

STAMP: Subtype-Aware Spatio-Temporal Atlases for Personalized Prediction of Longitudinal Alzheimer’s Disease Progression

John Kalkhof^[0000–0001–7316–1903], Marco Lorenzi^[0000–0003–0521–2881]

Inria Center at University Côte d’Azur
Sophia Antipolis, France
john.kalkhof@inria.fr

Abstract. Automatic disentanglement of progression subtypes in neurodegenerative disorders is critical for patient stratification in clinical management and trial enrichment. Existing disease subtyping methods operate exclusively on low-dimensional scalar biomarkers, discarding the fine-grained spatial patterns necessary for biologically meaningful subtype discovery and patient-level personalization. Conversely, continuous-time models that directly operate on longitudinal 3D brain MRI recover a single shared population trajectory and perform no subtyping, critically limiting personalized prediction and clinical interpretability. To fill this gap, we present STAMP, a longitudinal progression subtyping framework that, integrates automatic subtype discovery directly into the optimization of a continuous-time spatio-temporal atlas from longitudinal 3D T1 brain MRI. STAMP jointly learns subtype-specific temporal placements, dense image deformation fields, and region-wise progression functions, enabling probabilistic subtype assignment and personalized trajectory prediction for new subjects. Evaluated on longitudinal ADNI-1/GO/2 data, STAMP achieves the lowest forecasting error in 16 of 18 anatomical regions and statistically significant improvements in 11 regions compared to single-trajectory and post-hoc subtyping baselines, while recovering clinically interpretable subtype structure.

Keywords: Longitudinal Disease Progression · Subtyping · Spatio-temporal Atlases

1 Introduction

The automatic disentanglement of disease progression subtypes in neurodegenerative disorders such as Alzheimer’s disease is critical for patient stratification in clinical management and trial enrichment strategies [6]. Disease heterogeneity has been repeatedly identified as a key factor underlying variable treatment response, including in recent anti-amyloid trials, highlighting the need for models that explicitly account for multiple progression patterns within a population [4].

A first line of work addresses this challenge through trajectory-based subtyping methods operating on low-dimensional biomarkers. The seminal Subtype and Stage Inference (SuStaIn) framework models phenotypic and temporal

heterogeneity jointly by inferring discrete disease subtypes and subject-specific stages from scalar biomarker measurements [25]. Subsequent extensions, including Bayesian continuous-trajectory clustering approaches, further advanced probabilistic modeling of subtype structure and disease staging [17, 24]. However, these methods rely on predefined scalar summaries rather than learning progression directly in high-dimensional image space, thereby discarding fine-grained spatial patterns relevant to subtype differentiation.

In parallel, continuous-time models have been developed directly on longitudinal 3D brain MRI. For example, recent approaches combine variational autoencoders or latent diffusion models with longitudinal trajectory modeling to generate spatiotemporally consistent predictions of future brain states at the individual level [16, 20, 22]. These methods demonstrate strong performance for prognostic imaging and anatomical synthesis. However, existing image-based progression models typically learn a single shared population trajectory, with inter-subject variability captured through continuous latent representations. While such latent variability encodes subject-specific deviations, it does not explicitly disentangle multiple structured progression subtypes with distinct disease timelines, nor does it provide interpretable subtype assignments.

Taken together, current approaches either (i) perform subtype discovery using low-dimensional biomarkers, or (ii) model full spatiotemporal image trajectories without explicitly accounting for heterogeneous progressions. To our knowledge, a unified framework that jointly learns subtype structure and continuous-time progression directly in high-dimensional image space remains lacking.

In this work, we introduce STAMP, a longitudinal progression subtyping framework that integrates automatic subtype discovery into the optimization of a continuous-time spatiotemporal atlas from longitudinal 3D T1-weighted brain MRI. STAMP jointly estimates: (1) subtype membership, (2) subject-specific temporal positioning along a subtype-specific disease timeline, and (3) dense 3D deformation fields encoding subtype-specific morphological trajectories. The model can be interpreted as a mixture of continuous-time atlas trajectories, where each subtype is associated with distinct temporal dynamics while preserving the geometric interpretability of deformation-based atlas models. Unlike post-hoc clustering strategies, subtype structure is learned end-to-end within the progression model itself. Given a new subject, STAMP infers subtype membership and generates a personalized future trajectory conditioned on the inferred subtype. Experiments on longitudinal ADNI-1/GO/2 data show that STAMP achieves the lowest forecasting error as compared to the baselines, with statistically significant improvements compared to single-trajectory baselines, while recovering clinically interpretable subtype structure.

We make our framework available under gitlab.inria.fr/jkalkhof/stamp.

2 Methodology

Given time-series of structural brain images $(x_{t_1}^j, \dots, x_{t_{N_j}}^j)$ for individuals $j = 1, \dots, N$ and time points $t_1, \dots, t_{N_j}$ our goal is to model subtypes of longitu-




Fig. 1. (A) Baseline clinical scores and relative regional volumes are used to infer subtype weights; (B) subtype-specific progression functions define regional disease dynamics along the shared time axis; (C) the resulting mixed progression signal conditions the deformation module to predict subject-specific anatomical change; (D) time-binned balancing constrains subtype usage across disease stages; and (E) entropy regularization enforces confident subtype separation.

dinal spatio-temporal image progressions in a continuous-time atlas space. Our framework is designed to account for temporal changes at multiple spatial scales: i) dense longitudinal deformation fields at the voxel-level are parameterized by ii) temporal trajectories of regional brain changes associated to different disease subtypes. Finally, for each individual, STAMP predicts a personalized image evolution based on the associated most likely cluster. Therefore, our full framework jointly estimates time-axis alignment, subtype-specific progression functions, and subject-level subtype assignments.

2.1 Atlas Backbone

We adopt a continuous-time spatio-temporal atlas backbone similar to SMART [9] and extend it to model heterogeneous subtype-specific trajectories. Concretely, we consider a multi-scale formulation of longitudinal progression defined by (i) a shared disease time axis, (ii) subtype-specific regional progression functions, and (iii) local spatio-temporal deformation field.

Estimating progression time axis and regional progressions. For each brain region k and subtype s , STAMP learn continuous progression functions $g^s(t) \in \mathbb{R}_{\geq 0}^K$, along the disease time axis, where $g_k^s(t)$ denotes the instantaneous rate of change for region k at disease time t for subtype $s = 1, \dots, C$.

Similarly to previous disease progression modeling approaches [12, 24], the estimation of the progressions $g^s(t)$, parameterized as neural ODEs [3], is performed by optimizing individualized time-shift parameters v^j to account for individual variability in progression timing. In particular, the time-shift parameter is here estimated from input relative regional volumes (Figure 1B). STAMP therefore assumes that each disease subtype evolves according to the specific trajectory $g^s(t)$, where individualized progressions are captured by the subject-specific time-shift v^j . The neural ODE parameterization allows these subtype-specific progression functions to capture non-linear temporal dynamics.

For a forecasting interval δt , the regional change for individual j and cluster s is therefore computed as

$$I^j(\delta t; s, v^j) = \int_{v^j}^{v^j + \delta t} g^s(t) dt. \quad (1)$$

Estimating local spatio-temporal deformations. The mixed regional progression signal I^j is projected via the baseline segmentation to form a dense growth map, which conditions a deformation model based on M3D-NCA [8, 10]. Neural Cellular Automata (NCA) [13] are local, recurrent convolutional models that iteratively update voxel states, enabling pixel-precise and spatially conditioned deformation modeling [21]. The network predicts a stationary velocity field representing subject-specific anatomical evolution over δt (Figure 1C), which is integrated to obtain a diffeomorphic displacement.

Gating estimator. To effectively estimate progression subtypes, gating weights $\theta^j \in \Delta^{C-1}$ are inferred for each individual j from the input baseline anatomy and clinical information (Figure 1A). The output of the gating network are subject-specific weights producing the final individualized convex combination of regional progression signal:

$$I^j = \sum_{c=1}^C \theta_c^j I^{j,(c)}. \quad (2)$$

Within this formulation, we thus let the mixed progression signal I^j to condition the deformation module to produce subject-specific anatomical forecasts. All components, including subtype-specific trajectories, time-axis alignment, and gating weights, are trained jointly end-to-end.

2.2 Learning Stable Subtype Separation

Joint optimization of subtype trajectories and gating weights can lead to trivial solutions in which all subjects collapse onto a single trajectory or subtype assignments fluctuate across training. We therefore introduce regularization terms that encourage stable and meaningful separation.




Fig. 2. Qualitative example of a subject evaluated under both subtype trajectories, showing localized differences in predicted regional growth along the disease time axis.

Entropy regularization. Joint optimization of subtype trajectories and gating weights can yield diffuse assignments in which subjects are softly shared across clusters, reducing effective specialization. To discourage this, we penalize the entropy of the gating distribution $\theta^j \in \Delta^{C-1}$:

$$\mathcal{L}_{\text{ent}} = \frac{1}{B} \sum_{j=1}^B H(\theta^j), \quad H(\theta^j) = - \sum_{c=1}^C \theta_c^j \log \theta_c^j. \quad (3)$$

Minimizing $\mathcal{L}_{\text{ent}}$ encourages confident subtype assignments and sharper separation between trajectories.

Time-binned balance constraint. A trivial solution partitions subjects by stage while keeping progression functions identical. We prevent this stage-based collapse by enforcing balanced subtype usage within bins along the estimated disease axis. Let τ^j denote the estimated disease stage of subject j and $\mathcal{S}_b$ the set of subjects whose stage falls into bin b . The empirical subtype usage within bin b is defined as

$$u_{b,c} = \frac{1}{|\mathcal{S}_b|} \sum_{j \in \mathcal{S}_b} \theta_c^j, \quad (4)$$

where $\theta^j \in \Delta^{C-1}$ are the gating weights. To discourage collapse, subtype usage is constrained to remain within predefined bounds $u_{b,c} \in [\alpha, 1 - \alpha]$, implemented via a smooth barrier penalty

$$\mathcal{L}_{\text{bin}} = \frac{1}{B} \sum_{b=1}^B \sum_{c=1}^C \phi(u_{b,c}), \quad (5)$$

where $\phi(\cdot)$ penalizes violations of the interval. To reduce minibatch noise, $u_{b,c}$ is estimated using an exponential moving average across training iterations.

Curriculum scheduling. Early in training, subtype trajectories are poorly formed, and strong entropy regularization can prematurely force separation. We

therefore employ a curriculum strategy: the weight of $\mathcal{L}_{\text{bin}}$ is initially emphasized to prevent collapse, then gradually reduced to allow differentiation, while the weight of $\mathcal{L}_{\text{ent}}$ is increased later to consolidate confident subtype specialization.

The overall objective combines the forecasting loss with $\mathcal{L}_{\text{ent}}$ and $\mathcal{L}_{\text{bin}}$.

3 Experimental Setup and Results

We evaluate STAMP on longitudinal MRI from ADNI-1/GO/2 [7, 15] to assess its ability to model heterogeneous disease progression and improve subtype-aware forecasting. STAMP is implemented in PyTorch [14] and optimized using Adam [11]. We compare against BrLP [19], CounterSynth [16], and Conditional VoxelMorph [5, 1] for followup predictions. All baselines are reproduced using their official public implementations and trained with matched preprocessing, resolution, and data splits to ensure fair comparison.

3.1 Dataset

We evaluate STAMP on three longitudinal MRI cohorts from ADNI-1, ADNI-GO, and ADNI-2, comprising 446 subjects and 1,041 scans. Only subjects with at least two timepoints were included (mean follow-up: 2.95 years). We retain scans with available cognitive assessments (ADAS11, FAQ, MMSE, CDRSB, MOCA) and restrict the cohort to subjects with confirmed amyloid pathology ($A\beta < 192$ pg/mL) [23]. All scans were skull-stripped, affinely aligned using turboprep [18], and resampled to a common voxel grid of $128 \times 160 \times 128$. Brain tissue segmentation was obtained using SynthSeg 2.0 [2]. Data were split at the patient level into 80% training, 5% validation, and 15% testing sets.

3.2 Subtype Interpretation and Personalization

Figure 3 visualizes the learned subtype-specific progression functions. Across both subtypes, the magnitude of the growth signal reflects plausible anatomical change: lateral ventricles show the strongest expansion, followed by the hippocampus and inferior lateral ventricles, consistent with established Alzheimer’s-related atrophy patterns. Figure 2 further illustrates this separation by projecting a single subject along both subtype trajectories, revealing localized differences in predicted regional growth under identical temporal alignment. To interpret the gating mechanism, we analyze associations with clinical and biomarker variables. Subtype 1 correlates with disease severity (ADAS13: $p < 10^{-13}$) and amyloid- β levels ($p < 10^{-6}$), whereas Subtype 2 aligns with greater pathology, reflected by elevated TAU ($p < 10^{-6}$) and increased APOE4 carrier frequency ($p < 10^{-2}$). Diagnostic prevalence is consistent with this pattern: CN subjects are primarily assigned to Subtype 1 (95%), Dementia cases to Subtype 2 (100%), and MCI subjects span both trajectories. The estimated individual time shift shows only moderate correlation with ADAS13 (Spearman $\rho = 0.43$, $p < 10^{-3}$), while subtype assignment is not correlated ($p > 0.05$).




Fig. 3. Subtype-specific regional progression functions. The horizontal axis (x) denotes reparameterized disease time, and the vertical axis (y) the predicted regional growth signal conditioning deformation.

with time shift, and both subtypes span the full estimated disease time axis. Thus, the gate does not simply partition subjects by stage, but captures distinct progression regimes beyond temporal severity. We therefore interpret the recovered groups as data-driven progression regimes rather than definitive biological disease subtypes. Concerning the assessment of the personalization capabilities, we note that inverting the gating assignments reduces predictive accuracy (Table 2, inv. Gate). This result demonstrates the importance of subtyping for improving individualized forecasting accuracy.

3.3 Quantitative Results

Table 1 reports region-wise MAE for STAMP, BrLP, CounterSynth and Conditional VoxelMorph. We focus the quantitative evaluation on image-based longitudinal forecasting. Averaged across regions, STAMP achieves Dice 0.907 and Hausdorff95 1.010, comparable to CounterSynth (0.907 / 1.018) and superior to Conditional VoxelMorph (0.856 / 1.065) and BrLP (0.665 / 2.207). STAMP yields the lowest MAE in 16 of 18 regions and is statistically superior in 11 ($p < 0.05$, two-tailed paired t-test) compared to CounterSynth. Gains are most pronounced in medial temporal and ventricular structures, including the hippocampus and lateral ventricles, as well as cortical and global atrophy markers (e.g., cerebral cortex, CSF), indicating improved longitudinal forecasting through subtype-specific trajectory modeling.

Method	Hip ↓	Amyg ↓	L-Vent ↓	C. WM ↓	C. Cortex ↓	Inf.-L. Vent ↓
STAMP (ours)	0.013 ± 0.011*	0.006 ± 0.006	0.140 ± 0.135*	0.388 ± 0.240*	0.356 ± 0.315*	0.013 ± 0.018*
CounterSynth [16]	0.019 ± 0.015	0.007 ± 0.007	0.290 ± 0.238	0.440 ± 0.300	0.434 ± 0.377	0.020 ± 0.025
BrLP [19]	0.033 ± 0.028	0.020 ± 0.018	0.369 ± 0.514	0.802 ± 0.593	3.259 ± 0.933	0.040 ± 0.044
Cond VXM [5]	0.034 ± 0.035	0.008 ± 0.007	0.195 ± 0.181	2.102 ± 1.984	0.701 ± 0.438	0.025 ± 0.029
Method	Cereb. WM ↓	Cereb. Cortex ↓	Thalamus ↓	Caudate ↓	Putamen ↓	Pallidum ↓
STAMP (ours)	0.032 ± 0.024	0.109 ± 0.172*	0.012 ± 0.010*	0.035 ± 0.030	0.010 ± 0.009*	0.005 ± 0.005
CounterSynth [16]	0.030 ± 0.024*	0.126 ± 0.200	0.017 ± 0.013	0.011 ± 0.011*	0.013 ± 0.010	0.005 ± 0.004
BrLP [19]	0.311 ± 0.172	0.335 ± 0.278	0.054 ± 0.052	0.058 ± 0.046	0.048 ± 0.037	0.024 ± 0.015
Cond VXM [5]	0.172 ± 0.177	0.124 ± 0.155	0.018 ± 0.019	0.033 ± 0.032	0.047 ± 0.048	0.015 ± 0.015
Method	Accumbens ↓	VentralDC ↓	Brain Stem ↓	3rd Vent ↓	4th Vent ↓	CSF ↓
STAMP (ours)	0.003 ± 0.002*	0.009 ± 0.007	0.015 ± 0.012	0.007 ± 0.006*	0.005 ± 0.004	0.497 ± 0.370*
CounterSynth [16]	0.004 ± 0.003	0.009 ± 0.007	0.015 ± 0.012	0.009 ± 0.007	0.005 ± 0.004	0.601 ± 0.430
BrLP [19]	0.011 ± 0.008	0.060 ± 0.034	0.100 ± 0.083	0.028 ± 0.022	0.039 ± 0.027	2.699 ± 1.413
Cond VXM [5]	0.005 ± 0.004	0.021 ± 0.027	0.062 ± 0.071	0.014 ± 0.017	0.028 ± 0.029	1.601 ± 1.924

Table 1. Mean absolute error (MAE) of regional volumes as percentage of total brain volume (mean ± standard deviation); * $p < 0.05$ (pairwise test).

Method	Hip ↓	L-Vent ↓	Inf.-L. Vent ↓	C. Cortex ↓	Amyg ↓	AVG (All 18) ↓
STAMP	0.013 ± 0.011	0.140 ± 0.135*	0.013 ± 0.018*	0.356 ± 0.315	0.006 ± 0.006	0.092 ± 0.076
Personalisation						
inv. Gate	0.013 ± 0.012	0.164 ± 0.176	0.015 ± 0.021	0.375 ± 0.344	0.007 ± 0.006	0.094 ± 0.082
Ablation						
w/o Subtypes	0.019 ± 0.015	0.217 ± 0.204	0.020 ± 0.025	0.433 ± 0.377	0.007 ± 0.007	0.107 ± 0.090
w/o Scheduler	0.012 ± 0.010	0.136 ± 0.150	0.013 ± 0.017	0.884 ± 0.600	0.008 ± 0.006	0.135 ± 0.098
w/o TimeBin	0.013 ± 0.012	0.164 ± 0.175	0.014 ± 0.020	0.682 ± 0.506	0.007 ± 0.006	0.120 ± 0.092
w/o Entropy	0.013 ± 0.011	0.142 ± 0.157	0.013 ± 0.020	0.784 ± 0.578	0.007 ± 0.006	0.125 ± 0.097

Table 2. STAMP compared to inverted gates and ablations without subtypes, scheduling, time-binned balancing, or entropy regularization; * $p < 0.05$ (pairwise test).

3.4 Ablation

The ablation in Table 2 quantifies the contribution of each component of STAMP. Removing subtype modeling (w/o Subtypes) increases error consistently across regions (AVG: 0.107 vs. 0.092), confirming the benefit of modeling heterogeneous trajectories. Without time-binned balancing, cortical error rises from 0.356 to 0.682 and the global average increases to 0.120, consistent with collapse toward a dominant trajectory. Omitting entropy regularization similarly elevates cortical error (0.784), indicating reduced subtype separation. The largest degradation occurs without curriculum scheduling, where cortical error increases to 0.884, reflecting unstable trajectory specialization during training.

4 Conclusion

We presented STAMP, a subtype-aware continuous-time spatiotemporal atlas model that integrates disease subtyping directly into deformation-based progression modeling. By jointly learning subtype-specific trajectories, temporal alignment, and gated routing within a unified optimization framework, STAMP captures heterogeneous progression dynamics and enables subtype-aware personalized forecasting of longitudinal anatomical change. Ablation and personalization analyses show that explicit subtype separation and regularised routing are necessary for improved predictive performance. These findings demonstrate that

modeling disease progression as a structured mixture of continuous trajectories improves longitudinal forecasting while preserving the geometric interpretability of atlas-based models. Finally, evaluation on independent external cohorts remains an important direction for future validation of our framework.

Acknowledgments. This work was supported by the grants ANR-22-FAI1-0003 (TRAIN) and ANR-24-IAS2-0001 (Fed-Ops), by the PEPR Santé Numérique program through the grant ANR-22-PESN-0016 (REWIND), and by the 3IA Côte d’Azur initiatives, including ANR-19-P3IA-0002 (3IA Côte d’Azur) and ANR-23-IACL-0001 (3IA Côte d’Azur 2030).

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging* **38**(8), 1788–1800 (2019)
2. Billot, B., Greve, D.N., Puonti, O., Thielscher, A., Van Leemput, K., Fischl, B., Dalca, A.V., Iglesias, J.E., et al.: Synthseg: Segmentation of brain mri scans of any contrast and resolution without retraining. *Medical image analysis* **86**, 102789 (2023)
3. Chen, R.T., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. *Advances in neural information processing systems* **31** (2018)
4. Cummings, J., Apostolova, L., Rabinovici, G., Atri, A., Aisen, P., Greenberg, S., Hendrix, S., Selkoe, D., Weiner, M., Petersen, R., et al.: Lecanemab: appropriate use recommendations. *The journal of prevention of Alzheimer’s disease* **10**(3), 362–377 (2023)
5. Dalca, A., Rakic, M., Guttag, J., Sabuncu, M.: Learning conditional deformable templates with convolutional networks. *Advances in neural information processing systems* **32** (2019)
6. Habes, M., Grothe, M.J., Tunc, B., McMillan, C., Wolk, D.A., Davatzikos, C.: Disentangling heterogeneity in alzheimer’s disease and related dementias using data-driven methods. *Biological psychiatry* **88**(1), 70–82 (2020)
7. Jack Jr, C.R., Bernstein, M.A., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P.J., L. Whitwell, J., Ward, C., et al.: The alzheimer’s disease neuroimaging initiative (adni): Mri methods. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* **27**(4), 685–691 (2008)
8. Kalkhof, J., Ihm, N., Köhler, T., Gregori, B., Mukhopadhyay, A.: Med-nca: Bio-inspired medical image segmentation. *Medical Image Analysis* p. 103601 (2025)
9. Kalkhof, J., Lorenzi, M.: SMART: A Flexible, Interpretable, and Scalable Spatio-temporal Brain Atlas from High-Resolution Imaging Data (Feb 2026), <https://hal.science/hal-05489861>
10. Kalkhof, J., Mukhopadhyay, A.: M3d-nca: Robust 3d segmentation with built-in quality control. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 169–178. Springer (2023)
11. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)

12. Li, D., Iddi, S., Thompson, W.K., Raffi, M.S., Aisen, P.S., Donohue, M.C., Initiative, A.D.N.: Bayesian latent time joint mixed-effects model of progression in the alzheimer’s disease neuroimaging initiative. *Alzheimer’s & Dementia: Diagnosis, Assessment & Disease Monitoring* **10**(1), 657–668 (2018)
13. Mordvintsev, A., Randazzo, E., Niklasson, E., Levin, M.: Growing neural cellular automata. *Distill* **5**(2), e23 (2020)
14. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)
15. Petersen, R.C., Aisen, P.S., Beckett, L.A., Donohue, M.C., Gamst, A.C., Harvey, D.J., Jack Jr, C., Jagust, W.J., Shaw, L.M., Toga, A.W., et al.: Alzheimer’s disease neuroimaging initiative (adni) clinical characterization. *Neurology* **74**(3), 201–209 (2010)
16. Pombo, G., Gray, R., Cardoso, M.J., Ourselin, S., Rees, G., Ashburner, J., Nachev, P.: Equitable modelling of brain imaging by counterfactual augmentation with morphologically constrained 3d deep generative models. *Medical Image Analysis* **84**, 102723 (2023)
17. Poulet, P.E., Durrleman, S.: Mixture modeling for identifying subtypes in disease course mapping. In: *International Conference on Information Processing in Medical Imaging*. pp. 571–582. Springer (2021)
18. Puglisi, L.: Turboprep. <https://github.com/LemuelPuglisi/turboprep> (2023), accessed: 2025-10-08
19. Puglisi, L., Alexander, D.C., Ravi, D.: Enhancing spatiotemporal disease progression models via latent diffusion and prior knowledge. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 173–183. Springer (2024)
20. Puglisi, L., Alexander, D.C., Ravi, D.: Brain latent progression: Individual-based spatiotemporal disease progression on 3d brain mris via latent diffusion. *Medical Image Analysis* (2025)
21. Ranen, A., Kalkhof, J., Mukhopadhyay, A.: Ncadapt: Dynamic adaptation with domain-specific neural cellular automata for continual hippocampus segmentation. In: *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 3834–3843. IEEE (2025)
22. Sauty, B., Durrleman, S.: Progression models for imaging data with longitudinal variational auto encoders. In: *International conference on medical image computing and computer-assisted intervention*. pp. 3–13. Springer (2022)
23. Shaw, L.M., Vanderstichele, H., Knapik-Czajka, M., Clark, C.M., Aisen, P.S., Petersen, R.C., Blennow, K., Soares, H., Simon, A., Lewczuk, P., et al.: Cerebrospinal fluid biomarker signature in alzheimer’s disease neuroimaging initiative subjects. *Annals of neurology* **65**(4), 403–413 (2009)
24. Viani, A., Custo, A., d’Angremont, E., Garibotto, V., Frisoni, G.B., Gutman, B.A., Lorenzi, M.: Disease progression modeling and stratification for detecting sub-trajectories in the natural history of pathologies: Application to alzheimer’s disease trajectory modeling. *Imaging Neuroscience* **3**, IMAG-a (2025)
25. 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., et al.: Uncovering the heterogeneity and temporal complexity of neurodegenerative diseases with subtype and stage inference. *Nature communications* **9**(1), 4273 (2018)