UNIVERSITY OF AMSTERDAM

qurAI

AMLAB Amsterdam Machine Learning Lab

Longitudinal Flow Matching for Trajectory Modeling

Mohammad Mohaiminul Islam Thijs P. Kuipers Sharvaree Vadgama Coen de Vente
Afsana Khan Clara I. Sánchez Erik J. Bekkers

Visit our paper and GitHub!

Standard Flow Matching

- Two distributions: Source $X_0 \rightarrow$ Target X_1
- Requires a path construction from source to target (usually **linear interpolation**)
- Learn dynamics along this path
- Usually, we care about **only endpoints**.
- But works fine for **two marginals**

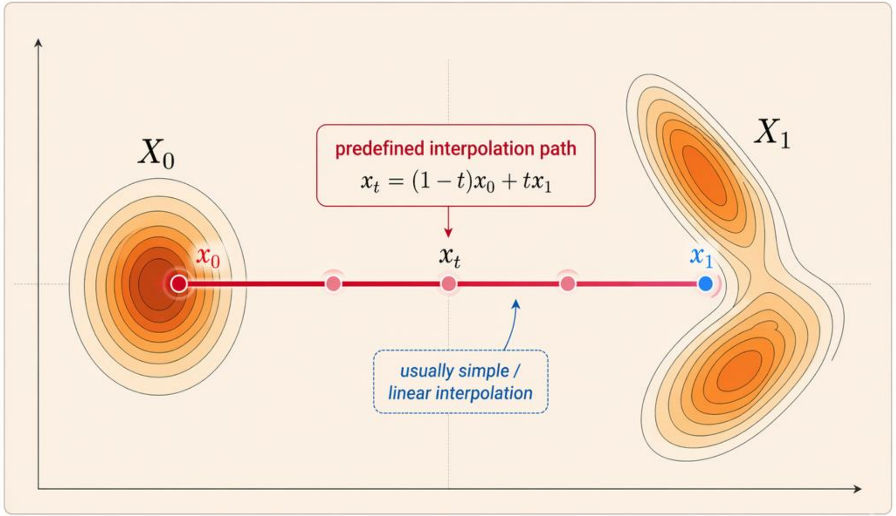

The Problem

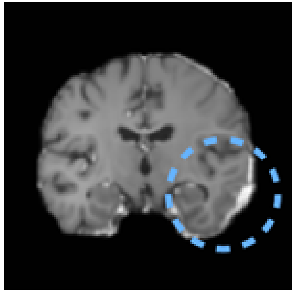

x_0

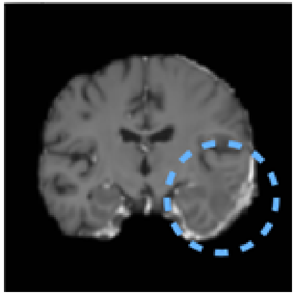

x_1

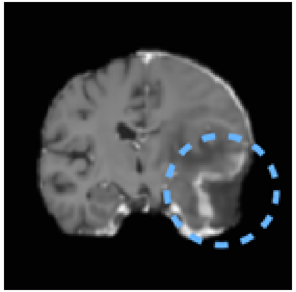

x_2

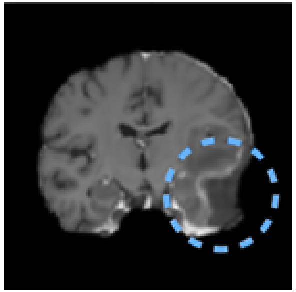

x_3

1. **Longitudinal data** → sparse, irregular, high-dimensional

The Problem

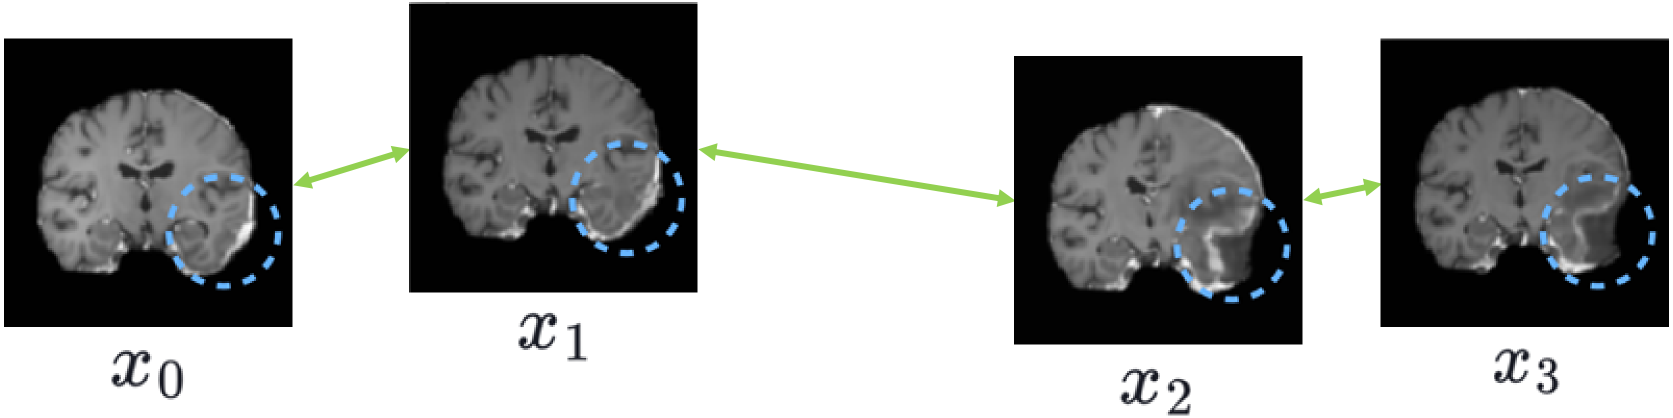

1. **Longitudinal data** → sparse, irregular, high-dimensional
2. **Pairwise transitions** → no global consistency and temporal alignment

The Problem

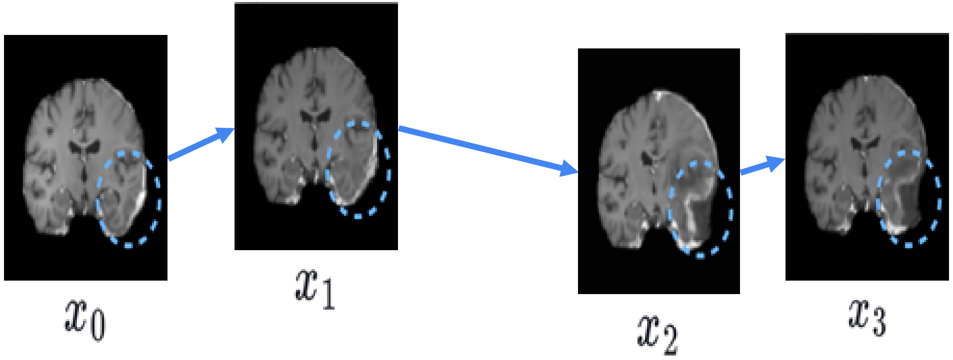

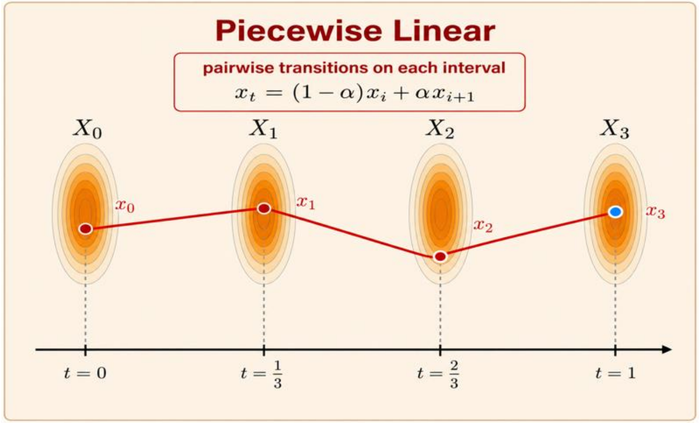

1. **Longitudinal data** → sparse, irregular, high-dimensional
2. **Pairwise transitions** → no global consistency and temporal alignment
3. **Piecewise linear Interp.** → Sharp transitions, bad behavior near intermediate marginals

The Problem

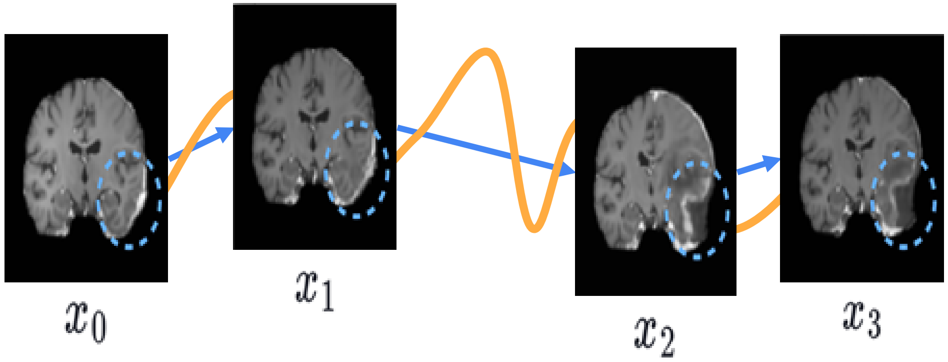

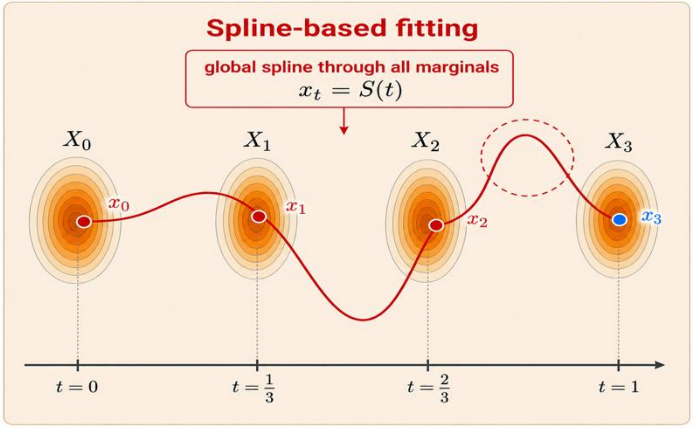

1. **Longitudinal data** → sparse, irregular, high-dimensional
2. **Pairwise transitions** → no global consistency and temporal alignment
3. **Piecewise linear Interp.** → Sharp transitions, bad behaviour near intermediate marginals
4. High-dimensional **spline** fitting → can be expensive, noisy, and unstable

Interpolative Multi-Marginal Flow Matching (IMMFM)

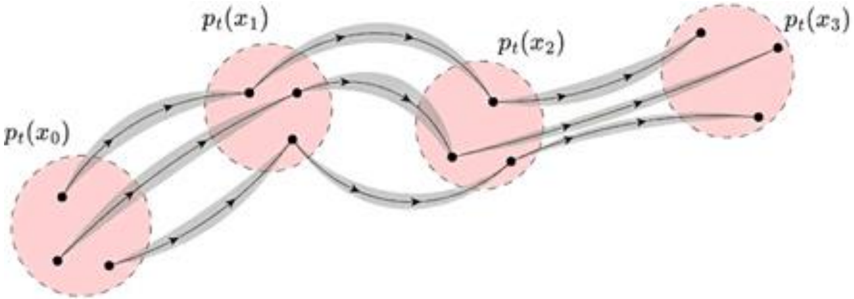

Quadratic Interp. with look-ahead

$$\mu_t(z) = x_{t_i} + v_i(t - t_i) + \frac{1}{2}\,\alpha_t(v_i - v_{i+1})(t - t_i)$$

- **Resolves** unstable, expensive spline fitting via quadratic interpolation with look-ahead.

Interpolative Multi-Marginal Flow Matching (IMMFM)

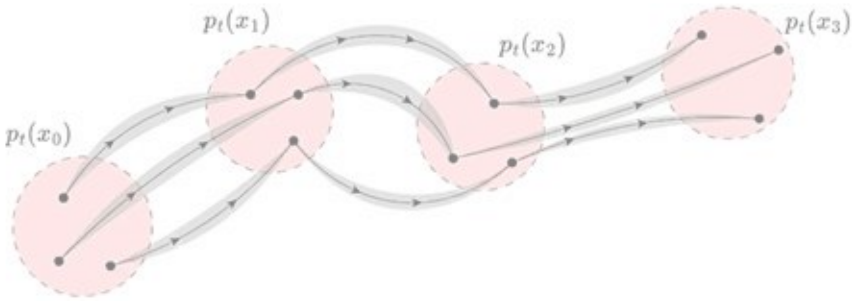

Quadratic Interp. with look-ahead

$$\mu_t(z) = x_{t_i} + v_i(t - t_i) + \frac{1}{2}\,\alpha_t(v_i - v_{i+1})(t - t_i)$$

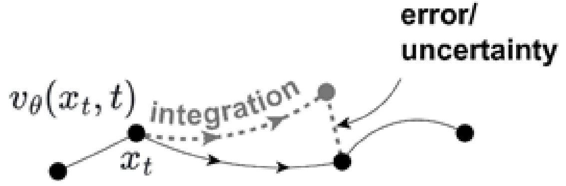

- **Resolves** unstable, expensive spline fitting via quadratic interpolation with look-ahead.
- **Absorbs** complex volatility ignored by the smooth drift → learned diffusion coefficient

Interpolative Multi-Marginal Flow Matching (IMMFM)

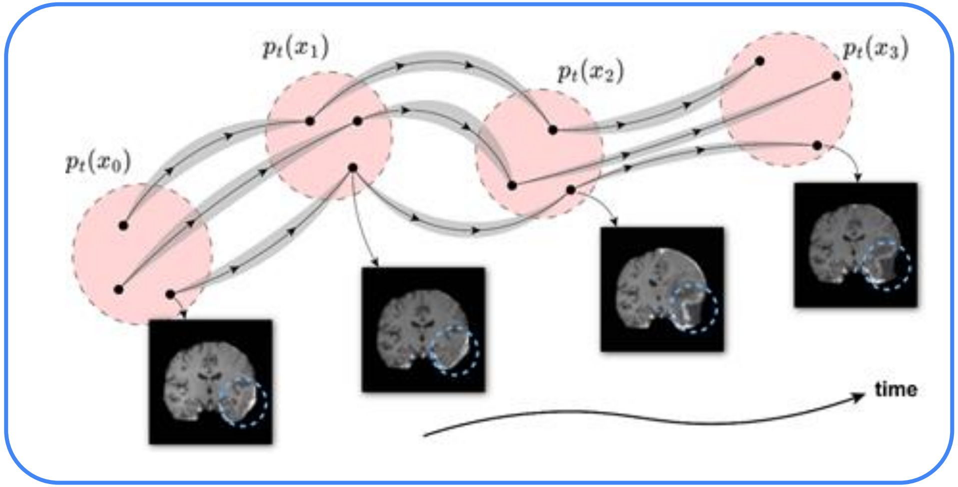

Quadratic Interp. with look-ahead

$$\mu_t(z) = x_{t_i} + v_i(t - t_i) + \frac{1}{2}\,\alpha_t(v_i - v_{i+1})(t - t_i)$$

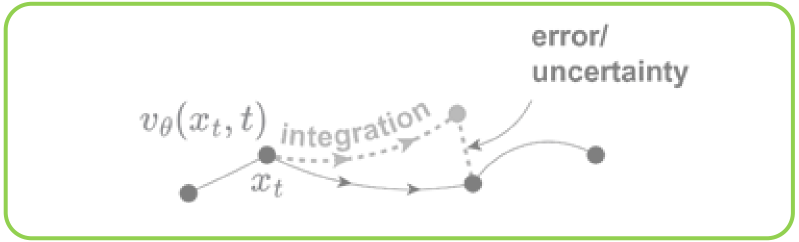

- **Resolves** unstable, expensive spline fitting via quadratic interpolation with look-ahead.
- **Absorbs** complex volatility ignored by the smooth drift → learned diffusion coefficient
- **Learns** globally consistent, stochastic dynamics from noisy sparse data.

Interpolative Multi-Marginal Flow Matching (IMMFM)

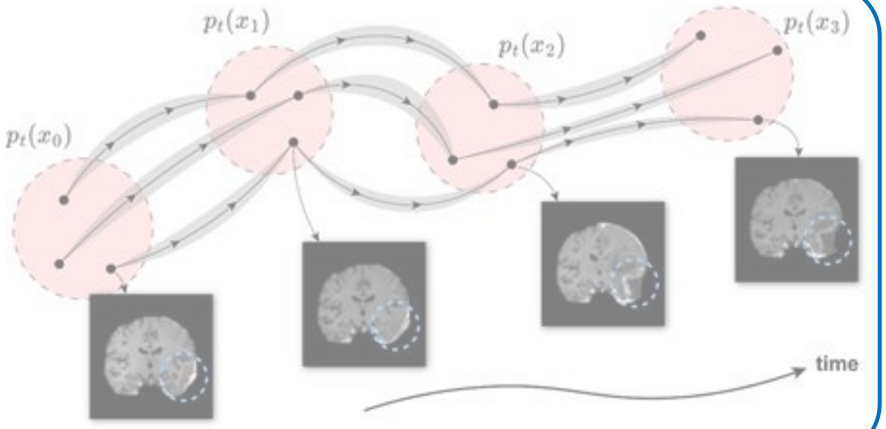

Quadratic Interp. with look-ahead

$$\mu_t(z) = x_{t_i} + v_i(t - t_i) + \frac{1}{2}\alpha_t(v_i - v_{i+1})(t - t_i)$$

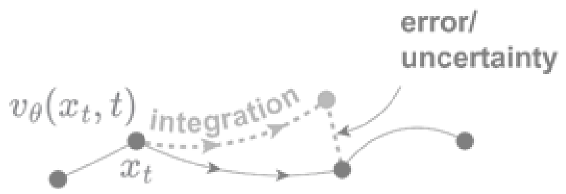

- **Resolves** unstable, expensive spline fitting via quadratic interpolation with look-ahead.
- **Absorbs** complex volatility ignored by the smooth drift → learned diffusion coefficient
- **Learns** globally consistent, stochastic dynamics from noisy sparse data.
- **Practical conditioning** → via observed marginals, no explicit conditioning required!

Training Objective

Optimize **three** decoupled components: **1)** Flow drift v_θ , **2)** Score s_θ , **3)** Diffusion g_θ^2 .

$$\mathcal{L}_{\text{IMMFM}}(\theta) = \mathbb{E}_{t,x}\left[\|v_\theta - u_t^\circ(x|z)\|_2^2\right] +$$
→ O-IMMFM: Deterministic ODE

$$\lambda(t)^2\mathbb{E}_{t,x}\left[\|s_\theta - \nabla_x \log p_t(x|z)\|_2^2\right] +$$
→ S-IMMFM: Stochastic SDE

$$\beta\mathbb{E}_{t,x}\left[\|g_\theta^2 - r_\theta^2\|_2^2\right]$$
→ SU-IMMFM: SDE + diffusion

Training Objective

Optimize **three** decoupled components: **1)** Flow drift v_θ , **2)** Score s_θ , **3)** Diffusion g_θ^2 .

$$\mathcal{L}_{\text{IMMFM}}(\theta) = \mathbb{E}_{t,x}\left[\|v_\theta - u_t^\circ(x|z)\|_2^2\right] +$$

→ O-IMMFM: Deterministic ODE

$$\lambda(t)^2\mathbb{E}_{t,x}\left[\|s_\theta - \nabla_x \log p_t(x|z)\|_2^2\right] +$$

→ S-IMMFM: Stochastic SDE

$$\beta\mathbb{E}_{t,x}\left[\|g_\theta^2 - r_\theta^2\|_2^2\right]$$

→ SU-IMMFM: SDE + diffusion

We prove: learning the residual term jointly **does not** bias the drift and score

Results on Synthetic Data

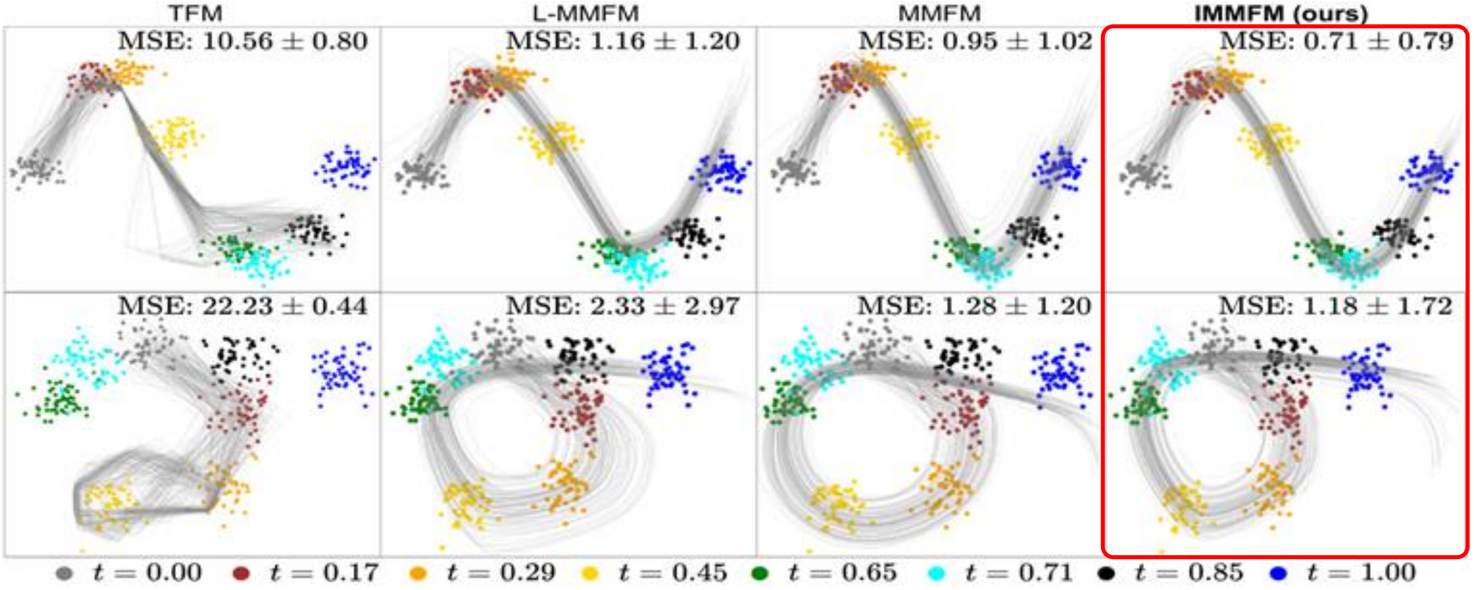

- **8 time-points** , each of the **marginals** is given by a Gaussian.
- **TMF** → Rolling window based, **L-MMFM** →P. Linear interp. MMFM → Spline

More Results

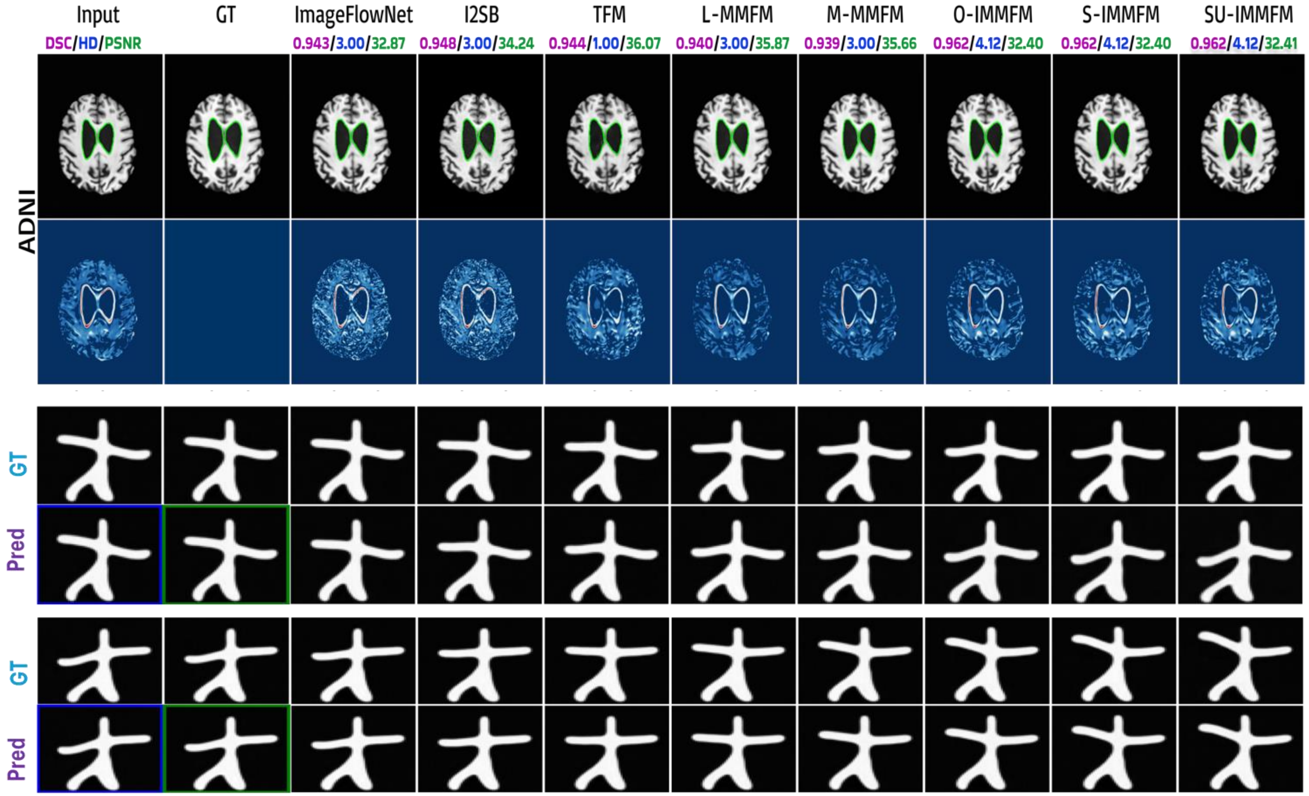

Visit our project page to watch these simulations in action !!

Early Alzheimer’s Detection

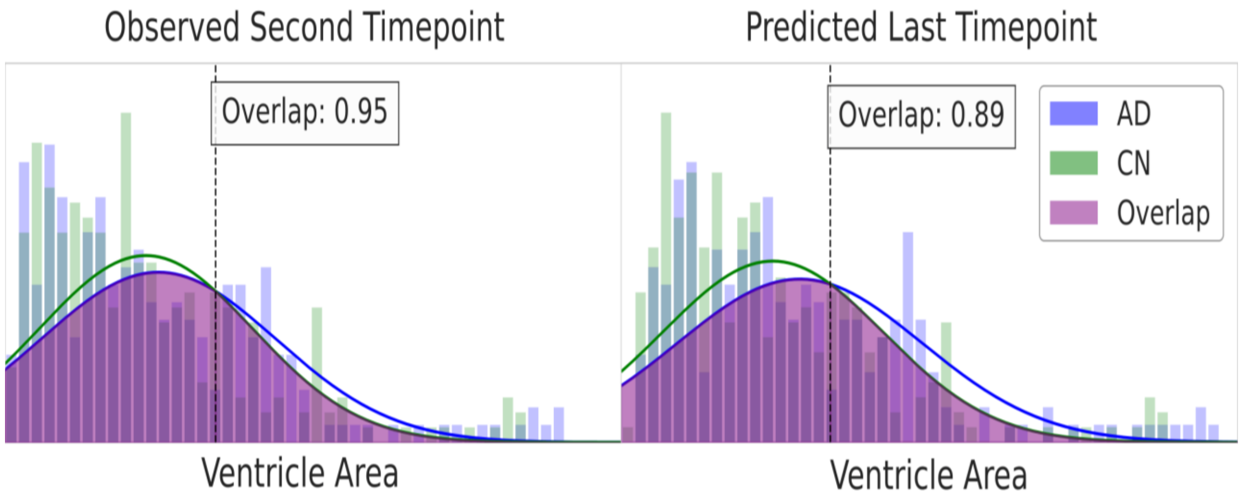

- Classification based on just one biomarker, **the ventricle area.**
- **+9.1%** on **early Alzheimer's** (with 18-month lead time) Classification.

Thank You !!

Drop by our
poster #180 this afternoon to learn more

Visit our paper and GitHub!