abnormal. Each biomarker (brain region, CSF, cognitive test or FDG-PET) is coloured according to the proportion of the population in which it has become abnormal by a particular stage along the central ordering. This proportion is estimated for the event-based model by the number of MCMC samples (figure 1), and for the generalised Mallows model by the probability (calculated using the central ordering π and spread θ) of an event appearing at or before a particular stage. This proportion ranges from 0 in yellow to 1 in red. Regions not included in the model are shown in grey. At each stage yellow biomarkers can be interpreted as being normal, red biomarkers as being abnormal, and orange biomarkers as varying in whether they have become abnormal across the population.

regional volume loss is a noisy measure on a cross-sectional level, i.e. subjects whose regional volume looks mildly abnormal compared to the rest of the population but have no other biomarker abnormalities. The variance θ_c is greater for the clusters of the Dirichlet process mixture model than the generalised Mallows model (as shown by an increase in the number of orange regions in figure 5 compared to figure 3), likely because each cluster only contains a proportion of the population, meaning that there are fewer subjects to fit the model to, and due to the uncertainty in the cluster assignment of each subject. Our Gibbs sampling technique returns samples of all of the model parameters. For example, we are




Fig. 4: Estimate of the uncertainty in a subject’s disease stage obtained by using Gibbs sampling to fit the generalised Mallows event-based model. We show an estimate of the probability of each stage for an example cognitively normal subject (green), mild cognitive impairment subject (blue), and Alzheimer’s disease subject (red). Each stage corresponds to the number of biomarkers in the sequence that have become abnormal.

able to estimate the uncertainty in the disease stage of each subject for both models, and the cluster assignment of each subject from the Dirichlet process mixture, producing a similar diagram to figure 4.

6 Conclusions

We have proposed a generalised family of event-based models that relax the assumption of a common event sequence over the population in different ways. We formulated these models so that they work for a large multi-modal set of sporadic disease biomarkers, and developed a Gibbs sampler that provides an estimate of the uncertainty on each model parameter. We have fitted this family of models to regional imaging data from ADNI to determine the sequence of regional volume loss in sporadic Alzheimer’s disease, incorporating a much more extensive, multi-modal set of biomarkers than has been seen previously. We find that the generalised Mallows model estimates a similar event sequence to the original event-based model, but with a larger variation across subjects. Fitting a Dirichlet process model detects subgroups of the population with different event sequences.

Many other possible extensions to this family of models are interesting to consider in the future. In their current form these new models do not provide much additional clinical utility, however they have the potential to provide much richer information than the original event-based model. Future work will extend the Dirichlet process model to incorporate data from different disease datasets, allowing the automatic extraction of biomarker orderings in different diseases and mixed pathology. The sampling techniques described naturally extend to incorporate multiple time points within an individual [10]. Extending these models to include longitudinal data will provide richer datasets for characterising the heterogeneity in individual event sequences. The family of models described here




Fig. 5: As figure 3 but for the clusters identified by the Dirichlet process mixture of generalised Mallows event-based models, with clusters 1 to 3 displayed from top to bottom.

have a large number of parameters, developing hierarchical models [15] will reduce the required number of parameters by capturing sections of the sequence that are common or distinct amongst subjects to allow more robust fitting to cross-sectional data.

The family of disease progression models we describe are potentially applicable to any disease or developmental process. The multiple orderings of events described by these models have potential use for outlier detection, differential diagnosis and to characterise disease subtypes, e.g. genetic subtypes, for improved patient stratification.

7 Acknowledgements

The EPSRC support this work with grant EP/J020990/01, EP/M020533/1 and EP/M006093/1. This work was supported by the NIHR Queen Square Dementia Biomedical Research Unit and the Leonard Wolfson Experimental Neurology Centre. The Dementia Research Centre is an ARUK coordination centre. D. Cash receives funding from an anonymous charitable foundation and support from the Brain Research Trust. S. Ourselin is partially funded by the EU-FP7 project VPH- DARE@IT (FP7-ICT-2011-9-601055) and the National Institute for Health Research University College London Hospitals Biomedical Research Centre (NIHR BRC UCLH/UCL High Impact Initiative). Data collection and sharing for this project was funded by the Alzheimer’s Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2- 0012).

References

1. C.R. Jack, D.S. Knopman, W.J. Jagust, L.M. Shaw, P.S. Aisen, M.W. Weiner, R.C. Petersen, and J.Q. Trojanowski. Hypothetical model of dynamic biomarkers of the Alzheimer’s pathological cascade. *Lancet neurology*, 9(1):119–28, 2010.
2. G. McKhann, D. Drachman, M. Folstein, R. Katzman, D. Price, and E.M. Stadlan. Clinical diagnosis of Alzheimer’s disease: report of the NINCDS-ADRDA Work Group under the auspices of Department of Health and Human Services Task Force on Alzheimer’s Disease. *Neurology*, 34:939–944, 1984.
3. V.L. Villemagne, S. Burnham, P. Bourgeat, B. Brown, K.A. Ellis, O. Salvado, C. Szoeke, S.L. Macaulay, R. Martins, P. Maruff, D. Ames, C.C. Rowe, and C.L. Masters. Amyloid β deposition, neurodegeneration, and cognitive decline in sporadic Alzheimer’s disease: a prospective cohort study. *Lancet neurology*, 12(4):357–67, 2013.
4. C.R. Jack, D.S. Knopman, W.J. Jagust, R.C. Petersen, M.W. Weiner, P.S. Aisen, L.M. Shaw, P. Vemuri, H.J. Wiste, S.D. Weigand, T.G. Lesnick, V.S. Pankratz, M.C. Donohue, and J.Q. Trojanowski. Tracking pathophysiological processes in Alzheimer’s disease: an updated hypothetical model of dynamic biomarkers. *Lancet neurology*, 12(2):207–16, 2013.
5. C.R. Jack and D.M. Holtzman. Biomarker Modeling of Alzheimer’s Disease. *Neuron*, 80(6):1347–1358, 2013.
6. H.M. Fonteijn, M. Modat, M.J. Clarkson, J. Barnes, M. Lehmann, N.Z. Hobbs, R.I. Scallan, S.J. Tabrizi, S. Ourselin, N.C. Fox, and D.C. Alexander. An event-based model for disease progression and its application in familial Alzheimer’s disease and Huntington’s disease. *NeuroImage*, 60(3):1880–9, 2012.
7. A.L. Young, N.P. Oxtoby, P. Daga, D.M. Cash, N.C. Fox, S. Ourselin, J.M. Schott, and D.C. Alexander. A data-driven model of biomarker changes in sporadic Alzheimer’s disease. *Brain*, 137:2564–2577, 2014.
8. M.A. Fligner and J.S. Verducci. Distance Based Ranking Models. *JSTOR: Journal of the Royal Statistical Society. Series B (Methodological)*, 48:359–369, 1986.
9. C.E. Antoniak. Mixtures of Dirichlet Processes with Applications to Bayesian Nonparametric Problems, 1974.

10. J. Huang and D.C. Alexander. Probabilistic Event Cascades for Alzheimer’s disease. In *Advances in Neural Information Processing Systems 25*, pages 3104–3112, 2012.
11. M. Meila and H. Chen. Dirichlet Process Mixtures of Generalized Mallows Models. In *Uncertainty in Artificial Intelligence (UAI)*, pages 285–294, 2010.
12. M. Meila and L. Bao. Estimation and Clustering with Infinite Rankings. In *Uncertainty in Artificial Intelligence (UAI)*, pages 393–402, 2008.
13. M.J. Cardoso, R. Wolz, M. Modat, N.C. Fox, D. Rueckert, and S. Ourselin. Geodesic information flows. *Medical image computing and computer-assisted intervention: MICCAI*, 15:262–70, 2012.
14. A. Klein, A. Worth, J. Tourville, B. Landman, T. Dal, C. Satrajit, and S.G. David. An interactive tool for constructing optimal brain colormaps. 16:84358, 2010.
15. C. Meek and M. Meila. Recursive Inversion Models for Permutations. *Advances in Neural Information Processing Systems 27*, pages 631–639, 2014.