Discovery of Latent Clinical Phenotypes Through Optimal Transport of Brain Dynamics

Luca Taffarello[1], Federico Calesella[1], Giacomo Barzon[1], Benedetta Mariani[1], Antonio Gallo[2], Manuela Altieri[2], Alessandro Pasquale De Rosa[2], Alessandro D'Ambrosio[2], Fabrizio Esposito[2], Stefano De Marchi[1,3], Antonino Vallesi[1,4], Michele Allegra[1,5], Alessandra Bertoldo[1,6], Samir Suweis[1,5], Marco Zorzi[7,8*]

[1]Padova Neuroscience Center, University of Padova, Padova, Italy.
[2]Department of Advanced Medical and Surgical Sciences, Advanced MRI Research Center, University of Campania Luigi Vanvitelli, Naples, Italy.
[3]Department of Medicine, University of Padova, Padova, Italy.
[4]Department of Neuroscience, University of Padova, Padova, Italy.
[5]Department of Physics and Astronomy Galileo Galilei, University of Padova, Padova, Italy.
[6]Department of Information Engineering, University of Padova, Padova, Italy.
[7]Department of General Psychology, University of Padova, Padova, Italy.
[8]IRCCS San Camillo Hospital, Venice, Italy.

*Corresponding author(s). E-mail(s): marco.zorzi@unipd.it;

Abstract

The increasing availability of large clinical neuroimaging datasets has enabled the discovery of objective, data-driven clinical phenotypes. However, current approaches typically rely on structural or static functional metrics, thereby overlooking the rich information embedded in brain dynamics. To address this limitation, we developed a novel analytical pipeline integrating Gaussian Hidden Markov Modeling (HMM) with Optimal Transport (OT) theory. By representing individual brain dynamics as probability distributions over a shared latent state space, our approach quantifies the cost of transitioning between subjects' dynamic profiles, providing a sensitive metric for inter-individual variability that captures temporal aspects of brain function. We applied this framework to multiple sclerosis (MS), a condition where substantial within-group heterogeneity challenges traditional clinical classification. Analyzing resting-state functional Magnetic Resonance Imaging data from 122 relapsing-remitting MS patients and 97 matched healthy controls, HMM identified three latent states with distinct activation and connectivity patterns. While patients and controls differed in fractional occupancies and dwell times, we focused on quantifying heterogeneity specifically within the MS cohort. Using the Schrödinger bridge — a regularized optimal transport cost — we computed pairwise dissimilarities between all patients based on their dynamic profiles. Unsupervised clustering of the resulting dissimilarity matrix revealed two distinct patient subgroups that showed significant differences in clinical disability and structural pathology markers. Crucially, this stratification was not detectable when applying clustering procedures to static functional connectivity matrices, demonstrating the added value of dynamical features. Our results establish a generalizable methodological framework for extracting meaningful subtypes from temporal brain dynamics, with potential applications across neurological and psychiatric conditions characterized by phenotypic heterogeneity.

1 Introduction

Traditionally, the clinical phenotyping of neurological and psychiatric diseases relied heavily on behavioral markers. The advent of large-scale neuroimaging has since enabled the delineation of phenotype boundaries based on anatomical and functional correlates of behavioral or cognitive impairment [1, 2]. While this paradigm has yielded significant insights into psychiatric [3, 4], neurodegenerative [5–7], and neurodevelopmental diseases [8], it typically relies on either structural markers or static, time-averaged functional connectivity. By focusing on

static features, these approaches overlook the rich information embedded in dynamic activity patterns, which are known to be essential for explaining variance in behavioral traits and higher cognition [9–13]. Despite the documented alterations in brain dynamics across various conditions [14–16], the use of whole-brain dynamical markers for clinical stratification remains relatively underexplored [17].

Here, we propose a principled framework to characterize inter-individual variability in whole-brain dynamics for subject stratification. Unlike clustering methods based on engineered features, our approach provides a geometrically grounded distance metric rooted in optimal transport theory. We first employ Gaussian Hidden Markov Modeling (HMM) [18, 19] to identify a set of latent states common to all subjects. Compared to sliding-window functional connectivity [14, 20, 21], HMM offers more robust inference with lower dependence on free parameters [22] and naturally represents dynamics as Markov transitions between discrete states, yielding probability distributions suitable for transport-theoretic analysis. While latent states are identified at the group level, individual dynamics are characterized by subject-specific transition matrices. For each pair of subjects, we use the *Schrödinger bridge* method [23] to determine the *transport cost* required to map one individual's dynamics onto another's. This cost quantifies the perturbation of a subject's transition matrix needed to match another individual's state distribution, and admits a control-theoretic interpretation as the smallest input amplitude required to realize this transformation [24], situating our approach within the emerging field of control-theoretic neuroscience [25]. While previous studies applied the Schrödinger bridge technique to estimate, within individuals, the transition cost between resting and task-evoked brain activity distributions [26, 27], we extend this approach to quantify inter-individual distances using the costs to construct a pairwise dissimilarity matrix. This enables fully unsupervised clustering and the embedding of subjects within a low-dimensional dynamic feature space.

To demonstrate the clinical utility of this framework, we applied it to a large dataset of patients with relapsing-remitting multiple sclerosis (RRMS) and healthy controls. Multiple sclerosis (MS) is a chronic inflammatory, demyelinating, and degenerative disease of the central nervous system [28]. While MS is traditionally classified into three phenotypes based on symptomatology (relapsin-remitting, primary progressive, and secondary progressive), significant within-group heterogeneity highlights the need to identify biologically grounded endophenotypes [29–35]. Whereas previous dynamic connectivity studies identified reliable patient-control differences, they did not reveal pathological subtypes [36–38]. Using our optimal transport framework, we demonstrate that RRMS patients stratify into two dynamically distinct subgroups showing significant differences in clinical disability and structural pathology, a stratification invisible to static functional connectivity analysis.

2 Materials and methods

2.1 Dataset

MRI and clinical data were obtained from 122 patients with relapsing-remitting multiple sclerosis (RRMS) and 97 age- and sex-matched healthy controls (HCs), collected at the University of Campania Luigi Vanvitelli (Naples, Italy) [35]. All participants underwent neurological and neuropsychological evaluation, including the Expanded Disability Status Scale (EDSS) and the Symbol Digit Modalities Test (SDMT). MRI scans were acquired on a 3T GE Signa HDxt scanner (GE Healthcare, Milwaukee, USA) equipped with an eight-channel head coil. The imaging protocol included high-resolution 3D T1-weighted images (voxel size $= 1 \times 1 \times 1.2\ mm^3$). Resting-state fMRI data were acquired using an echo-planar imaging (EPI) sequence (TR = 1500 ms; voxel size $= 4 \times 4 \times 4\ mm^3$; 240 volumes; duration $\approx$ 6 min). Participants were instructed to remain still, awake, and with eyes closed during the rs-fMRI acquisition. Full details on participant recruitment, inclusion criteria, and MRI acquisition parameters are reported in [35].

2.2 Preprocessing of MRI data

Functional and anatomical MRI data were preprocessed using fMRIPrep 24.1.1 [39, 40]. A custom anatomical module was incorporated to generate a high-quality, template-informed brain mask. T1-weighted images were first corrected for intensity inhomogeneities, then linearly and nonlinearly aligned to the MNI152 template. The resulting transformations were inverted to project a standardized template brain mask into native space, yielding a robust subject-specific mask. T1-weighted structural images also underwent surface reconstruction and tissue segmentation with FreeSurfer. Spatial normalization to MNI152NLin6Asym and MNI152NLin2009cAsym templates was performed using nonlinear ANTs registration. Functional preprocessing included creation of a reference volume, motion correction with MCFLIRT, and boundary-based co-registration of BOLD data to the T1w reference. Several confound time series were computed, including framewise displacement (FD), global signals, and both anatomical and temporal CompCor components. Surface and grayordinate representations were generated by resampling BOLD data into fsLR space. Post-processing of fMRIPrep output was performed using the eXtensible Connectivity Pipeline-DCAN (XCP-D) [41]. BOLD data were despiked with AFNI's 3dDespike. Denoising followed the aCompCor strategy, incorporating the top five WM and CSF components, 12 motion parameters (6 motion parameters and their derivatives), and high-pass cosine regressors. High-motion volumes

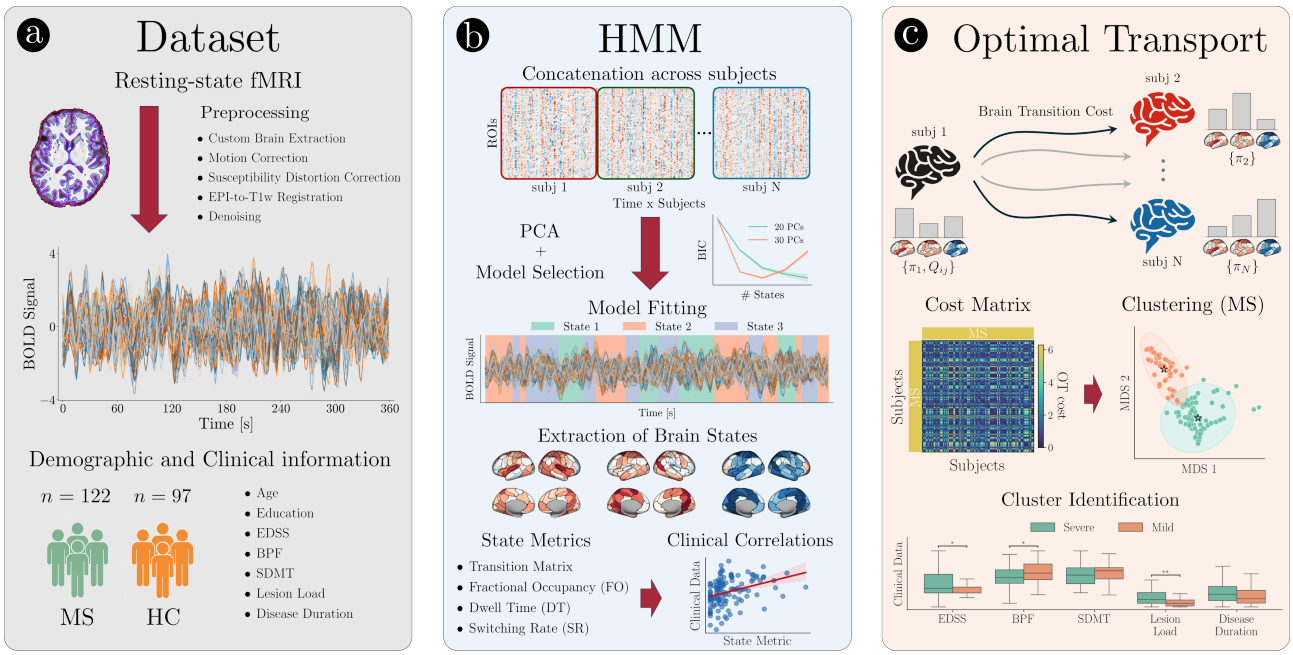

Fig. 1 **Workflow** a) *Dataset and preprocessing.* Resting-state fMRI data were collected from 122 patients with multiple sclerosis (MS) and 97 healthy controls (HC). Preprocessing included custom brain extraction, motion and distortion correction, EPI-to-T1w registration, and denoising. Demographic, clinical (EDSS, SDMT), and MRI-derived measures (BPF, lesion load) were also acquired for all participants. b) *Hidden Markov Model (HMM) inference.* ROI-wise BOLD time series were standardized and concatenated across subjects. Dimensionality was reduced via PCA, and a Gaussian HMM was fitted to the group-level data to identify recurrent brain states. Each state captured distinct spatial patterns of activity and covariance, while subject-specific state time courses were used to compute dynamic metrics including fractional occupancy, dwell time, and switching rate. These metrics were then correlated with clinical and imaging variables to evaluate their associations with disease and cognitive measures. c) *Optimal transport analysis.* For each subject, the distribution of fractional occupancies across HMM states and the empirical transition probabilities were used to estimate pairwise transition costs via the Schrödinger bridge formulation of entropy-regularized optimal transport. The resulting cost matrix quantified the dissimilarity between subjects' brain state dynamics and was used for clustering within the MS cohort. Clinical and imaging variables were subsequently compared across the identified subgroups to characterize their differences.

(FD $>$ 0.5 mm) were interpolated prior to band-pass filtering (0.01–0.1 Hz). The cleaned BOLD data were smoothed using Connectome Workbench with a 6 mm FWHM kernel. Parcel-wise time series were then extracted using the 100-parcels version of the Schaefer atlas.

2.3 Hidden Markov Model

To characterize the temporal dynamics of large-scale brain activity, we employed a Gaussian Hidden Markov Model (HMM) [18, 42]. This probabilistic framework models fMRI time series as transitions between a finite set of latent brain states. The HMM simultaneously estimates these latent states, each fully characterized by a state-specific spatial pattern of mean BOLD activity and covariance, along with the probabilities governing their temporal transitions. This approach provides a dynamic description of large-scale functional organization, capturing temporal fluctuations that are not accessible through static connectivity analyses. More explicitly, let $\mathbf{x}_t \in \mathbb{R}^D$ denote the vector of BOLD signals at time t from D regions of interest (ROIs), and let $z_t \in \{1, \ldots, K\}$ indicate the latent state active at that time. According to the HMM framework, the observed data are modeled as sequences of transitions among K hidden states, where each state defines a distinct multivariate Gaussian distribution of regional activity. The temporal evolution of states follows a first-order Markov process:

$$p(\mathbf{x}_{1:T}, z_{1:T}) = p(z_1)\, p(\mathbf{x}_1 \mid z_1) \prod_{t=2}^{T} p(z_t \mid z_{t-1})\, p(\mathbf{x}_t \mid z_t), \tag{1}$$

with $p(z_t \mid z_{t-1}) = A_{z_{t-1}, z_t}$ denoting the transition probabilities between states, and the observation model for each state k is given by

$$p(\mathbf{x}_t \mid z_t = k) = \mathcal{N}(\mathbf{x}_t \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \tag{2}$$

where $\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}_k$ represent the mean activation and covariance associated with state k. Although the HMM can be estimated separately for each participant, the limited length of typical resting-state fMRI time series provides insufficient samples to reliably estimate all model parameters. Therefore, the model was fitted on temporally concatenated data across subjects to obtain a group-level representation of the latent states. In this formulation, the parameters $(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ were estimated jointly at the group level, while the inferred state time courses, reflecting

the posterior probability of each state being active at each time point, remained subject-specific. To construct the group-level data used for model fitting, each subject's ROI-wise time series was z-scored across time (zero mean and unit variance per ROI) and then temporally concatenated across subjects to form a data matrix $\mathbf{X} \in \mathbb{R}^{(T\times S)\times D}$, where T is the number of time points per subject, S the number of subjects, and D the number of ROIs. Dimensionality reduction was performed using principal component analysis (PCA), keeping the first M components. This step reduces the number of free parameters by operating in a lower-dimensional subspace ($M < D$), improving the stability of covariance estimation and the statistical efficiency of model inference [18]. We selected $M = 30$, which explained approximately 80% of the total variance. The model parameters were estimated using the Expectation–Maximization (EM) algorithm implemented in the *Dynamax* library [43]. During the E-step, the algorithm computes the posterior probabilities of the latent states given the current parameters, and during the M-step it updates the transition matrix and the parameters $(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$. These steps are iterated until convergence of the marginal log-likelihood

$$\mathcal{L}(\Theta) = \log \sum_{z_{1:T}} p(\mathbf{x}_{1:T}, z_{1:T} \mid \Theta), \tag{3}$$

which quantifies the overall model fit. The number of hidden states (K) was determined by minimizing the Bayesian Information Criterion (BIC) across a range of values $K \in \{2, 3, 4, 5, 6\}$, yielding a shallow minimum around $K \in \{3, 4\}$ (see Fig. A1). We selected $K = 3$, as it was the most consistently optimal solution when re-fitting the HMM over a grid of (K, M) configurations, with $K \in \{2, 3, 4, 5, 6\}$ and $M \in \{10, 20, 25, 30, 35, 40, 45, 60, 80\}$, each fitted using 10 random initializations. For each model, the BIC was computed from the final log-likelihood and the number of free parameters of the Gaussian HMM in M dimensions. From the estimated model, we derived subject-specific metrics describing the temporal organization of brain-state activity: fractional occupancy (FO), mean dwell time (DT), and switching rate (SR). The FO of each state quantifies the proportion of time spent in that state and was defined as

$$\mathrm{FO}_k = \frac{1}{T} \sum_{t=1}^{T} \gamma_t(k), \tag{4}$$

where $\gamma_t(k) = p(z_t = k \mid \mathbf{x}_{1:T})$ denotes the posterior probability of state k being active at time t. For the computation of DT and SR, each time point was assigned to the most probable state,

$$\hat{z}_t = \arg\max_k \, p(z_t = k \mid \mathbf{x}_{1:T}), \tag{5}$$

and the DT of state k was defined as the average duration of time intervals during which the state remained active:

$$\mathrm{DT}_k = \frac{1}{N_k} \sum_{n=1}^{N_k} d_{k,n}, \tag{6}$$

where $d_{k,n}$ denotes the duration of the n-th interval in state k, and N_k is the total number of such intervals. The SR quantified the frequency of transitions between distinct states:

$$\mathrm{SR} = \frac{1}{T-1} \sum_{t=2}^{T} \mathbb{I}\big[\hat{z}_t \neq \hat{z}_{t-1}\big], \tag{7}$$

where $\mathbb{I}[\cdot]$ is the indicator function. Together, these measures summarize the temporal organization of the latent-state dynamics inferred by the HMM and were used to compare group-level and individual differences in subsequent analyses.

2.4 Transition cost between brain-state configurations via Schrödinger bridge

To quantify the dissimilarity between subjects in terms of their HMM-derived brain dynamics, we computed a transition cost between their respective brain state configurations using the Schrödinger bridge framework [26, 27, 44]. Each subject was characterized by a vector of fractional occupancies $\boldsymbol{\pi} = [\mathrm{FO}_1, \ldots, \mathrm{FO}_K]$, representing the empirical distribution of time spent in each latent state. This vector defines a low-dimensional representation of each subject's brain state configuration. Furthermore, for each subject, we derived the empirical joint probability matrix of consecutive states, Q_{ij}, from the sequence of most-likely states inferred by the HMM:

$$Q_{ij} = p(z_{t-1} = i, z_t = j), \tag{8}$$

thus capturing the subject-specific transition statistics. These Q_{ij} matrices were used as priors characterizing the intrinsic (or spontaneous) dynamics of each subject. Within this framework, the Schrödinger bridge formalism provides a probabilistic approach to model how the distribution of latent states of one subject $\boldsymbol{\pi}^A$ (source subject) can be transformed into another $\boldsymbol{\pi}^B$ (target subject), by optimally modulating the source's intrinsic dynamics Q_{ij}. In other words, it identifies the most likely stochastic evolution (controlled path) that connects $\boldsymbol{\pi}^A$ to $\boldsymbol{\pi}^B$, while remaining as close as possible to the spontaneous (uncontrolled) transition probabilities of the source subject. Among all possible controlled transitions, the optimal one, corresponding to the Schrödinger bridge, minimizes the Kullback–Leibler divergence between the probability of the modulated dynamics and that of the original uncontrolled dynamics. Therefore, we can define its transition cost as

$$\text{Cost} = \min_{P_{ij} \in \mathcal{U}} D_{\text{KL}}(P_{ij} \,||\, Q_{ij}), \tag{9}$$

where $\mathcal{U} = \{P_{ij} > 0 \,|\, \sum_j P_{ij} = \pi_i^A,\ \sum_i P_{ij} = \pi_j^B\}$ is the set of admissible joint transition probabilities (couplings) whose marginals match the source and target configurations. This formulation is equivalent to an entropy-regularized optimal transport problem [45]:

$$\text{Cost} = \min_{P_{ij} \in \mathcal{U}} \sum_{i,j} T_{ij} P_{ij} - H(P_{ij}), \tag{10}$$

where $T_{ij} = -\log Q_{ij}$ defines the transport cost matrix derived from the prior transition probabilities, and $H(P_{ij}) = -\sum_{i,j} P_{ij} \log P_{ij}$ denotes the entropy term. The optimal coupling P_{ij}^* was computed using the Sinkhorn algorithm [46, 47], which efficiently solves the entropy-regularized optimal transport problem. Intuitively, this optimization identifies the most likely stochastic transformation that maps one subject's brain-state configuration into another, while remaining consistent with the transition constraints inferred from the HMM. The resulting transition cost quantifies how much the intrinsic dynamics of the source subject must be altered to reproduce the target configuration: higher costs indicate that the two subjects exhibit distinct dynamical regimes in the latent state space, whereas lower costs reflect more comparable organizations of brain-state dynamics. The outcome of this framework is a cost matrix C_{ij}, where each entry represents the transition cost between the brain-state configurations of subject i and subject j, obtained as in Eq. 10.

2.5 Statistical analysis

All statistical tests were two-sided with a significance level of $\alpha = 0.05$.
State-specific functional connectivity differences (Fig. 2c) were assessed using permutation tests (10,000 resamples), using the difference in mean correlation within each network block as the test statistic, with Bonferroni correction across network pairs.
Group differences in HMM-derived metrics between MS patients and healthy controls (Fig. 3a) were evaluated using Mann-Whitney U-tests, with false discovery rate (FDR) correction across all performed comparisons. Associations between HMM-derived metrics and clinical or MRI variables within the MS cohort (Fig. 3b) were assessed using Spearman rank correlations with FDR correction.
Clinical and structural differences between OT-derived MS subgroups (*Mild* vs *Severe*; Fig. 4c) were assessed after testing normality using the Shapiro-Wilk test. Welch's t-tests were used for approximately normal variables, Mann-Whitney U-tests for non-normal or ordinal variables, and Fisher's exact test for categorical variables. FDR correction was applied across the four primary clinical outcomes (EDSS, BPF, lesion load, SDMT), whereas potential confounders (age, sex, education, disease duration) were tested separately.

3 Results

3.1 Characterization of hidden brain states

We applied Gaussian Hidden Markov Modeling (HMM) to analyze fMRI recordings of 122 RRMS patients and 97 age- and sex-matched healthy controls. The Gaussian HMM can capture group-level recurrent configurations of large-scale brain dynamics, each defined by a distinct pattern of mean activation and functional connectivity. Following common practice in the literature [48–52], we trained a single model on patients and controls to yield a common state space that faclilitates direct group comparison. Model selection identified $K = 3$ as the optimal number of states. To visualize these patterns in the original brain space, the mean activation vectors ($\boldsymbol{\mu}_k$) and covariance matrices ($\boldsymbol{\Sigma}_k$), originally estimated in a PCA-reduced space, were projected back to the original ROI space. State-specific FC matrices were then obtained by normalizing the reconstructed covariances, providing a representation of how regions co-fluctuate within each state (Fig. 2a,b). To quantitatively characterize how the connectivity patterns of each state differed from the overall functional organization across participants, we

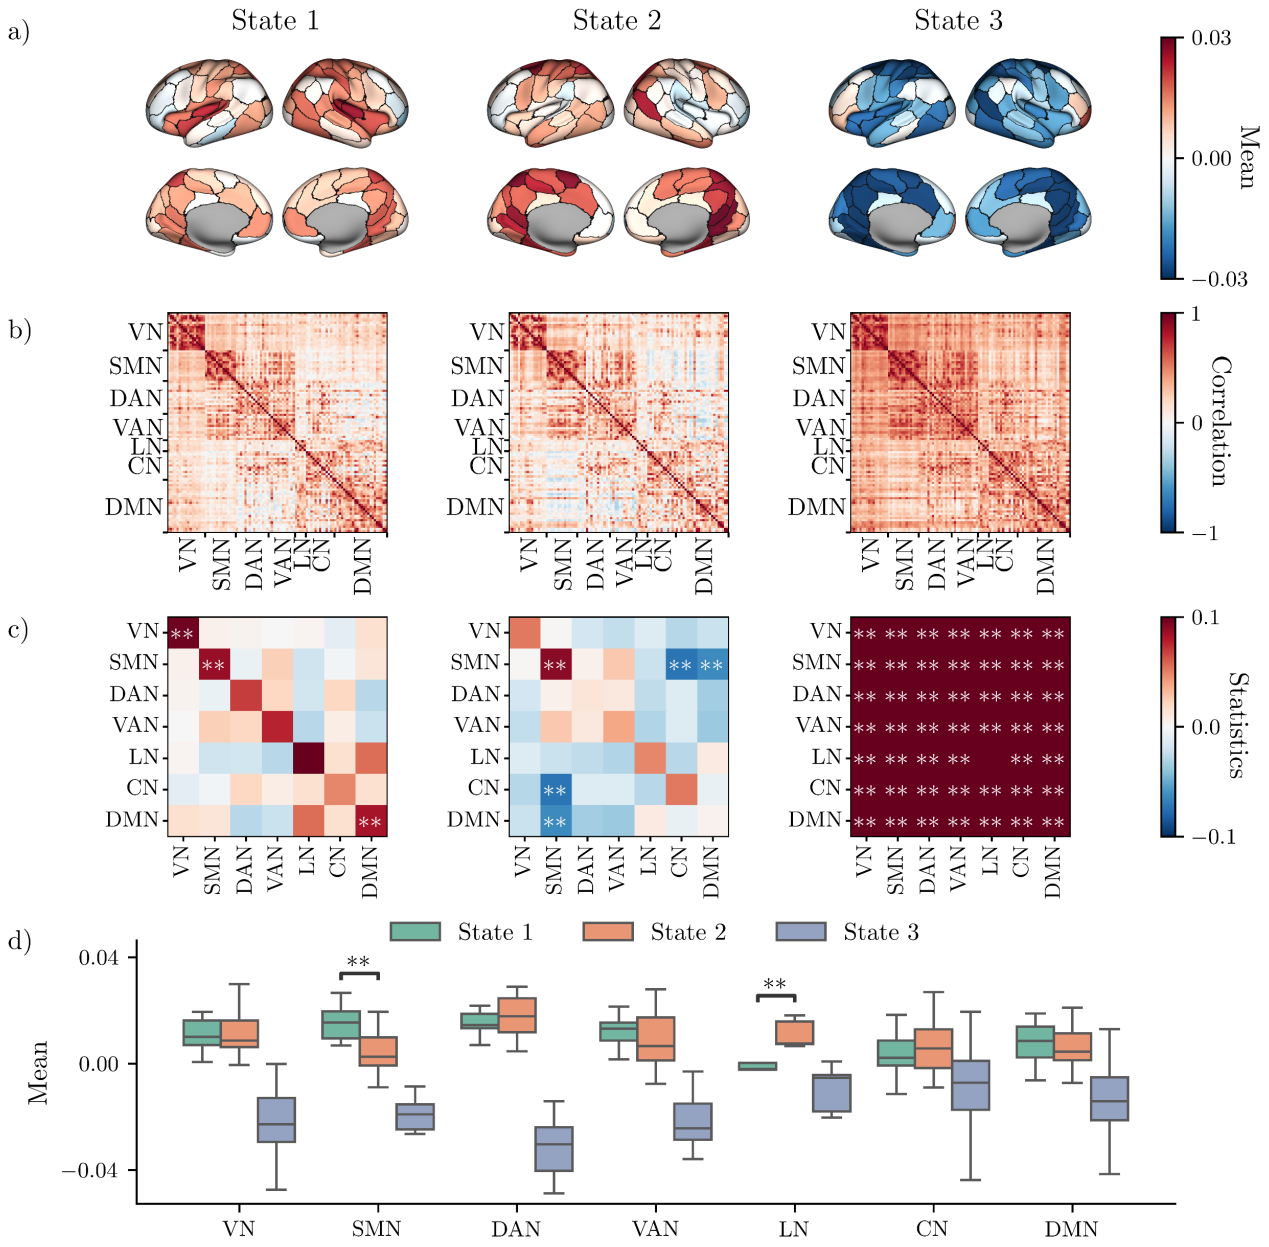

Fig. 2 HMM-derived brain states. a) *Mean activation patterns.* Group-level Gaussian HMM identified three recurrent latent states ($K = 3$), each defined by a distinct spatial pattern of mean BOLD activation ($\boldsymbol{\mu}_k$) projected back to the original ROI space. b) *State-specific functional connectivity (FC).* Covariance matrices ($\boldsymbol{\Sigma}_k$) were reconstructed and normalized to obtain FC maps describing region-to-region co-fluctuations within each state. c) *State-specific connectivity differences.* Differences between each state's FC matrix and the average FC across subjects are shown at the network level. Statistical significance was assessed using permutation tests (10,000 resamples; difference in mean connectivity strength as test statistic), with Bonferroni correction across network pairs. Positive values indicate stronger connectivity relative to the group-average FC, while negative values indicate weaker coupling. d) *Mean activation by network.* Average state activations were computed for each canonical resting-state network. Statistical comparisons (Mann–Whitney tests, FDR correction) were focused on States 1 and 2, revealing significant differences in the somatomotor (SMN) and limbic (LN) networks. Statistical significance is denoted as $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

compared the FC matrix of each state with the mean FC across all time points (Fig. 2c). For every pair of resting-state networks, we tested whether the distribution of correlation values within each block differed significantly from the corresponding block in the mean FC matrix using a permutation test (10, 000 resamples, Bonferroni correction). This procedure allowed us to identify network interactions that were significantly stronger or weaker in a given state relative to the average connectivity pattern across participants. The results revealed distinct network coupling patterns across states. State 1 showed higher within-network connectivity in the visual (VN), somatomotor (SMN), and default mode (DMN) networks compared to the mean FC, indicating a configuration with stronger functional integration within these networks. State 2 showed stronger connectivity within the SMN and reduced interactions between the SMN and higher-order networks such as the DMN and control (CN), indicating tighter intra-SMN integration together with inter-network segregation. State 3 was characterized

by broadly positive correlations across networks, consistent with a synchronized pattern and reduced network differentiation. Mean activation patterns also differed across states (Fig. 2a,d). States 1 and 2 showed overall positive activations, while State 3 displayed negative mean activations across all networks. Because States 1 and 2 also showed significant group differences in HMM-derived metrics (see Fig. 3), statistical comparisons of mean activation were focused on these two states, revealing higher activation in the SMN and lower activation in the limbic (LN) network for State 1 relative to State 2. Overall, these results highlight three distinct large-scale brain states that formed the basis for the subsequent analyses.

3.2 Group-level differences in brain-state dynamics

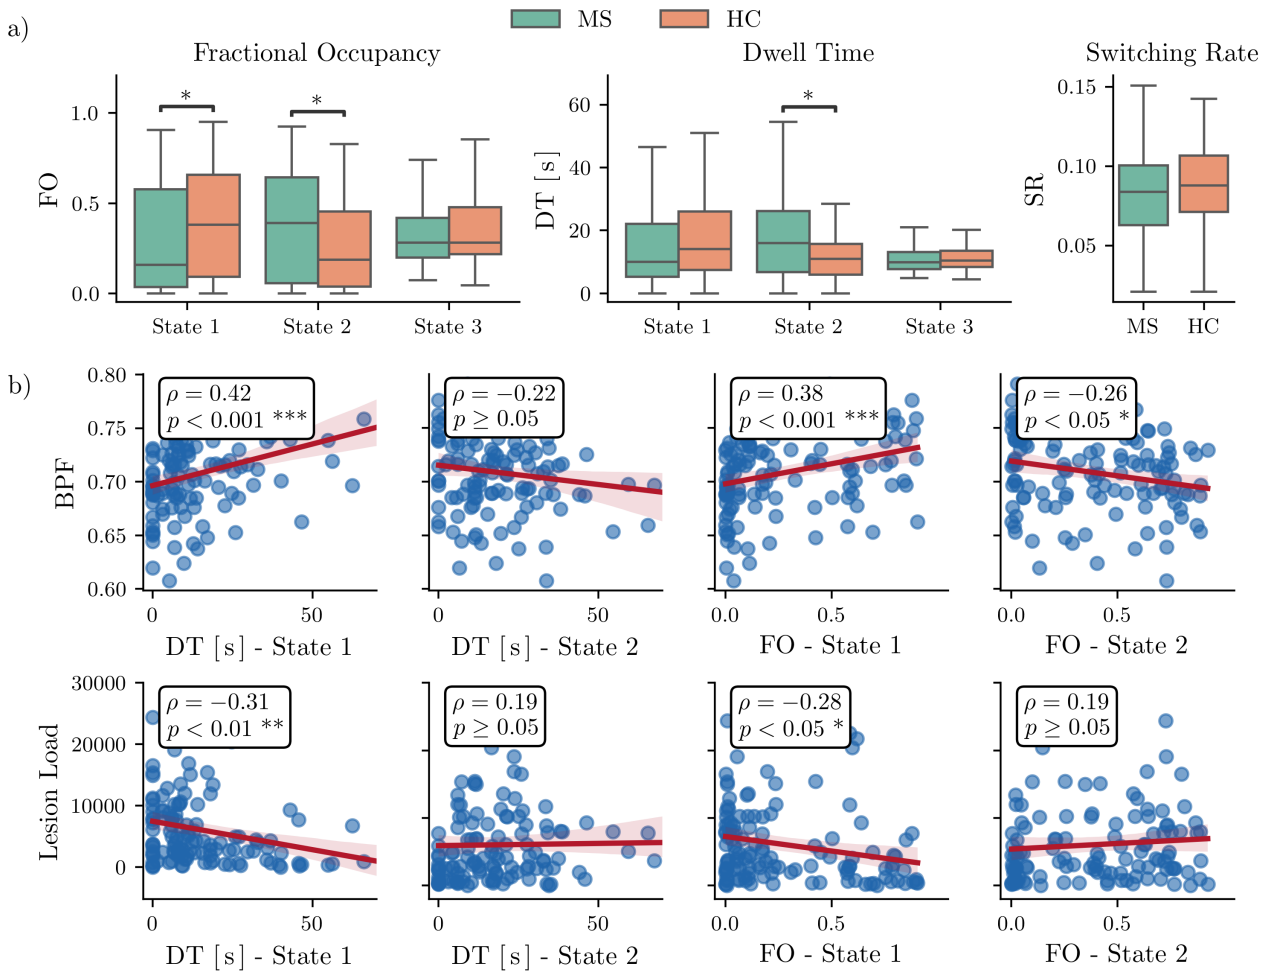

Fig. 3 HMM-derived metrics and clinical correlations in MS. a) *Group comparison of HMM-derived metrics.* Fractional occupancy (FO), dwell time (DT), and switching rate (SR) were compared between multiple sclerosis (MS) patients and healthy controls (HC). MS patients showed higher FO and DT in State 2 and lower FO and DT in State 1, while no group differences were observed for State 3 or SR. b) *Clinical correlations in MS.* Spearman correlations between HMM-derived metrics and clinical variables (EDSS, SDMT, BPF, and lesion load) were computed within the MS group, with statistical significance corrected for multiple comparisons (FDR-corrected p-values < 0.05). Higher FO and DT in State 1 correlated positively with BPF and negatively with lesion load, whereas FO in State 2 correlated negatively with BPF. Statistical significance is denoted as $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

We next compared the HMM-derived metrics between multiple sclerosis (MS) patients and healthy controls (HC) (Fig. 3). Group differences were mainly observed in States 1 and 2, while State 3 showed comparable values across groups. MS patients spent more time in State 2, as reflected by higher fractional occupancy (FO) and longer dwell time (DT), and less time in State 1, which showed lower FO and DT relative to HCs (all Mann–Whitney tests, FDR-corrected p-values < 0.05). Switching rate (SR) did not differ significantly between groups, indicating that the overall frequency of transitions between states was comparable. Considering the properties of these states (Fig. 2), this pattern indicates that MS patients tend to remain longer in a configuration of sensorimotor segregation (State 2), characterized by stronger connectivity within the somatomotor network and weaker interactions with higher-order systems such as the default mode and control networks. In

contrast, healthy controls more frequently occupy State 1, which shows stronger functional integration within visual, somatomotor, and default-mode networks. These results suggest that MS is associated with a shift toward more functionally segregated brain dynamics. To assess the clinical relevance of these alterations, we examined correlations between HMM-derived metrics and clinical measures in MS patients, including the Expanded Disability Status Scale (EDSS), Symbol Digit Modalities Test (SDMT), brain parenchymal fraction (BPF), and lesion load (Fig. 3b). We considered only correlations that remained statistically significant after correction for multiple comparisons (FDR-corrected p-values < 0.05). Higher dwell time and fractional occupancy in State 1 were positively correlated with BPF ($\rho = 0.42$, $p < 0.001$; $\rho = 0.38$, $p < 0.001$) and negatively correlated with lesion load ($\rho = -0.31$, $p < 0.01$; $\rho = -0.28$, $p < 0.05$). Conversely, fractional occupancy in State 2 correlated negatively with BPF ($\rho = -0.26$, $p < 0.05$). These correlations support the interpretation that spending less time in State 1 and more time in State 2 reflects greater structural impairment in MS, consistent with lower BPF and higher lesion load, while State 1 represents a healthier configuration more typical of controls. No significant relationships were found for State 3. Altogether, these findings support the conclusion that reduced structural integrity in MS is associated with an altered balance between integrated and sensorimotor-related brain configurations, consistent with previous evidence of disrupted large-scale network integration in MS. Overall, these dynamic markers provide a concise description of individual brain organization and form the basis for the subsequent subject-level stratification analysis.

3.3 Identification of clinical phenotypes

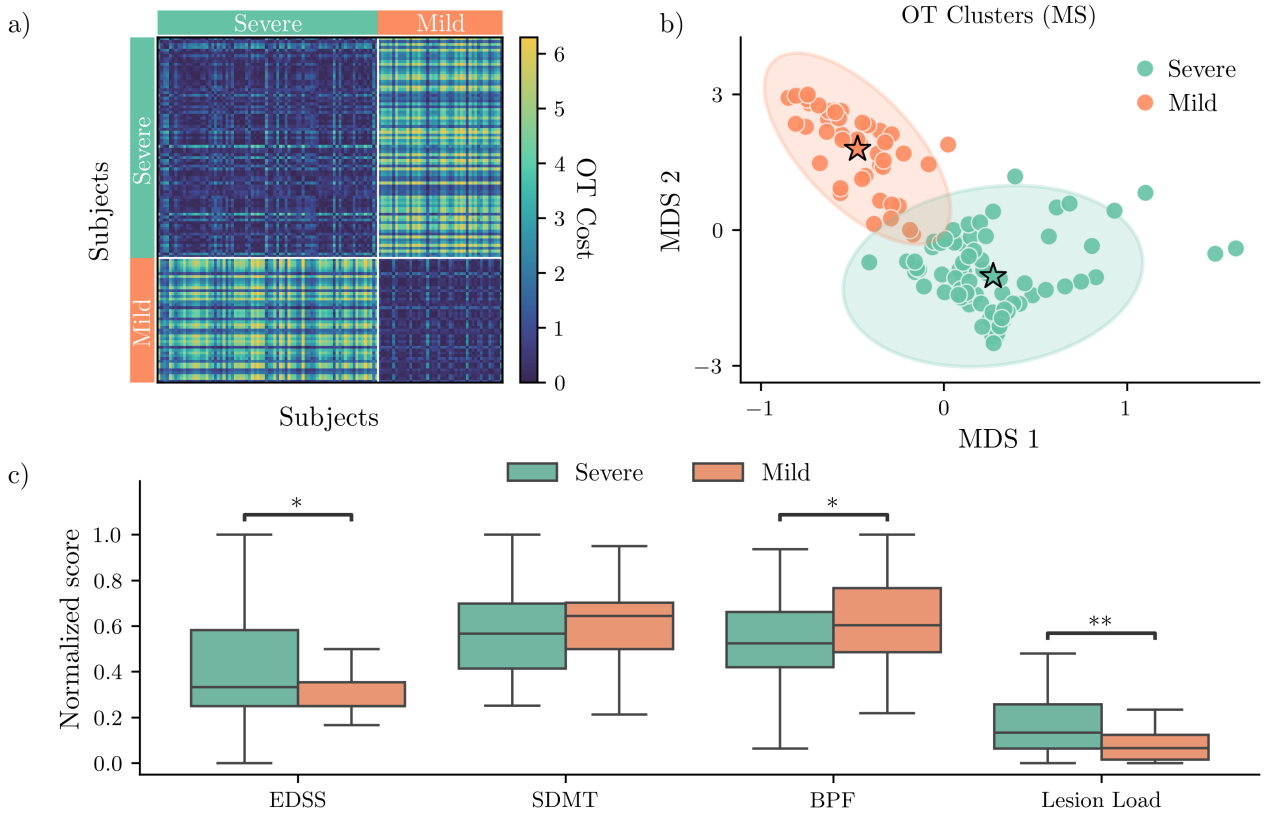

Fig. 4 Data-driven stratification of MS patients based on brain-state transition costs. a) *OT cost matrix.* Optimal Transport (OT) cost matrix derived from the Schrödinger bridge framework, representing pairwise transition costs between MS subjects based on HMM-derived brain dynamics. b) *OT-derived MS clusters.* Agglomerative clustering applied to the symmetrized OT distance matrix identifies two patient groups, visualized in a low-dimensional space using multidimensional scaling (MDS). c) *Clinical and structural group comparison.* Compared to the *Mild* group, the *Severe* group shows higher EDSS, lower brain parenchymal fraction (BPF), and greater lesion load. Statistical significance is denoted as $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$ (FDR-corrected p-values). In (a) and (b) cluster labels (*Mild* and *Severe*) are assigned a posteriori based on the differences observed in (c).

We next applied the optimal transport (OT) framework to quantify subject-level dissimilarities in brain state dynamics (Fig. 4). A pairwise cost matrix was computed using the Schrödinger bridge formulation (see Methods), which identifies the minimum modulation of HMM dynamics required to transform one subject's fractional occupancies into another's. Because these OT costs are inherently directional ($C_{ij} \neq C_{ji}$), we derived

a symmetric dissimilarity measure by averaging reciprocal entries, i.e., $\tilde{C}_{ij} = (C_{ij} + C_{ji})/2$, and used $\tilde{C}$ as input to agglomerative clustering. Although this procedure necessarily discards directional information, adopting alternative symmetrization schemes based on $\min(C_{ij}, C_{ji})$ and $\max(C_{ij}, C_{ji})$, corresponding to best- and worst-case reciprocal dissimilarities, respectively, led to identical subject partitions (ARI=1.0; see Sec. A.2). This indicates that the resulting subdivision reflects a robust global structure of OT-derived distances rather than a consequence of the specific symmetrization choice. The two-cluster solution achieved the highest silhouette score (0.69; compared with 0.60 and 0.55 for three and four clusters, respectively) and was visualized using multidimensional scaling (MDS; Fig. 4b). The resulting dimensions of MDS embedding (MDS1 and MDS2) showed strong associations with HMM-derived dynamical metrics (see Fig. A2), consistent with the interpretation that the OT embedding captures meaningful inter-individual differences in latent brain-state dynamics, while providing a low-dimensional geometric representation that integrates multiple aspects of the underlying temporal organization.

While the two groups were comparable in terms of demographic factors and disease duration, they presented different clinical and structural profiles, as revealed by key clinical measures (EDSS, SDMT) and MRI-derived structural indices (BPF, lesion load). Accordingly, the clusters were labeled *Mild* and *Severe*. Between-cluster comparisons were statistically assessed as described in Methods, with FDR correction across the four primary outcomes (EDSS, BPF, lesion load, SDMT). The *Severe* group showed higher EDSS scores ($p_{\mathrm{MW}} = 0.039$), lower BPF ($p_t = 0.022$), and greater lesion load ($p_{\mathrm{MW}} = 0.008$), while no significant difference was found in SDMT ($p_{\mathrm{MW}} = 0.11$). The two groups were comparable in age ($p_t = 0.39$), sex distribution ($p_{\mathrm{Fisher}} = 0.16$), and education level ($p_{\mathrm{MW}} = 0.33$). The *Severe* group showed a non-significant tendency toward longer disease duration ($p_{\mathrm{MW}} = 0.08$), suggesting that disease chronicity alone does not explain the observed differences in clinical and structural measures. Crucially, parallel clustering applied to standard static functional connectivity embeddings (obtained from PCA or UMAP) completely failed to yield clinically or structurally distinct subgroups (see Sec. A.4, Fig. A3). This directly underscores the added value of our OT framework, demonstrating that latent brain-state dynamics harbor critical, clinically-relevant heterogeneity that is invisible to conventional static analyses.

4 Discussion

In this study, we developed a novel framework to characterize inter-individual variability in resting-state brain dynamics and applied it to derive data-driven subtypes within relapsing-remitting multiple sclerosis (RRMS). Hidden Markov Models (HMMs) allow the identification of recurrent latent brain states as probabilistic configurations of mean activity and covariance, inferred at the group level from temporally concatenated data. This approach avoids the need to select an arbitrary window length or to cluster windowed connectivity patterns, which are well-known limitations of sliding-window dynamic functional connectivity methods [14, 21]. Moreover, HMMs provide a fully probabilistic description of brain-state dynamics, in which state representations and transition probabilities are jointly inferred within a single model that explicitly accounts for temporal dependencies [18, 42].

While HMMs have been widely used to identify recurrent brain states and characterize their temporal organization [18, 19], previous approaches have typically relied on summary dynamic metrics (e.g., fractional occupancy, dwell time, and switching rate) for group comparisons or prediction. Here, we moved beyond this feature-based description by combining a Gaussian HMM with an optimal transport (OT) formulation based on the Schrödinger bridge problem. Within this framework, we identified the minimal perturbation required to an individual's brain transition matrix in order to match another individual's brain state distribution. The resulting cost can be interpreted as a subject-level *dynamic distance*, capturing differences in the global organization of brain-state dynamics. This formulation enables subject embedding and unsupervised stratification based on inter-individual differences in latent-state transition dynamics. Our approach is consistent with emerging control-theoretic views of brain dynamics, in which transitions between large-scale brain configurations are interpreted in terms of energetic or probabilistic costs [24, 25]. More broadly, dynamic analysis pipelines have been proposed as a key approach to capture the complex nature of behavior and brain disorders [53], and the present approach contributes to this emerging landscape by providing a control-theoretic tool for dynamic phenotyping at the subject level.

Application of this framework to RRMS revealed systematic alterations in the temporal organization of brain dynamics. Compared to matched controls, patients exhibited increased occupancy of a state characterized by stronger somatomotor segregation and weaker coupling with higher-order networks, and reduced time in a more integrated configuration involving visual, somatomotor, and default-mode systems. These shifts correlated with markers of structural pathology — brain parenchymal fraction and lesion load — suggesting that white matter damage may constrain the repertoire of accessible brain states, biasing dynamics toward less integrated configurations. Such constraints on state transitions may reflect reduced network flexibility, a phenomenon that has been linked to cognitive impairment and functional disability across multiple neurological conditions [54, 55].

The observed relationship between dynamic alterations and structural damage further supports the notion that macroscale functional dynamics provide a sensitive window into the consequences of distributed tissue injury on large-scale network organization.

These findings converge with previous dynamic functional connectivity (dFC) studies in MS, which have reported alterations in dynamic network organization associated with disability and cognitive decline. Hidalgo de la Cruz et al. observed reduced dFC strength in frequently occurring states and increased connectivity within sensorimotor and visual networks in less frequent states, with progressive MS phenotypes showing more pronounced deviations and stronger associations with disability [36]. Similarly, d'Ambrosio et al. found that cognitively impaired MS patients show reduced global dynamism — characterized by fewer transitions between meta-states and shorter dwell times in highly connected configurations — that correlated with structural damage and cognitive deficits [37]. Tozlu et al. demonstrated that dFC features outperform both static functional connectivity and structural connectivity in predicting disability status [38], highlighting the added clinical value of dynamic information. The present findings extend this literature by demonstrating that similar alterations in the balance between network integration and segregation can be quantified through a transport-theoretic framework. Rather than relying on state-wise summary statistics, our approach distills these dynamic alterations into inter-individual dissimilarity metrics, providing a geometrically grounded representation of the phenomena previously linked to disability and cognitive impairment.

Crucially, the OT-based dissimilarity metric enabled the identification of two clinically distinct patient subgroups through unsupervised clustering. These subgroups differed markedly in disability severity (EDSS), brain atrophy (brain parenchymal fraction), and lesion load, yet shared the same clinical diagnosis (relapsing-remitting) and comparable demographic characteristics. This stratification was completely absent when clustering was applied to the embeddings (linear or nonlinear) of static functional connectivity matrices, demonstrating that temporal dynamics capture clinically meaningful heterogeneity invisible to time-averaged approaches. This finding has important implications for clinical research and practice. First, it suggests that dynamic brain signatures may identify patients at higher risk for more aggressive disease progression, potentially informing treatment decisions, clinical trial enrollment, or the timing of therapeutic escalation. Second, the observation that dynamically defined subtypes align with established markers of structural pathology suggests that these patterns reflect genuine neurobiological differences rather than measurement artifacts. While previous MRI-based subtyping efforts based in MS have relied primarily on structural imaging [31–35], our results indicate that dynamic functional markers provide a complementary and potentially more sensitive axis for patient stratification.

Several limitations should be considered when interpreting the present findings and in view of future developments of this framework. First, the relatively short duration of the resting-state acquisitions (approximately 6 minutes) constrains the precision of subject-specific estimates of state occupancies and transition probabilities, and consequently the accuracy of the Schrödinger-bridge distances. Higher temporal resolution acquisitions and longer resting-state runs, such as those with shorter repetition times and extended scan durations, would enhance the reliability of subject-specific transition estimates [56] and likely increase the sensitivity of this approach. Although such protocols are often challenging to implement in clinical settings, they provide an important benchmark for assessing the reliability and practical utility of dynamic phenotyping methods. Second, the cross-sectional, single-center nature of the dataset restricts the generalizability of both the inferred latent states and the derived subject embeddings. Longitudinal and multicenter studies will be essential to assess the temporal stability of dynamic distances, their sensitivity to disease progression, and their potential modulation by treatment. From a broader perspective, recent work on brain-state dynamics and dynamic computational phenotyping highlights the importance of reliability, identifiability, and trait-like stability of dynamic measures [57, 58]. Future studies should therefore explicitly quantify the test–retest reliability and psychometric properties of OT-based dynamic distances.

In summary, this study shows that combining HMM-based state estimation with a Schrödinger-bridge formulation yields a subject-level distance in brain dynamics that is sensitive to clinically and structurally relevant heterogeneity in RRMS. While several methodological and practical challenges remain, particularly concerning optimal scan duration, test-retest reliability, and generalizability, our results establish the feasibility and clinical utility of transport-theoretic dynamic phenotyping. This framework is broadly applicable to any neurological or psychiatric condition characterized by heterogeneous clinical presentation and may prove particularly valuable for stratifying patients in clinical trials, predicting disease trajectories, or identifying targets for personalized intervention. As neuroimaging datasets continue to grow in size and temporal resolution, control-theoretic approaches to brain dynamics offer a promising avenue to refine our understanding of individual differences in brain function and their relationship to clinical phenotypes.

Funding Information

Work supported by #NEXTGENERATIONEU (NGEU) and funded by the Ministry of University and Research (MUR), National Recovery and Resilience Plan (NRRP), project MNESYS (PE0000006) – "A Multiscale integrated approach to the study of the nervous system in health and disease" (DN. 1553 October 11, 2022).

Conflict of Interest Statement

The authors declare no competing interests.

Data Availability Statement

Data used in this study are available from the authors upon reasonable request. The code used for the analyses is available at https://github.com/LTaffarello/hmm-ot

References

[1] Greene, A.S., Constable, R.T.: Clinical Promise of Brain-Phenotype Modeling: A Review. JAMA Psychiatry **80**(8), 848 (2023) https://doi.org/10.1001/jamapsychiatry.2023.1419

[2] Wen, J., Antoniades, M., Yang, Z., Hwang, G., Skampardoni, I., Wang, R., Davatzikos, C.: Dimensional Neuroimaging Endophenotypes: Neurobiological Representations of Disease Heterogeneity Through Machine Learning. Biological Psychiatry **96**(7), 564–584 (2024) https://doi.org/10.1016/j.biopsych.2024.04.017

[3] Goodkind, M., Eickhoff, S.B., Oathes, D.J., Jiang, Y., Chang, A., Jones-Hagata, L.B., Ortega, B.N., Zaiko, Y.V., Roach, E.L., Korgaonkar, M.S., Grieve, S.M., Galatzer-Levy, I., Fox, P.T., Etkin, A.: Identification of a Common Neurobiological Substrate for Mental Illness. JAMA Psychiatry **72**(4), 305–315 (2015) https://doi.org/10.1001/jamapsychiatry.2014.2206

[4] Miranda, L., Paul, R., Pütz, B., Koutsouleris, N., Müller-Myhsok, B.: Systematic Review of Functional MRI Applications for Psychiatric Disease Subtyping. Frontiers in Psychiatry **12** (2021) https://doi.org/10.3389/fpsyt.2021.665536

[5] Chen, P., Yao, H., Tijms, B.M., Wang, P., Wang, D., Song, C., Yang, H., Zhang, Z., Zhao, K., Qu, Y., Kang, X., Du, K., Fan, L., Han, T., Yu, C., Zhang, X., Jiang, T., Zhou, Y., Lu, J., Han, Y., Liu, B., Zhou, B., Liu, Y.: Four Distinct Subtypes of Alzheimer's Disease Based on Resting-State Connectivity Biomarkers. Biological Psychiatry **93**(9), 759–769 (2023) https://doi.org/10.1016/j.biopsych.2022.06.019

[6] Zhou, C., Wang, L., Cheng, W., Lv, J., Guan, X., Guo, T., Wu, J., Zhang, W., Gao, T., Liu, X., Bai, X., Wu, H., Cao, Z., Gu, L., Chen, J., Wen, J., Huang, P., Xu, X., Zhang, B., Feng, J., Zhang, M.: Two distinct trajectories of clinical and neurodegeneration events in Parkinson's disease. npj Parkinson's Disease **9**(1), 111 (2023) https://doi.org/10.1038/s41531-023-00556-3

[7] Young, A.L., Marinescu, R.V., Oxtoby, N.P., Bocchetta, M., Yong, K., Firth, N.C., Cash, D.M., Thomas, D.L., Dick, K.M., Cardoso, J., van Swieten, J., Borroni, B., Galimberti, D., Masellis, M., Tartaglia, M.C., Rowe, J.B., Graff, C., Tagliavini, F., Frisoni, G.B., Laforce, R., Finger, E., de Mendonça, A., Sorbi, S., Warren, J.D., Crutch, S., Fox, N.C., Ourselin, S., Schott, J.M., Rohrer, J.D., Alexander, D.C.: Uncovering the heterogeneity and temporal complexity of neurodegenerative diseases with Subtype and Stage Inference. Nature Communications **9**(1), 4273 (2018) https://doi.org/10.1038/s41467-018-05892-0

[8] Hong, S.-J., Valk, S.L., Di Martino, A., Milham, M.P., Bernhardt, B.C.: Multidimensional Neuroanatomical Subtyping of Autism Spectrum Disorder. Cerebral Cortex **28**(10), 3578–3588 (2018) https://doi.org/10.1093/cercor/bhx229

[9] Liégeois, R., Li, J., Kong, R., Orban, C., Van De Ville, D., Ge, T., Sabuncu, M.R., Yeo, B.T.T.: Resting brain dynamics at different timescales capture distinct aspects of human behavior. Nature Communications **10**(1), 2317 (2019) https://doi.org/10.1038/s41467-019-10317-7

[10] Jia, H., Hu, X., Deshpande, G.: Behavioral Relevance of the Dynamics of the Functional Brain Connectome. Brain Connectivity **4**(9), 741–759 (2014) https://doi.org/10.1089/brain.2014.0300

[11] Vidaurre, D., Llera, A., Smith, S.M., Woolrich, M.W.: Behavioural relevance of spontaneous, transient brain network interactions in fMRI. NeuroImage **229**, 117713 (2021) https://doi.org/10.1016/j.neuroimage.2020.117713

[12] Fong, A.H.C., Yoo, K., Rosenberg, M.D., Zhang, S., Li, C.-S.R., Scheinost, D., Constable, R.T., Chun, M.M.: Dynamic functional connectivity during task performance and rest predicts individual differences in attention across studies. NeuroImage **188**, 14–25 (2019) https://doi.org/10.1016/j.neuroimage.2018.11.057

[13] Cabral, J., Vidaurre, D., Marques, P., Magalhães, R., Silva Moreira, P., Miguel Soares, J., Deco, G., Sousa, N., Kringelbach, M.L.: Cognitive performance in healthy older adults relates to spontaneous switching between states of functional connectivity during rest. Scientific Reports **7**(1), 5135 (2017) https://doi.org/10.1038/s41598-017-05425-7

[14] Preti, M.G., Bolton, T.A., Van De Ville, D.: The dynamic functional connectome: State-of-the-art and perspectives. NeuroImage **160**, 41–54 (2017) https://doi.org/10.1016/j.neuroimage.2016.12.061

[15] Valsasina, P., Hidalgo de la Cruz, M., Filippi, M., Rocca, M.A.: Characterizing Rapid Fluctuations of Resting State Functional Connectivity in Demyelinating, Neurodegenerative, and Psychiatric Conditions: From Static to Time-Varying Analysis. Frontiers in Neuroscience **13**, 618 (2019) https://doi.org/10.3389/fnins.2019.00618

[16] Filippi, M., Spinelli, E.G., Cividini, C., Agosta, F.: Resting State Dynamic Functional Connectivity in Neurodegenerative Conditions: A Review of Magnetic Resonance Imaging Findings. Frontiers in Neuroscience **13** (2019) https://doi.org/10.3389/fnins.2019.00657

[17] Chen, J., Iraji, A., Fu, Z., Duda, M., Andrés-Camazón, P., Thapaliya, B., Liu, J., Calhoun, V.D.: Dynamic Fusion of Genomics and Functional Network Connectivity in UK Biobank Reveals Schizophrenia-Related SNP Manifolds. medRxiv (2025). https://doi.org/10.1101/2024.01.09.24301013

[18] Vidaurre, D., Smith, S.M., Woolrich, M.W.: Brain network dynamics are hierarchically organized in time. Proceedings of the National Academy of Sciences **114**(48), 12827–12832 (2017) https://doi.org/10.1073/pnas.1705120114

[19] Vidaurre, D., Abeysuriya, R., Becker, R., Quinn, A.J., Alfaro-Almagro, F., Smith, S.M., Woolrich, M.W.: Discovering dynamic brain networks from big data in rest and task. NeuroImage **180**, 646–656 (2018) https://doi.org/10.1016/j.neuroimage.2017.06.077

[20] Calhoun, V.D., Miller, R., Pearlson, G., Adalı, T.: The Chronnectome: Time-Varying Connectivity Networks as the Next Frontier in fMRI Data Discovery. Neuron **84**(2), 262–274 (2014) https://doi.org/10.1016/j.neuron.2014.10.015

[21] Lurie, D.J., Kessler, D., Bassett, D.S., Betzel, R.F., Breakspear, M., Kheilholz, S., Kucyi, A., Liégeois, R., Lindquist, M.A., McIntosh, A.R., Poldrack, R.A., Shine, J.M., Thompson, W.H., Bielczyk, N.Z., Douw, L., Kraft, D., Miller, R.L., Muthuraman, M., Pasquini, L., Razi, A., Vidaurre, D., Xie, H., Calhoun, V.D.: Questions and controversies in the study of time-varying functional connectivity in resting fMRI. Network Neuroscience **4**(1), 30–69 (2020) https://doi.org/10.1162/netn_a_00116

[22] Vidaurre, D.: A new model for simultaneous dimensionality reduction and time-varying functional connectivity estimation. PLOS Computational Biology **17**(4), 1008580 (2021) https://doi.org/10.1371/journal.pcbi.1008580

[23] Chen, Y., Georgiou, T.T., Pavon, M.: On the Relation Between Optimal Transport and Schrödinger Bridges: A Stochastic Control Viewpoint. Journal of Optimization Theory and Applications **169**(2), 671–691 (2016) https://doi.org/10.1007/s10957-015-0803-z

[24] Kamiya, S., Kawakita, G., Sasai, S., Kitazono, J., Oizumi, M.: Optimal Control Costs of Brain State Transitions in Linear Stochastic Systems. Journal of Neuroscience **43**(2), 270–281 (2023) https://doi.org/10.1523/JNEUROSCI.1053-22.2022 . Chap. Research Articles

[25] Tang, E., Bassett, D.S.: Colloquium: Control of dynamics in brain networks. Reviews of Modern Physics **90**(3), 031003 (2018) https://doi.org/10.1103/RevModPhys.90.031003

[26] Kawakita, G., Kamiya, S., Sasai, S., Kitazono, J., Oizumi, M.: Quantifying brain state transition cost via

Schrödinger Bridge. Network Neuroscience **6**(1), 118–134 (2022) https://doi.org/10.1162/netn_a_00213

[27] Barzon, G., Ambrosini, E., Vallesi, A., Suweis, S.: EEG microstate transition cost correlates with task demands. PLOS Computational Biology **20**(10), 1012521 (2024) https://doi.org/10.1371/journal.pcbi.1012521

[28] Dobson, R., Giovannoni, G.: Multiple sclerosis - a review. European Journal of Neurology **26**(1), 27–40 (2019) https://doi.org/10.1111/ene.13819

[29] Chard, D.T., Alahmadi, A.A.S., Audoin, B., Charalambous, T., Enzinger, C., Hulst, H.E., Rocca, M.A., Rovira, À., Sastre-Garriga, J., Schoonheim, M.M., Tijms, B., Tur, C., Gandini Wheeler-Kingshott, C.A.M., Wink, A.M., Ciccarelli, O., Barkhof, F.: Mind the gap: From neurons to networks to outcomes in multiple sclerosis. Nature Reviews Neurology **17**(3), 173–184 (2021) https://doi.org/10.1038/s41582-020-00439-8

[30] Pitt, D., Lo, C.H., Gauthier, S.A., Hickman, R.A., Longbrake, E., Airas, L.M., Mao-Draayer, Y., Riley, C., De Jager, P.L., Wesley, S., Boster, A., Topalli, I., Bagnato, F., Mansoor, M., Stuve, O., Kister, I., Pelletier, D., Stathopoulos, P., Dutta, R., Lincoln, M.R.: Toward Precision Phenotyping of Multiple Sclerosis. Neurology: Neuroimmunology & Neuroinflammation **9**(6), 200025 (2022) https://doi.org/10.1212/NXI.0000000000200025

[31] Pontillo, G., Penna, S., Cocozza, S., Quarantelli, M., Gravina, M., Lanzillo, R., Marrone, S., Costabile, T., Inglese, M., Morra, V.B., Riccio, D., Elefante, A., Petracca, M., Sansone, C., Brunetti, A.: Stratification of multiple sclerosis patients using unsupervised machine learning: A single-visit MRI-driven approach. European Radiology **32**(8), 5382–5391 (2022) https://doi.org/10.1007/s00330-022-08610-z

[32] Eshaghi, A., Young, A.L., Wijeratne, P.A., Prados, F., Arnold, D.L., Narayanan, S., Guttmann, C.R.G., Barkhof, F., Alexander, D.C., Thompson, A.J., Chard, D., Ciccarelli, O.: Identifying multiple sclerosis subtypes using unsupervised machine learning and MRI data. Nature Communications **12**(1), 2078 (2021) https://doi.org/10.1038/s41467-021-22265-2

[33] Smets, I., Goris, A., Vandebergh, M., Demeestere, J., Sunaert, S., Dupont, P., Dubois, B.: Quantitative MRI phenotypes capture biological heterogeneity in multiple sclerosis patients. Scientific Reports **11**(1), 1573 (2021) https://doi.org/10.1038/s41598-021-81035-8

[34] de Boer, A., van den Bosch, A.M.R., Mekkes, N.J., Fransen, N.L., Dagkesamanskaia, E., Hoekstra, E., Hamann, J., Smolders, J., Huitinga, I., Holtman, I.R.: Disentangling the heterogeneity of multiple sclerosis through identification of independent neuropathological dimensions. Acta Neuropathologica **147**(1), 90 (2024) https://doi.org/10.1007/s00401-024-02742-w

[35] De Rosa, A.P., d'Ambrosio, A., Bisecco, A., Altieri, M., Cirillo, M., Gallo, A., Esposito, F.: Functional gradients reveal cortical hierarchy changes in multiple sclerosis. Human Brain Mapping **45**(6), 26678 (2024) https://doi.org/10.1002/hbm.26678

[36] Hidalgo de la Cruz, M., Valsasina, P., Sangalli, F., Esposito, F., Rocca, M.A., Filippi, M.: Dynamic Functional Connectivity in the Main Clinical Phenotypes of Multiple Sclerosis. Brain Connectivity **11**(8), 678–690 (2021) https://doi.org/10.1089/brain.2020.0920

[37] d'Ambrosio, A., Valsasina, P., Gallo, A., De Stefano, N., Pareto, D., Barkhof, F., Ciccarelli, O., Enzinger, C., Tedeschi, G., Stromillo, M.L., Arévalo, M.J., Hulst, H.E., Muhlert, N., Koini, M., Filippi, M., Rocca, M.A.: Reduced dynamics of functional connectivity and cognitive impairment in multiple sclerosis. Multiple Sclerosis Journal **26**(4), 476–488 (2020) https://doi.org/10.1177/1352458519837707

[38] Tozlu, C., Jamison, K., Gauthier, S.A., Kuceyeski, A.: Dynamic Functional Connectivity Better Predicts Disability Than Structural and Static Functional Connectivity in People With Multiple Sclerosis. Frontiers in Neuroscience **15** (2021) https://doi.org/10.3389/fnins.2021.763966

[39] Esteban, O., Markiewicz, C.J., Blair, R.W., Moodie, C.A., Isik, A.I., Erramuzpe, A., Kent, J.D., Goncalves, M., DuPre, E., Snyder, M., Oya, H., Ghosh, S.S., Wright, J., Durnez, J., Poldrack, R.A., Gorgolewski, K.J.: fMRIPrep: A robust preprocessing pipeline for functional MRI. Nature Methods **16**(1), 111–116 (2019) https://doi.org/10.1038/s41592-018-0235-4

[40] Markiewicz, C.J., Esteban, O., Goncalves, M., Poldrack, R.A., Gorgolewski, K.J.: fMRIPrep: A Robust Preprocessing Pipeline for Functional MRI. Zenodo (2026). https://doi.org/10.5281/zenodo.18248358

[41] Mehta, K., Salo, T., Madison, T.J., Adebimpe, A., Bassett, D.S., Bertolero, M., Cieslak, M., Covitz, S., Houghton, A., Keller, A.S., Lundquist, J.T., Luo, A., Miranda-Dominguez, O., Nelson, S.M., Shafiei, G., Shanmugan, S., Shinohara, R.T., Smyser, C.D., Sydnor, V.J., Weldon, K.B., Feczko, E., Fair, D.A., Satterthwaite, T.D.: XCP-D: A robust pipeline for the post-processing of fMRI data. Imaging Neuroscience **2**, 2–00257 (2024) https://doi.org/10.1162/imag_a_00257

[42] Baker, A.P., Brookes, M.J., Rezek, I.A., Smith, S.M., Behrens, T., Probert Smith, P.J., Woolrich, M.: Fast transient networks in spontaneous human brain activity. eLife **3**, 01867 (2014) https://doi.org/10.7554/eLife.01867

[43] Linderman, S.W., Chang, P., Harper-Donnelly, G., Kara, A., Li, X., Duran-Martin, G., Murphy, K.: Dynamax: A Python package for probabilistic state space modeling with JAX. Journal of Open Source Software **10**(108), 7069 (2025) https://doi.org/10.21105/joss.07069

[44] Gabriel, P., Marco, C.: Computational Optimal Transport with Applications to Data Sciences. Foundations and Trends in Machine Learning **11**(5-6), 355–607 (2019) https://doi.org/10.1561/2200000073

[45] Léonard, C.: A Survey of the Schrödinger Problem and Some of Its Connections with Optimal Transport. arXiv (2013). https://doi.org/10.48550/arXiv.1308.0215

[46] Sinkhorn, R.: A Relationship Between Arbitrary Positive Matrices and Doubly Stochastic Matrices. The Annals of Mathematical Statistics **35**(2), 876–879 (1964) https://doi.org/10.1214/aoms/1177703591

[47] Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. Advances in neural information processing systems **26** (2013)

[48] Moretto, M., Silvestri, E., Zangrossi, A., Corbetta, M., Bertoldo, A.: Unveiling whole-brain dynamics in normal aging through Hidden Markov Models. Human Brain Mapping **43**(3), 1129–1144 (2022) https://doi.org/10.1002/hbm.25714

[49] Moretto, M., Silvestri, E., Facchini, S., Anglani, M., Cecchin, D., Corbetta, M., Bertoldo, A.: The dynamic functional connectivity fingerprint of high-grade gliomas. Scientific Reports **13**(1), 10389 (2023) https://doi.org/10.1038/s41598-023-37478-2

[50] Tessadori, J., Boscolo Galazzo, I., Storti, S.F., Pini, L., Brusini, L., Cruciani, F., Sona, D., Menegaz, G., Murino, V.: Linking dynamic connectivity states to cognitive decline and anatomical changes in Alzheimer's disease. NeuroImage **320**, 121448 (2025) https://doi.org/10.1016/j.neuroimage.2025.121448

[51] Zhang, X., Yang, L., Lu, J., Yuan, Y., Li, D., Zhang, H., Yao, R., Xiang, J., Wang, B.: Reconfiguration of brain network dynamics in bipolar disorder: A hidden Markov model approach. Translational Psychiatry **14**(1), 507 (2024) https://doi.org/10.1038/s41398-024-03212-3

[52] Jespersen, K.V., Stevner, A., Kringelbach, M., Van Someren, E., Vidaurre, D., Vuust, P.: Reduced switching between brain states in insomnia: Evidence from modeling of fMRI brain dynamics. Cerebral Cortex **35**(11), 314 (2025) https://doi.org/10.1093/cercor/bhaf314

[53] Bolton, T.A.W., Morgenroth, E., Preti, M.G., Van De Ville, D.: Tapping into Multi-Faceted Human Behavior and Psychopathology Using fMRI Brain Dynamics. Trends in Neurosciences **43**(9), 667–680 (2020) https://doi.org/10.1016/j.tins.2020.06.005

[54] Polverino, A., Troisi Lopez, E., Minino, R., Liparoti, M., Romano, A., Trojsi, F., Lucidi, F., Gollo, L., Jirsa, V., Sorrentino, G., Sorrentino, P.: Flexibility of Fast Brain Dynamics and Disease Severity in Amyotrophic Lateral Sclerosis. Neurology **99**(21), 2395–2405 (2022) https://doi.org/10.1212/WNL.0000000000201200

[55] Sorrentino, P., Rucco, R., Baselice, F., De Micco, R., Tessitore, A., Hillebrand, A., Mandolesi, L., Breakspear, M., Gollo, L.L., Sorrentino, G.: Flexible brain dynamics underpins complex behaviours as observed in Parkinson's disease. Scientific Reports **11**(1), 4051 (2021) https://doi.org/10.1038/s41598-021-83425-4

[56] Fang, X., Marxen, M.: Test-retest reliability of dynamic functional connectivity parameters for a two-state model. Network Neuroscience **9**(1), 371–391 (2025) https://doi.org/10.1162/netn_a_00437

[57] Lee, K., Ji, J.L., Fonteneau, C., Berkovitch, L., Rahmati, M., Pan, L., Repovš, G., Krystal, J.H., Murray, J.D., Anticevic, A.: Human brain state dynamics reflect individual neuro-phenotypes. bioRxiv,

2023–0918557763 (2024) https://doi.org/10.1101/2023.09.18.557763

[58] Schurr, R., Reznik, D., Hillman, H., Bhui, R., Gershman, S.J.: Dynamic computational phenotyping of human cognition. Nature Human Behaviour **8**(5), 917–931 (2024) https://doi.org/10.1038/s41562-024-01814-x

Appendix A Supplementary Materials

A.1 Model selection

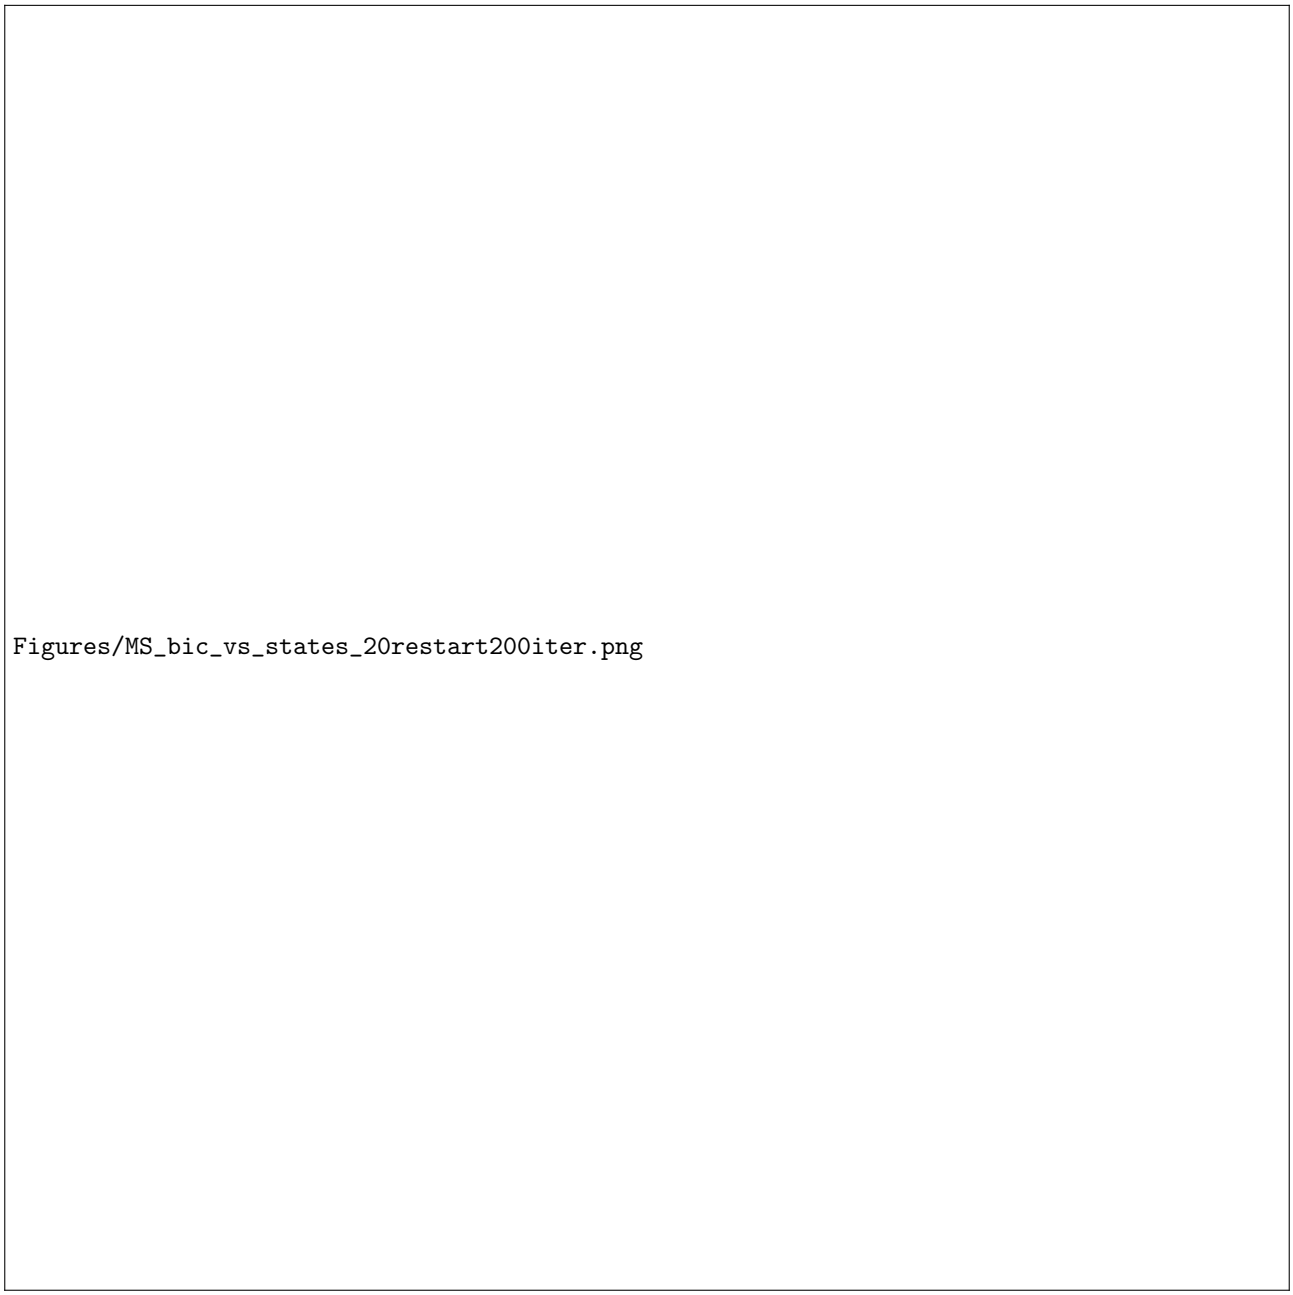

Fig. A1 Model Selection Bayesian Information Criterion (BIC) as a function of the number of hidden states (K) for different numbers of retained PCA components (M). Gaussian HMMs were fitted for $K \in \{2, 3, 4, 5, 6\}$ and $M \in \{10, 20, 25, 30, 35, 40, 45, 60, 80\}$, with 10 random initializations per configuration. Points indicate the mean BIC across initializations, and shaded areas denote $\pm$ one standard deviation.

A.2 Robustness of OT-based clustering to matrix symmetrization

Because Schrödinger bridge OT costs are directional, the resulting subject-by-subject cost matrix is asymmetric. To quantify the degree of asymmetry, we computed the symmetry degree

$$1 - \frac{\|C - C^\top\|_F}{\|C\|_F},$$

where $\|\cdot\|_F$ denotes the Frobenius norm. In MS patients this yielded a value of 0.689, indicating modest asymmetry. Prior to clustering, we therefore derived a symmetric dissimilarity matrix by averaging reciprocal entries,

$\tilde{C}_{ij} = (C_{ij} + C_{ji})/2$. To assess whether the clustering outcome depended on the specific symmetrization strategy, we repeated the agglomerative clustering using two alternative symmetric constructions, namely $\min(C_{ij}, C_{ji})$ (best-case) and $\max(C_{ij}, C_{ji})$ (worst-case) reciprocal dissimilarities. Across all three approaches (mean, min, max), silhouette analysis consistently selected a two-cluster solution (all ≈ 0.69), and subject assignments were identical (Adjusted Rand Index = 1.0). Despite this identical partitioning, local geometric properties differed across symmetrizations. Specifically, the maximum relative change in individual pairwise distances reached 63%, and only 61% (mean vs. max) and 42% (mean vs. min) of subjects preserved identical sets of their five nearest neighbors, indicating substantial local rearrangements of neighborhood structure. In contrast, the global geometry of the distance matrices was highly preserved, with correlations between upper-triangular entries exceeding 0.98 for mean vs. max and mean vs. min (and 0.96 for max vs. min). Because average-linkage agglomerative clustering is primarily driven by average inter-cluster distances rather than by local neighborhood relations, this preserved global structure explains why final partitions were obtained across symmetrization schemes. Consistently, inter-cluster mean distances were rescaled but not structurally altered (mean: 2.97; max: 3.38; min: 2.57). Together, these analyses demonstrate that the identified MS subgroups are driven by stable large-scale organization in OT-derived dissimilarities and are robust to the specific choice of symmetrization, despite modest directional asymmetry and local geometric variations.

A.3 Interpretation of OT-MDS dimensions

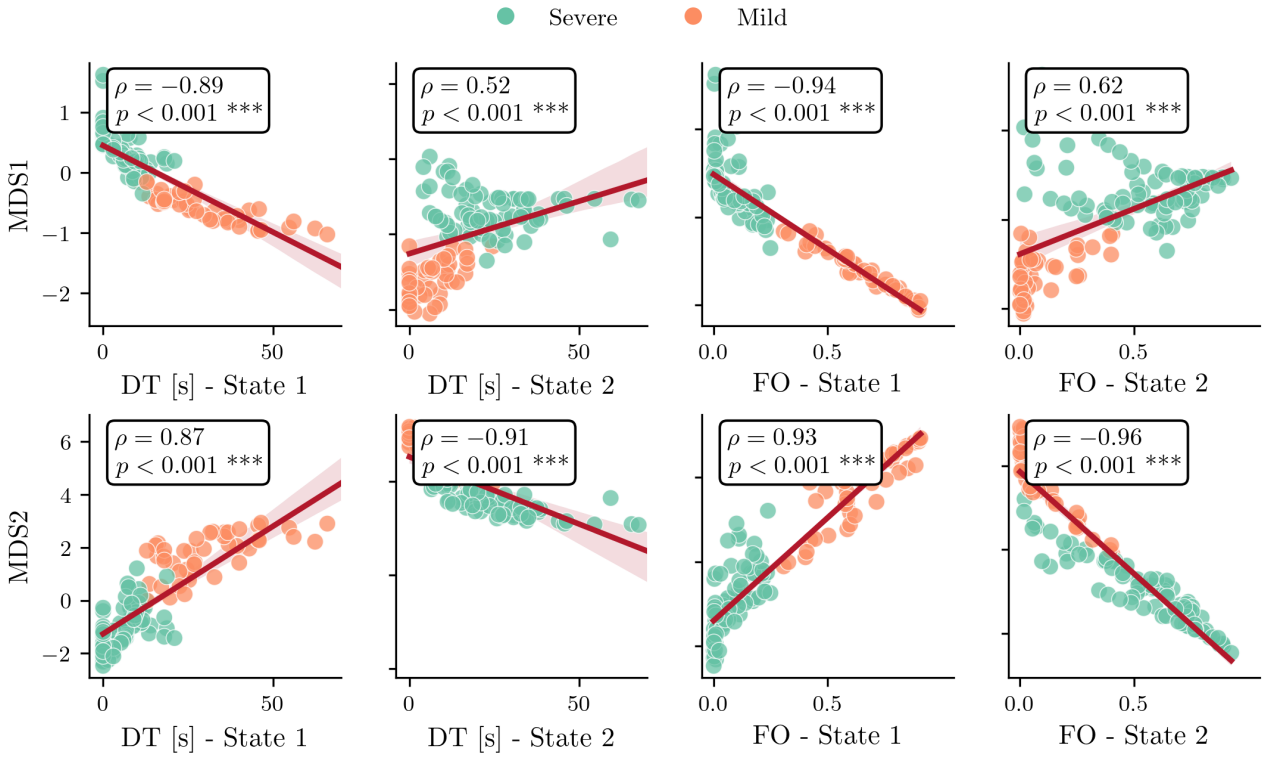

Fig. A2 Associations between OT-derived MDS dimensions and HMM-derived dynamical metrics in MS. Spearman correlations between the first two multidimensional scaling dimensions (MDS1 and MDS2), obtained from the OT-based subject dissimilarity matrix and HMM-derived dynamical metrics, within the MS group. Only associations surviving multiple-comparison correction (FDR-corrected p-values < 0.05) are shown. Both MDS dimensions exhibited strong correlation with FO and DT in States 1 and 2, indicating that the low-dimensional OT embedding captures meaningful inter-individual variability in latent brain-state dynamics. Subjects are color-coded based on their cluster assignment (Mild vs Severe). Statistical significance is denoted as $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

A.4 FC-based clustering

As a comparison to the OT-based stratification, we repeated the clustering analysis using static functional connectivity (FC). For each MS patient, whole-brain FC matrices were vectorized and assembled into a subject-by-feature matrix, followed by dimensionality reduction using principal component analysis (PCA). Agglomerative clustering was then applied to the PCA representations, and a two-cluster solution was selected for consistency with the OT-based analysis and guided by silhouette scores. The resulting PCA-derived clusters

were compared in terms of clinical and structural measures, and no significant group differences were observed. Comparable results were obtained across a range of PCA dimensionalities, suggesting that the limited clinical separation was not driven by a specific choice of embedding dimension. We further evaluated a nonlinear embedding of static FC features using UMAP, followed by agglomerative clustering into two groups. Because UMAP embeddings can depend on initialization, we also used a consensus strategy in which UMAP embedding and clustering were repeated across multiple random runs; the resulting co-association matrix was used to derive stable consensus labels. Also in this case, group comparisons did not reveal significant differences in clinical or structural measures. Overall, these analyses suggest that clustering based on static FC provides weaker clinical separation than approaches that incorporate the temporal organization of brain-state dynamics.

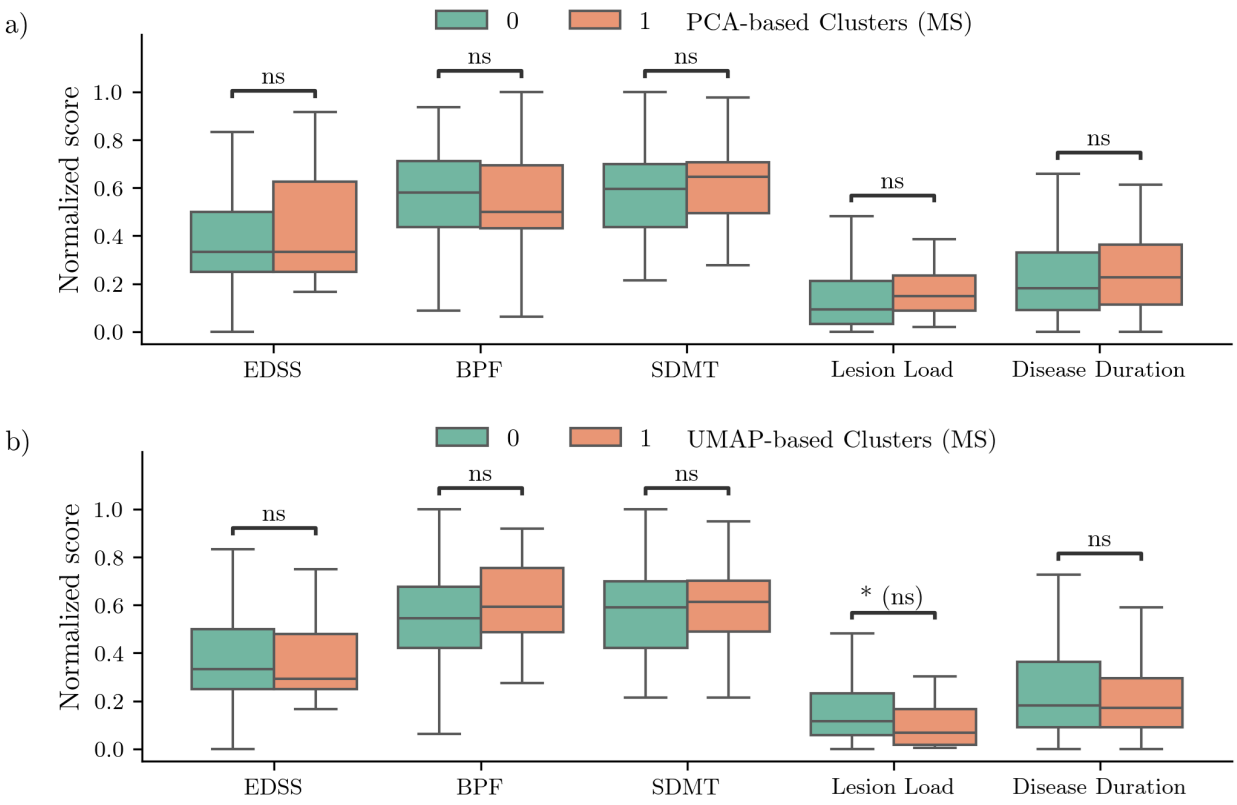

Fig. A3 Static FC-based clustering of MS patients. Two-group agglomerative clustering based on PCA (a) and UMAP embeddings (b) of static FC features in MS patients. Boxplots summarize clinical and structural measures across clusters. No comparisons remain significant after FDR correction.Asterisks (if present) denote uncorrected effects, whereas (ns) indicates lack of significance after FDR correction. Statistical significance is denoted as $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.