Longitudinal Flow Matching for Trajectory Modeling

UNIVERSITY OF AMSTERDAM

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

qurAI   AMLAB Amsterdam Machine Learning Lab

Problem

- Longitudinal data → **sparse, irregularly sampled, and high-dimensional**.
- Modeling dynamics through **pairwise / two-marginal transitions** or rolling windows → miss **global consistency and temporal alignment**.
- Multi-marginal **spline fitting** in high dimensions → **fails** due to overfitting, **noisy target paths**, and **numerical instability (Runge phenomenon)**.

Solution

- ***Interpolative Multi-Marginal Flow Matching*** (IMMFM) learns globally consistent, **continuous stochastic dynamics** directly from sparse multi-timepoint data.
- IMMFM avoids sharp transitions from piecewise linear paths and unstable spline fitting via a tractable **quadratic look-ahead interpolation path**.
- IMMFM incorporates a **learned data-driven diffusion coefficient** to absorb complex local volatility ignored by the smooth drift.

Key Contributions & Methodology

1. Continuous Framework via ODE/SDEs: Formulates individual multi-marginal trajectories using a **continuous target Probability Density Function**, solving the multi-timepoint sparsity mapping.

2. Time-Aware Quadratic Flow: Introduces a **tractable, time-aware quadratic interpolation path** that avoids sharp transitions caused by piecewise linear interpolation near intermediate marginals, while remaining more reliable and practical than spline fitting in high-dimensional settings.

3. Learned Diffusion for Stochastic Volatility: Instead of uniform noise, our framework characterizes intrinsic system variability by **actively learning a localized, data-driven diffusion** (σ) alongside the drift field.

4. Multi-Marginal Coupling via Pairwise OT: Connects observed marginals through task-specific **pairwise OT couplings**, unifying settings with known identities, image registration costs, or independent cross-sectional samples into a common trajectory-construction framework.

5. Practical Conditioning: Uses the **observed marginals** themselves as the conditioning signal, enabling subject-specific trajectory generation from available observations **without requiring explicit class labels** or auxiliary conditioning variables.

Visit our paper and GitHub!

Method Overview

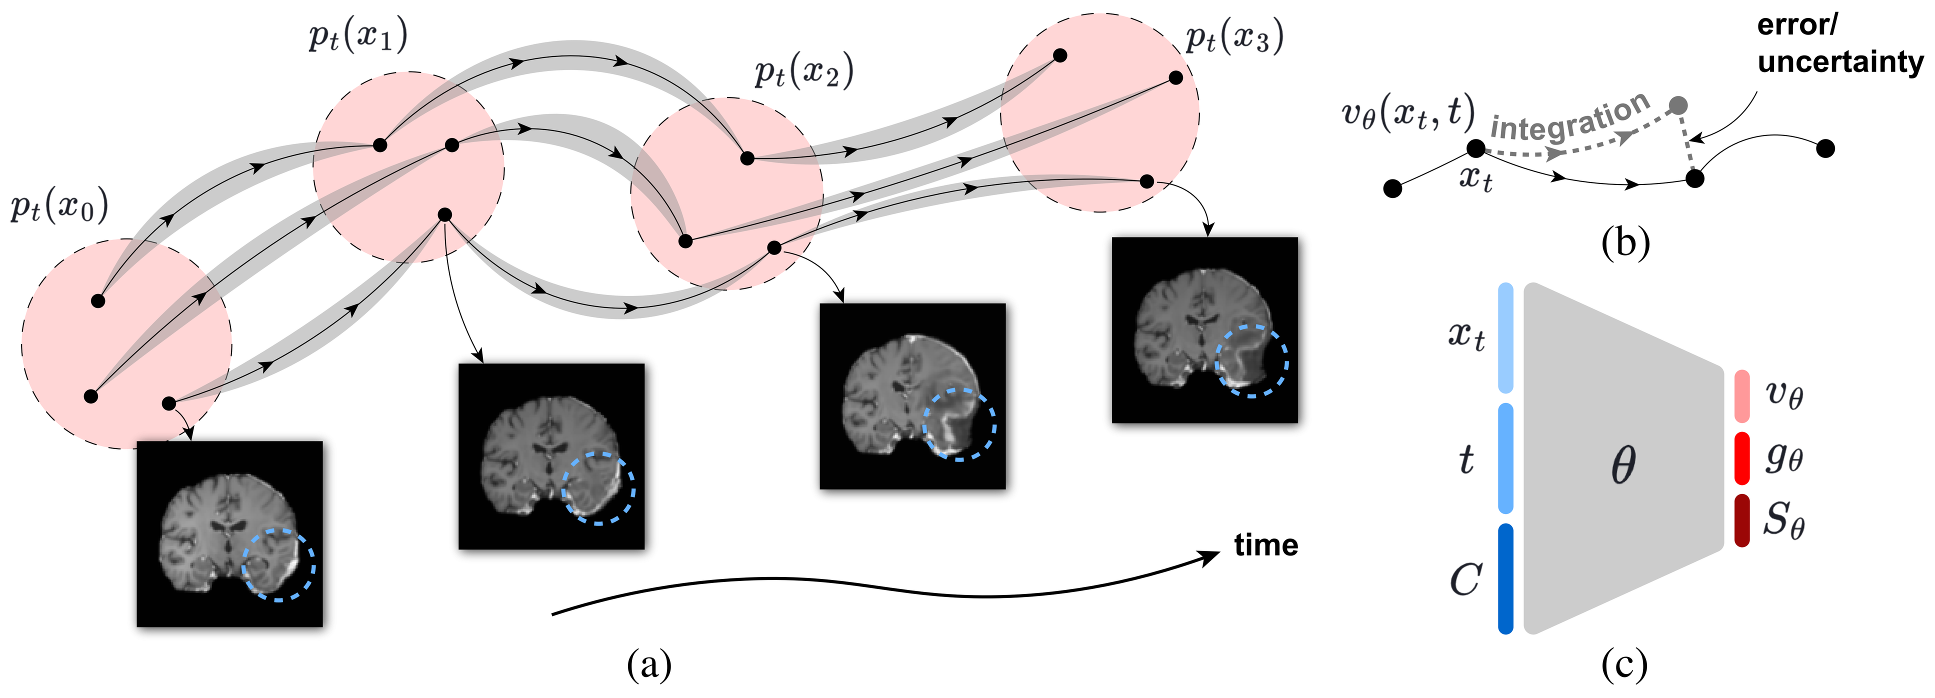

Figure 1. Generates continuous trajectories by bridging isolated observations to a flexible, explicitly-modeled smooth conditional path.

Simulation-Free Trajectory Matching Algorithm

Step 1: Multi-Marginal Prior Alignment (Optimal Transport) : Construct temporal couplings between observed marginals using pairwise OT with a task-specific cost k, inducing valid pseudo-trajectories:

$$C(x_{t_0}, \ldots, x_{t_M}) = \sum_{i=0}^{M-1} k(x_{t_i}, x_{t_{i+1}})$$

Step 2: Quadratic Interpolation with Look-ahead : Form a tractable quadratic target path that avoids unstable spline fitting by using next-segment velocity as a look-ahead correction:

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

Step 3: Flow and Score Matching : Sample $t \sim U(0,1)$ and $x_t \sim p_t(x \mid z)$ along the constructed conditional bridge, then train the drift and score (Tong et al., 2023) directly from analytical targets.

Step 4: Learning Data-Driven Diffusion : The diffusion factor g_ϕ handles the irreducible residual variance naturally discarded by the smoothed drift curve. Learned jointly by matching its squared magnitude to the empirical spatial residuals ($\mathcal{L}_{\text{unc}}$).

Training Objective

We optimize three decoupled components, the flow drift v_θ, the score s_θ, and the diffusion g_θ, using a joint conditional formulation:

$$\mathcal{L}_{\text{IMMFM}}(\theta) = \underbrace{\mathbb{E}_{t,x}\left[\|v_\theta - u_t^\circ(x|z)\|_2^2\right]}_{\text{Base: O-IMMFM}} + \underbrace{\lambda(t)^2\mathbb{E}_{t,x}\left[\|s_\theta - \nabla_x \log p_t(x|z)\|_2^2\right]}_{\text{+ Score } (\to \text{ S-IMMFM})} + \underbrace{\beta\mathbb{E}_{t,x}\left[\|g_\theta^2 - r_\theta^2\|_2^2\right]}_{\text{+ Uncertainty } (\to \text{ SU-IMMFM})}$$

The IMMFM Family: O-IMMFM: Deterministic ODE | S-IMMFM: Stochastic SDE | SU-IMMFM: SDE + Learned Diffusion

Unbiased Joint Optimization: We prove that learning the uncertainty term jointly does not bias the drift/score learning; the stationary points remain identical to the drift/score-only case.

Results at a Glance

- **Fidelity gains:** +4.4% Dice; +2.2 dB PSNR.
- **3D (SU-IMMFM):** DSC +2.9%; PSNR +3.6 dB; HD +3.7.
- **Clinical utility:** +9.1% AD vs CN (36-month forecasts).
- **Quadratic path:** PSNR +2.10 dB; DSC +3.7%; HD +2.16.
- **Data diffusion:** PSNR +0.47 dB; DSC +1.5%; HD +6.29.

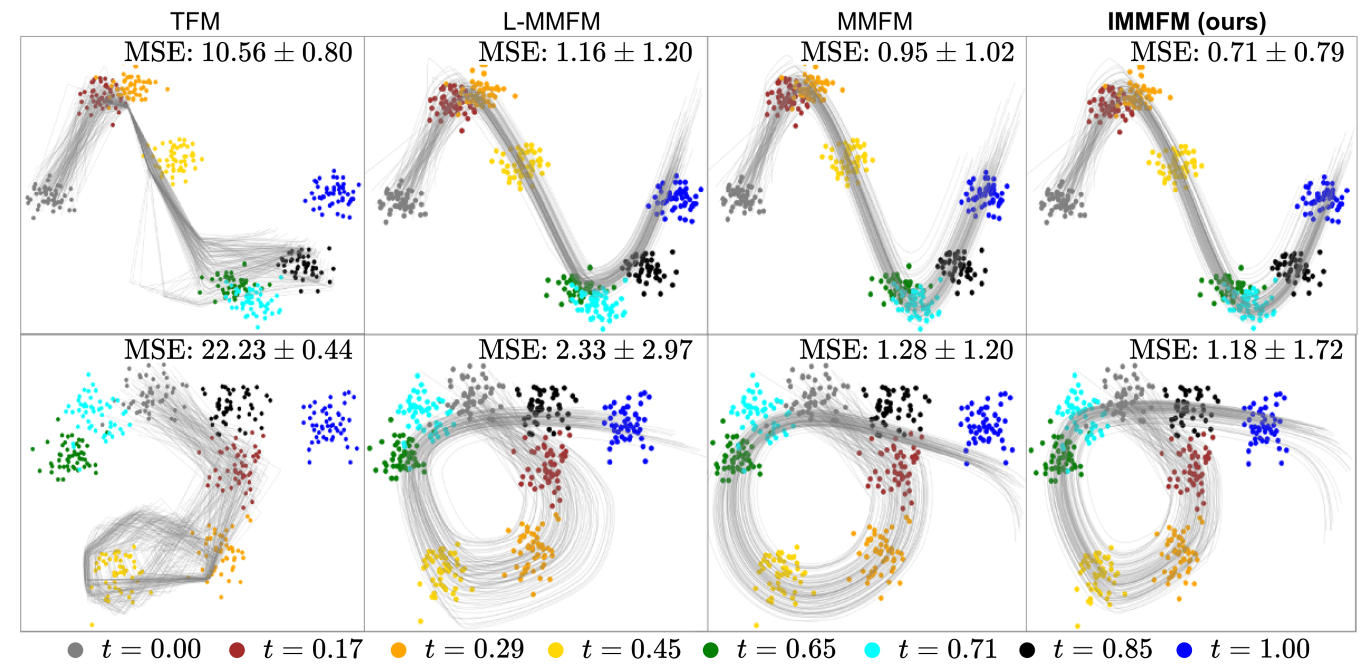

Figure 2. IMMFM preserves smooth, continuous trajectories in curved dynamic scenarios.

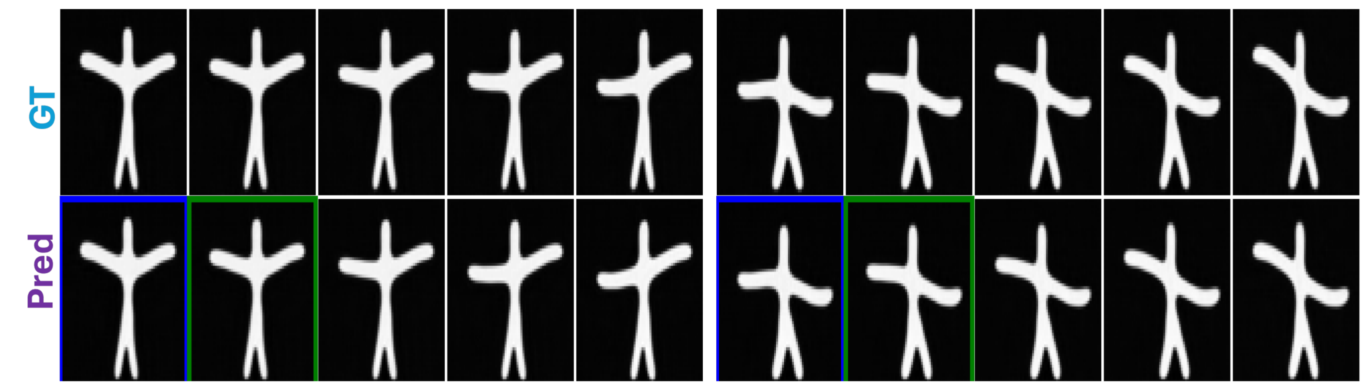

Figure 3. Motion Simulation on Starmen Dataset.

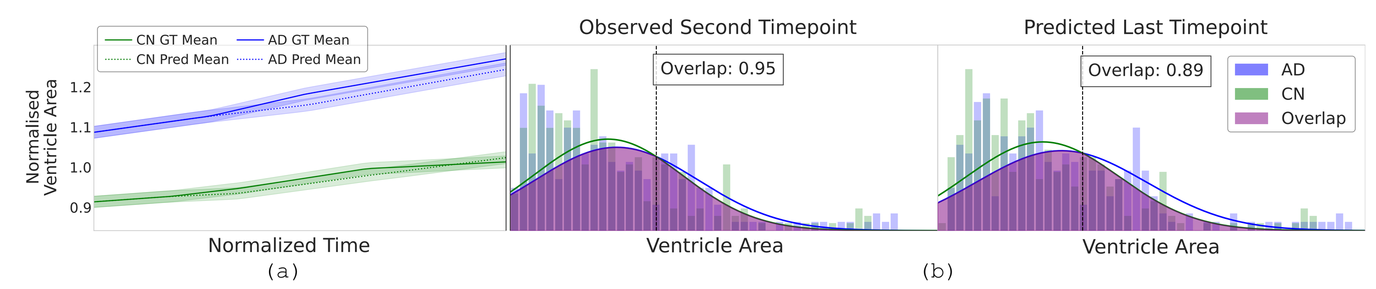

Figure 4. Ventricle biomarker trajectories from 18→36 month forecasts improve AD vs CN accuracy from 71.7% to 80.8% (+9.1%), enabling an ∼18-month lead time for possible early diagonosis.

Temporal Generalization & Stability

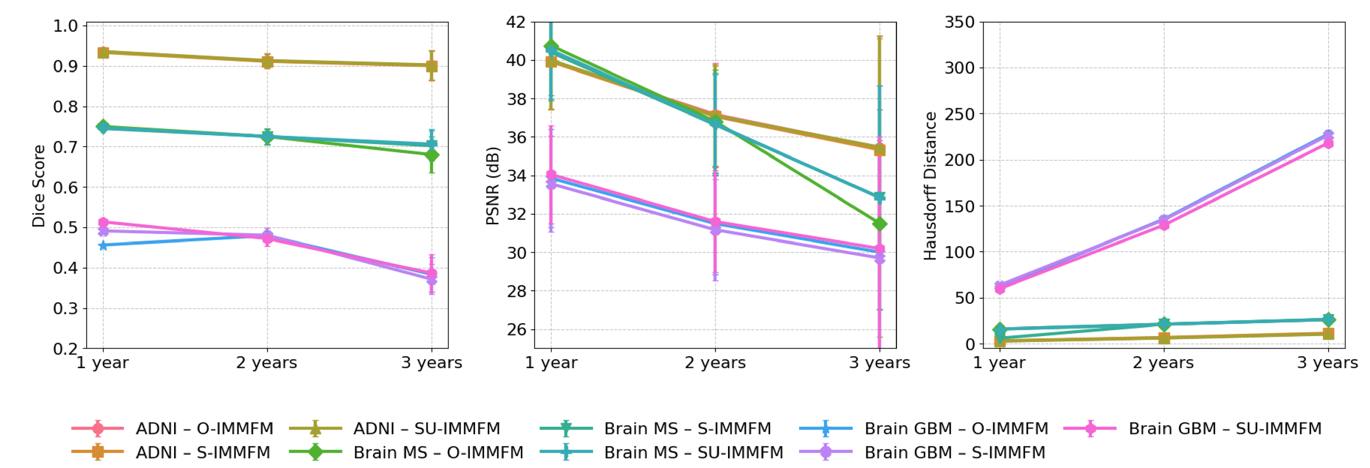

Figure 5. Predictive stability >3 yr horizons: IMMFM preserves high-fidelity Dice/PSNR/HD.