NeuroImage 235 (2021) 117980

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Investigating hypotheses of neurodegeneration by learning dynamical systems of protein propagation in the brain

Sara Garbarino*, Marco Lorenzi, For the Alzheimer's Disease Neuroimaging Initiative[1]

Université Côte d'Azur, Inria, Epione Research Project, France

ARTICLE INFO

Keywords:
Neurodegeneration
Causal model
Dynamical systems
Protein propagation
Gaussian process
Brain connectivity

ABSTRACT

We introduce a theoretical framework for estimating, comparing and interpreting mechanistic hypotheses on long term protein propagation across brain networks in neurodegenerative disorders (ND). The model is expressed within a Bayesian non-parametric regression setting, where mechanisms of protein dynamics are inferred by means of gradient matching on dynamical systems (DS). The Bayesian formalism, combined with stochastic variational inference, naturally allows for model comparison via assessment of model evidence, while providing uncertainty quantification of causal relationship underlying protein progressions. When applied to in–vivo AV45-PET brain imaging data measuring topographic amyloid deposition in Alzheimer's disease (AD), our model identified the mechanisms of accumulation, clearance and propagation as the best suited DS for bio-mechanical description of amyloid dynamics in AD, enabling realistic and accurate personalized simulation of amyloidosis.

1. Introduction

It is common hypothesis that while under normal conditions an efficient clearance process allows the brain to control the accumulation of toxic misfolded proteins, in neurodegenerative disorders (ND) the equilibrium between accumulation and clearance is broken (Bateman et al., 2006). In this case, misfolded proteins may aggregate in plaques, eventually propagating between regions, leading to cellular dysfunction, disruption of synaptic connections, and neuronal loss (Soto and Pritzkow, 2018). Despite the involvement of distinct proteins in different ND, the process of protein misfolding is thought to remain similar, although still not completely understood (Editorial, 2018; Sweeney et al., 2017): proteins aggregates could self–propagate and subsequentially spread the pathology between cells and tissues along pathways largely overlapping with functional or structural brain networks (Brettschneider et al., 2015; Fornito and Bullmore, 2015; Jucker and Walker, 2013).

While a comprehensive model describing protein dynamics on the whole brain and along the whole span of the diseases is still missing, the current dominant paradigm for the study of proteinopathies in–vivo consists in associations analysis. For example, the relationship between pathological states and protein burden in the brain is usually quantified by means of correlation or regression models (Hansson et al., 2018; Melzer et al., 2019; Müller et al., 2019; Näslund et al., 2000; Palmqvist et al., 2014; Rentz et al., 2010). Nevertheless, findings issued from these studies shed rather limited insights about the pathophysiological mechanisms underlying the observations.

Orthogonal analysis paradigms are based on simulations, and have been proposed to investigate the dynamics of disease progression from a mechanistic point of view. The ambition of these approaches consists in discerning the generative spatio–temporal physical processes underlying the observations. These models thus offer the possibility to investigate the pathology from a causal perspective formalized through bio–mechanical relationships among observations (Iturria-Medina et al., 2017; Zhou et al., 2012). For example, a variety of generative models based on dynamical systems (DS) has been proposed for describing the kinetics governing the dynamical processes of accumulation, clearance and propagation of proteins (Cauda et al., 2018; Iturria-Medina et al., 2014; Oxtoby et al., 2017; Raj et al., 2015; 2012; Weickenmeier et al., 2018). DS define protein kinetics through (typically very large) systems of non–linear differential equations (ODE) which encode the underlying bio–mechanical processes in a set of *kinetic parameters* governing the system.

Many DS models define the propagation dynamics through diffusion equations (Cauda et al., 2018; Raj et al., 2015; 2012). This modelling choice allows to reduce the number of parameters to be estimated, but comes at the expenses of an oversimplification of the dynamics governing the dynamical processes of protein propagation. Indeed, while the pathological kinetics may be assimilated to diffusive processes in

* Corresponding author.
E-mail addresses: sara.garbarino@inria.fr (S. Garbarino), marco.lorenzi@inria.fr (M. Lorenzi).

[1] Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report.

https://doi.org/10.1016/j.neuroimage.2021.117980.
Received 7 January 2021; Received in revised form 20 February 2021; Accepted 12 March 2021
Available online 3 April 2021.

short term observations, the long term evolution of NDs are unlikely to have diffusive properties. For example, the asymptotically constant behaviour of NDs may not be described by the stationary and constant rate of change specified by diffusion equations. With the aim of overcoming this drawback, more recent works describe the processes of protein misfolding by combining classical population dynamics equations with anisotropic diffusion to describe constant and simultaneous propagation and aggregation (Weickenmeier et al., 2018), or by developing DS with protein concentration-dependent effects (Garbarino et al., 2019a).

Prediction and personalisation stem naturally from the ordinary differential equation structure of DS, which allows extrapolation from known initial conditions and prescribed kinetic parameters. When studying brain protein measurements in–vivo, the possibility of comparing models issued from different DS is of paramount importance for investigating the bio–mechanical hypotheses underlying disease progression. This is a crucial aspect, for observational data can be equally well described by conceptually different mechanistic hypotheses, and is especially important in the clinical setting, when direct intervention on the brain system is ultimately not possible. Unfortunately, since DS are usually identified by a large number of kinetic parameters and different relationship across variables, a consistent theoretical framework for DS comparison is often missing.

Another source of complexity in the modeling of protein propagation in–vivo arises from the lack of a precise definition of the time axis, thus preventing models of disease dynamics to reproduce patterns of protein propagation compatible across the entire history of ND. A well–defined temporal reference is indeed not available in typical longitudinal data sets of ND: the patient's time of onset is generally unknown, while the rate of pathological progression is highly variable across individuals. Moreover, ND are commonly studied through observational studies focusing on specific clinical stages, collecting individual measurements over a limited time–span, resulting in few time–points collected for each subject. To overcome the problem of modeling short term measurements with undefined temporal reference, several data–driven disease progression models (DPM) have been proposed (Donohue et al., 2014; Garbarino et al., 2019b; Jedynak et al., 2012; Lorenzi et al., 2017; Marinescu et al., 2019b; Schiratti et al., 2015; Villemagne et al., 2013; Young et al., 2014), based on the concept of self–modeling regression (Lawton et al., 1972). These approaches allow to reconstruct biomarkers trajectories along an ideal *long term* disease progression by optimally "stitching" together *short term* individual measurements. Each subject is then characterized by specific time parameters quantifying their pathological stage with respect to the estimated long term group–wise evolution. These models provide a description of temporal progression across biomarkers, usually expressed in terms of regression curves, or event orderings, without however elucidating the kinetics and relationships across them. This aspect indicates a limited ability of DPM in providing an understanding of the pathological mechanisms.

To date, no modeling framework allows inference and comparison of mechanistic hypotheses of protein processes across the long term evolution of ND. The problem is challenging since it requires to simultaneously account for short term observations to reconstruct the long term disease progression, and to estimate the kinetic parameters specified by high dimensional dynamical systems.

In this study we tackle this issue by formulating Gaussian process progression modeling for dynamical systems (GPPM–DS): a framework for specifying hypothetical models of protein dynamics, fitting them to the data and comparing their evidence through Bayesian model comparison. As a result, we can identify the most plausible dynamics underlying the measurements, and provide an interpretable model of the potential causal relationship across variables. This framework enables to formalize and test novel hypotheses on disease dynamics, and to challenge them against alternative ones. This idea introduces a paradigm shift in the analysis of neurodegeneration in neuroimaging data analogous to the one brought by Dynamic Causal Modeling Friston et al. (2003) for the hypothesis–driven analysis in functional neuroimaging. Compared to this latter framework, here we tackle the additional problem of the lack of a well–defined time axis for the pathology, and present a completely novel theoretical framework for scalable inference.

Overall, in this work we provide a framework for i) testing hypothesis–driven models of protein dynamics against longitudinal data, ii) simulating protein propagation along the whole disease progression time span, and iii) predicting individual protein deposition in unseen data. We test GPPM–DS on a variety of synthetic data and evaluate its performances in estimating protein evolution as well as kinetic and individual time parameters, for each proposed DS, as compared to standard DPM based on monotonic constraints (Lorenzi et al., 2017). Results on synthetic data can be found in the Supplementary Material.

We then demonstrate GPPM–DS on the modeling of the evolution of cerebral amyloid protein accumulation from AV45–PET data of Alzheimer's disease (AD) subjects from the ADNI data set. We show that GPPM–DS identifies the non–linear process of accumulation, clearance and propagation as the best suited DS for bio–mechanical description of amyloid dynamics.

2. Methods

GPPM–DS is described as a Bayesian non–parametric constrained regression problem (Section 2.1), where constraints are imposed on the dynamics of protein evolution and are given by a DS for protein dynamics (Section 2.2), while the long term data are reconstructed from short term observation by using time reparametrization techniques typical of the DPM framework.

Bayesian inference techniques are used to solve the GPPM–DS problem (Section 2.3). Specifically, we determine a lower bound for the generally intractable model posterior. We resort to optimizing an approximation of the marginal function, and imposing sparsity constraints on the model parameters. Efficient optimisation is performed through stochastic gradient descent via backpropagation and mini–batch optimisation.

2.1. Bayesian non–parametric constrained regression framework

Let us assume to have S subjects ($s = 1, \dots, S$), each one with associated measurements of protein concentrations $\mathbf{Y}^s$, at different brain regions and at different time–points over a short term time span $\mathbf{t}^s = \{t^s_1, \dots, t^s_{T^s}\}$, where T^s denotes the total number of measurements for a given subject s. We then assume that each individual measurement is obtained with respect to an absolute time–frame τ through a function $t = g^s(\tau)$ modeling the time–reparameterization with respect to the common group–wise evolution. This means assuming that the short term time–points are obtained through reparametrization of some unknown time–points $\tau^s_k \in I \subseteq \mathbb{R}$ over a long term span, such that:

$$\begin{aligned} g^s &: I \to J^S \subset I \\ &\tau^s_k \mapsto t^s_k = g^s\left(\tau^s_k\right) \end{aligned} \tag{1}$$

Finally, we assume that the measurements $\mathbf{Y}^s$ are realizations of an unknown group–wise process $\boldsymbol{f}$, describing the temporal evolution of protein concentrations. According to these assumptions, we model the observations $\mathbf{Y}^s$ via mixed–effect regression with time reparametrization:

$$\mathbf{Y}^s(g^s(\tau)) = \boldsymbol{f}(g^s(\tau)) + \boldsymbol{\nu}^s(g^s(\tau)) + \epsilon. \tag{2}$$

A variety of models has been proposed to tackle this kind of regression problem, either via parametric or non–parametric techniques. Here we use GP Progression Model (Lorenzi et al., 2017), which is a non–parametric Bayesian DPM describing $\boldsymbol{f}$ as a Gaussian process (GP), and g^s as a translation parameterized by an individual time–shift: $g^s(\tau^s_k) = t^s_k + d^s$. Finally, $\boldsymbol{\nu}^s$ are assumed to be Gaussian random effects $\mathcal{N}(0, \phi^s)$.

We introduce constraints on the dynamics of the model, enforcing the concentrations' evolution to a general DS, governed by a functional $\mathcal{H}_\theta$, which depends on M *kinetics parameters* $\{\theta_i\}_{m=1}^M = \boldsymbol{\theta}$. This means

specifying a family of admissible functions whose derivatives evaluated at points $\mathbf{t}$ satisfy the DS:

$$\mathscr{H} = \{\boldsymbol{f}(t) : \dot{\boldsymbol{f}}(t) = \mathcal{H}_{\theta}(\boldsymbol{f}(t), t)\}. \tag{3}$$

The GPPM–DS model is finally described as Eq. (2) subject to constraints (3).

2.2. Dynamical systems for modeling misfolded proteins

We introduce three different DS accounting for protein dynamics: i) diffusion (Diff): a purely diffusive model of constant propagation, based on the work presented in Raj et al. (2012); ii) reaction–diffusion (RD): a model where aggregation and propagation of proteins are simultaneous, constant and opposite, and the total process eventually reaches a plateau, to reproduce the DS proposed by Weickenmeier et al. (2018); and iii) accumulation clearance and propagation (ACP): a model where propagation is triggered when the aggregation saturates, and then both aggregation and propagation eventually reach a plateau (Garbarino et al., 2019a). The three chosen models offer a range of varying complexity in terms of understanding of the underlying pathology: from the simplest one (Diff), which allows to estimate purely diffusive effects of an hypothetical constant propagation of proteins, to the most complex (ACP), which incorporates the effects of multivariate propagation, aggregation saturation and clearance. In this context, the ACP model is related to the highest range of parameters describing protein propagation, and so potentially offers the highest capacity in describing the pathological progression. We aim at identifying, amongst these three models, the one that best explain the dynamics of the data, while at the same time accounting for model complexity.

Here we consider the brain as a system of N interconnected regions, where each region i ($i = 1, \dots, N$) is characterized by its concentration of proteins $f_i(\tau)$ along the whole long term time interval I. We describe the network of brain connections as a graph, whose adjacency matrix is denoted by A, its degree matrix by D, and the corresponding Laplacian matrix by $H = D - A$. In the following, $\odot$ denotes the Hadamard product between either matrices or vectors.

The three DS are defined as follows:

$$\begin{aligned}\textbf{Diff}: \dot{\boldsymbol{f}}(t) &= \mathcal{H}_{\theta_D}(\boldsymbol{f}(\tau), \tau) \\ &= \left(A \odot K_{\theta_D}\right) H \boldsymbol{f}(\tau).\end{aligned} \tag{4}$$

Here the matrix K_{θ_D} encodes region–wise, constant rates of protein propagation across adjacent regions and is described as $\left(K_{\theta_D}\right)_{ij} = k_{ij}$, where k_{ij} is the (symmetrical) rate of propagation between regions i and j. The Diff model is linear and produces patterns of protein concentration that accumulate indefinitely over time. The set of kinetic parameters for the Diff model is $\boldsymbol{\theta}_D = \left(k_{ij}\right)$.

$$\begin{aligned}\textbf{RD}: \dot{\boldsymbol{f}}(t) &= \mathcal{H}_{\theta_{RD}}(\boldsymbol{f}(\tau), \tau) \\ &= \left(A \odot K_{\theta_{RD}}\right) H \boldsymbol{f}(\tau) + R_{\theta_{RD}} \boldsymbol{f}(\tau) \odot (\boldsymbol{\upsilon} - \boldsymbol{f}(\tau)).\end{aligned} \tag{5}$$

This model includes both propagation and aggregation mechanisms, which are constant and simultaneous. We assume that no aggregation nor propagation occur in healthy conditions, while protein plaques aggregation develops when the accumulation–clearance equilibrium breaks. The model is comprised of two terms: a standard diffusion term $\left(A \odot K_{\theta_{RD}}\right) H \boldsymbol{f}(\tau)$ for describing constant protein propagation (see above) and a reaction term $R_{\theta_{RD}} \boldsymbol{f}(t) \odot (\boldsymbol{\upsilon} - \boldsymbol{f}(\tau))$ for describing protein aggregation. The second term balances constant propagation with a constant aggregation term encoded by the matrix $R_{\theta_{RD}} = k_t I$, where k_t is the rate of total aggregation and is constant across regions. This term eventually reaches a plateau when protein concentration get to a maximal concentration threshold $\boldsymbol{\upsilon}$. This equation is known as the Fisher–Kolmogorov equation (Adomian, 1995). The set of kinetic parameters for the RD model is $\boldsymbol{\theta}_{RD} = \left(k_{ij}, k_t, \boldsymbol{\upsilon}\right)$.

$$\begin{aligned}\textbf{ACP}: \dot{\boldsymbol{f}}(\tau) &= \mathcal{H}_{\theta_{ACP}}(\boldsymbol{f}(\tau), \tau) \\ &= \left(A \odot K_{\theta_{ACP}}(\boldsymbol{f}, \tau)\right) H \boldsymbol{f}(\tau) + R_{\theta_{ACP}}(\boldsymbol{f}, \tau) \boldsymbol{f}(\tau)\end{aligned} \tag{6}$$

The ACP model describes the three processes of accumulation, clearance and propagation of proteins allowing for non–constant effects. We assume again no aggregation nor propagation to occur in healthy conditions, and aggregation to develop when the accumulation–clearance equilibrium breaks. The propagation term $K_{\theta_{ACP}}$ is concentration–dependent: the toxic protein concentration in each region saturates when reaching a first critical threshold $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_N)$, and subsequently triggers propagation towards the connected regions. Propagation also reaches a plateau when passing a second critical threshold $\boldsymbol{\eta} = (\eta_1, \dots, \eta_N)$. This can be modeled by setting $\left(K_{\theta_{ACP}}\right)_{ij} = k_{ij_{ACP}}(\boldsymbol{f}(\tau), \boldsymbol{\gamma}, \boldsymbol{\eta})$, where $k_{ij_{ACP}}$ are sigmoid functions of the constant propagation rates k_{ij}

$$k_{ij_{ACP}}(f_j(\tau), \gamma_j, \eta_j) = \frac{k_{ij}}{\left(1 + e^{-l_1(f_j(\tau) - \gamma_j)}\right)\left(1 + e^{l_2(f_j(\tau) - \eta_j)}\right)}. \tag{7}$$

For the aggregation term $R_{\theta_{ACP}}$, we hypothesize a similar concentration–dependent threshold above which the total aggregation process reaches a plateau. In this context, $\left(R_{\theta_{ACP}}\right)_{ij}$ is re–written as

$$\left(R_{\theta_{ACP}}\right)_{ij} = \begin{cases} k_{t_{ACP}}(\boldsymbol{f}(\tau), \boldsymbol{\eta}) & \text{if} \quad i = j \\ 0 \quad \text{otherwise;} \end{cases} \tag{8}$$

where

$$k_{t_{ACP}}(f_i(\tau), \eta_i) = \frac{k_t}{1 + e^{l_2(f_i(\tau) - \eta_i)}}. \tag{9}$$

The key difference with the previous RD model, is that the ACP model does not assume aggregation and propagation to be simultaneous nor constant, but rather hypothesizes the existence of critical values of protein concentrations, different for each region, at which the regional concentrations saturate and subsequently trigger propagation. The set of kinetic parameters for the ACP model is $\boldsymbol{\theta}_{ACP} = \left(k_{ij}, k_t, \boldsymbol{\gamma}, \boldsymbol{\eta}\right)$.

Fig. 1 shows examples of Diff, RD and ACP models with varying key parameters on a toy data set of $N = 3$ synthetic brain regions. In Fig. 1A), for the linear Diff model, the highest the rates of propagation k_{ij}, the fastest the proteins travel amongst the simulated brain regions, whose protein concentration tend to balance over time. Fig. 1B)-C) show trajectories obtained by varying the k_{ik} (1B) and υ (1C) parameters for the RD model (while the other parameter is kept constant). When k_{ij} increases, the model predicts a behaviour similar to the Diff model in 1A), with fastest propagation amongst region and concentration trajectories that balance over time. In Fig. 1C) we note how accumulation reaches a plateau once the maximal concentration threshold parameter $\boldsymbol{\upsilon}$ is hit. Fig. 1D)-F) show protein trajectories predicted by the ACP model with varying parameters. At increasing k_{ij} again correspond fastest propagation amongst regions (1D), while when varying the l_1 and l_2 parameters, the trajectories display an increasingly sigmoidal shape (1E). Finally, when the parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\eta}$ increase, the concentrations saturate, triggering propagation, and reach a plateau, later in the progression (1F).

2.3. Inference

Solving the GPPM–DS system of Eqs. (2) and (3) requires inference of the kinetic parameters $\boldsymbol{\theta}$, the group–wise dynamics $\boldsymbol{f}$, and, for each subject s, of the parameters of the individual random effects $\boldsymbol{\phi}^s$ and of the time–shift parameters d^s. We first define $\boldsymbol{F}^s$ as the realization of $\boldsymbol{f}$ at times $\boldsymbol{t}^s$, and $\dot{\boldsymbol{F}}^s$ as the set of realizations of $\boldsymbol{f}$ and of its derivatives at (generally different) times $\boldsymbol{u}^s$. To ease the notation, we assume both $\boldsymbol{f}$ and its derivatives to have realizations on the same time–points $\boldsymbol{t}^s$, but computation extends easily to more complex scenarios. We also indicate by $\boldsymbol{F}$, $\boldsymbol{\nu}$, $\dot{\boldsymbol{F}}$, $\boldsymbol{d}$ and $\boldsymbol{\phi}$ the collections of $\boldsymbol{F}^s$, $\boldsymbol{\nu}^s$, $\dot{\boldsymbol{F}}^s$, d^s and ϕ^s for all

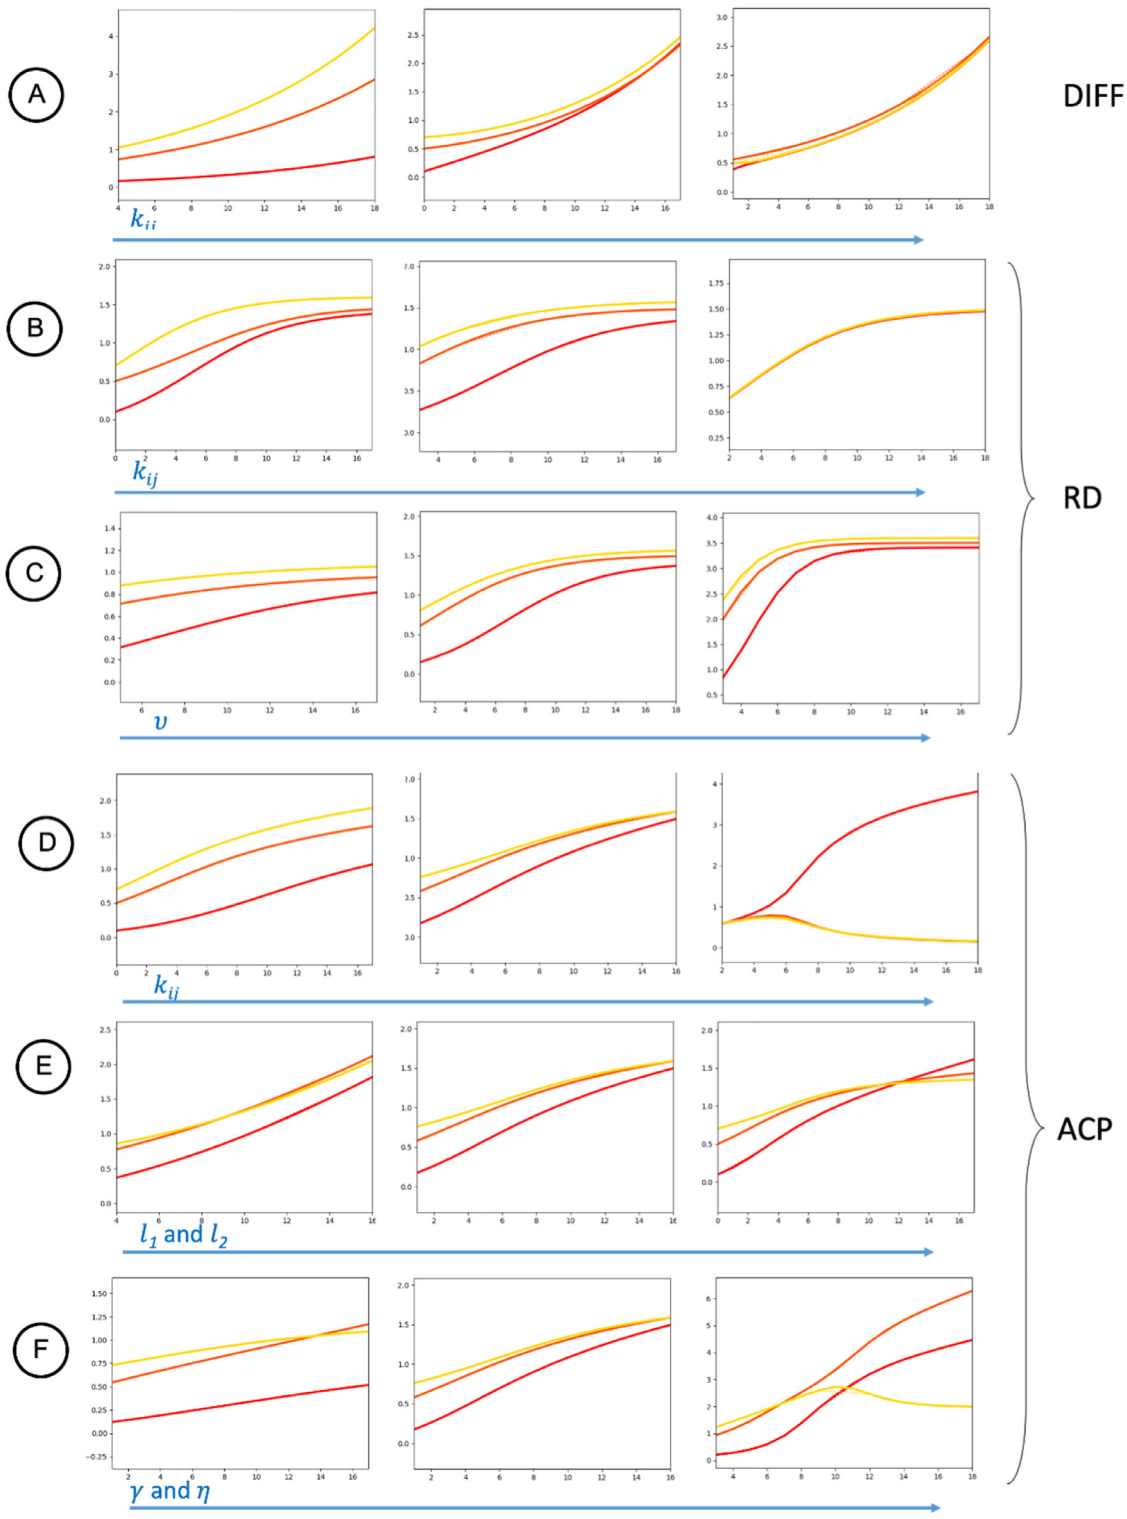

Fig. 1. Examples of Diff, RD and ACP models with varying key parameters on a toy data set of $N = 3$ synthetic brain regions. A) Diff model with varying rates of propagation k_{ij}. For each biomarker, k_{ij} is a random realisation from a Gaussian centered in (from left to right): 0.05, 0.5 and 1. B) RD model with varying rates of propagation k_{ij}. For each biomarker, k_{ij} is a random realisation from a Gaussian centered in (from left to right): 0.05, 0.5 and 1. The threshold parameter $\boldsymbol{v}$ is set at 1. C) RD model with varying threshold parameter $\boldsymbol{v}$. For each biomarker, $\boldsymbol{v}$ is a random realisation from a Gaussian centered in (from left to right): 0.5, 1 and 3. The k_{ij} parameters are set at 0.5. D) ACP model with varying rates of propagation k_{ij}. For each biomarker, k_{ij} is a random realisation from a Gaussian centered in (from left to right): 0.05, 0.5 and 1. The threshold parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\eta}$ ares set at 0.6 and 0.9 respectively. The sigmoid parameters l_1 and l_2 are set at 3. E) ACP model with varying sigmoid parameters l_1 and l_2. For each biomarker, l_1 and l_2 are a random realisation from a Gaussian centered in (from left to right): $(0.5, 0.5)$, $(3, 3)$ and $(5, 5)$. The k_{ij} parameters are set at 0.5. The threshold parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\eta}$ are set at $(0.6, 0.9)$. F) ACP model with varying threshold parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\eta}$. For each biomarker, $\boldsymbol{\gamma}$ and $\boldsymbol{\eta}$ are a random realisation from a Gaussian centered in (from left to right): $(0.1, 0.3)$, $(0.6, 0.9)$ and $(2, 3)$. The k_{ij} parameters are set at 0.5. The sigmoid parameters l_1 and l_2 are set at 3.

the subjects ($s = 1, \dots, S$). Following Lorenzi and Filippone (2018) and Cutajar et al. (2017), we describe the constrained regression problem in a Bayesian setting and solve the inference problem for $\boldsymbol{\theta}$ and $\boldsymbol{F}$ by determining a lower bound for the marginal

$$p(\mathbf{Y}, \mathcal{H}|\boldsymbol{\phi}, \boldsymbol{d}, \mathbf{t}, \epsilon, \zeta) = \int p(\mathbf{Y}|\boldsymbol{F}, \boldsymbol{\phi}, \boldsymbol{d}, \mathbf{t}, \epsilon) p\big(\mathcal{H}|\dot{\boldsymbol{F}}, \boldsymbol{\theta}, \boldsymbol{d}, \mathbf{t}, \zeta\big) \\ p\big(\boldsymbol{F}, \dot{\boldsymbol{F}}|\boldsymbol{\phi}, \boldsymbol{d}, \mathbf{t}\big) p(\boldsymbol{\theta}) d\boldsymbol{F}\, d\dot{\boldsymbol{F}}\, d\boldsymbol{\theta}, \tag{10}$$

where

$$p(\boldsymbol{F}, \dot{\boldsymbol{F}}|\boldsymbol{\phi}, \boldsymbol{d}, \mathbf{t}) d\boldsymbol{F} = p(\dot{\boldsymbol{F}}|\boldsymbol{F}) p(\boldsymbol{F}|\boldsymbol{\phi}, \boldsymbol{d}, \mathbf{t}). \tag{11}$$

Analogously to Lorenzi and Filippone (2018), we assume the two likelihood terms (the data fidelity term $p(\mathbf{Y}|\boldsymbol{F}, \boldsymbol{\phi}, \mathbf{t}, \boldsymbol{d}, \boldsymbol{\epsilon})$ and the constraint term $p(\mathscr{H}|\dot{\boldsymbol{F}}, \boldsymbol{\theta}, \mathbf{t}, \boldsymbol{d}, \boldsymbol{\zeta})$) to be respectively Gaussian and Student–t with respective variance and scale parameters ϵ and ζ.

Due to the general intractable form of (10), we approximate the GP $\boldsymbol{F}$ via random features (RF) expansion (Rahimi and Recht, 2008). Specifically, a GP with radial basis function covariance can be expressed as $\boldsymbol{F} \approx h(\mathbf{t}\boldsymbol{\Omega})\mathbf{W}$, where $\boldsymbol{\Omega}$ is a linear projection of the input $\mathbf{t}$ into the RF space specified by trigonometric activation functions $h(\cdot) = (\cos(\cdot), \sin(\cdot))$, and $\mathbf{W}$ are the regression parameters (Cutajar et al., 2017). Such approximation extends to the derivatives of the GP thanks to the chain rule (Lorenzi and Filippone, 2018). As a result, both GP realization and its derivatives can be identified by the same parameters $\mathbf{W}$ and $\boldsymbol{\Omega}$.

The RF approximation allows to replace inference on $\boldsymbol{F}$ with inference on $\mathbf{W}$ and $\boldsymbol{\Omega}$. Following Cutajar et al. (2017), we optimize (10) through variational inference of $\mathbf{W}$ and $\boldsymbol{\theta}$, and assume $\boldsymbol{\Omega}$ to be sampled from the prior with fixed randomness. This leads to the optimization of the following lower bound of the log–marginal (or evidence lower bound, ELBO):

$$log(p(\mathbf{Y}, \mathcal{H}|\boldsymbol{\phi}, \boldsymbol{d}, \mathbf{t}, \epsilon, \zeta)) \geq E_{q(\mathbf{W})}\big[\log\left(p(\mathbf{Y}|\boldsymbol{\Omega}, \mathbf{W}, \phi, \boldsymbol{d}, \mathbf{t}, \epsilon)\right)\big] \\ + E_{q(\mathbf{W})q(\boldsymbol{\theta})}\big[\log\left(p(\mathcal{H}|\boldsymbol{\Omega}, \mathbf{W}, \boldsymbol{\theta}, \boldsymbol{d}, \mathbf{t}, \zeta)\right)\big] \\ - DKL(q(\mathbf{W})|p(\mathbf{W})) - DKL(q(\boldsymbol{\theta})|p(\boldsymbol{\theta})). \tag{12}$$

Here, we perform variational inference on the kinetic parameters $\boldsymbol{\theta}$ and group–wise dynamics $\boldsymbol{F}$ (and so $\mathbf{W}$), while finding a maximum likelihood estimate of time shift parameters $\boldsymbol{d}$ and of the individual random effects $\boldsymbol{\phi}$. The $DKL(q|p)$ term represents the Kullback Leibler divergence between the prior p and the variational approximation q, and we assume $q(\mathbf{W})$ and $q(\boldsymbol{\theta})$ to be Gaussians $\mathcal{N}(\mu_\theta, \sigma_\theta)$ and $\mathcal{N}(\mu_\mathbf{W}, \sigma_\mathbf{W})$. In this way the last two terms of Eq. (12) have closed forms (Cutajar et al., 2017). Thanks to separability and analytical forms of likelihood and DKL terms, the ELBO (12) can be optimized through stochastic gradient descent via backpropagation and mini–batch (Kingma et al., 2015). Full optimization is performed with stochastic gradient descent with adaptive moment estimation (Adam) (Kingma and Ba, 2014), through the alternate optimization of

i) the approximated posterior over $\mathbf{W}$, GP parameters and random effects: $q(\mathbf{W})$, ϵ and ϕ;
ii) the individual time–shifts parameters $\boldsymbol{d}$;
iii) the approximated posterior over the kinetic parameters and likelihood parameters of the constraints: $q(\boldsymbol{\theta})$ and ζ.

2.4. Sparsity on the kinetic parameters

A well–known limitation of fitting non–linear DS concerns the identifiability of the associated kinetic parameters $\boldsymbol{\theta}$. Here, inspired by previous work on dropout and variational dropout (Antelmi et al., 2019; Kingma et al., 2015; Molchanov et al., 2017), we impose parsimonious representations of the model kinetic parameters through sparsity constraints (sparse GPPM–DS). As we will see, sparse GPPM–DS sets to zero the kinetic parameters which are associated with high estimated variance - parameters which can be considered un–identifiable, given the observed data.

It has been shown Kingma et al. (2015) that a way to impose sparsity on Bayesian neural networks (i.e. regularize their weights w) is to use the so–called local reparametrization trick: re–parametrize the Gaussian approximation $q(w) \sim \mathcal{N}(\mu_w, \sigma_w)$ of the weights of the approximated posterior $p(w)$ as $\tilde{q}(w) \sim \mathcal{N}(\mu_w, \alpha\mu_w^2)$, where α is the *dropout rate*. The authors in Molchanov et al. (2017) proposed a form for the prior on w consistent with the optimization of the ELBO associated with new $\tilde{q}$: the log–scale uniform function $p(|w|) \propto \frac{1}{|w|}$. In our case, we are interested in imposing sparsity constraints on the kinetic parameters $\boldsymbol{\theta}$. For this reason, we apply the local re–parametrization trick to the Gaussian approximation of the posterior distribution of the kinetic parameters $q(\boldsymbol{\theta}) \sim \mathcal{N}(\mu_\theta, \sigma_\theta)$, obtaining $\tilde{q}(\boldsymbol{\theta}) \sim \mathcal{N}(\mu_\theta, \alpha\mu_\theta^2)$. This modification promotes sparsity, as large values of θ_i correspond to large uncertainty expressed by the variance term $\alpha\mu_{\theta_i}^2$, indicating that the associated kinetic parameter θ_i can be set to zero. In what follows, when using sparse GPPM–DS, the associated $D_{KL}(\tilde{q}(\boldsymbol{\theta})|p(\boldsymbol{\theta}))$ is computed according to the numerical approximation as described in Molchanov et al. (2017).

2.5. Simulation results

We tested GPPM–DS in both sparse and full versions on a variety of synthetic data sets and compared its performances in recovering kinetic parameters, simulated evolution and time reparametrization parameters, for each of the three proposed DS, as compared to standard DPM.

As reference DPM, we implement a GPPM based on monotonic constraints (GPPM–mono) (Lorenzi et al., 2017): it is, as the model proposed in this work, a Gaussian process–based random effect modelling of longitudinal progressions, and simply assumes a steady temporal evolution of biomarkers from normal to pathological values, enforcing the biomarker trajectories to follow a monotonic behaviour. GPPM–mono provides us with a "null progression" hypothesis, in which no propagation parameters are accounted for, while regions are simply assumed to steadily evolve from normal to pathological values. Specifically (full details can be found in the Supplementary Material, Section S1):

1. we generated ground truth data using each of the proposed DS and then attempt reconstruction using every GPPM–DS, demonstrating, in terms of both RMSE and ELBO, that the candidate model is generally optimal when it corresponds to the ground truth DS (Supplementary Section S1.1).
2. We tested the variational dropout scheme for identifying zero–valued kinetic parameters and demonstrate that the sparse model has good accuracy in identifying such parameters (Supplementary Section S1.2).
3. We analyzed the performance of the sparse models in predicting unseen data as compared to the full models, demonstrating that they perform similarly in terms of RMSE while requiring only a fraction of parameters (Supplementary Section S1.3).

Moreover, in order to probe whether the mechanistic inference is truly driven by biological dynamics, we performed a set of simulations in which the ground truth data, generated using one of the propagation models, have been shuffled 100 times, thus obtaining 100 data sets with no biological signal (Supplementary Section S1.5). Then, we reconstructed with every GPPM–DS and showed that, in terms of ELBO, the three models best reconstruct the data roughly the same number of times. These results show that under the null hypothesis of no biological signal underlying the spatial data, the models are roughly equivalent from the point of view of Bayesian model comparison.

3. Results

GPPM–DS relies on the gradient-matching approach presented in Lorenzi and Filippone (2018) that overcomes the issue of using

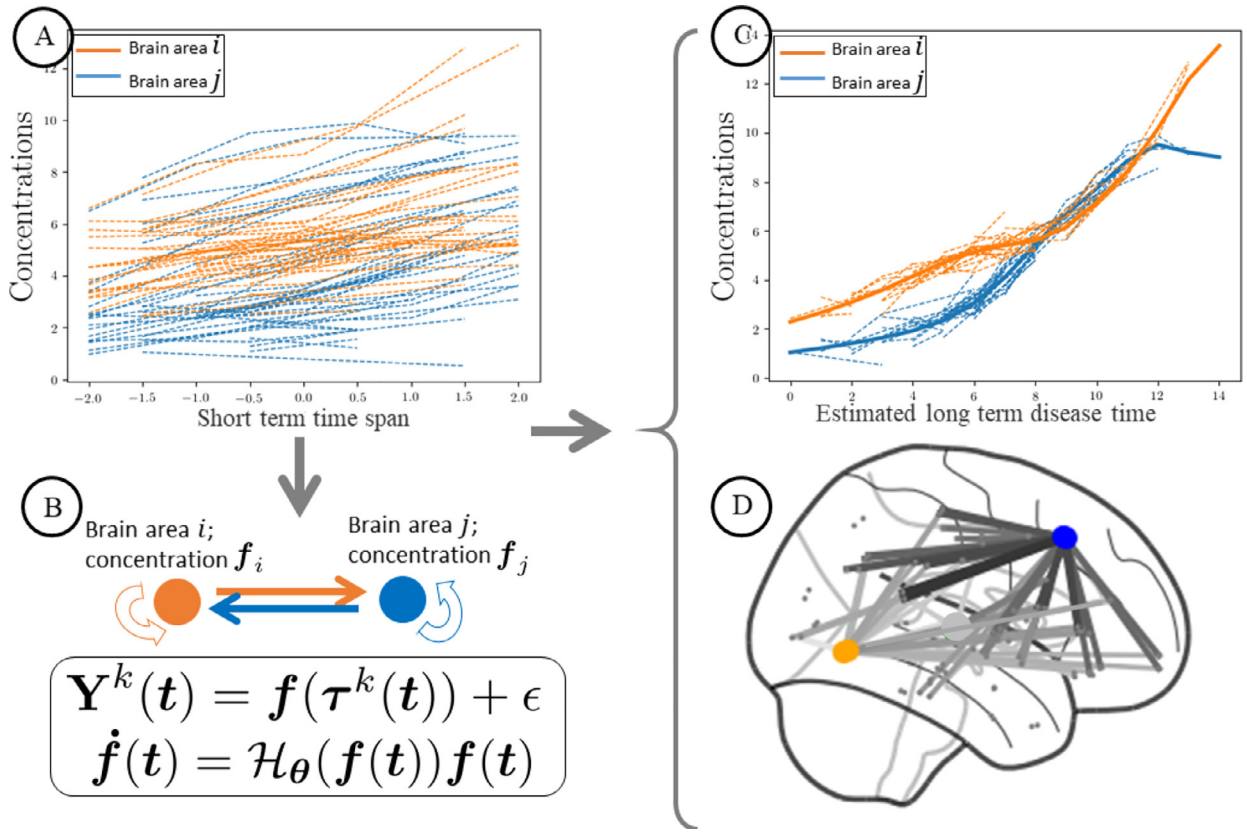

Fig. 2. Schematic representation of the proposed GPPM–DS framework. Regional protein concentrations f_i and f_j are collected for k subjects over a short term time span t, encoded in a measurement array $\mathbf{Y}^k(t)$ (A). The dynamics of such concentrations are described in terms of a functional $\mathcal{H}_\theta$, with unknown parameters θ, encoded in a dynamical system for the vector of concentrations f (B). The proposed framework estimates such parameters encoding the strength of propagation (D) and the long term protein concentrations with respect to the estimated long term time axis (C).

costly numerical integration procedures for inferring ODE parameters (Macdonald and Husmeier, 2015), thus allowing scalable inference of protein progression jointly with associated kinetic parameters. The proposed framework is formulated as a constrained regression problem in a Bayesian non–parametric setting, where the protein progression is modeled as a Gaussian process (GP), while bio–mechanical processes are defined as constraints on the protein dynamics through DS expressed by ODE. The Bayesian setting allows for uncertainty quantification of the protein dynamics while, to achieve tractability, the inference problem is solved via stochastic variational inference. The framework also provides a principled theory for model comparison via assessment of model evidence: here we investigate use of the evidence lower bound (ELBO) as a surrogate for model comparison. Fig. 2 shows a schematic representation of our framework.

We propose GPPM–DS by comparing three different DS of amyloid protein dynamics in AD. We investigate a purely diffusive model of constant propagation (Diff), based on the preliminary work of Raj et al. (2012); a reaction-diffusion model (RD), where aggregation and propagation of proteins are simultaneous, constant and opposite, and the total process eventually reaches a plateau, to reproduce the DS proposed by Weickenmeier et al. (2018); a non–linear accumulation clearance and propagation model (ACP) (Garbarino et al., 2019a), where propagation is triggered when aggregation reaches saturation, and then both aggregation and propagation reach a plateau. For each DS, GPPM–DS provides a complete description of the protein dynamics, which can be subsequently used for simulating and predicting proteins changes over time through forward integration. To address the ill–posed nature of the identification of the DS parameters (Saccomani, 2015), we further leverage on variational dropout techniques (Antelmi et al., 2019; Kingma et al., 2015; Molchanov et al., 2017) for imposing parsimonious representations of the model parameters through sparsity constraints (sparse GPPM–DS).

We show first that GPPM–ACP model outperforms the others in terms of ELBO and predictive accuracy of amyloid deposition in unseen follow–up data. This is confirmed by the analysis of the null scenario case where the model is optimised with respect to the sole monotonic constraint of the trajectories (GPPM–mono) (Lorenzi et al., 2017). Full results on GPPM–mono are shown in the Supplementary Material, Figures S6 and S10 A). We then show that GPPM–ACP model allows biomechanical interpretation of amyloid dynamics in AD, while providing plausible description of the pathological evolution.

3.1. Data acquisition and preprocessing

We analyzed AV45–PET brain imaging data of 770 subjects from ADNI, with a total of 1477 longitudinal data points. The neuronal connections along which protein propagation occurs is approximated by estimating the average structural connectomes from 24 young and healthy volunteers of the HCP data set. In this case, the substrate for protein propagation is assumed to be unaffected by the pathology.

HCP data: Data used in the preparation of this work were obtained from the MGH-USC Human Connectome Project database. We collected 3D T1w and DTI of 24 age and gender–matched subjects. Details on the selected subjects and the pipeline for structural connectome generation are described in Oxtoby et al. (2017). We averaged the 24 connectomes and obtained an average young, healthy connectome on 82 FreeSurfer regions (Desikan et al., 2006). For computational reasons, and in order to remove weak connections, we averaged left and right

Table 1
Socio demographic and clinical information for the ADNI study cohort at baseline (770 subjects). ADAS13: Alzheimer's Disease Assessment Scale-cognitive subscale, 13 items; FAQ: Functional Assessment Questionnaire; MMSE: Mini-Mental State Exam; RAVLT learning: Rey Auditory Verbal Learning Test, learning item.

Group	CN-	CN+	MCI	Dementia
N (female)	173(87)	133(76)	307(124)	157(65)
age (std)	73.2(6.5)	75.3(6.4)	73.0(7.3)	74.6(8.1)
years education (std)	16.7(2.7)	16.3(2.5)	16.2(2.8)	15.8(2.6)
APOE4+ (% positive)	26(15%)	56(42%)	190(62%)	105(67%)
ADAS13 (std)	8.6(4.3)	9.2(4.3)	16.2(6.8)	28.8(9.2)
FAQ (std)	0.3(1.3)	0.3(0.8)	3.1(3.8)	11.8(7.4)
MMSE(std)	29.0(1.2)	29.0(1.2)	27.9(1.8)	23.8(2.5)
RAVLT learning (std)	6.2(2.1)	5.7(2.5)	4.4(2.6)	2.0(1.8)

hemispheres (remaining with 41 regions) and we set to 0 all the weights below the estimated average weights across nodes, and to 1 the weight above.

ADNI data: The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. The primary goal of ADNI has been to test whether serial Magnetic Resonance Imaging (MRI), Positron Emission Tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of MCI and early AD. For up-to-date information, see www.adni-info.org.

We collected clinical, demographic and AV45-PET SUVr data from the adnimerge package (adni.bitbucket.io (The Alzheimer's Disease Neuroimaging Initiative, 2019)) for subjects with diagnosis label "Dementia", Mild Cognitive Impairment" (MCI) or "Cognitively Normal" (CN). As we are interested in modeling amyloid evolution in well phenotyped Alzheimer's trajectories, we selected, among the MCI or Dementia subjects, only the one effectively displaying amyloid accumulation. On the other hand, we selected the whole cohort of CN (both amyloid positive: CN+ and amyloid negative: CN-), in order to maximize the time–span of the data. The threshold for amyloid positivity was established at amyloid level in the cerebrospinal fluid < 192 ng/l (Almdahl et al., 2017). The information was extracted from the adnimerge R package (MEDIAN field of the upennbiomkmaster table). ADNI AV45-PET SUVr data are already computed on a brain anatomical parcellation estimated with the software Freesurfer (Desikan et al., 2006), and normalized against cerebellar uptake. For each region we discarded white matter, ventricular and cerebellar regions, remaining with 82 regions, and averaged the SUVr of each region from both hemispheres, remaining with 41 regions. Demographic and clinical details are shown in Table 1. Finally we split the data set in two parts. A training data set D_1 was defined by including all the longitudinal data for each subject up to the second-to-last time points. The remaining time–points were included in a testing data set D_2. Subjects with one measurement only were included in D_1. The resulting data set D_1 includes 984 longitudinal measures; D_2 contains 493 cross-sectional measures. We run the models on D_1, estimating protein dynamics, GP parameters and individual time–shifts, and used D_2 to validate model predictions.

3.2. GPPM–DS for modeling amyloid deposition in Alzheimer's disease

The proposed GPPM–DS framework can be effectively used to compare different hypotheses of protein kinetics. We analyze AV45–PET data with three different kinetic models, whose mathematical description can be found in the Method Section 4.2: Diff model of Eq. (4), RD model of Eq. (5) and ACP of Eq. (6). We use the ELBO as surrogate for model comparison (as for the synthetic data, see Supp. Mat.): we compared the ELBO while training of D_1, and the RMSE for prediction on the unseen data (D_2). We computed the RMSE for each individual prediction (averaging over the biomarkers), and bootstrapped over the kinetic parameters 100 times, obtaining RMSE distributions. Individual predictions on D_2 are performed by forward integration of the associated dynamical system with the estimated kinetic parameters, from given initial conditions (the individual baseline values) for the sparse model (i.e. having removed the kinetic parameters with dropout threshold of $p < 0.65$, see synthetic data results, Supp. Mat.).

GPPM–ACP model has best performances according to both ELBO and predictive accuracy on follow-up data (Table 2).

The GPPM–Diff model produces the worst RMSE for prediction, which may be due to the unrealistic patterns of indefinite accumulation that the underlying dynamical system allows for. Table 2 also reports RMSE for prediction for each disease status. Results are consistent across diagnosis for GPPM–ACP and GPPM–RD, while for GPPM–Diff predictions deteriorate sensibly for MCI and AD subjects.

The last line of Table 2 reports results for the null model (GPPM–mono), which produces the worst ELBO. Moreover, by not estimating any dynamical system, it does not provide individual predictions on D_2 via forward integration, and so the RMSE on prediction is not provided.

Fig. 3 shows long term trajectories estimated with the three models on three sample regions of interest, precuneus, supramarginal and lingual, over the estimated long term time axis (roughly 25 years). The three regions of interest were selected to be representative of, respectively, early, intermediate and late disease progression (times in which regions reach abnormality are shown in Fig. 5). We note that while the estimated long term trajectories may appear similar across DS, the underlying dynamic properties are different by construction. Fig. 3 also shows individual measurements, colored according to the diagnosis and shifted according to the estimated time reparametrization parameters. These parameters are similar across DS, and produce similar diagnostic separation (Supp. Fig. S10). The long term trajectories of the whole set of regions are in Supp. Figs. S7,–S9 (S6 for the GPPM–mono model). We note that most regions show an initial plateau at 10 years, and only few are still accumulating amyloid in the time-span estimated by the models. Plots of the RMSE distributions for the individual predictions of the GPPM–DS models across biomarkers are in Supp. Figures S11, S12 and S13.

3.3. GPPM–ACP: kinetic parameters and propagation pathways

For each region i, the GPPM–ACP model estimates saturation (γ_i) and plateau (η_i) thresholds with associated variability (Eq. (7, Method Section 4.2). The associated times to reach saturation and plateau can be computed as t_{γ_i} such that $f_i(t_{\gamma_i}) = \gamma_i$, and t_{η_i} such that $f_i(t_{\eta_i}) = \eta_i$.

Fig. 4 A) shows GP functions along with the estimated γ_i and η_i for three brain regions (precuneus, supramarginal and lingual). The estimated times t_{γ_i} and t_{η_i} vary across regions and reflect the respective role in the process of amyloid deposition in AD. For instance, we note how the precuneus region saturates quite early, triggering propagation to connected regions. This result is compatible with previous findings in histo–pathological and imaging studies of amyloid deposition in AD (Braak and Braak, 1991; Fantoni et al., 2018; Grothe et al., 2017; Murray et al., 2015; Thal et al., 2018; 2002). The GPPM–ACP model also estimates $k_{ij_{ACP}}$, the region–wise time–dependent propagation parameters describing amyloid propagation among connected regions - see Eq. (7), Method Section 4.2. With reference to the red/purple/light blue vertical bars in Fig. 4A), Fig. 4B) shows the connectomes where the edges' colors are proportional to the values of the estimated kinetic parameters. The three regions of interest are highlighted, and we can observe that each region's propagation is triggered when the region saturates, gets stronger while propagation is ongoing, and finally reaches a plateau for the regions that reach a plateau. Fig. 4C) shows the same results as in Fig. 4B), but the connectome is visualised as an adjacency matrix, where gray/black cells correspond to present/missing connections amongst regions, while colors for the lingual, precuneus and supramarginal connections are proportional to the values of the estimated kinetic parameters. An animated version of Fig. 4 is available as Supp. data.

Table 2
RMSE mean (and std) for individual prediction on unseen data set D_2 and ELBO of the models while training on D_1. Units for the RMSE is normalised AV45-PET UPTAKE.

Model	RMSE on prediction					negative
	CN-	CN+	MCI	AD	average	ELBO
GPPM–Diff	0.38(0.20)	0.42(0.20)	0.43(0.22)	0.44(0.23)	0.42(0.22)	-10122
GPPM–RD	0.16(0.10)	0.17(0.10)	0.17(0.11)	0.17(0.11)	0.17(0.11)	-11002
GPPM–ACP	**0.13(0.09)**	**0.13(0.09)**	**0.13(0.10)**	**0.13(0.10)**	**0.13(0.10)**	**-11874**
GPPM–mono	–	–	–	–	–	-9975

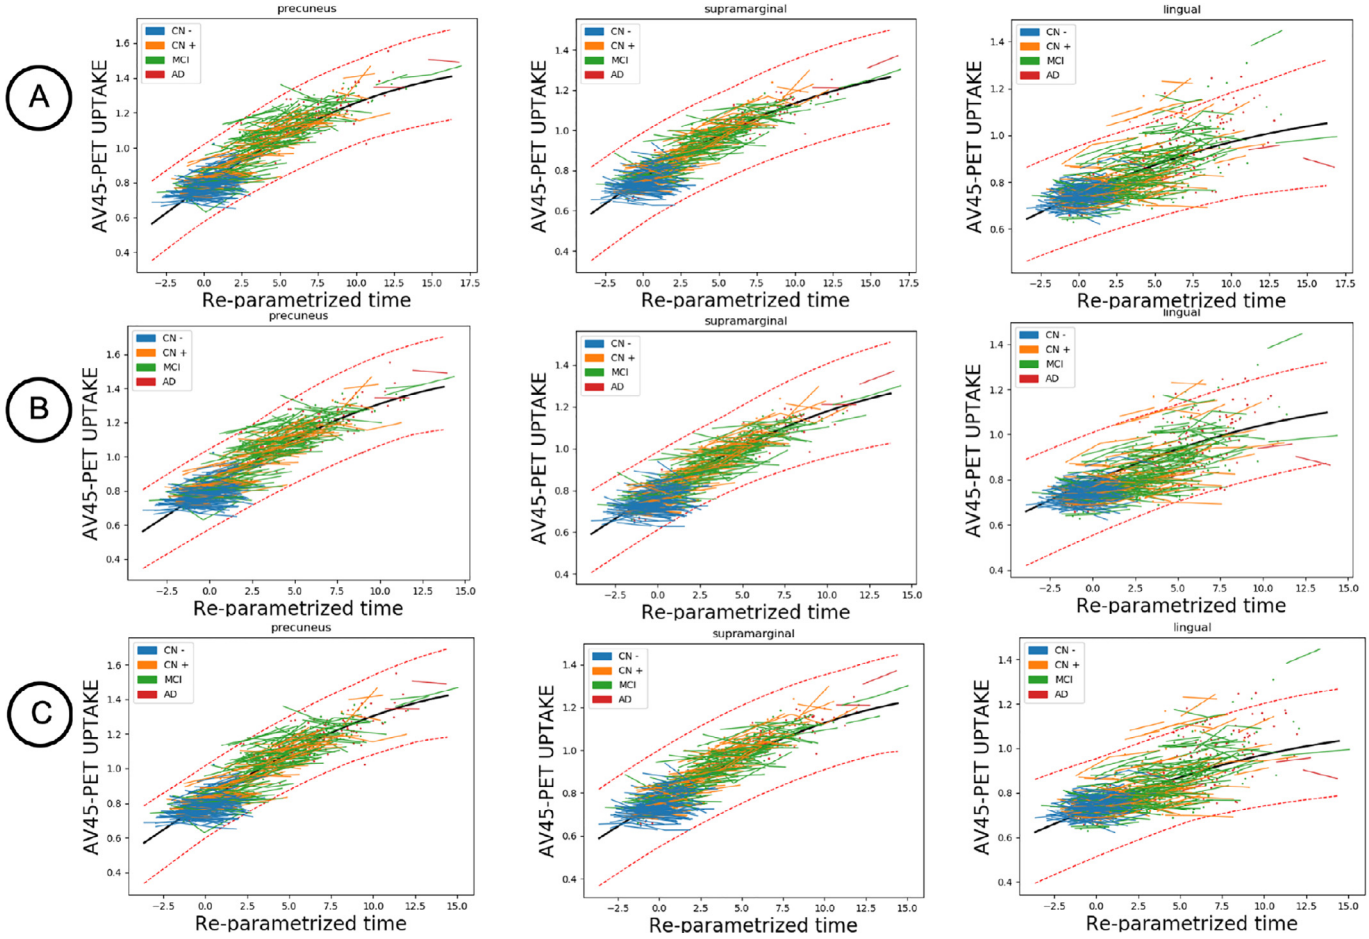

Fig. 3. Modeling amyloid deposition in Alzheimer's disease. Long term trajectories and individual short term measurements estimated with the three GPPM–DS models on three sample regions of interest: precuneus, supramarginal and lingual, selected to be representative of early, intermediate and late progression. Specifically: in A) results for GPPM–Diff of Eq. (4); B) results for GPPM–RD of Eq. (5); C) results for GPPM–ACP of Eq. (6). The black solid line represents the average trajectory, while the red dashed ones ± 3 standard deviations. In blue CN- subjects, in orange CN+ subjects, in green MCI subjects and in red AD subject.

3.4. GPPM–ACP: saturation thresholds as a proxy of time-to-abnormality

The estimated kinetic parameters can be interpreted in term of disease progression, for example in terms of ordering in which regions become abnormal. Fig. 5A) shows the regional ordering induced by the time to reach abnormality, measured as the time at which maximal separation between CN- and AD subjects was measured for that region. Fig. 5B) shows the regional ordering induced by the estimated time to reach saturation t_{γ_i}. The two orderings share a pattern compatible with previous findings in amyloid deposition, reporting frontal and parietal areas as the first regions involved in amyloid deposition (Mawuenyega et al., 2010; Rodrigue et al., 2012). This result highlights that the regions that first reaches saturation and start propagation are also the first one to show abnormality. The last involved regions are the subcortical ones, such as the thalamus (Irvine et al., 2008).

3.5. GPPM–ACP: personalisation of protein dynamics

For each subject, the DS structure of GPPM–ACP allows to integrate the dynamics over time given an initial condition (the individual baseline measurements) thus obtaining a vector field governing forward and backward evolution in time associated with the individuals. Fig. 6A) shows the estimated vector field for the dynamics associated with two regions of interest: precuneus and lingual. The field was obtained by integrating the associated ACP dynamical system while setting the other biomarkers constant to their mean values. We can appreciate the non–linear dynamics of the ACP model. Given the baseline values, we follow the predicted dynamics over time associated with the variability on the kinetic parameters. By plotting the subject's follow-up data of the testing data set D_2, we note that the model can reliably predict the unseen point. Specifics on the selected subject can be found in Supp. Table

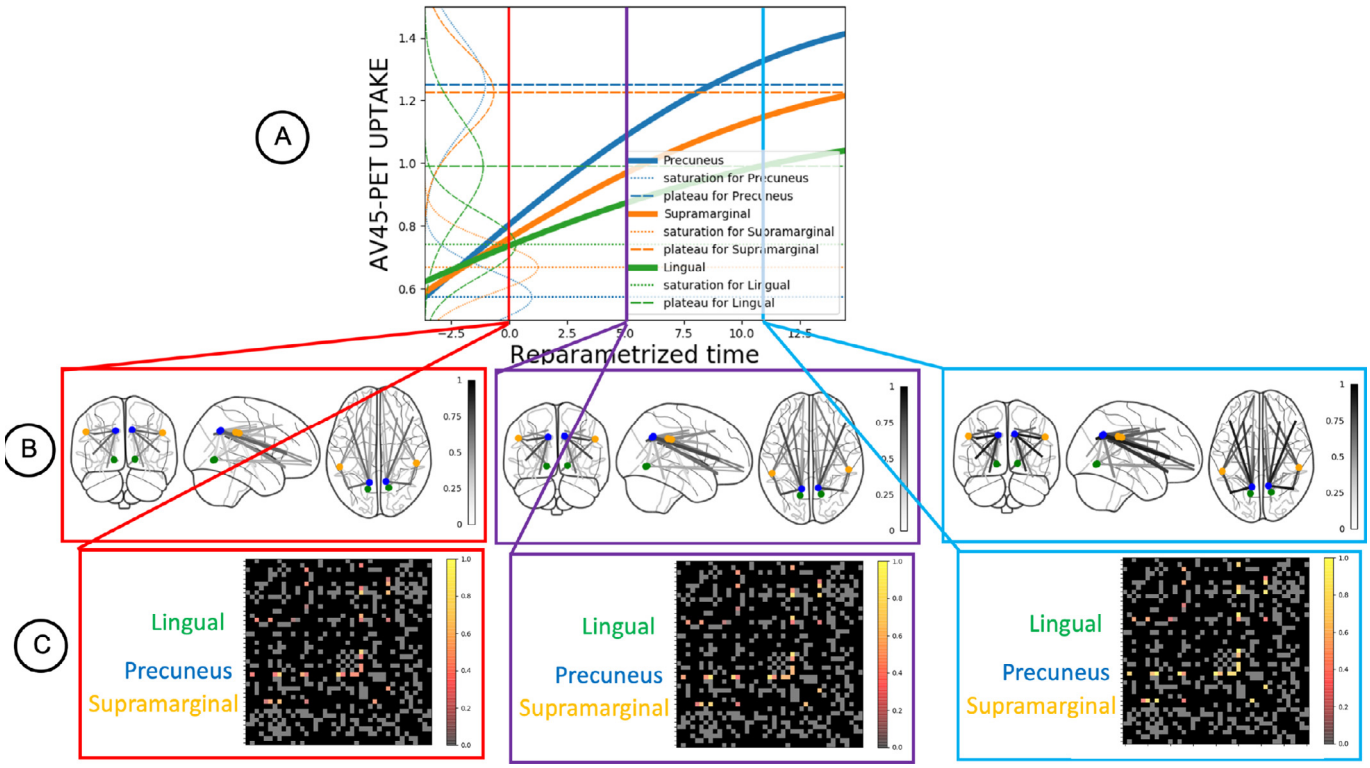

Fig. 4. Kinetic parameters and propagation pathway. A) Long term trajectories estimated by GPPM-ACP; mean saturation threshold parameters (horizontal dashed lines); mean plateau threshold parameters (horizontal dotted lines) and corresponding distributions on the y–axis; for 3 regions of interest: precuneus (in blue), supramarginal (in orange) and lingual (in green). B) Estimated time–dependent propagation parameters to/from the 3 regions of interest along the associated anatomical connections, sampled at 3 times (corresponding to the red/purple/light blue vertical bars shown in A). The colors of the edges of the connectome are proportional to the values of the estimated kinetic parameters $k_{ij_{ACP}}$, which is the region–wise time–dependent propagation parameters describing amyloid propagation among connected regions. Glass brain images obtained with Nilearn Abraham et al. (2014) (available at https://nilearn.github.io/). C) Estimated time–dependent propagation parameters to/from the 3 regions of interest overlaid on the binary adjacency matrix of the connectome, sampled at 3 times (corresponding to the red/purple/light blue vertical bars shown in A). The adjacency matrix shows black/gray entries for missing/present connections amongst regions. The colors of the entries (i, j) describing connections to/from the three regions of interest are proportional to the values of $k_{ij_{ACP}}$.

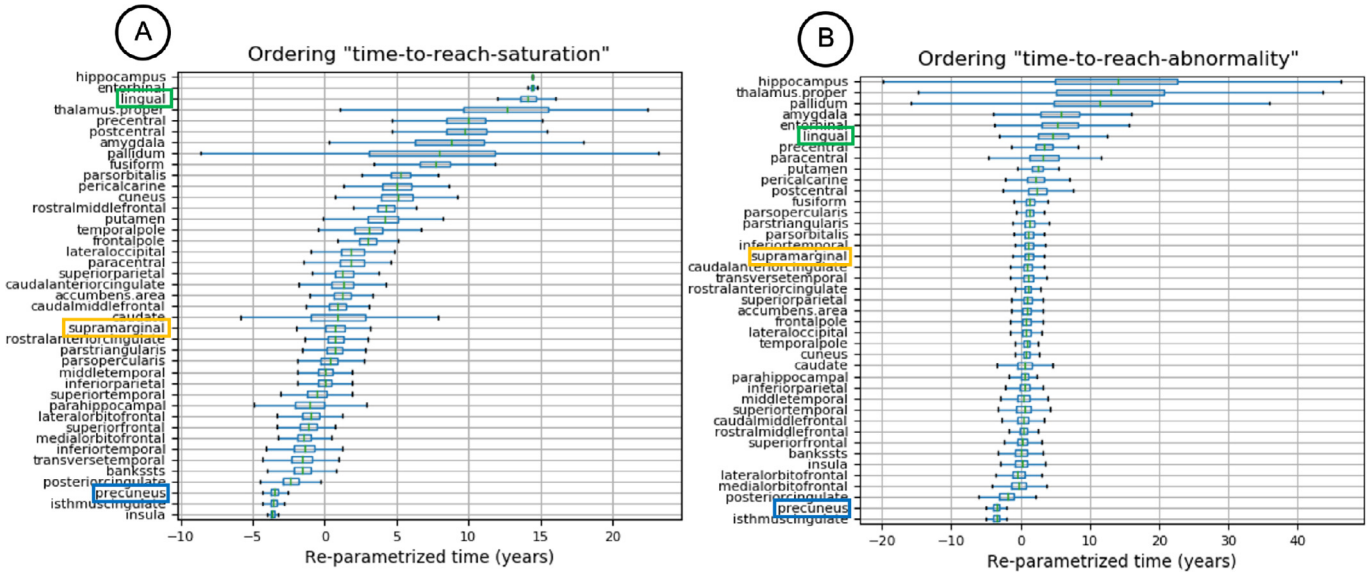

Fig. 5. Saturation thresholds as a proxy of time-to-abnormality. A) Regional ordering, according to the time to reach abnormality, measured as the time at which maximal separation between CN- and AD subjects was measured for each region, with associated variability. B) Regional ordering, according to the time to reach saturation t_{y_i} and associated variability - see Eq. (7), Method Section 4.2. In both panels, the 3 regions of interest - precuneus, supramarginal and lingual, are highlighted in blue, orange and green, respectively.

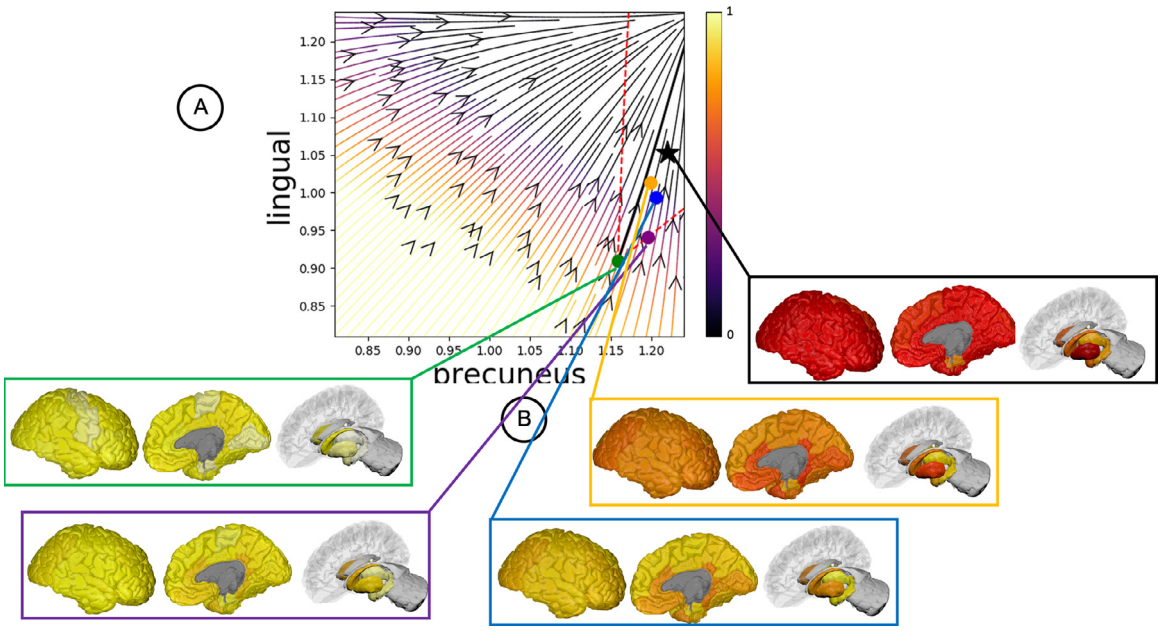

Fig. 6. Personalisation of protein dynamics. For each subject, the DS structure of GPPM–ACP allows to integrate the dynamics over time given an initial condition (the individual baseline measurements), thus obtaining a vector field governing forward and backward evolution in time associated with the individuals. A) Streamlines for the estimated amyloid deposition dynamics of precuneus and lingual regions for a sample healthy individual converted to MCI at 8 years from baseline, and subsequently to AD at 10 years. Black lines: predicted dynamics. Red dashed lines: associated variability. Colored dots: 4 observed time–points at 4, 6, 8, 10 years. Star shaped point: unseen follow–up at 11 years; B) Predicted cumulative amyloid deposition for the time-points highlighted in A). Brain images obtained with BrainPainter (Marinescu et al., 2019a) (available at brainpainter.csail.mit.edu/).

S10. Fig. 6B) shows the predicted cumulative amyloid deposition in the whole brain. An animated version of Fig. 6 is available as Supp. data.

4. Discussion

We presented a framework for spatio–temporal modeling of protein dynamics over brain networks from short term imaging data, which enables the investigation of bio–mechanical hypotheses governing disease progression via Bayesian model comparison. The framework leverages on Bayesian non–parametric regression techniques, and is coupled with a variational inference approach for scalable inference, in this way providing uncertainty quantification of the kinetic parameters governing protein dynamics and of the long term ND progression. Constraints on the protein dynamics are enforced via DS, naturally allowing prediction of protein changes over time, and so ultimately enabling realistic personalized simulation of pathological evolution.

When applied to AV45-PET brain imaging data, our framework provides new insights on the mechanisms of amyloid deposition in AD, indicating the ACP model as the most accurate dynamical system for bio–mechanical interpretation of amyloid dynamics. The model provides plausible simulations of protein propagation, and achieves accurate predictions of individual protein deposition in unseen data.

Our results further indicate that fundamentally different propagation mechanisms can be associated with similar progression patterns. This is a natural consequence of the ill-posed nature of the modeling problem tackled in this work. For this reason, the diagnostic separation along the disease progression axis is mildly affected from the choice of the underlying mechanistic process. This is confirmed by the analysis of the limit case where the model is optimized with respect to the sole monotonic constraint of the trajectories (GPPM–mono). Although not associated with any specific mechanistic hypothesis, the diagnostic separation obtained with this simple model is still very similar.

As implemented in our framework, Bayesian model comparison through the inspection of the ELBO accounts for model complexity. Indeed, highly parameterized models are automatically penalised by the Kullback-Leibler divergence term, which scales with the number of parameters. This is a direct consequence of the fact that the lower bound is a surrogate of the model evidence. To further mitigate the ill–posedness nature of the modelling problem, we also introduced a sparsity constraint through variational dropout, which we have shown to enhance model parsimony by considerably reducing the number of effective parameters. For all these reasons, our theoretical setup provided us with formal and practical guarantees against overfit. As a demonstration of this aspect, we also showed that the ELBO is associated with out of sample predictive accuracy.

We observe that, while we have implemented mechanistic models attempting to explain amyloid spread in humans based on common and ackowledged hypotheses (Bateman et al., 2006; Soto and Pritzkow, 2018), alternative hypotheses have been presented (Whittington et al., 2018), describing the distribution of amyloid in Alzheimer's disease as the result of heterogeneous regional carrying capacities.

We note that a limitation in the use of DS concerns the identifiability of the kinetic parameters, whose analysis typically requires symbolic computation (Chis et al., 2011) and can become prohibitive as the system size and non–linearity increase. In particular, while for linear dynamical systems there are well–established frameworks for establishing either structural, local, or data–driven identifiability, relying on either the power series method (Pohjanpalo, 1978) or the Laplace transform (Delbary et al., 2016), proof of identifiability for general non–linear dynamical systems is still missing (Saccomani, 2015). A number of methods has been proposed in recent years (Thomaseth and Saccomani, 2018; Villaverde et al., 2016), for analyzing identifiability of sub–classes of non–linear dynamical systems: to our knowledge however, no method can deal with non–linear system with time–varying parameters, which is, for instance, the case of the ACP model. Further, the proposed framework uniquely combines the DS with a disease progression models for estimating long term protein trajectories from short term data. This adds a second layer of parameters for which identifiability analysis is required: the individual time parameters and the parameters for the protein trajectories fit. Although we do not provide here formal investigation of iden-

tifiability of the GPPM–DS framework, the proposed Bayesian framework naturally allows for uncertainty quantification on the model parameters. Also, we leverage on variational dropout techniques (Antelmi et al., 2019; Kingma et al., 2015; Molchanov et al., 2017) for imposing parsimonious representations of the model parameters through sparsity constraints by means of the sparse GPPM–DS. From a biological perspective, imposing sparsity on the kinetic parameters implies retaining only the brain connections along which protein propagation is strong. Since brain networks have "small–world" organization, segregation and integration are balanced in such a way that connections are somehow redundant - in this way, by retaining only connections along which propagation is strong, we may be discarding weak paths while the flow of propagation could still able to reach the other regions.

We assume here propagation to take place along the structural connectome (Crossley et al., 2014; Iturria-Medina et al., 2014; Prusiner, 2012; Raj et al., 2015; Warren et al., 2013). Anatomical connectivity networks are a natural choice for propagation models as they estimate physical connections between brain regions, rather than the correlations estimated in functional (Ogawa et al., 1990) and gray matter structural covariance networks (Alexander-Bloch et al., 2013). To estimate the structural connectomes we rely on tractography, which is prone to false positive and negative connections (Maier-Hein et al., 2017; Thomas et al., 2014). Nevertheless, here we take an average connectome over multiple young and healthy subjects, which we believe can provide an accurate anatomical reference, although not accounting for the effects of the disease on the connectome itself (Oxtoby et al., 2017). Future models could approximate connectivity impairment arising from white matter damage commensurate with amyloid deposition. Further, limitations of current diffusion MRI tractography techniques makes it impossible to consider non symmetrical connectivity matrices, and as a consequence anterograde or retrograde protein propagation processes, which could potentially present different dynamics. Nonetheless, one major advantage of our framework is that it naturally allows the use of different or complementary networks that provide directionality, as fMRI or EEG, by introducing a non–symmetrical graph adjacency matrix A and a non–symmetrical propagation matrix. K_θ. The caveat is that this would increase the total number of parameters to be estimated by the model. Finally, we note that the chosen ADNI dataset is suboptimal for longitudinal amyloid data analysis, as longitudinal amyloid measurement is more stable when using a white matter reference region (Brendel et al., 2015; Chen et al., 2015; Landau et al., 2015) and using a longitudinal pipeline (Reuter and Fischl, 2011). Similarly, the chosen dataset precludes an investigation as to how resolution (i.e. number of ROIs)/atlas affects model fitting.

The ideas we propose here may apply to a much larger range of neurological diseases with proteinopathies and alternative models of propagation. The proposed framework will be used in the future to investigate a model of the natural history of AD, integrating a comprehensive panel of biomarkers: imaging, cognitive, genetic and demographic data. To this end, we aim at developing DS jointly accounting for biomarkers of protein dynamics (including AV1451–PET for Tau accumulation in AD), hypometabolism and atrophy (Abi-Nader et al., 2019). This extension would require the definition of appropriate mechanistic constraints within- and across- different protein spreading processes. Of course, the experimental validation of this extension requires the availability of sufficient longitudinal samples. Also, the GPPM framework rests on the assumption of consistent regional propagation across individuals, which is the case for amyloid protein (Grothe et al., 2017), while tau as well as MRI-derived atrophy patterns are far more heterogeneous (Murray et al., 2011; Tam et al., 2019). Future work could be devoted to study protein kinetics on more homogeneous subgroups or to explore the variability among different subgroups.

Finally, the application of the model in clinical trials, although of certain interest, falls beyond the scope of the current study. A future validation of this framework with clinical partners is currently under study.

GPPM–DS availability

GPPM–DS source code is available at https://gitlab.inria.fr/epione/GP_progression_model_V2 Data used in the preparation of this work are from ADNI. ADNI is a public-private partnership. All ADNI data are shared without embargo through the LONI Image and Data Archive (https://ida.loni.usc.edu/login.jsp) a secure research data repository. Interested scientists may obtain access to ADNI imaging, clinical, genomic, and biomarker data for the purposes of scientific investigation, teaching, or planning clinical research studies. Access is contingent on adherence to the ADNI Data Use Agreement. For up-to-date information please see http://adni.loni.usc.edu/wp-content/uploads/how-to-apply/ADNI-DSP-Policy.pdf.

Declaration of Competing Interest

The authors declare no competing interests.

Credit authorship contribution statement

Sara Garbarino: Conceptualization, Formal analysis, Data curation, Writing - original draft. **Marco Lorenzi:** Conceptualization, Resources, Writing - original draft.

Acknowledgements

SG acknowledges financial support from the French government managed by L'Agence Nationale de la Recherche under Investissements d'Avenir UCA JEDI (ANR-15-IDEX-01) through the project "AtroPro-Dem: A data-driven model of mechanistic brain Atrophy Propagation in Dementia". ML was supported by the French government, through the UCA JEDI and 3IA Côte d'Azur Investments in the Future project managed by the National Research Agency (ANR-19-P3IA-0002, ANR-19-CE45-0006).

This work is supported by the Inria Sophia Antipolis - Méditerranée, "NEF" computation cluster.

Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer's Association; Alzheimer's Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol- Myers Squibb Company; CereSpir, Inc.; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research and Development, LLC.; Johnson and Johnson, Pharmaceutical Research and Development LLC.; Lumosity; Lundbeck; Merck and Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing fuND to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Disease Cooperative Study at the University of California, San Diego. ADNI data are disseminated by the Laboratory for NeuroImaging at the University of Southern California.

The HCP project (Principal Investigators: Bruce Rosen, M.D., Ph.D., Martinos Center at Massachusetts General Hospital; Arthur W. Toga, Ph.D., University of California, Los Angeles, Van J. Weeden, MD, Martinos Center at Massachusetts General Hospital) is supported by the National Institute of Dental and Cranio-facial Research (NIDCR), the National Institute of Mental Health (NIMH) and the National Institute of

Neurological Disorders and Stroke (NIND). Collectively, the HCP is the result of efforts of co-investigators from the University of California, Los Angeles, Martinos Center for Biomedical Imaging at Massachusetts General Hospital (MGH), Washington University, and the University of Minnesota.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at 10.1016/j.neuroimage.2021.117980

References

Abi-Nader, C., Ayache, N., Robert, P., Lorenzi, M., 2019. Monotonic Gaussian process for spatio-temporal disease progression modeling in brain imaging data. Neuroimage. 2019

Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., Varoquaux, G., 2014. Machine learning for neuroimaging with scikit-learn. Front. Neuroinform. 8, 14.

Adomian, G., 1995. Fisher-Kolmogorov equation. Appl. Math. Lett. 8 (2), 51–52.

Alexander-Bloch, A., Raznahan, A., Bullmore, E., Giedd, J., 2013. The convergence of maturational change and structural covariance in human cortical networks. J. Neurosci. 33 (7), 2889–2899.

Almdahl, I.S., Lauridsen, C., Selnes, P., Kalheim, L.F., Coello, C., Gajdzik, B., et al., 2017. Cerebrospinal fluid levels of amyloid beta 1–43 mirror 1–42 in relation to imaging biomarkers of Alzheimer's disease. Front. Aging Neurosci. 9, 9.

Antelmi, L., Ayache, N., Robert, P., Lorenzi, M., 2019. Sparse multi-channel variational autoencoder for the joint analysis of heterogeneous data. June.

Bateman, R.J., Munsell, L.Y., Morris, J.C., Swarm, R., Yarasheski, K.E., Holtzman, D.M., 2006. Human amyloid-beta synthesis and clearance rates as measured in cerebrospinal fluid in vivo. Nat. Med. 12 (7), 856–861.

Braak, H., Braak, E., 1991. Neuropathological stageing of alzheimer-related changes. Acta Neuropathol. 82 (4), 239–259.

Brendel, M., Högenauer, M., Delker, A., Sauerbeck, J., Bartenstein, P., Seibyl, J., Rominger, A., Alzheimer's Disease Neuroimaging Initiative, 2015. Improved longitudinal [18f]-AV45 amyloid PET by white matter reference and VOI-based partial volume effect correction. Neuroimage 108, 450–459.

Brettschneider, J., Del Tredici, K., Lee, V.M.Y., Trojanowski, J.Q., 2015. Spreading of pathology in neurodegenerative diseases: a focus on human studies. Nat. Rev. Neurosci. 16 (2), 109.

Cauda, F., Nani, A., Manuello, J., Premi, E., Palermo, S., Tatu, K., Duca, S., Fox, P.T., Costa, T., 2018. Brain structural alterations are distributed following functional, anatomic and genetic connectivity. Brain 141 (11), 3211–3232.

Chen, K., Roontiva, A., Thiyyagura, P., Lee, W., Liu, X., Ayutyanont, N., Protas, H., Luo, J.L., Bauer, R., Reschke, C., Bandy, D., 2015. Improved power for characterizing longitudinal amyloid-beta PET changes and evaluating amyloid-modifying treatments with a cerebral white matter reference region. J. Nucl. Med. 56 (4), 560–566.

Chis, O.T., Banga, J.R., Balsa-Canto, E., 2011. Structural identifiability of systems biology models: a critical comparison of methods. PLoS One 6 (11), e27755.

Crossley, N.A., Mechelli, A., Scott, J., Carletti, F., Fox, P.T., McGuire, P., Bullmore, E.T., 2014. The hubs of the human connectome are generally implicated in the anatomy of brain disorders. Brain 137 (8), 2382–2395.

Cutajar, K., Bonilla, E.V., Michiardi, P., Filippone, M., 2017. Random feature expansions for deep Gaussian processes. In: Proceedings of the 34th International Conference on Machine Learning, 70, pp. 884–893.

Delbary, F., Garbarino, S., Vivaldi, V., 2016. Compartmental analysis of dynamic nuclear medicine data: models and identifiability. Inverse Probl. 32 (12), 125010.

Desikan, R.S., Segonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., Albert, M.S., Killiany, R.J., 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. Neuroimage 31, 968–980.

Donohue, M.C., Jacqmin-Gadda, H., Le Goff, M., Thomas, R.G., Raman, R., Gamst, A.C., Beckett, L.A., Jack Jr, C.R., Weiner, M.W., Dartigues, J.F., Aisen, P.S., 2014. Estimating long-term multivariate progression from short-term data. Alzheimer's Dementia 10 (5), S400–S410.

Editorial, 2018. Focus on neurodegenerative disease. Nat. Neurosci. 21, 1293.

Fantoni, E.R., Chalkidou, A., O'Brien, J.T., Farrar, G., Hammers, A., 2018. A systematic review and aggregated analysis on the impact of amyloid PET brain imaging on the diagnosis, diagnostic confidence, and management of patients being evaluated for Alzheimer's disease. J. Alzheimers Dis. 63 (2), 783–796.

Fornito, A., Bullmore, E.T., 2015. Connectomics: a new paradigm for understanding brain disease. Eur. Neuropsychopharmacol. 25 (5), 733–748.

Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modelling. Neuroimage 19 (4), 1273–1302.

Garbarino, S., Lorenzi, M., Alzheimer's Disease Neuroimaging Initiative, 2019. Modeling and inference of spatio-temporal protein dynamics across brain networks. In: International Conference on Information Processing in Medical Imaging. Springer, Cham, pp. 57–69.

Garbarino, S., Lorenzi, M., Oxtoby, N., Vinke, E.J., Marinescu, R.V., Eshaghi, A., Ikram, M.A., Niessen, W.J., Ciccarelli, O., Barkhof, F., Schott, J.M., Vernooij, M.W., Alexander, D.C., 2019. Differences in topological progression profile among neurodegenerative diseases from imaging data. Elife 8, e49298.

Grothe, M.J., Barthel, H., Sepulcre, J., Dyrba, M., Sabri, O., Teipel, S.J., Alzheimer's Disease Neuroimaging Initiative, 2017. In vivo staging of regional amyloid deposition. Neurology 89 (20), 2031–2038.

Hansson, O., Seibyl, J., Stomrud, E., Zetterberg, H., Trojanowski, J.Q., Bittner, T., Lifke, V., Corradini, V., Eichenlaub, U., Batrla, R., Buck, K., 2018. CSF biomarkers of Alzheimer's disease concord with amyloid-beta PET and predict clinical progression: a study of fully automated immunoassays in bioFINDER and ADNI cohorts. Alzheimer's Dementia 14 (11), 1470–1481.

Irvine, G.B., El-Agnaf, O.M., Shankar, G.M., Walsh, D.M., 2008. Protein aggregation in the brain: the molecular basis for Alzheimer's and Parkinson's diseases. Mol. Med. 14 (7–8), 451–464.

Iturria-Medina, Y., Carbonell, F.M., Sotero, R.C., Chouinard-Decorte, F., Evans, A.C., 2017. Multifactorial causal model of brain (dis)organization and therapeutic intervention: application to alzheimer's disease. Neuroimage 152, 60–77.

Iturria-Medina, Y., Sotero, R.C., Toussaint, P.J., Evans, A.C., Alzheimer's Disease Neuroimaging Initiative, 2014. Epidemic spreading model to characterize misfolded protein propagation in aging and associated neurodegenerative disorders. PLoS Comput. Biol. 10 (11), e1003956.

Jedynak, B.M., Lang, A., Liu, B., Katz, E., Zhang, Y., Wyman, B.T., Raunig, D., Jedynak, C.P., Caffo, B., Prince, J.L., 2012. A computational neurodegenerative disease progression score: method and results with the Alzheimer's disease neuroimaging initiative cohort. Neuroimage 63 (3), 1478–1486.

Jucker, M., Walker, L.C., 2013. Self-propagation of pathogenic protein aggregates in neurodegenerative diseases. Nature 501 (7465), 45.

Kingma, D. P., Ba, J., Adam: a method for stochastic optimization. arXiv preprint arXiv:1412.6980.

Kingma, D.P., Salimans, T., Welling, M., 2015. Variational dropout and the local reparameterization trick. In: Advances in Neural Information Processing Systems, pp. 2575–2583.

Landau, S.M., Fero, A., Baker, S.L., Koeppe, R., Mintun, M., Chen, K., Reiman, E.M., Jagust, W.J., 2015. Measurement of longitudinal beta-amyloid change with 18f-florbetapir PET and standardized uptake value ratios. J. Nucl. Med. 56 (4), 567–574.

Lawton, W.H., Sylvestre, E.A., Maggio, M.S., 1972. Self modeling nonlinear regression. Technometrics 14 (3), 513–532.

Lorenzi, M., Filippone, M., 2018. Constraining the Dynamics of Deep Probabilistic Models. In: Proceedings of the 35th International Conference on Machine Learning, Vol. 80, pp. 3233–3242.

Lorenzi, M., Filippone, M., Frisoni, G.B., Alexander, D.C., Ourselin, S., 2017. Probabilistic disease progression modeling to characterize diagnostic uncertainty: application to staging and prediction in Alzheimer's disease. Neuroimage 190, 56–68.

Macdonald, B., Husmeier, D., 2015. Gradient matching methods for computational inference in mechanistic models for systems biology: a review and comparative analysis. Front. Bioeng. Biotechnol. 3, 180. 2015

Maier-Hein, K.H., Neher, P.F., Houde, J.C., Côté, M.A., Garyfallidis, E., Zhong, J., Chamberland, M., Yeh, F.C., Lin, Y.C., Ji, Q., Reddick, W.E., 2017. The challenge of mapping the human connectome based on diffusion tractography. Nat Commun 8 (1), 1349.

Marinescu, R. V., Alexander, D. C., Golland, P., Brainpainter: a software for the visualisation of brain structures, biomarkers and associated pathological processes. arXiv preprint arXiv:1905.08627.

Marinescu, R.V., Eshaghi, A., Lorenzi, M., Young, A.L., Oxtoby, N.P., Garbarino, S., Crutch, S.J., Alexander, D.C., Alzheimer's Disease Neuroimaging Initiative, 2019. DIVE: a spatiotemporal progression model of brain pathology in neurodegenerative disorders. Neuroimage 192, 166–177.

Mawuenyega, K.G., Sigurdson, W., Ovod, V., Munsell, L., Kasten, T., Morris, J.C., Yarasheski, K.E., Bateman, R.J., 2010. Decreased clearance of CNS b-amyloid in alzheimer's disease. Science 330 (6012). 1774–1774

Melzer, T.R., Stark, M.R., Keenan, R.J., Myall, D.J., MacAskill, M.R., Pitcher, T.L., Livingston, L., Grenfell, S., Horne, K.L., Young, B., Pascoe, M., 2019. Beta amyloid deposition is not associated with cognitive impairment in Parkinson's disease. Front. Neurol. 10, 391.

Molchanov, D., Ashukha, A., Vetrov, D., 2017. Variational dropout sparsifies deep neural networks. In: Proceedings of the 34th International Conference on Machine Learning, 70, pp. 2498–2507. August, JMLR. org.

Müller, E.G., Edwin, T.H., Stokke, C., Navelsaker, S.S., Babovic, A., Bogdanovic, N., Knapskog, A.B., Revheim, M.E., 2019. Amyloid-beta PET'correlation with cerebrospinal fluid biomarkers and prediction of alzheimer' s disease diagnosis in a memory clinic. PLoS One 14 (8).

Murray, M.E., Graff-Radford, N.R., Ross, O.A., Petersen, R.C., Duara, R., Dickson, D.W., 2011. Neuropathologically defined subtypes of Alzheimer's disease with distinct clinical characteristics: a retrospective study. Lancet Neurol. 10 (9), 785–796.

Murray, M.E., Lowe, V.J., Graff-Radford, N.R., Liesinger, A.M., Cannon, A., Przybelski, S.A., Rawal, B., Parisi, J.E., Petersen, R.C., Kantarci, K., Ross, O.A., 2015. Clinicopathologic and 11c-pittsburgh compound b implications of ThAl amyloid phase across the Alzheimer's disease spectrum. Brain 138 (5), 1370–1381.

Näslund, J., Haroutunian, V., Mohs, R., Davis, K.L., Davies, P., Greengard, P., Buxbaum, J.D., 2000. Correlation between elevated levels of amyloid beta-peptide in the brain and cognitive decline. JAMA 283 (12), 1571–1577.

Ogawa, S., Lee, T.M., Kay, A.R., Tank, D.W., 1990. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proc. Natl. Acad. Sci. 87 (24), 9868–9872.

Oxtoby, N.P., Garbarino, S., Firth, N.C., Warren, J.D., Schott, J.M., Alexander, D.C., Alzheimer's Disease Neuroimaging Initiative, 2017. Data-driven sequence of changes to anatomical brain connectivity in sporadic Alzheimer's disease. Front. Neurol. 8, 580.

Palmqvist, S., Zetterberg, H., Blennow, K., Vestberg, S., Andreasson, U., Brooks, D.J., Owenius, R., Hägerström, D., Wollmer, P., Minthon, L., Hansson, O., 2014. Accuracy of brain amyloid detection in clinical practice using cerebrospinal fluid beta-amyloid

42: a cross-validation study against amyloid positron emission tomography. JAMA Neurol. 71 (10), 1282–1289.
Pohjanpalo, H., 1978. System identifiability based on the power series expansion of the solution. Math. Biosci. 41 (1–2), 21–33.
Prusiner, S.B., 2012. A ing role for prions in neurodegenerative diseases. Science 336 (6088), 1511–1513.
Rahimi, A., Recht, B., 2008. Random features for large-scale kernel machines. In: Advances in Neural Information Processing Systems, pp. 1177–1184.
Raj, A., LoCastro, E., Kuceyeski, A., Tosun, D., Relkin, N., Weiner, M., (Alzheimer's Disease Neuroimaging Initiative (ADNI, 2015. Network diffusion model of progression predicts longitudinal patterns of atrophy and metabolism in alzheimer's disease. Cell Rep. 10 (3), 359–369.
Raj, A., Kuceyeski, A., Weiner, M., 2012. A network diffusion model of disease progression in dementia. Neuron 73 (6), 1204–1215.
Rentz, D.M., Locascio, J.J., Becker, J.A., Moran, E.K., Eng, E., Buckner, R.L., Sperling, R.A., Johnson, K.A., 2010. Cognition, reserve, and amyloid deposition in normal aging. Ann. Neurol. 67 (3), 353–364.
Reuter, M., Fischl, B., 2011. Avoiding asymmetry-induced bias in longitudinal image processing. Neuroimage 57 (1), 19–21.
Rodrigue, K.M., Kennedy, K.M., Devous, M.D., Rieck, J.R., Hebrank, A.C., Diaz-Arrastia, R., Mathews, D., Park, D.C., 2012. B-amyloid burden in healthy aging: regional distribution and cognitive consequences. Neurology 78 (6), 387–395.
Saccomani, M.P., 2015. Identifiability of nonlinear ODE models in systems biology: results from both structural and data-based methods. In: Ortuño, F., Rojas, I. (Eds.), Bioinformatics and Biomedical Engineering. IWBBIO 2015. Lecture Notes in Computer Science, 9043. Springer, Cham.
Schiratti, J.B., Allassonniere, S., Colliot, O., Durrleman, S., 2015. Learning Spatiotemporal Trajectories from Manifold-valued Longitudinal Data. In: Advances in Neural Information Processing Systems, pp. 2404–2412.
Soto, C., Pritzkow, S., 2018. Protein misfolding, aggregation, and conformational strains in neurodegenerative diseases. Nat. Neurosci. 21 (10), 1332–1340.
Sweeney, P., Park, H., Baumann, M., Dunlop, J., Frydman, J., Kopito, R., McCampbell, A., Leblanc, G., Venkateswaran, A., Nurmi, A., Hodgson, R., 2017. Protein misfolding in neurodegenerative diseases: implications and strategies. Transl. Neurodegener. 6 (1), 6.
Tam, A., Dansereau, C., Iturria-Medina, Y., Urchs, S., Orban, P., Sharmarke, H., Breitner, J., Bellec, P., Alzheimer's Disease Neuroimaging Initiative, 2019. A highly predictive signature of cognition and brain atrophy for progression to Alzheimer's dementia. Gigascience 8 (5), giz055.
Thal, D.R., Beach, T.G., Zanette, M., Lilja, J., Heurling, K., Chakrabarty, A., Ismail, A., Farrar, G., Buckley, C., Smith, A.P., 2018. Estimation of amyloid distribution by [18 f] flutemetamol PET predicts the neuropathological phase of amyloid beta-protein deposition. Acta Neuropathol. 136 (4), 557–567.
Thal, D.R., Rub, U., Orantes, M., Braak, H., 2002. Phases of ab-deposition in the human brain and its relevance for the development of AD. Neurology 58 (12), 1791–1800.
The Alzheimer's Disease Neuroimaging Initiative, 2019. R package version 0.0.1 of ADNIMERGE: Alzheimer's disease neuroimaging.
Thomas, C., Frank, Q.Y., Irfanoglu, M.O., Modi, P., Saleem, K.S., Leopold, D.A., Pierpaoli, C., 2014. Anatomical accuracy of brain connections derived from diffusion MRI tractography is inherently limited. Proc. Natl. Acad. Sci. 111 (46), 16574–16579.
Thomaseth, K., Saccomani, M.P., 2018. Local identifiability analysis of nonlinear ODE models: how to determine all candidate solutions. IFAC-PapersOnLine 51 (2), 529–534.
Villaverde, A.F., Barreiro, A., Papachristodoulou, A., 2016. Structural identifiability of dynamic systems biology models. PLoS Comput. Biol. 12 (10), e1005153.
Villemagne, V.L., Burnham, S., Bourgeat, P., Brown, B., Ellis, K.A., Salvado, O., et al., 2013. Amyloid beta deposition, neurodegeneration, and cognitive decline in sporadic Alzheimer's disease: a prospective cohort study. Lancet Neurol. 12 (4), 357–367.
Warren, J.D., Rohrer, J.D., Schott, J.M., Fox, N.C., Hardy, J., Rossor, M.N., 2013. Molecular nexopathies: a new paradigm of neurodegenerative disease. TreND Neurosci. 36 (10), 561–569.
Weickenmeier, J., Kuhl, E., Goriely, A., 2018. The multiphysics of prion-like diseases: progression and atrophy. Phys. Rev. Lett. 121, 158101.
Whittington, A., Sharp, D.J., Gunn, R.N., Alzheimer's Disease Neuroimaging Initiative, 2018. Spatiotemporal distribution of beta-amyloid in alzheimer disease is the result of heterogeneous regional carrying capacities. J. Nucl. Med. 59 (5), 822–827.
Young, A.L., Oxtoby, N.P., Daga, P., Cash, D.M., Fox, N.C., Ourselin, S., Schott, J.M., Alexander, D.C., 2014. A data-driven model of biomarker changes in sporadic Alzheimer's disease. Brain 137 (9), 2564–2577.
Zhou, J., Gennatas, E.D., Kramer, J.H., Miller, B.L., Seeley, W.W., 2012. Predicting regional neurodegeneration from the healthy brain functional connectome. Neuron 73 (6), 1216–1227.