# Unscrambling disease progression at scale: fast inference of event permutations with optimal transport

Peter A. Wijeratne (p.wijeratne@sussex.ac.uk)
Daniel C. Alexander (d.alexander@ucl.ac.uk)

NEURAL INFORMATION PROCESSING SYSTEMS

## Factor of >1000x speed-up enables scaling to large feature sets, providing new utility

Disease progression models infer group-level temporal trajectories of change in patients' features, providing unique insight into disease biology and staging systems. Discrete models consider disease progression as a latent permutation of events, where each event corresponds to a feature becoming probabilistically abnormal. However, permutation inference using traditional maximum likelihood approaches becomes prohibitive due to combinatoric explosion, severely limiting model dimensionality and utility.

- We leverage optimal transport (OT) to derive a new generative model of disease progression, the variational event-based model (vEBM).
- We use synthetic data to demonstrate the vEBM's improved speed and robustness to noise over appropriate baselines (see [1] for details).
- We use the vEBM with data from Alzheimer's disease (AD) and age-related macular degeneration (AMD), revealing, for the first time, pixel-level disease progression events in the brain and eye, respectively.

## Variational event-based model (vEBM)

### Reframing disease progression using optimal transport

We describe disease progression as a latent doubly-stochastic permutation matrix of events, $S$, which defines the OT coupling between distributions of normality and abnormality for each feature (Figure 1). We use entropy-regularised OT [2] and define a differentiable variational evidence lower bound (ELBO) using a Gumbel-Sinkhorn prior. Note that the model infers the sequence from a single set of features per individual ("snapshots").

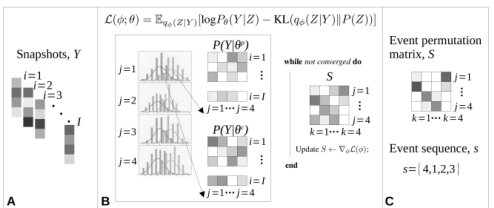


**Figure 1.** Schematic of the variational event-based model for a toy 4-feature dataset. **A.** The dataset contains snapshots from I individuals, with $j, k = \{1, 2, 3, 4\}$ features and latent events; the features can be of any type and can be incomplete. **B.** Before inference, probabilistic models of normality and abnormality are fit to the dataset, giving the likelihood look-up tables $P(Y|\theta p, c)$; these are fixed throughout inference, as denoted by the inner box outside the training loop. To infer the permutation matrix $S$, the ELBO is optimised and $S$ is updated each iteration using the Sinkhorn-Knopp algorithm. **C.** The resulting hard permutation, $s$, i.e., the disease event sequence, is obtained from $S$, using the Hungarian algorithm.

## Pixel-level disease events in the brain and eye

### First application of vEBM to imaging data in AD and AMD

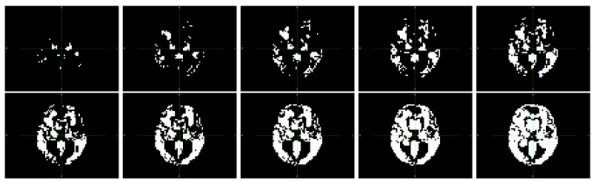

**Figure 2.** Pixel-level disease progression sequence in AD obtained by the vEBM. White pixels correspond to events that have occurred by the corresponding point of the sequence. The figure shows 10 sequence positions at uniform steps of 100 across the total of 1344, with the top left figure corresponding to position 50 (the first 50 events have occurred) and the bottom right to position 950. Images were made from the vEBM output using 3D Slicer (https://www.slicer.org/).

We apply the vEBM separately to magnetic resonance imaging (MRI) tensor-based morphometry data from the Alzheimer's Disease Neuroimaging Initiative dataset (Figure 2); and optical coherence tomography (OCT) data from the Duke University Opthalmology dataset (Figure 3). The vEBM provides pixel-level visualisation of the order of pathology appearance and new insights into disease progression.

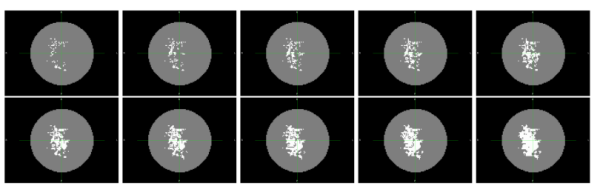

**Figure 3.** Pixel-level disease progression sequence in AMD obtained by the vEBM. White pixels correspond to events that have occurred by the corresponding point of the sequence. We have selected 10 sequence positions at uniform steps of 50 across the total of 537 in the full sequence, with the top left figure corresponding to position 80 and the bottom right to position 530. Images were made from the vEBM output using 3D Slicer (https://www.slicer.org/).

To evaluate our ADNI pixel-level model with respect to previous analyses that have used segmented regional brain volumes, we map the vEBM pixel-level events post hoc to pixel-level labels obtained from the FreeSurfer (https://surfer.nmr.mgh.harvard.edu/) segmentation of the reference template (Figure 4). Our findings are in broad agreement with previous results, but our model provides more fine-grained insights; we now obtain continuous trajectories of change, which capture interesting non-linearities

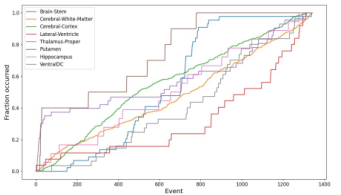


Figure 4. Trajectories of regional brain volumes in our ADNI cohort, obtained by mapping the vEBM pixel-level events to pixel-level labels obtained from the FreeSurfer segmentation of the reference template. The horizontal axis shows the event number (from 0 – 1344), and the vertical axis shows the fraction of pixel-events that have occurred in each regional brain volume at the corresponding event number, as defined by the vEBM event sequence.

## Broader impact

### Low-compute disease progression modelling at scale

The vEBM enables disease progression modelling at scale in multiple areas of medical imaging, not only the modalities demonstrated here, e.g., diffusion weighted imaging, microstructure modelling, connectivity; other imaging modalities, e.g., positron emission tomography, computed tomography, X-rays, ultrasound; and non-radiological imaging modalities, e.g., microscopy. Furthermore, the vEBM can run quickly on a relatively low-spec computer without the need for GPU infrastructure, making it accessible to research labs – and potentially clinics – that have limited resources, while further minimising its carbon impact by reducing compute time.The vEBM code is available here: https://github.com/pawij/vebm

### References

[1] Peter A. Wijeratne & Daniel C. Alexander (2024). "Unscrambling disease progression at scale: fast inference of event permutations with optimal transport". In: Advances in Neural Information Processing Systems (NeurIPS) 38. arXiv:2410.14388

[2] Marco Cuturi (2013). "Sinkhorn distances: lightspeed computation of optimal transport distances". In: Advances in Neural Information Processing Systems (NeurIPS) 26. arXiv:1306.0895

# Unscrambling disease progression at scale: fast inference of event permutations with optimal transport

## Factor of >1000x speed-up enables scaling to large feature sets, providing new utility

Disease progression models infer group-level temporal trajectories of change in patients' features, providing unique insight into disease biology and staging systems. Discrete models consider disease progression as a latent permutation of events, where each event corresponds to a feature becoming probabilistically abnormal. However, permutation inference using traditional maximum likelihood approaches becomes prohibitive due to combinatoric explosion, severely limiting model dimensionality and utility.

- We leverage optimal transport (OT) to derive a new generative model of disease progression, the variational event-based model (vEBM).
- We use synthetic data to demonstrate the vEBM's improved speed and robustness to noise over appropriate baselines (see [1] for details).
- We use the vEBM with data from Alzheimer's disease (AD) and age-related macular degeneration (AMD), revealing, for the first time, pixel-level disease progression events in the brain and eye, respectively.

## Pixel-level dis

First application of vEI

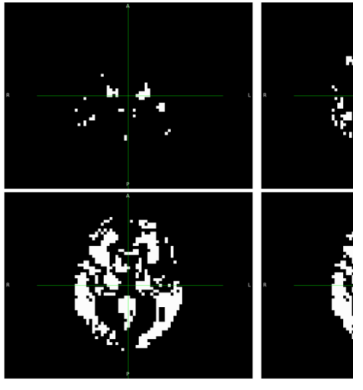

**Figure 2.** Pixel-level disease pro
have occurred by the correspond
100 across the total of 1344, with
bottom right to position 950.

We apply the vEBM se
tensor-based morphor

# Variational event-based model (vEBM)

Reframing disease progression using optimal transport

We describe disease progression as a latent doubly-stochastic permutation matrix of events, $S$, which defines the OT coupling between distributions of normality and abnormality for each feature (Figure 1). We use entropy-regularised OT [2] and define a differentiable variational evidence lower bound (ELBO) using a Gumbel-Sinkhorn prior. Note that the model infers the sequence from a single set of features per individual (“snapshots”).

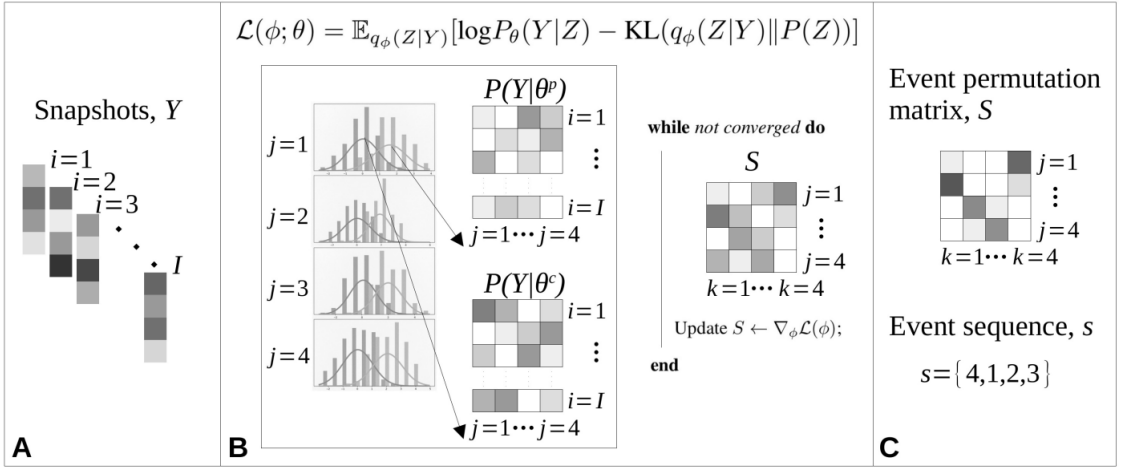


**Figure 1.** Schematic of the variational event-based model for a toy 4-feature dataset. **A.** The dataset contains snapshots from I individuals, with *j, k* = {1, 2, 3, 4} features and latent events; the features can be of any type and can be incomplete. **B.** Before inference, probabilistic models of normality and abnormality are fit to the dataset, giving the likelihood look-up tables $P(Y|\theta p,c)$; these are fixed throughout inference, as denoted by the inner box outside the training loop. To infer the permutation matrix *S*, the ELBO is optimised and *S* is updated each iteration using the Sinkhorn-Knopp algorithm. **C.** The resulting hard permutation, *s*, i.e., the disease event sequence, is obtained from *S*, using the Hungarian algorithm.

# Pixel-level disease events in the brain and eye

First application of vEBM to imaging data in AD and AMD

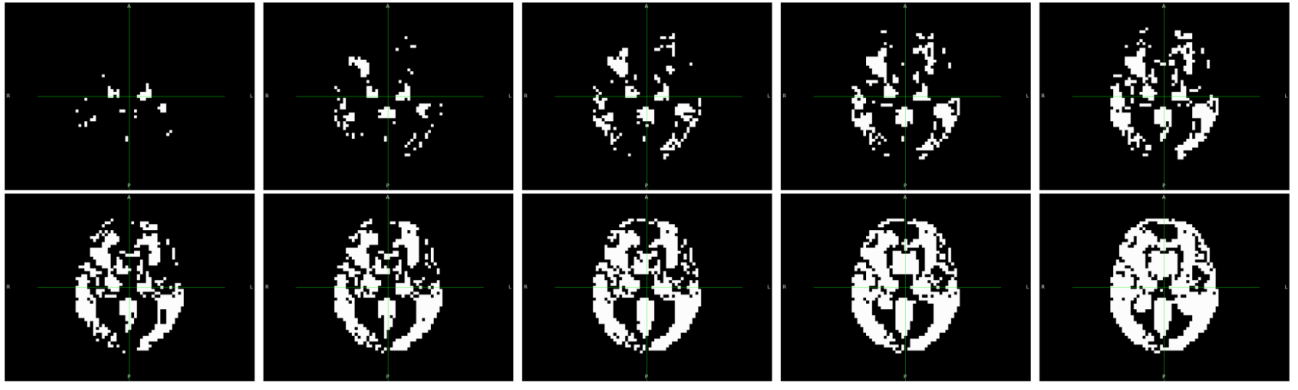

**Figure 2.** Pixel-level disease progression sequence in AD obtained by the vEBM. White pixels correspond to events that have occurred by the corresponding point of the sequence. The figure shows 10 sequence positions at uniform steps of 100 across the total of 1344, with the top left figure corresponding to position 50 (the first 50 events have occurred) and the bottom right to position 950. Images were made from the vEBM output using 3D Slicer (https://www.slicer.org/).

We apply the vEBM separately to magnetic resonance imaging (MRI) tensor-based morphometry data from the Alzheimer's Disease Neuroimaging Initiative dataset (Figure 2); and optical coherence tomography (OCT) data from the Duke University Opthalmology dataset (Figure 3). The vEBM provides pixel-level visualisation of the order of pathology appearance and new insights into disease progression.

Figure 4.
pixel-lev
event num

Broa

Low-co

ussex.ac.uk)

@ucl.ac.uk)

NEURAL INFORMATION PROCESSING SYSTEMS

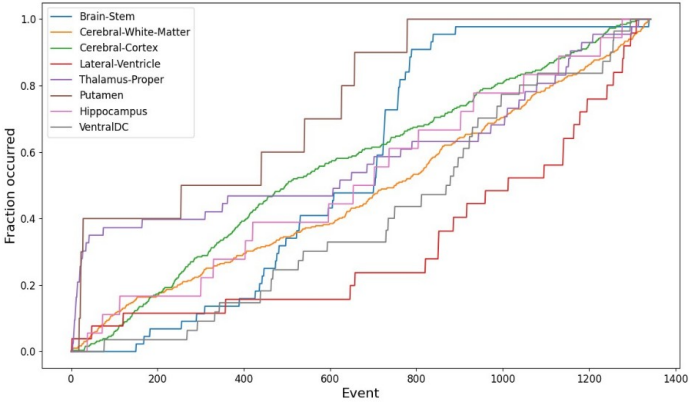


Figure 4. Trajectories of regional brain volumes in our ADNI cohort, obtained by mapping the vEBM pixel-level events to pixel-level labels obtained from the FreeSurfer segmentation of the reference template. The horizontal axis shows the event number (from 0 – 1344), and the vertical axis shows the fraction of pixel-events that have occurred in each regional brain volume at the corresponding event number, as defined by the vEBM event sequence.

We apply the vEBM separately to magnetic resonance imaging (MRI) tensor-based morphometry data from the Alzheimer's Disease Neuroimaging Initiative dataset (Figure 2); and optical coherence tomography (OCT) data from the Duke University Opthalmology dataset (Figure 3). The vEBM provides pixel-level visualisation of the order of pathology appearance and new insights into disease progression.

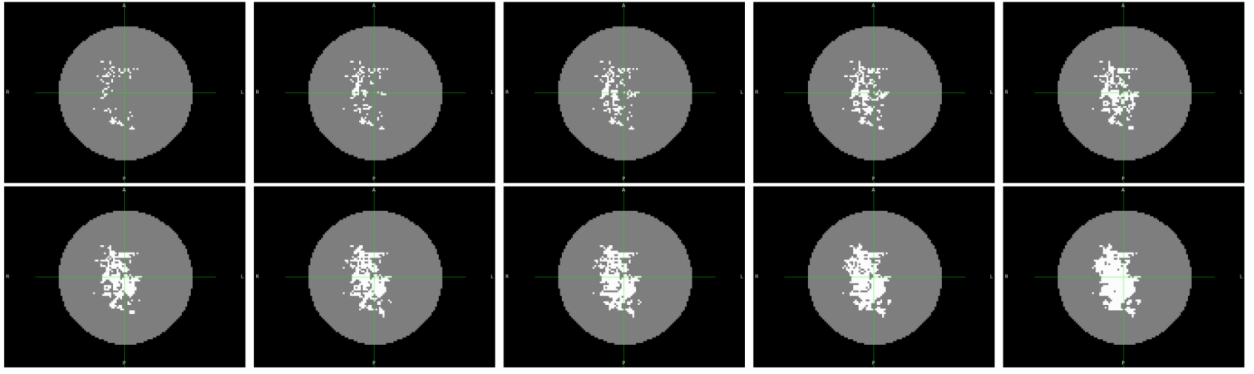

**Figure 3.** Pixel-level disease progression sequence in AMD obtained by the vEBM. White pixels correspond to events that have occurred by the corresponding point of the sequence. We have selected 10 sequence positions at uniform steps of 50 across the total of 537 in the full sequence, with the top left figure corresponding to position 80 and the bottom right to position 530. Images were made from the vEBM output using 3D Slicer (https://www.slicer.org/).

To evaluate our ADNI pixel-level model with respect to previous analyses that have used segmented regional brain volumes, we map the vEBM pixel-level events post hoc to pixel-level labels obtained from the FreeSurfer (https://surfer.nmr.mgh.harvard.edu/) segmentation of the reference template (Figure 4). Our findings are in broad agreement with previous results, but our model provides more fine-grained insights; we now obtain continuous trajectories of change, which capture interesting non-linearities

# Broader impact

## Low-compute disease progression modelling at scale

The vEBM enables disease progression modelling at scale in multiple areas of medical imaging, not only the modalities demonstrated here, e.g., diffusion weighted imaging, microstructure modelling, connectivity; other imaging modalities, e.g., positron emission tomography, computed tomography, X-rays, ultrasound; and non-radiological imaging modalities, e.g., microscopy. Furthermore, the vEBM can run quickly on a relatively low-spec computer without the need for GPU infrastructure, making it accessible to research labs – and potentially clinics – that have limited resources, while further minimising its carbon impact by reducing compute time.The vEBM code is available here: https://github.com/pawij/vebm

## References


[1] Peter A. Wijeratne & Daniel C. Alexander (2024). “Unscrambling disease progression at scale: fast inference of event permutations with optimal transport”. In: Advances in Neural Information Processing Systems (NeurIPS) 38. arXiv:2410.14388

[2] Marco Cuturi (2013). “Sinkhorn distances: lightspeed computation of optimal transport distances”. In: Advances in Neural Information Processing Systems (NeurIPS) 26. arXiv:1306.0895