A STATISTICAL FRAMEWORK FOR THE MODELING OF THE CONTINUOUS PROGRESSION OF ALZHEIMER'S DISEASE.

BY ANAMIKA AGRAWAL[1,6], VICTORIA M. RACHLEFF[2,4], KYLE J. TRAVAGLINI[2], SHUBHABRATA MUKHERJEE[5], PAUL K. CRANE[5], MICHAEL HAWRYLYCZ[2], C. DIRK KEENE[4], ED LEIN[2], GONZALO E. MENA[3,a] AND MARIANO I. GABITTO[2,7,b]

[1]*Center for Data-Driven Discovery for Biology, Allen Institute*

[2]*Human Cell Types Department, Allen Institute*

[3]*Department of Statistics & Data Science, Carnegie Mellon University*

[4]*Department of Laboratory Medicine and Pathology, University of Washington*

[5]*Department of Medicine, University of Washington*

[6]*Department of Physiology and Biophysics, University of Washington*

[7]*Department of Statistics, University or Washington,* [a]*gmena@andrew.cmu.edu;* [b]*mariano.gabitto@alleninstitute.org*

Throughout an organism's life, a multitude of biological systems transition through complex biophysical processes. These processes serve as indicators of the underlying biological states. Inferring these latent unobserved states is a key problem in modern biology and neuroscience. Unfortunately, in many experimental setups we can at best obtain snapshots of the system at different times for different individuals, and one major challenge is the one of reconciling those measurements. This formalism is particularly relevant in the study of Alzheimer's Disease (AD) progression, in which we observe in brain donors the aggregation of pathological proteins but the underlying disease state is unknown. The progression of AD can be modeled by assigning a latent score - termed pseudotime - to each pathological state, creating a pseudotemporal ordering of donors based on their pathological burden.

This paper proposes a hierarchical Bayesian framework to model AD progression using detailed quantification of multiple AD pathological proteins from the Seattle AD Brain Cell Atlas consortium (SEA-AD). Inspired by biophysical models, we model pathological burden as an exponential process. Theoretical properties of the model are studied, by using linearization to reveal convergence and identifiability properties. We provide Markov chain Monte Carlo estimation algorithms, and show the effectiveness of our approach with multiple simulation studies across data conditions. Applying the methodology to SEA-AD brain data, we infer pseudotime for each donor and order them by pathological burden. Finally, we analyze the information within each pathological feature and utilize it to refine the model by focusing on the most informative pathologies. This lays the groundwork for suggesting future experimental design approaches.

1. Introduction. Modeling how cells or individuals change over time, due to endogenous or exogenous effects, is a fundamental problem in the natural and social sciences. At the heart of this problem lies the longstanding goal of learning a longitudinal trajectory delineating the dynamic evolution of each system. A wealth of research has cast this problem as learning a latent unobserved variable -termed pseudotime- dictating the progression of events. Under this framework, relevant measured quantities are expressed as functions of

Keywords and phrases: Hierarchical Bayesian Inference, Repulsive Prior, Identifiable, Alzheimer's Disease, Trajectory Inference.

this variable. Pseudotime inference was first motivated by the famous metaphor of Conrad Waddington in which undifferentiated cells progress along a continuum developmental trajectory to reach their committed cell fates (Waddington, 1940, 1957). Inference of such models is challenging because of two reasons. Firstly, the lack of a universal definition of time renders the problem unidentifiable unless constraints are imposed. Secondly and most critically, the successful inference of the parameters governing the system's evolution requires the inference of the pseudo time variable as well, adding an extra layer of complexity.

This is particularly true in modern biological applications, where measurements describing progression can only be collected once per individual, providing only a cross-sectional snapshot of the system's state. Aligning these cross-sections along the pseudotime axis poses an additional inferential challenge, potentially impacting our ability to recover underlying dynamics compared to scenarios with multiple time measurements per individual or when pseudotimes are known beforehand.

Although our methodology is general enough to accommodate similar problems, here we focus on the study of neurodegeneration in Alzheimer's Disease (AD). We consider biophysically-inspired models for protein pathology accumulation and the pseudo-temporal dynamics of disease progression and infer them by leveraging a recently published dataset profiling multiple brain donors spanning the entire AD spectrum. Crucially, these cross-sectional measurements encompass a diverse set of pathophysiological markers that provide information on different stages of disease progression.

1.1. *Relation to prior work.* Pseudotime or ordering methods are widely applied due to their practical significance across various domains. While both share the goal of inferring data's inherent order, pseudotime goes a step further than ordination by attributing significant importance to the pseudotime latent variable, quantifying the system's progression along a dynamic trajectory. Here, we offer a concise overview of related methods. In the analysis of cellular development using single-cell RNA-seq data, pseudotimes are classically derived by computing a low-dimensional representation of each cells, inferring principal curves, followed by computing suitable minimum spanning trees (Magwene, Lizardi and Kim, 2003; Trapnell et al., 2014; Street et al., 2018). Opposed to these non-parametric methods, a line of methods states the problem from a model-based, parametric perspective. For example, in the context of ecological niches and species modeling (Hui et al., 2015; O'Hara and van der Veen, 2024; Hui et al., 2023; Popovic, Hui and Warton, 2022; van der Veen et al., 2023; Roberts, 2020; Hoegh and Roberts, 2020), measurements of relative abundances of species are sorted along sites of greater abundance. There, generalized latent variable models are used to represent ordination as latent variables, and no true ordination occurs, but instead, multivariate latent variable ordination is used to represent underlying gradients that influenced species composition. In Political Science, the related problem of sorting latent preferences of agents (such as judges) given observed votes appears under the name of 'ideal point' models (Gelman and Hill, 2006) and this latent individual ordering is inferred by stating random-effects variables.

Model-based approaches facilitate the simultaneous inference of the pseudotimes as well as any other relevant parameters. For example, closest to our work is the method by Campbell and Yau (2018) that, in the context of single-cell omics, states the pseudotime inference as a matrix factorization-like problem where parameters are inferred through Bayesian methods. Then, it is possible to characterize both the timing of events as well as the progression of each feature (genes). While their setup is similar to ours, their work focuses on modelling, and it doesn't study in-depth the interplay of the different variables (as the number of samples and dimension of cross-sections), the effects of pooling or shrinkage, and it does not provide rigorous consistency and performance guarantees as we do here.

Finally, a growing line of research applies pseudotime-based analysis to the context of aging (Pierson et al., 2019) and disease progression analysis, notably, to the study of cancer progression (Gupta and Bar-Joseph, 2008; Huang et al., 2023) and neural degeneration, including AD (Mukherjee et al., 2020). However, the majority of cited studies fail to address a fundamental degeneracy present in latent variable inference problems, wherein pseudotime solutions can be scaled or affine transformed while still producing a valid solution—a mathematical indeterminacy inherent in many proposed frameworks. Our work attempts to resolve such degeneracy by imposing relevant mathematical constraints and is unique in that it addresses both the theoretical and applied aspects, and we expect it will help better inform the deployment of pseudotime inference methods for disease progression analysis.

1.2. *Organization.* This paper is organized as follows. First, in section 2 we introduce the study of neurodegeneration and protein aggregation caused by AD pathology, and the data set in which we based our work. In section 3, we introduce a Biophysically-inspired model to describe AD neuropathology aggregation. Biophysical considerations justify the use of generalized linear models with an exponential link function tying the unobserved timing of events to observations. We supplement this likelihood specification with priors that constrain the parameter space and represent complex dependency structures across measurements. As a result, we obtain a family of Hierarchical Bayesian models that we infer using modern computational tools. In section 4 we provide a mathematical analysis of a simplified linear model under which we are able to make precise statements regarding the identifiability, consistency and convergence rates of our model as the number of individuals and measurements increase. Chief to this analysis is the observation that the underlying inferential target can be understood at the most fundamental level as an instance of noisy matrix factorization. These theoretical results are supplemented with extensive simulations in Section 5. These simulations show that in most aspects the theoretical results presented in the linear case also manifest in the generalized-linear setup. Simulations also show a few distinctive phenomena of generalized linear models that are not well captured by theoretical results on the linear case. In section 6, we deploy our method to the real-world dataset, and building on previous sections and specialized biological knowledge, we demonstrate that we are able to estimate pseudotime and progression parameters reliably. In section 6.2, we extend our modelling framework to inform future experimental design. Finally, in section 7, we discuss the significance and limitations of our work, and we sketch future directions.

2. SEA-AD dataset. This work is motivated by the need to understand the impact of pathological protein aggregation on cellular vulnerability across individuals spanning the complete spectrum of AD pathology. Quantitative neuropathological analysis was performed on the middle temporal gyrus (MTG) (Fig. 1A), a region of the cerebral cortex organized in a layered structured (from layers 1 to 6), of $d = 84$ postmortem brain donors included in the Seattle Alzheimer's Disease Brain Cell Atlas (SEA-AD) study (Gabitto et al., 2023). Donors from the SEA-AD cohort were drawn from two ongoing longitudinal studies: the Adult Changes in Thought (ACT) and the University of Washington Alzheimer's Disease Research Center (ADRC). We note that the SEA-AD cohort consists of donors with age distribution skewed towards advanced age (average age at death 88, SD=8) (Fig. 1B), which suggest a priori that high AD pathology values might be overrepresented in the cohort. The publicly available quantitative neuropathology dataset (accessible at sea-ad.org) comprises $M = 35$ total measurements derived from 6 immunohistochemical stains targeting pathological proteins (hyperphosphorylated tau (pTau) and amyloid beta (A-β), along with additional markers selected for examining cellular changes in neurons, astrocytes, and microglia. Table 1 presents key markers used in this study and their corresponding extracted features.

Immunostains for each protein and cellular marker are processed by machine learning algorithms and creates masks for each protein pathology, segments positive areas, and quantify the burden of each pathology (Fig. 1A). Next, we provide a brief overview of general AD pathophysiology and elucidate how these pertinent variables contribute to understanding AD progression.

AD is the most common cause of age-related dementia and is characterized neuropathologically by the stereotyped spatiotemporal progression of two pathological proteins, intracellular hyperphosphorylated tau (pTau) in the form of neurofibrillary tangles (NFTs) and amyloid beta ($A\beta$) in the form of extracellular plaques. In AD, pTau and $A\beta$ follow opposite progression patterns throughout the brain. Neuropathologists have mapped the progression of these pathological proteins and established discrete staging scales for each of them. These staging protocols order donors based on the presence or absence of pathological proteins in defined brain regions - Braak stages (I-VI) semiquantitatively assess pTau NFTs (Braak and Braak, 1991), while Thal Phasing (1-5) semiquantitatively assesses $A\beta$ plaques (Thal et al., 2002). A higher score in any scale indicates the presence of the relevant pathology across a greater number of brain regions. The AD neuropathologic change score (ADNC) (Hyman et al., 2012), attempts to integrate both scores into a general scale of AD progression. The two hallmark AD pathological proteins spatially overlap for the first time in the temporal lobe of the cerebral cortex, specifically the inferior and middle temporal gyrus (MTG). The MTG a key transition zone which delineates normal aging and preclinical AD from advanced stages of AD associated with widespread neocortical spread of pathology (Braak and Braak, 1991).

TABLE 1
Measurements used named by markers and the measured quantity

Feature name	Biological measurement	Type of marker
percent 6e10 positive area	amyloid-beta immunoreactive area	pathological peptide
number of 6e10 positive objects per area	number of amyloid-beta plaques	pathological peptide
average 6e10 positive object area	size of amyloid-beta plaques	pathological peptide
number of AT8 positive cells per area	number of pTau positive cells	pathological peptide
percent AT8 positive area	pTau immunoreactive area	pathological peptide
average AT8 positive cell area	pTau immunoreactive area	pathological peptide
number of NeuN strong positive cells per area	NeuN immunoreactive area (mature neurons)	cellular
percent NeuN positive area	NeuN immunoreactive area (mature neurons)	cellular
percent of Iba1 and 6e10 positive co-localized objects	Iba1 immunoreactive area (microglia)	cellular-pathological colocalization
number of Iba1 and 6e10 positive co-localized objects per area	Iba1 immunoreactive area (microglia)	cellular-pathological colocalization
average Iba1 positive process area per cell	Iba1 immunoreactive area (microglia)	cellular-pathological colocalization
percent GFAP positive area	GFAP immunoreactive area (astrocytes)	cellular marker

In addition to pathology measurements, our dataset encompasses a range of supplementary metadata pertaining to each brain donor, which includes clinical, cognitive, and demographic information. Established pathological assessments, previously described, such as Braak staging, Thal phasing, and ADNC are part of the repertoire of variables included, together with the last Cognitive Abilities Screening Instrument (CASI) (Teng et al., 1994a) test, administered during the donor's lifetime to evaluate cognitive function (Fig. 1C). We used these

FIG 1. *Seattle Alzheimer's Disease Brain Cell Atlas (SEA-AD) data sets profiling neuropathology in AD brain donors.* (A) Left, image of the middle temporal gyrus of a representative brain donor in which pathology data was stained and imaged over L = 5 cortical layers of MTG (middle box). Right, amyloid-beta aggregates stained with 6e10 stain are segmented by using ML algorithms (green masks) and quantified. (B) Histogram depicting age distribution of donors at death, average age at death = 88, median age = 90. (C) Histogram depicting the distribution of ADNC and CASI scores for SEA-AD donors. (D1-D4) Neuropathological measurements, percentage AT8 (measuring phosphorylated Tau pathology) and 6e10 (measuring amyloid β pathology) positive stain area, for each MTG cortical layer organizing donors by (1) BRAAK stages, (2) THAL phases, (3) ADNC scores, and (4) CASI scores.

per-donor-metrics to provide initial descriptions of the data and performed data-entry quality controls. Illustrative quantitative neuropathology measurements, categorized by staging and cognitive criteria, are depicted in Fig. 1D1-D4. This figure illustrates a noticeable increment in pathological burden as metrics increase in value. While the full dataset consists of $M = 35$ pathological markers, we use a subset of those in our model. We considered $M = 12$ markers (described in table 1), removing noisy biomarkers that have low signal-to-noise ratio or that were sparse, i.e. data was available just for a few donors.

3. Model Description. In a broader context, we aim to elucidate the challenge of modeling the inherent variability in a biological process by using a latent variable that represents pseudotime. The latent variable is named pseudotime because it not only orders observations but also governs the progression of the biological process. While we specifically focus on AD pathology aggregation (see section 2), this framework has broader applications, such as investigating aging-related phenotypes or studying different cell types in cellular development.

In modeling AD progression, we assume donors traverse a single disease trajectory spanning from no disease to the end stage, yet this trajectory remains unknown. Donors are randomly selected across the disease spectrum, but we lack prior knowledge of the sampling distribution, $p(t)$. For each of the donors, we collect multiple measurements reflecting the state along disease trajectory, including varying levels of prionic proteins. Lastly, observations are corrupted by observational noise (Fig. 2). Notable previous work attempting to solve similar problem in an aging context, discovered that to ensure identifiability of the link function between the latent pseudotime and the observational process, the link function (represented by $f(t)$ in Fig. 2) should be monotonic and injective (Pierson et al., 2019). However, in that case, time was a known variable, compared to our case in which nothing is known about latent pseudotime.

In the next sections we construct a model for cross-sectional SEA-AD donor data, $X_{m,l}^d$, which comprises quantitative neuropathological measurements from $d = 1 \ldots D = 84$ donors, encompassing $m = 1 \ldots M = 12$ distinct protein measurements taken across $l = 1 \ldots L = 5$ cortical layers in MTG. These measurements reflect the accumulation of various pathological proteins within each donor. We define a donor-specific one-dimensional latent variable, referred to as $t^d \in [0, 1]$. This variable represents the pseudo-temporal score that organizes donors according to their varying disease burdens, from low ($t_d = 1$) to high ($t_d = 1$). Inspired by biophysical principles of protein aggregation, and the cited aging-related studies (Pierson et al., 2019), we will find that a monotonic and injective exponential link function best accommodates our data. We will infer all model parameters through Bayesian inference.

FIG 2. *Schematic representing the modeling of a general biological process by a latent variable, t, termed pseudotime. In this work, pseudotime represents a scale that describe the trajectory of Alzheimer's Disease along the entire disease spectrum. Donors are sampled along AD trajectory from distribution $p(t)$. Then, multiple observed biophysical processes are measured from each donor. These measurements can be linked to the underlying pseudotime by a mean function, $f(t)$, and when observed, are corrupted by noise (described by the distribution $p(x|t)$ of observational noise.*

3.1. *Deterministic Biophysical model of Pathology Aggregation.* Alzheimer's Disease is characterized by the proliferation of misfolded protein aggregates, often called prions (Zhou and Liu (2013)). As mentioned before, amyloid-Beta ($A\beta$) aggregates and Tau neurofibrillary tangles (τ-NFT) are two examples of well-studied protein aggregates that have been widely used as pathological markers to study the spread and severity of AD (Braak and Braak, 1991; Thal et al., 2002). Several mechanistic models, chiefly rooted in polymer kinetics, have been used to study the aggregation and seeding dynamics of protein aggregates (Vaquer-Alicea and Diamond, 2019; Davis and Sindi, 2016). In our work, we will leverage the Nucleated Polymerization Model (NPM), which is the most widely accepted model used to describe protein aggregates propagation dynamics (Masel, Jansen and Nowak, 1999), to describe the temporal progression of AD-associated pathology markers.

The NPM model describes a cycle of aggregate growth followed by multiplication, leading to the dominance of toxic protein aggregates growth. The growth of such aggregates occurs through attachment with surrounding 'healthy' monomeric units, which is then followed by a fragmentation process that multiplies the number of 'seeds' capable of growing into many more misfolded aggregates (Fig. 3A). Once the monomers are incorporated, they become a part of the "unhealthy" or misfolded population. The misfolded aggregates also tend to be more stable than their healthy counterparts (Tompa et al., 2001). In an aggregation model for a spatially-limited domain (e.g. a single anatomical region of the brain), the important dynamical parameters to consider are the initial seeding levels, denoted by β_0, and the rate of multiplication through aggregation and division, denoted by β_1 ((Fig. 3B). In the 'early' time regime, i.e. when the available pool of healthy protein monomers is not limiting, we will describe the evolution of a pathological protein $X_{m,l}$ for measurement m in layer l with the following dynamics arising from the NPM biophysical model (Masel, Jansen and Nowak, 1999):

$$X_{m,l}(t) = e^{\beta_0^{m,l} + \beta_1^{m,l} t} \tag{1}$$

The details of the kinetic model of pathology propagation, and the simplifying assumptions used to arrive to such functional form, are included in the supplement section 1.

3.2. *Modeling the Generative Distribution of Observed Pathology.* The above biophysical considerations justify using a trajectory that is an exponential transformation of an affine function of the latent pseudotime variable. To account for observation noise, we couple this latent trajectory with a sampling model where observed counts $X_{m,l}^d \in \mathbf{N}_0$ describing pathological proteins follow a Negative Binomial distribution. Specifically, we consider a count parameter $A_{m,l} \in \mathbf{N}^+$, and a probability of success in the (negative) logit scale $\rho \in \mathbf{R}$ such that the likelihood of $X_{m,l}^d$ given $A_{m,l}$ and $\rho = \beta_0^{m,l} + \beta_1^{m,l} t^d$ writes as

$$p(X_{m,l}^d = x | A_{m,l}, \beta_1^{m,l}, \beta_0^{m,l}, t^d) = \binom{x + A_{m,l} - 1}{x} \left(\frac{1}{1 + e^{-\rho}} \right)^x \left(\frac{1}{1 + e^{\rho}} \right)^{A_{m,l}}. \tag{2}$$

The mean and variance of the above distribution are given by $\mu = A_{m,l} e^\rho$, $\sigma^2 = \mu + \frac{\mu^2}{A_{m,l}}$. This implies that the conditional mean of $X_{m,l}^d$ given $\beta_0^{m,l}, \beta_1^{m,l}$ and t^d is exactly the quantity $X_{m,l}(t)$ defined in Eq. (1), up to the $A_{m,l}$ factor that can be absorbed into $\beta_0^{m,l}$. Consequently, the Negative Binomial model can be understood as providing a noisy version of the underlying dynamics. As customary, we understand A as an overdispersion parameter A (Gelman, Hill and Vehtari, 2021) modulating different levels of observation noise whereby in the limit $A \to \infty$ we recover the usual Poisson regression. Eq. (2) is one instance of a generalized linear model with a logarithmic link function. Still, other choices could be made if observations were of different nature. For example, logistic regression could be used if observations were binary.

FIG 3. *We assume that AD protein pathology aggregation can be described by an exponential process. A) Schematic representing protein aggregation process. Due to protein misfolding, an initial seeding event is triggered which later makes the aggregate to grow. Randomly, protein aggregates break down multiplying in number. The conjunction of growths and multiplications events can be described by exponential dynamics. B) scatterplot depicting measurement for $A\beta$ burden, the percentage area stained positive by 6e10 antibody, in each cortical layer ordered according to our model. For each protein pathology the model fits an initial seeding variable, β_0, and a growth aggregation rate β_1.'*

3.2.1. *Priors.* In addition to the the likelihood in Eq. (2) we must specify priors for the model parameters $\Theta = \{t, \beta_0, \beta_1, A\}$. First, we assume that times are sampled independently and follow a Beta distribution with parameters t_a, t_b:

$$(3) \qquad t^i \sim Beta(t_a, t_b).$$

In particular, $t^d \in [0, 1]$. The decision to choose such a bounded interval is two-fold. At the conceptual level, biological systems are constrained in space and time. This implies in our application to AD that protein propagation and disease trajectory do not last indefinitely. When looking at a particular brain region, the accumulation of pathological proteins will range from 0 to a high number, with the last case implying that the pathological protein has covered the entire region. At the mathematical level, as we will later show in Section 4, choosing a bounded interval enables identifiability of the entire system, a desired property in our setup where our latent variable represents not only ordination but also pseudo time. The case $t_a = t_b = 1$ corresponds to a uniform prior. Although this prior can seem unbiased, it encodes strong beliefs about the phenotypical diversity of the sample; for example, that the average sampled time is 0.5. This justifies the use of a more general $Beta$ prior to encode information on the sampling process. We will see in Section 4 that we can benefit from suitable choices of this prior.

We must also specify priors on $\beta_0^{m,l}, \beta_1^{m,l}$, parameters that represent the baseline level of pathology, and the exponential rate of pathological growth of each of the markers. We will consider two models, one with measurement and layer-independent priors on $\beta_1^{m,l}$ and the second with hierarchical priors pooling information across layers. In the former case, we are assuming that only the protein aggregation rate (through its parameter β_1) can be related across layers, and for each measurement, we will infer a covariance matrix across layers. Finally, we consider priors for the over-dispersion parameters $A^{m,l}$ since they are not revealed to us. In this case, we impose half-normal distribution. In summary, the plate diagrams of our hierarchical and fully factorized models are depicted in figure 4, and we detail next the entire generative process of our model:

1. For each donor, d
   a) Choose $t_d \sim \text{Beta}(1, 1)$
2. For each measurement, m
   a) **If** Hierarchical
      Then $\beta_1^m \sim \text{N}(0, 1)$
         $L_\Sigma^m \sim \text{LKJ}(M, 1)$
         $d_\sigma^m \sim \text{HCauchy}(1)$
         $\Sigma^m = \sqrt{d_\sigma^m} L_\Sigma^m$
      Else Pass
   b) For each layer, l
      i. Choose $A^{m,l} \sim \mathcal{H}\text{N}(0, 1)$
      ii. Choose $\beta_0^{m,l} \sim \text{N}(0, 1)$
      iii. **If** Hierarchical
         Then Choose $\beta_1^{m,l} \sim \text{N}(\beta_1^m, \Sigma_m)$
         Else Choose $\beta_1^{m,l} \sim \text{N}(0, 1)$

FIG 4. *Plate diagram depicting a model in which parameters are inferred independently in each measurement and each layer (A), and model with hierarchical priors pooling information across layers in each measurement (B). In both models, pathological measurements are depicted by $X_d^{m,l}$, where $m = 1 \ldots M$ index each of the measurements and $l = 1 \ldots L$ each layer. Latent pseudotime variable is denoted by t_d, where $d = 1 \ldots D$ index the brain donor. Measurement and layer-specific initial-seeding, growth-rate, and total-counts parameters are denoted by $\beta_0^{m,l}$, $\beta_1^{m,l}$, and $A^{m,l}$. In the case of the hierarchical model in (B), we pool information across layers and infer covariance matrices Σ_m, assuming that each layer has its own initial seeding events but their growth rate are similar.*

3.3. *Posterior Bayesian Inference through HMC sampling.* To perform posterior inference from our models we used numpyro version 0.13.0 (Bingham et al. (2019); Phan, Pradhan and Jankowiak (2019)), a Python based probabilistic programming environment. Similar to the more traditional probabilistic language, Stan (Carpenter et al. (2017)), numpyro is equipped with effective posterior sampling algorithms such as Hamiltonian Monte Carlo (HMC, Duane et al. (1987); Neal et al. (2011)). In particular, we utilized the implementation of the No-U-Turn Sampler (NUTS) (Hoffman et al. (2014)) kernel to obtain posterior

samples, and monitored convergence using R-hat and effective sample size metrics (Gelman et al. (2021)).

To infer posterior parameters of our models (Fig. 4) using the SEA-AD dataset, we sampled from the posterior with 5 HMC chains run sequentially for 5000 warm-up iterations, and then 1000 samples were collected in each chain. We randomly rotate the seed for the pseudo-random number generator in each chain. To ensure convergence, we discarded the chains that had an R-hat below 1.05 or had an effective sample size below 50. The results from these fits are described in Section 6. For experiments with synthetic data, we used different parameters to the ones described above (see Supplement Section 2 for details).

4. Model analysis. Our problem departs from the classical regression setup in that we not only estimate the regression parameters β but also must infer the nuisance parameter t. Here, we delve into the inferential implications of the joint estimation of t and β; we study the fundamental statistical limitations to our problem and argue that our modeling choices are well-justified in that they enable efficient use of all available sources of information to counter these limitations.

Results in this section are statements regarding conditions for identifiability and the best possible convergence rates for the mean squared error (MSE) of estimators of β and t. Although we keep our Bayesian estimators in mind, this analysis is rather concerned with the properties of the problem itself than with our particular method. We will later show in Section 5 that our Bayesian estimators (understood as posterior means) exhibit the behavior of the best possible estimators.

Instead of the generalized linear model in Eq. (2), here we consider the simpler linear model where a rich theoretical understanding of statistical limits is readily available:

$$(4) \qquad X_d^n = \beta_0^n + \beta_1^n t_d + \epsilon_d^n.$$

We understand X as a $D \times N$ matrix by vectorizing layer and feature information and treating them as a unique column 'feature' indexed by $n = 1, \ldots, N$ where $N = L \times M$. Here, ϵ_d^n are i.i.d Gaussian centered variables with variance ϕ^2. These noise variances play an analog role to the overdispersion parameter A in Eq. (2).

The key observation is that Eq. (4) corresponds to rank-two noisy factorization problem, so that at the most elementary level, the estimation limits of β and t are the limits of this low-rank factorization. Although the theory is general enough to enable the analysis of (4), for simplicity will neglect the β_0^n term and study the rank-one noise matrix factorization

$$(5) \qquad X = \beta t^\top + \epsilon.$$

To apply existing statistical results, we divide the estimation of t and β into the estimation of the normalized parameters $\beta' = \beta/||\beta||_2$, $t' = t/||t||_2$ and the individual signal strengths $||t||_2^2$ and $||\beta||^2$. To state consistency results, we conceive limiting experiments consisting of sequences of true $\{t\}_d$, $\{\beta\}_n$ (for example, obtained by sampling from their population distributions), and of matrices X and ϵ that are revealed as either D or N (or both) grow. For our asymptotic statements on D or N we must assume that $\bar{t}_\infty^2 := \lim_D \bar{t}_D^2$ and $\bar{\beta}_\infty^2 := \lim_N \bar{\beta}_N^2$ exist, where $\bar{t}_D^2 := ||t||_2^2/D$ and $\bar{\beta}_N^2 := ||\beta||_2^2/N$, respectively.

4.1. *Subspace estimation.* We first study the ability to estimate the normalized parameters t' and β'. We use the standard Frobenius metric for comparing subspaces U_1 and U_2 (see Appendix for details), which in the one-dimensional case reduces to

$$D_F(U_1, U_2) := \min\{||U_1 - U_2||_2, ||U_1 + U_2||_2\}.$$

To obtain estimation bounds, we turn to specialized, best-known risk analysis for subspace estimation Cai and Zhang (2018) based on variants of the well-known Wedin's $\sin \Theta$ perturbation theorem Wedin (1972). Let $\hat{\beta}$ and $\hat{t}$ be the first left and right singular vectors of X (in particular, they have unit norm).

PROPOSITION 1. *As a consequence of Theorems 3 and 4 in Cai and Zhang (2018), we have*

$$\frac{1}{N} ED_F(\hat{\beta}, \beta')^2 \lesssim \frac{\phi^2}{ND} \frac{\left(\overline{\beta}_N^2 \overline{t}_D^2 + \frac{\phi^2}{N}\right)}{\left(\overline{\beta}_N^2 \overline{t}_D^2\right)^2} \wedge 1$$

$$\frac{1}{D} ED_F(\hat{t}, t')^2 \lesssim \frac{\phi^2}{ND} \frac{\left(\overline{\beta}_N^2 \overline{t}_D^2 + \frac{\phi^2}{D}\right)}{\left(\overline{\beta}_N^2 \overline{t}_D^2\right)^2} \wedge 1$$

These estimators are minimax in the sense that any other estimators of β' and t' have risk lower bounded by the right sides above (up to a constant)

We can use Proposition 1 to understand how estimation performance for β' and t' improve as we increase either D or N (or both). We observe that for fixed N, the average MSE in the estimation of β' decreases at the rate D if $\overline{t}_D^2$ is assumed to remain bounded away from zero (we are excluding sparse non-informative sequences of donors), which is the same rate that would be achieved if D where known. Therefore, adding more donors is helpful in the estimation of β'. Likewise, if D is fixed, then the average MSE over t' decreases at the rate $1/N$ if $\overline{\beta}_N^2$ (i.e. we are not adding too many non-informative features), implying that adding more features help. Also, larger magnitudes $\overline{\beta}_N^2$ are beneficial for the estimation of t' and larger magnitudes of $\overline{t}_D^2$ are beneficial for the estimation of β'.

4.2. *Signal strength estimation.* Proposition 1 suggests that the spaces spanned by t' and β' can be estimated as easily as in the usual regression case (t known). We now turn to the problem of estimating the signal strengths $||\beta||$ and $||t||$. The ability to estimate these parameters consistently is critical since otherwise, we are at risk of inflating the significance of our estimators of β: for example, if $||\beta||$ is grossly overestimated, we may come to falsely detect many β_n as significant features explaining the progression of the disease.

Estimating the signal strengths is more delicate since identifiability issues might be at play. The main result is that by imposing the constrain that times belong to $[0, 1]$ (or any other compact interval defined beforehand), we can identify $||\beta||$ and $||t||$ if both N and D simultaneously grow to infinity.

PROPOSITION 2. *Suppose that t_d are sampled from a distribution $\mathbf{t}$ supported on $[0, 1]$. Suppose that both t_∞ and β_∞ exist. Then, there are estimators of t_D and β_N converging almost surely to t_∞ and β_∞ as N and D grow to infinity.*

The mathematical details are presented in the Appendix. At a very high level, the proof consists in first showing that we can first identify the product $\beta_\infty \times t_\infty$. Then, we can use the information on the support of $\mathbf{t}$ to re-scale the principal vector $\hat{t}$ so that it will lie in $[0, 1]$. This re-scaling enables the identifiability of each of the terms β_∞ and t_∞. The statement in

Proposition (2) holds regardless of whether or not ϕ^2 is known: in the latter case, we can estimate it incurring in an asymptotically negligible bias.

Proposition 2 gives an optimistic message that the entire system can be identified, but the main downside is that we are forced to consider simultaneous asymptotics in N and D. Unfortunately, as we will show, this heavy requirement cannot be relaxed, and the estimation of $||\beta||$ will be biased. In other words, the estimation of $||\beta||$ will be biased even if a cohort of infinite donors were available. We will comment on strategies to overcome this limitation. Even if these biases persist, the overall message of Proposition 2 is that if multiple informative features β are available, they should all be used to inform inferences.

5. Simulation studies.   The previous results provide us with tools to reason about inferential capabilities in our model. In this section, we complement these results with a comprehensive set of simulations. These simulations not only give an empirical validation to theory in Section 4; more importantly, we use them to understand our procedure's behavior in cases not covered by the theory.

This gap arises in three aspects: first, results in Section 4 are concerned with the estimation problem and not with our particular estimator. Second, they concern a simplified model that doesn't incorporate the effect of a nonlinear link function and a non-Gaussian likelihood, and where coefficients β are not assumed to possess a hierarchical structure that we could possibly benefit from. From the experiments in this section, we are able to make the following claims

- Our Bayesian estimator exhibits the behavior anticipated by Propositions 1 and 2 in that simplified setup (Sections 5.1, 5.3)
- The use of suitable priors on t can be helpful to overcome the identifiability issues described in Section 4 (Section 5.2)
- Beyond the simplified linear model, the qualitative behavior of estimators still matches the phenomena described in the previous section, although some additional new phenomena emerge:
  - In non-linear/non-Gaussian models, hierarchical models for β are beneficial if such structure is present in the data. We don't observe a benefit from using hierarchical models in linear Gaussian models (Section 5.4)
  - Non-linearities in the link function imply that different features can be differentially informative at different stages of the progression. This implies the need for careful choice of features and donor experimental design (Section 5.5)

Here, we only describe the main features of each experiment but defer a more comprehensive explanation of each setup and parameter and sampler values to the Appendix.

5.1. *Elementary simulations in the linear model.*   We start by studying the empirical performance of our Bayesian estimators in the setup of the linear model (4) with a single feature ($M = 1$). Unlike the full hierarchical model of Section 3, we consider a flat prior over β and a uniform distribution in $[0, 1]$ as prior on t_d and assume that true times were selected from a uniform grid in $[0, 1]$. We study estimators of β comparing against the regression estimator that would be obtained if the times t_d were assumed to be known. If that was the case, $\beta = (\beta_0, \beta_1)$ could inferred via least squares with closed-form expressions.

Results are shown in Fig. 5; the need to infer times implies a sizable increase in the MSE of the posterior mean of β over the baseline where times are known. As expected, the MSE depends on the noise variance parameter ϕ^2 (assumed known). Although uncertainty quantification is not the main preoccupation of our work, we show that similar behavior occurs when we compare the posterior variance of our estimator with the posterior variance of the usual linear regression estimator (Fig. 5B).

FIG 5. **Elementary convergence study for the linear model (Eq. (4)) with a single ($N = 1$) feature**. *Left: MSE of $\beta = (\beta_0, \beta_1)$ for different noise variances ϕ^2 (colored lines) and number of sampled times (D, x-axis). Dashed lines correspond to the baseline where t_d's are known (i.e., solving a linear regression). Right: sum of posterior variances of β_0 and β_1.*

5.2. *Consequences of inconsistency: on the choice of time priors.* In light of Proposition 2, we expect that estimators of β will be biased even if D is arbitrarily large. This is consistent with Fig. 5, where the MSE of β eventually reaches a plateau representing the asymptotic (in D) error. We explore this phenomenon further by considering alternative choices of true times and time priors $p(t)$. Specifically, in addition to true times being sampled from an equispaced grid, we also consider the case where they are sampled from a $Beta(0.1, 0.1)$ distribution and from a "cubic" distribution; i.e., where t is proportional to t^3 on $[0, 1]$. Additionally, we consider two priors: a uniform distribution (as in the previous experiment) and the same $Beta(0.1, 0.1)$. Results are shown in Fig. 6: MSE of β is the closest to the regression baseline if the prior matches the true distribution of data. However, if there is a disagreement, the MSE can increase as D grows.

One pessimistic interpretation of these experiments is that the prior has a too sizable influence on the outcomes even for large donor sample size T. While this is true, at the same time, there is no way to get around this problem because of the inherent lack of identifiability when $M < \infty$. More optimistically, these results show that it is possible to encode our understanding of how donor sampling occurred if that knowledge is available.

5.3. *Using multiple features.* We now discuss the consequences of using multiple features, i.e., $M > 1$. In Fig. 7, we show the results for the case where $\beta_0^n = 0$ and $\beta_1^n = 0$; i.e., all features have the same coefficients β as in the $M = 1$ case. As we increase more features, the average MSE for each feature decreases, converging to the linear regression (known t_d) baseline. For many features (e.g., $M = 100$), divergences from this baseline are observed only for a large number of donors, where MSEs are already low. The main conclusion is that the addition of multiple features is beneficial for the estimation of each individual feature. This phenomenon can be understood in the light of our theoretical results: By virtue of Proposition (1), an increased number of features (with the same β) adding more features will lead to a more precise estimation of the (normalized) vector of times t. These more precise estimates will, in turn, lead to estimates $\hat{\beta}$ that are closer to the ones that would be obtained if times were known.

FIG 6. **Effect of choices of time priors in the linear model (Eq.** (4) **with a single feature (**$M = 1$**) on estimation error (MSE) of** $\beta = (\beta_0, \beta_1)$ *Priors* $p(t) \sim Uniform(0, 1)$ *and* $p(t) \sim Beta(0.1, 0.1)$ *are shown on different rows and population distribution of times* ***t*** *on different columns. Colors indicate noise levels* ϕ^2. *Dashed lines correspond to the baseline where* t_d*'s are known.*

In the previous experiment, we assumed that all features were equal. We also show that the same phenomenon still holds if features are not identical but are instead sampled from a distribution. To this end, we consider the scenario where $\beta_0^n = 0$ and β_1^n are sampled from a Gaussian distribution with mean $\beta_1 = 1$ and variance $\sigma^2 \in \{0, 10\}$. Results in Fig. 8 show that MSE for β is at a comparable scale independent of whether β_1^n is random or not ($\sigma^2 = 0$ v.s. $\sigma^2 = 10$).

However, unlike in the previous experiment, the added randomness in β_1^n is beneficial for time estimation as we observe that time estimates in the $\sigma^2 = 10$ case are consistently closer to true values. This may seem paradoxical as the $\sigma^2 = 10$ corresponds to a noisier regime, and we may reasonably expect noisier results. Again, we can understand this phenomenon in the light of Proposition 1: larger values of σ^2 translate into larger values of $||\beta||^2$, which in turn lead to a lower upper bound on the discrepancy between true and estimated times, up to a normalization constant.

Although we typically treat times as a nuisance parameter, if the pseudotime inference is a primary inferential goal, the above experiment highlights that diversity of features (i.e., larger variance in β) will enhance our capabilities to estimate times. We will come back to this point when we discuss non-linearities.

5.4. *Shrinkage in Gaussian and non-Gaussian models with hierarchical priors.* While previous experiments address the multiple features case and the effect of variability in feature coefficients β on our estimates, none of the Bayesian models in those experiments explicitly model the sampling structure in coefficients. One natural question is whether the James-Stein phenomenon would manifest here, i.e., whether the MSE can be lowered using shrinkage estimators. These estimators, in turn, are obtained by imposing a suitable hierarchical structure

FIG 7. *Effect of numbers of features N (colors) in the linear model (Eq. (4)) on estimation error (MSE) of $\beta = (\beta_0, \beta_1)$. Different rows correspond to different observation noise levels ϕ^2, and columns indicate the true distribution of times $\boldsymbol{t}$. Dashed lines correspond to the baseline where t_d's are known.*

on β. This question is not obvious since it depends on the interplay of β and the "nuisance" parameter t.

As in the previous experiment, we assume that β are sampled i.i.d from a superpopulation; i.e. $\beta_0^n \sim \mathcal{N}(\beta_0, \sigma^2)$, $\beta_1^n \sim \mathcal{N}(\beta_1, \sigma^2)$. We consider Bayesian estimators based on four models: i) the one from previous experiments; i.e. a fully factorized independence prior on β (Fig 4A), and three hierarchical models (similar, but simpler than Fig 4B) where each β_0^n, β_1^n are sampled from a Gaussian distribution centered at β_0, β_1) and model variance σ_{model}^2 where ii) $\sigma_{model}^2 = 0 \neq \sigma^2$, iii) $\sigma_{model}^2 = \sigma^2$ iv) σ_{model}^2 is a parameter for which we perform Bayesian inference from a half-normal hyperprior. As shown in the first row of Fig. 9, our estimators produce lower MSE (average over different t_d) for larger values of M and lower noise variance ϕ^2. Additionally, as shown in the second row of Fig. 9A, increasing the number of time samples has little impact on the average MSE.

Fig. 9A gives the seemingly disappointing message that it is impossible to profit from the shrinkage provided by the hierarchical model since the MSE of the fully factorized model is indistinguishable from the other hierarchical estimators. However, as shown in Fig. 10, this is only feature of the linear model (Eq. (4)): when we depart from this assumption and consider the Negative Binomial regression (Eq. (2)), there is a clear advantage of the hierarchical model. Nonetheless, this advantage again vanishes as the time samples D increases. The overall message is that the estimation of t benefits from shrinkage on β in non-linear models if the number of features is large compared to the number of time samples.

5.5. *Combining features with stage-specific information in non-linear models.* One main feature distinguishing linear and non-linear models is that the signal-to-noise ratio SNR depends nonlinearly on times, implying that individual features can be preferentially informative about particular stages of the disease progression. Indeed, if we define (see Hastie et al. (2009))

$$SNR = \frac{g^{-1}\left(Var(\rho)\right)}{Var(X|\rho)},$$

FIG 8. *Effect of variability in β on estimation σ^2 (colors) in linear model (Eq. (4)) with $N = 10$ features on parameter and pseudotime estimation (A) MSE of $\beta = (\beta_0, \beta_1)$. Dashed lines correspond to the baseline where the t_d's are known). (B) comparison of time estimates (averages over experiments)*

where l is the linear response ($\rho = \beta_0^n + \beta_1^n t^d$ in Eq. (2)) and g is the link function (logarithmic in Eq. (2)), we observe that if $\beta_1^n > 0$, the SNR can be much larger at large times compared to small times, with an opposite behavior if $\beta_1^n < 0$.

We illustrate this phenomenon with an example of a Negative Binomial model with two features and an overdispersion parameter $A = 100$. The first feature has parameters $\beta_0^1 = 4, \beta_1^1 = -4$, so the mean counts $E(X_t^1)$ rapidly decrease at earlier times and eventually stabilize (see the Appendix for examples of observations). The second feature has $\beta_0^2 = -4$ and $\beta_1^2 \in \{2, 3, 4, 5, 6\}$ so the behavior is inverted, and the overall strength of this feature can be smaller ($\beta_1^2 \in \{2, 3\}$), equal ($\beta_1^2 = 4$) or larger ($\beta_1^2 \in \{5, 6\}$) than the first feature.

We compare estimators based on each individual feature with the one based on both. Results are shown in Fig. 11. We observe that including both features leads to a lower MSE (on t and β) than any individual feature (Fig 11A), suggesting that that the cross-talk of information between different features is beneficial for the estimation of each of them as they are uniquely informative about different phases of the diseases's progression.

We can better understand this phenomenon by looking at time sequences inferred by each model. As shown in Fig. 11B, the model that only uses the first feature accurately infers earlier times but struggles in later-time inference, and the converse is true for the model that only uses the second feature. In contrast, the model that combines both accurately infers times along the entire progression. Fig. 11B also shows that this consensus will be achieved even if there is asymmetry (one feature stronger than the other) as long as there is enough signal in the system (i.e., if both β_1^1, β_1^2 are sufficiently large). These observations are also supported by Fig. 11C, where we report MSE on time inference by classifying each time sample as being on an 'early' ($t \leq 0.5$) or 'late' ($t \geq 0.5$) stage. Models with a single feature

FIG 9. ***Estimation error (MSE of times) of different hierarchical models for Eq.*** (4). *We considered different values of N (x-axis), number of time samples D, observation noise variances ϕ^2, and noise variance σ^2. Increasing the number of features decreases estimation error; however, the decrease doesn't depend on the chosen model*

FIG 10. ***Estimation error (MSE of times) of different hierarchical models for the linear-Gaussian model (Eq.*** (4)***) and the Negative Binomial model (Eq.*** (2)***)*** *more layers help and a hierarchical model is helpful for time estimation in the negative binomial case*

have the lowest MSE at the stages that the feature is informative about, and the MSE of models containing the second feature decreases with the strength of this feature in all stages.

5.6. *Additional experiments.* In the Appendix, we present more experiments showing that the phenomena described in linear models (Sections 5.1,5.2, 5.3, 4) also hold if we

use Negative Binominal likelihood (Eq. (2)). Additionally, we present supplementary experiments in for Sections 5.1,5.2, 5.3, Eq. (4) comparing the minimax optimal SVD estimator instead with our Bayesian estimator, showing they are typically close to each other in the linear case.

FIG 11. *Inferring multiple features in the Negative Binomial non-linear model (Eq.* (2)*) (A): Average MSE for times, and MSE for the strength of the first and second features (β_1^1 and β_1^2) as a function of the strength of the second feature β_1^2 in models that only use individual features as well as the one that uses both (different colors). (B) Comparison of true and inferred times (average over experiments) for different strengths of the second feature β_1^2. (C) Average MSE of inferred times for the early and late stages as well as their combination, for different strengths of the second feature β_1^2.*

6. Inferring latent pseudotime in the SEA-AD cohort. Finally, we apply our generative models (Fig. 4) to the problem of describing the progression of pathology aggregation in AD, and infer latent pseudotime for each donor in the SEA-AD cohort. As a reminder, our dataset consists of $D = 84$ donors in which $M = 12$ pathological markers are measured in each of $L = 5$ cortical layers of the MTG (as detailed in Section 2). After sampling from the models' posteriors (see Section 3.3), posterior pseudotime estimates, t_d, span the entire progression of disease (Fig. 12A), with an increased representations of later times, in accordance with the cohort distribution that is biased towards older ages and higher disease metrics (such as Braak, or ADNC, see Fig. 1B,C). Growth-rate parameter β_1 and the parameter quantifying the initial seed levels of pathology β_0, revealed that canonical AD protein markers follow an increasing exponential dynamics from small values, while, as expected, a marker quantifying the number of neuronal cells per unit area decreases as AD progresses (Fig. 12B,C).

Consistent with our predictions from simulation studies, we observed that for $T = 84$ donors, both the hierarchical model and the fully-factorized model result in similar posterior distributions of model parameters. This is due to the size of the dataset in which sufficient number of donors compensate for the extra information provided by the layer and measurement covariance incorporated by the hierarchical model. Quantitative metrics comparing both models include the expected log posterior predictive distribution (Vehtari, Gelman and Gabry, 2017) and the Watanabe Information Criteria (WAIC, Watanabe and Opper (2010)). These metrics indicate that the hierarchical model better fit the data but the gain at this data regime is marginal (Supplementary Figures 2 and 3).

FIG 12. *Posterior distributions of model parameters for SEAD-AD data set of pathological proteins. (A) Inferred pseudotimes t_d for donors sorted in ascending order. B) posterior distribution of β_1, growth-rate and (C) β_0, initial seeding model parameters. In each case, error bars denote distribution of posterior means inferred for each layer. (D) Posterior means (solid lines) plotted in conjunction with the data for 4 exemplary pathological markers. Color-codes indicate observations for each cortical layer.*

6.1. *Exploring the importance of pathological markers as model regressors.* We next investigated the contribution of each feature to our latent pseudotime. As our problem does not map to a regression problem, in which variable importance has been extensively studied, we resorted to multiple metrics to quantify each pathological marker contribution. To compare the effect of each pathological marker as an explanatory variable, making an analogy to

regression problems, we first look at the growth-rate parameter β_1. Parameter values differ considerably across all $M = 12$ features, with six having the highest absolute value (Fig. 13A). However, in our case, parameter magnitude could be dependent on the signal-to-noise ratio in each feature and be loosely related to pseudotime inference. To find a different way of quantifying associations between pseudotime and each pathological measurement, we resorted to the calculation of the mutual information (MI), understanding MI as a means to identify feature importance in nonlinear models (Beraha et al., 2019). MI identifies the 7 features as the most relevant ones to inform pseudotime inference (Fig. 13B).

As an additional feature importance computational experiment, we calculated pseudotime posterior estimates for models using each pathological feature in isolation. Then, we calculated the error in the posterior pseudotime parameter t_d inferred from only one feature compared to all features, ($|t_{d,M=12} - t_{d,M_i}|$ (Fig. 13 C) and the standard deviation in the posterior distribution of each pseudotime parameter t_d, to reveal posterior uncertainty on pseudotime estimation (Fig. 13D). Our previous simulation studies informed us that features with high growth-rate parameter β_1 would have a high-dynamic range of pathology and observations with high pathological values. This in turn is more informative for obtaining identifiable estimates of t_d. Our computational results are in accordance with our simulation studies, revealing that the first seven pathological variables (order according to decreasing MI) have small pseudotime error and variance (Fig. 13C-D). Taken together all these computational experiments reveal that the seven most informative variables are a more compact subset able to create reliable pseudotime posterior estimates.

FIG 13. *Multiple metrics describing the importance of each pathological feature to inform latent pseudotime inference. (A) Growth-rate parameter, β_1, for each pathological feature, ordered according to decreasing mutual information (MI). (B) MI between each pathological feature and posterior pseudotime. (C) Error in pseudotime inference by comparing posterior times inferred using one feature vs all $M = 12$ features. (D) The standard deviation of posterior estimate, t_d, for each single-feature model.*

Next, we evaluated our model through correlations of our posterior pseudotime estimates with established neuropathological assessments to independently validate our pseudotime estimates with orthogonal metadata not used to fit the model. We notice that the pseudo-progression parameter is highly correlated with neuropathological assessments Braak, Thal, and ADNC (positive correlation) and cognitive scores CASI (negative correlation Teng et al. (1994b)) (Fig. 14, despite the metrics using different pathological proteins to quantify progression (see Section 2).

FIG 14. *Inferred pseudoprogression correlates with existing metrics for pathology progression. (A) Spearman correlation between the associated progression metric and inferred pseudoprogression score t_d, as the number of measurements to infer t_d increase . Both models (fully factorized and hierarchical) show high correlations with these metrics. However only a few measurements are required to maximize the correlations. (B) All relevant metrics increase monotonically with the inferred pseudoprogression score inferred from our model.*

6.2. *Informing Future Experimental Design.* This section is driven by the aim of leveraging our models to guide the design of future experiments. While this aim could be framed as a Bayesian optimization problem, we opt for a qualitative approach guided by several principles. Firstly, the process of collecting, processing, and analyzing neuropathological measurements, as presented here, is labor-intensive. Our objective is to utilize the insights gained from the preceding section to identify a concise set of highly informative variables. Although reducing the feature set could be in detriment of the quality of fit (as described in our Simulated Studies Section 5), this will streamline the time required for future experiments. Secondly, the selection criteria for brain donors in the SEA-AD cohort resulted in a bias towards older donors and those at advanced pathological stages. Most donors are of old age, mean age = 88 years old, and exhibit an inferred pseudotime greater than 0.5. In the future, we would like to diversify the SEA-AD cohort by including donors that uniformly span the entire pseudotime spectrum (from 0 to 1). A balanced cohort, representative of all disease stages, facilitates downstream analysis of multiple disease epochs. Moreover, pseudotime being an inferred parameter poses challenges in its use for donor preselection and cohort balancing. To address this, we aim to enhance the model by incorporating other observable variables such as donor metadata or staging information. This approach will enable us to strategically sample donors and mitigate biases.

To achieve these goals, we investigate a model incorporating $M = 7$ pathological features, chosen for their high standardized beta coefficient, mutual information, and low error and standard deviation in time inference (see Fig. 12). These features constitute a minimal set as desired. To suggest future brain donors based on easily measurable variables, we augment the

model with categorical staging metrics (Braak, Thal, and ADNC). Let m represent a staging metric, d denote a donor, and k signify a one-hot encoded stage level then, the vector $s_k^{d,m}$ represents:

$$
(6) \qquad s_k^{d,m} = \begin{cases} 1, \text{if } k = \text{donor's stage} \\ 0, \text{if } k \neq \text{donor's stage} \end{cases}
$$

For example, Thal phases ranges from phase 0 to 5, in increasing phases indicating increasing number of brain regions containing amyloid-beta pathological burden. A donor exhibiting phase-2 Thal phase would be represented as $s^{d,Thal} = [0, 0, 1, 0, 0, 0]$. We incorporate staging metrics into our formalism and model them as:

$$
p_k^{d,m} = \alpha_{0,m}^k + \alpha_{1,m}^k t_d
$$

$$
(7) \qquad s^{d,m} = \text{DirichletMultinomial}(c^{d,m} = e^{p^{d,m}}, \text{count} = 1)
$$

Assuming that pseudotime $t_d = 0$ aligns with all staging metrics being zero and that their values increase thereafter, we assign prior parameters α_0 and α_1 for each stage k.

$$
\alpha_0^k \sim \mathcal{N}(-10k, 5)
$$

$$
(8) \qquad \alpha_1^k \sim \mathcal{N}(2k - 10, 5)
$$

FIG 15. ***Model augmented with staging metrics is able to effectively represent experimental data.*** *(A) Comparison of posterior time estimates predicted by models spanning all features and features plus staging metrics. (B) Posterior β_0 and β_1 distributions (concatenating values from all layers and samples) in both models (C) Posterior predictive estimates reveal effective modeling of staging metric (Braak, Thal and ADNC).*

Next, we perform Bayesian posterior inference and compare posterior estimates of the augmented model (Augmented, $M = 7$ + Staging metrics) against our model that includes all features (All Feats, $M = 12$) in Fig. 15 and Supplementary Fig. 8. Both models resulted in similar posterior estimates, revealed by comparing posterior pseudotime estimates for each donor and dynamic parameters (β_0, β_1) across measurements common to both settings (Fig. 15A-B). In both cases, neuropathological data used to infer posterior parameters results in concordant estimates and relatively small posterior standard deviations. We observed that staging information does not significantly impact pseudotime posterior estimates, both estimates concentrating around the same mean, when most informative neuropathological features are selected. Although speculative, this may be indicative of the fact that the staging information may be too coarse to provide any additional information on disease progression on top of the one already provided by neuropathological markers. At the same time, the augmented model is able to effectively model staging metrics, observed when comparing staging metrics' posterior predictive distributions and observed data (Fig. 15C). Taken together, these results highlights the ability of our augmented model together with the set of informative features to accurately model pseudotime and AD staging metrics in concert.

FIG 16. *Augmented Model incorporating staging information informs experimental design to balance SEA-AD cohort across disease stages (A) To balance SEA-AD cohort, we used kernel density estimates to describe the distribution of inferred pseudotime values. Next, we empirically calculate a pdf corresponding to the least probable time values. $N = 1000$ samples are drawn from this distribution and the resulting histogram is displayed (B) Pseudotime samples from (A) are then converted to the associated Braak, Thal and ADNC stage. Top plot shows the distribution of stages in the donor dataset, bottom plot shows the distribution of stages corresponding to the resampled times. (C) Heatmaps showing the Braak/ADNC and Thal/ADNC joint histograms. The associated distribution of stages can aid experimentalist to improved donor sampling to balance their cohort.*

Finally, we utilize our augmented model to guide experimental design, aiming to achieve balance (equal representation of donors across pseudotime) in the SEA-AD cohort. To

accomplish this, we first estimate the distribution of posterior pseudotimes, $p(t)$, by 1) aggregating mean posterior pseudotime estimates for each donor, 2) constructing a histogram, and then 3) applying kernel density estimation. Next, we approximate the unnormalized distribution of under-sampled pseudotimes in the SEA-AD cohort by computing $p_{under}(t) = 1 - p(t)$. This approximation enables us to generate samples of under-sampled pseudotimes ($n_{samples} = 1000$, see Fig. 16 A). These values can then be employed to compute posterior predictive staging metrics based on the posterior estimates of staging parameters (see Fig. 16B-C). These predictions inform the staging metrics of suggested new donors that should be included in the study to achieve cohort balance. Given the current distribution of our cohort, it is apparent that sampling from early and middle stages would be advantageous.

7. Discussion. We introduced a Bayesian framework to model Alzheimer's Disease progression, incorporating a latent pseudotime variable to depict disease evolution. We modeled observations utilizing biophysically-inspired assumptions that posit pathological aggregation follows exponential dynamics. Our inference method is useful for datasets restricted to cross-sectional data, consolidating observations and inferring a single and common disease trajectory. We demonstrated the utility of our framework using simulation studies and real datasets originating from the SEA-AD consortium, profiling AD neuropathological proteins in a brain region. Applying our model to the SEA-AD cohort dataset, consisting of 84 donors, allowed us to infer donor pseudotime and pathological protein dynamics. Pseudotime estimates are biased towards late values, in accordance with our observations that the SEA-AD cohort is biased towards older donors and high pathology. Posterior predictive checks validate the goodness of our fits. In addition, comparisons with existing neuropathological staging-based metrics reveal that pseudotime ordering aligns well with brain-wide neuropathological assessments, indicating that our model is able to map AD trajectory when detailed information from only one brain region is used (the Middle Temporal Gyrus).

Next, we assembled a collection of metrics to assess the impact of each pathological protein on pseudotime inference, identifying the most informative features. These features exhibited high growth rates, mutual information between their values and pseudotime, and individually exhibited precise predictions of pseudotime with low prediction error. Incorporating these highly informative variables alongside neuropathological assessments, we enhanced our model to deduce the latter. When integrated with staging data, our enhanced model can guide the identification of neuropathological stages for future donors to be included in the existing SEA-AD cohort, ensuring a balanced representation of donors across all pseudotimes and enhancing downstream analysis across the disease spectrum. While our work assumes that the pathology markers align donors along the same pseudotime trajectory, which is true for the MTG region of the brain, this assumption may not hold true for other brain regions. This would require approaches that can overcome low-data regimes to successfully leverage all regions in a combined overarching multi-regional pseudotime.

Our results demonstrated the success of our methodology, supported by our theoretical explorations. In our work, we derived theoretical results on a simplified linear model to understand the best possible performance. A stimulating recent line of work has provided algorithmic (Kidzinski et al., 2022) and statistical (Chen, Li and Zhang, 2020; Du, Wasserman and Roeder, 2023) guarantees for matrix factorization problems in generalized-linear setups, a setup akin to ours. However, these results by themselves are not sufficient to explain the richness of phenomena manifesting in our setup. Future studies should study the simplified model and extend Proposition 2 by deriving explicit rate bounds for the signal strength estimates, and also extend our theoretical results to a generalized linear model case. We made contributions in this direction through extensive simulations and showed that whether or not

we benefit from hierarchical models for pseudotime estimation depends on the type of model and error distribution, and benefits are only observed beyond linear models with Gaussian errors. Future studies should develop a theoretical framework to explain this phenomenon.

Acknowledgments. The authors would like to thank Sivaraman Balakrishnan for useful theoretical discussions, Bob Carpenter for insightful discussion on modeling and inference approaches, and the entire SEA-AD team for creating an open source resource for the study of Alzheimer's Disease.

Funding. A.A. was supported by the Shanahan Foundation Fellowship. M.I.G. was supported by National Institutes on Aging (NIA) Grant U19AG060909 and R01102-01-071-20. This work used Bridges-2 at Pittsburgh Supercomputing Center through allocation MTH230027 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services and Support (ACCESS) program, which is supported by NSF Grant 2138259, 2138286, 2138307, 2137603, and 213829.

REFERENCES

BERAHA, M., METELLI, A. M., PAPINI, M., TIRINZONI, A. and RESTELLI, M. (2019). Feature Selection via Mutual Information: New Theoretical Insights. In *2019 International Joint Conference on Neural Networks (IJCNN)* 1-9. https://doi.org/10.1109/IJCNN.2019.8852410

BINGHAM, E., CHEN, J. P., JANKOWIAK, M., OBERMEYER, F., PRADHAN, N., KARALETSOS, T., SINGH, R., SZERLIP, P. A., HORSFALL, P. and GOODMAN, N. D. (2019). Pyro: Deep Universal Probabilistic Programming. *J. Mach. Learn. Res.* **20** 28:1–28:6.

BRAAK, H. and BRAAK, E. (1991). Neuropathological stageing of Alzheimer-related changes. *Acta neuropathologica* **82** 239–259.

CAI, T. T. and ZHANG, A. (2018). Rate-Optimal Perturbation Bounds for Singular Subspaces with Applications to High-Dimensional Statistics. *The Annals of Statistics* **46** 60–89.

CAMPBELL, K. R. and YAU, C. (2018). Uncovering pseudotemporal trajectories with covariates from single cell and bulk expression data. *Nature communications* **9** 2442.

CARPENTER, B., GELMAN, A., HOFFMAN, M. D., LEE, D., GOODRICH, B., BETANCOURT, M., BRUBAKER, M. A., GUO, J., LI, P. and RIDDELL, A. (2017). Stan: A probabilistic programming language. *Journal of statistical software* **76**.

CHEN, Y., LI, X. and ZHANG, S. (2020). Structured latent factor analysis for large-scale data: Identifiability, estimability, and their implications. *Journal of the American Statistical Association* **115** 1756–1770.

DAVIS, J. K. and SINDI, S. S. (2016). A mathematical model of the dynamics of prion aggregates with chaperone-mediated fragmentation. *Journal of mathematical biology* **72** 1555–1578.

DU, J.-H., WASSERMAN, L. and ROEDER, K. (2023). Simultaneous inference for generalized linear models with unmeasured confounders. *arXiv preprint arXiv:2309.07261*.

DUANE, S., KENNEDY, A. D., PENDLETON, B. J. and ROWETH, D. (1987). Hybrid monte carlo. *Physics letters B* **195** 216–222.

GABITTO, M. I., TRAVAGLINI, K. J., RACHLEFF, V. M., KAPLAN, E. S., LONG, B., ARIZA, J., DING, Y., MAHONEY, J. T., DEE, N., GOLDY, J. et al. (2023). Integrated multimodal cell atlas of Alzheimer's disease. *Research Square*.

GELMAN, A. and HILL, J. (2006). *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press.

GELMAN, A., HILL, J. and VEHTARI, A. (2021). *Regression and other stories*. Cambridge University Press.

GELMAN, A., CARLIN, J. B., STERN, H. S., DUNSON, D. B., VEHTARI, A. and RUBIN, D. B. (2021). Bayesian Data Analysis Third edition (with errors fixed as of 6 April 2021). *Issue: April*.

GUPTA, A. and BAR-JOSEPH, Z. (2008). Extracting dynamics from static cancer expression data. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* **5** 172–182.

HASTIE, T., TIBSHIRANI, R., FRIEDMAN, J. H. and FRIEDMAN, J. H. (2009). *The elements of statistical learning: data mining, inference, and prediction* **2**. Springer.

HOEGH, A. and ROBERTS, D. W. (2020). Evaluating and presenting uncertainty in model-based unconstrained ordination. *Ecology and Evolution* **10** 59–69.

HOFFMAN, M. D., GELMAN, A. et al. (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15** 1593–1623.

HUANG, K., ZHANG, Y., GONG, H., QIAO, Z., WANG, T., ZHAO, W., HUANG, L. and ZHOU, X. (2023). Inferring evolutionary trajectories from cross-sectional transcriptomic data to mirror lung adenocarcinoma progression. *PLOS Computational Biology* **19** e1011122.

HUI, F. K. C., TASKINEN, S., PLEDGER, S., FOSTER, S. D. and WARTON, D. I. (2015). Model-based approaches to unconstrained ordination. *Methods in Ecology and Evolution* **6** 399-411. https://doi.org/10.1111/2041-210X.12236

HUI, F. K., WARTON, D. I., FOSTER, S. D. and HAAK, C. R. (2023). Spatiotemporal joint species distribution modelling: A basis function approach. *Methods in Ecology and Evolution* **14** 2150–2164.

HYMAN, B., PHELPS, C., BEACH, T., BIGIO, E., CAIRNS, N., CARRILLO, M., DICKSON, D., DUYCK-AERTS, C., FROSCH, M., MASLIAH, E., MIRRA, S., NELSON, P., SCHNEIDER, J., THAL, D., THIES, B., TROJANOWSKI, J., VINTERS, H. and MONTINE, T. (2012). National Institute on Aging-Alzheimer's Association guidelines for the neuropathologic assessment of Alzheimer's disease. *Alzheimers Dement.* **8**.

KIDZINSKI, L., HUI, F. K., WARTON, D. I. and HASTIE, T. J. (2022). Generalized Matrix Factorization: efficient algorithms for fitting generalized linear latent variable models to large data arrays. *Journal of Machine Learning Research* **23** 1–29.

MAGWENE, P. M., LIZARDI, P. and KIM, J. (2003). Reconstructing the temporal ordering of biological samples using microarray data. *Bioinformatics* **19** 842–850.

MASEL, J., JANSEN, V. A. and NOWAK, M. A. (1999). Quantifying the kinetic parameters of prion replication. *Biophysical chemistry* **77** 139–152.

MUKHERJEE, S., HEATH, L., PREUSS, C., JAYADEV, S., GARDEN, G. A., GREENWOOD, A. K., SIEBERTS, S. K., DE JAGER, P. L., ERTEKIN-TANER, N., CARTER, G. W. et al. (2020). Molecular estimation of neurodegeneration pseudotime in older brains. *Nature communications* **11** 5781.

NEAL, R. M. et al. (2011). MCMC using Hamiltonian dynamics. *Handbook of markov chain monte carlo* **2** 2.

O'HARA, R. B. and VAN DER VEEN, B. (2024). Hierarchical Ordination, A unifying framework for drivers of community processes. *bioRxiv* 2024–01.

PHAN, D., PRADHAN, N. and JANKOWIAK, M. (2019). Composable Effects for Flexible and Accelerated Probabilistic Programming in NumPyro. *arXiv preprint arXiv:1912.11554*.

PIERSON, E., KOH, P., HASHIMOTO, T., KOLLER, D., LESKOVEC, J., ERIKSSON, N. and LIANG, P. (2019). Inferring Multidimensional Rates of aging from Cross-sectional Data. *Proc. Mach. Learn. Res.* **89**.

POPOVIC, G. C., HUI, F. K. and WARTON, D. I. (2022). Fast model-based ordination with copulas. *Methods in Ecology and Evolution* **13** 194–202.

ROBERTS, D. W. (2020). Comparison of distance-based and model-based ordinations. *Ecology* **101** e02908.

STREET, K., RISSO, D., FLETCHER, R. B., DAS, D., NGAI, J., YOSEF, N., PURDOM, E. and DUDOIT, S. (2018). Slingshot: cell lineage and pseudotime inference for single-cell transcriptomics. *BMC genomics* **19** 1–16.

TENG, E., HASEGAWA, K., HOMMA, A., IMAI, Y., LARSON, E., GRAVES, A., SUGIMOTO, K., YAMAGUCHI, T., SASAKI, H., CHIU, D. and ET AL. (1994a). The Cognitive Abilities Screening Instrument (CASI): a practical test for cross-cultural epidemiological studies of dementia. *Int Psychogeriatr.* **6**.

TENG, E. L., HASEGAWA, K., HOMMA, A., IMAI, Y., LARSON, E., GRAVES, A., SUGIMOTO, K., YAMAGUCHI, T., SASAKI, H., CHIU, D. et al. (1994b). The Cognitive Abilities Screening Instrument (CASI): a practical test for cross-cultural epidemiological studies of dementia. *International psychogeriatrics* **6** 45–58.

THAL, D. R., RÜB, U., ORANTES, M. and BRAAK, H. (2002). Phases of Aβ-deposition in the human brain and its relevance for the development of AD. *Neurology* **58** 1791–1800.

TOMPA, P., TUSNÁDY, G., CSERZŐ, M. and SIMON, I. (2001). Prion protein: evolution caught en route. *Proceedings of the National Academy of Sciences* **98** 4431–4436.

TRAPNELL, C., CACCHIARELLI, D., GRIMSBY, J., POKHAREL, P., LI, S., MORSE, M., LENNON, N. J., LIVAK, K. J., MIKKELSEN, T. S. and RINN, J. L. (2014). The dynamics and regulators of cell fate decisions are revealed by pseudotemporal ordering of single cells. *Nature biotechnology* **32** 381–386.

VAN DER VEEN, B., HUI, F. K., HOVSTAD, K. A. and O'HARA, R. B. (2023). Concurrent ordination: Simultaneous unconstrained and constrained latent variable modelling. *Methods in Ecology and Evolution* **14** 683–695.

VAQUER-ALICEA, J. and DIAMOND, M. I. (2019). Propagation of protein aggregation in neurodegenerative diseases. *Annual review of biochemistry* **88** 785–810.

VEHTARI, A., GELMAN, A. and GABRY, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and computing* **27** 1413–1432.

WADDINGTON, C. H. (1940). *Organisers and Genes.* University Press; Cambridge.

WADDINGTON, C. H. (1957). *The strategy of genes.* Taylor and Francis Group.

WATANABE, S. and OPPER, M. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of machine learning research* **11**.

WEDIN, P.-Å. (1972). Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics* **12** 99–111.

ZHOU, J. and LIU, B. (2013). Alzheimer's disease and prion protein. *Intractable Rare Dis. Res.* **2** 35–44.

A

Middle Temporal Gyrus

Layer 1
Layer 2
Layer 3
Layer 4
Layer 5/6

Cortical Columnar Organization

Aβ

6e10 Immunostain
(Marking Aβ plaques)

Masks for Quantification
(Green: Aβ plaques masks)

B

Mean age
Median age

Age at death

C

ADNC
High   Not AD   Intermediate   Low

CASI
Not identified   Moderate   Mild   Severe

D 1

% AT8
+ area

Layer1
Layer2
Layer3
Layer4
Layer5-6

I   II   III   IV   V   VI
Braak stage

2

% 6e10
+ area

Layer1
Layer2
Layer3
Layer4
Layer5-6

0   1   2   3   4   5
Thal phase

3

% AT8
+ area

% 6e10
+ area

Layer1
Layer2
Layer3
Layer4
Layer5-6
Layer1
Layer2
Layer3
Layer4
Layer5-6

Not AD   Low   Intermediate   High
Adneurochange (ADNC)

4

% AT8
+ area

% 6e10
+ area

Layer1
Layer2
Layer3
Layer4
Layer5-6
Layer1
Layer2
Layer3
Layer4
Layer5-6

Not Identify   Mild   Moderate   Severe
CASI

Schematic Representation

Phenotypical Diversity
t
High AD
No AD

Hypothesized Sampling Process
$P(t)$
High AD
No AD

Measured Biological Process
$x = f(t)$

Measurements
$P(x) \sim NB(f(t), \varepsilon)$

AD Phenotype

t: Disease Severity

$P(t)$: Distribution of donors across disease severity

$f(t)$: Function linking observed measurement and disease severity

$P(x)$: Observational model describing measurement noise