DeepFRC: An End-to-End Deep Learning Model for Functional Registration and Classification

Siyuan Jiang†, Yihan Hu†, Wenjie Li, Pengcheng Zeng*

Institute of Mathematical Sciences, ShanghaiTech University

ICLR 2026

上海科技大学
ShanghaiTech University

* Corresponding e-mail: zengpch@shanghaitech.edu.cn

Outline

* Corresponding e-mail: zengpch@shanghaitech.edu.cn

IMS, ShanghaiTech University

Motivation

- Functional data (curves/trajectories) are ubiquitous in biomedicine, motion analysis, etc.
- Key challenges: phase variability (temporal misalignment) and classification.
- Traditional approaches treat registration and classification separately, which is inefficient.
- **Goal**: Jointly learn alignment and classification in an end-to-end framework.

*Corresponding e-mail: zengpch@shanghaitech.edu.cn IMS, ShanghaiTech University

Related Work

- **Functional registration**: landmark-based, metric-based, model-based methods – often costly and manual.
- **Functional classification**: fPCA, generalized regression – limited by handcrafted features.
- **Deep learning in FDA**: SrvfRegNet (registration only), TTN (joint but not smooth), Tang et al. (parametric, expensive).
- **Gap**: No existing method jointly learns diffeomorphic warping and classification with theoretical guarantees.

Contributions

1. **Unified architecture**: Neural deformation operator + spectral representation + class-aware contrastive loss.
2. **Theoretical foundations**: First approximation and generalization bounds for such a joint model.
3. **Empirical validation**: Outperforms SOTA on synthetic and real-world datasets; robust to noise, missing data, and scale.

* Corresponding e-mail: zengpch@shanghaitech.edu.cn
IMS, ShanghaiTech University

Overall Architecture

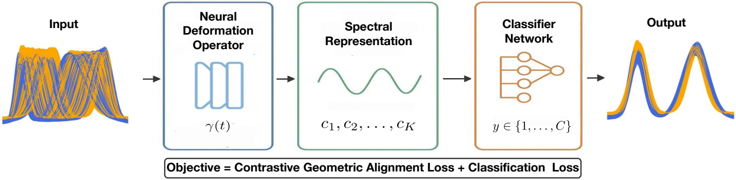

- Input: raw functional trajectories $x_i(t_i)$ with labels y_i.
- Neural deformation operator learns diffeomorphic warping $\gamma_i(t)$.
- Aligned signals are expanded in Fourier basis $\rightarrow$ spectral coefficients c_i.
- Classifier (MLP) predicts labels.
- Training with combined contrastive-geometric alignment loss and classification loss.

Neural Deformation Operator

- 1D CNN extracts temporal features $\tau(x_i(t_i))$.
- Construct diffeomorphic warping γ_i via cumulative sum of squared features and normalization:

$$\tilde{\gamma}_{ij} = \frac{\sum_{s=0}^{j} \tau_{is}^2}{\sum_{\nu=0}^{n} \tau_{i\nu}^2}, \quad \gamma_{ij} = \frac{\sum_{s=0}^{j} \tilde{\gamma}_{is}}{\sum_{\nu=0}^{n} \tilde{\gamma}_{i\nu}}.$$

- Guarantees $\gamma_i(0) = 0$, $\gamma_i(1) = 1$, and $\dot{\gamma}_i > 0$ (diffeomorphism).

* Corresponding e-mail: zengpch@shanghaitech.edu.cn IMS, ShanghaiTech University

Spectral Representation of Aligned Functions

- Warped curve $\tilde{x}_i(t)$ is obtained via linear interpolation.
- Instead of high-dimensional grid, expand in Fourier basis:

$$\tilde{x}_i(t) \approx \sum_{j=1}^{K} c_{ij}\phi_j(t).$$

- Coefficients c_{ij} estimated by least squares (closed-form).
- Yields compact, smooth, regularized embeddings.

Loss Functions

- **Contrastive geometric alignment loss** (SRVF space):

$$\mathcal{L}_1(\Theta_1) = \sum_{j=1}^{C} \frac{\sum_{i:y_i=j} \|Q_i(\gamma_i) - \bar{Q}^{(j)}\|}{N^{(j)}} + \alpha \sum_{1 \le u < v \le C} \|\bar{Q}^{(u)} - \bar{Q}^{(v)}\|^{-1}.$$

- **Classification loss** (cross-entropy):

$$\mathcal{L}_2(\Theta) = -\frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{C} y_{ij} \log \psi_{ij}.$$

- **Joint objective**:

$$\mathcal{L}(\Theta) = \mathcal{L}_1(\Theta_1) + \beta \mathcal{L}_2(\Theta).$$

Complete Model

Algorithm 1 Training DeepFRC

Require: Training data $\{(x_i(\boldsymbol{t}_i), y_i)\}_{i=1}^N$

1: **Set Hyperparameters**: Size of basis function K, loss-related $\{\alpha, \beta\}$, and training parameters $\boldsymbol{\eta}$ (epochs E, batch size, learning rate, etc.)
2: **Initialize Parameters** $\Theta = \Theta_{\text{initial}}$
3: **for** $e = 1$ to E **do**
4: **Forward Propagation**:
(1) Compute $\gamma_i(\boldsymbol{t}_i)$ for each $x_i(\boldsymbol{t}_i)$
(2) Warp $x_i(\boldsymbol{t}_i)$ to obtain $\tilde{x}_i(t)$, calculate its SRVF $Q_i(\boldsymbol{\gamma}_i)$, and extract coefficients $\tilde{\boldsymbol{c}}_i$
(3) Pass $\tilde{\boldsymbol{c}}_i$ through the classifier to compute $\psi(\tilde{\boldsymbol{c}}_i)$
(4) Compute the loss $\mathcal{L}(\Theta)$
5: **Backward Propagation**: Update Θ via AdamW optimizer using $\frac{\partial \mathcal{L}(\Theta)}{\partial \theta}, \theta \in \Theta$
6: **end for**
7: **Return** Trained DeepFRC with optimized parameters Θ^*

---

* Corresponding e-mail: zengpch@shanghaitech.edu.cn IMS, ShanghaiTech University

Theorem 1: Low Registration Error

Under standard regularity conditions (neural network capacity, smooth warping, bounded SRVF), for any $\epsilon > 0$, there exists an estimated warping $\hat{\gamma}$ such that

$$\Delta Q_{\text{reg}}(\gamma^*, \hat{\gamma}) < \epsilon.$$

- Shows neural networks can approximate optimal diffeomorphic warpings in SRVF space.
- Validated via controlled simulations with known ground truth.

* Corresponding e-mail: zengpch@shanghaitech.edu.cn IMS, ShanghaiTech University

Theorem 2: Low Generalization Error

Assume class separation in SRVF means and bounded probabilities. Then for weights $\hat{\Theta}$ estimated by the model,

$$\Delta R_{\text{gen}}(\hat{\Theta}) \lesssim \frac{T_0^{1-1/c_0}}{N},$$

where T_0 is number of steps, N sample size.

- Links alignment quality to classification performance.
- Guides hyperparameter selection (e.g., α, β).

*Corresponding e-mail: zengpch@shanghaitech.edu.cn

* Corresponding e-mail: zengpch@shanghaitech.edu.cn IMS, ShanghaiTech University

Real Data Application

Table 1: Quantitative comparison of registration and classification performance with state-of-the-art approaches across five real datasets. **Bold** indicates best results.

Model	Wave			Yoga			Symbol (2 classes)			Symbol (3 classes)			MotionSense		
	ATV	ACC	F_1-score	ATV	ACC	F_1-score	ATV	ACC	F_1-score	ATV	ACC	F_1-score	ATV	ACC	F_1-score
DeepFRC	**5.6**	**96.4%**	**0.965**	**16.2**	**89.8%**	**0.909**	**4.8**	**96.0%**	**0.959**	**3.2**	**96.3%**	**0.963**	**25.0**	**95.0%**	**0.952**
TTN	6.3	94.7%	0.948	57.7	89.4%	0.904	8.6	92.0%	0.918	4.5	93.3%	0.933	35.1	85.0%	0.857
SrvfRegNet+FCNN$_{\text{raw}}$	7.3	94.6%	0.947	136.0	81.0%	0.830	14.8	94.5%	0.942	6.5	94.7%	0.947	37.7	90.0%	0.909
SrvfRegNet+FCNN$_{\text{fourier}}$	7.3	94.9%	0.950	136.0	84.0%	0.852	14.8	95.0%	0.949	6.5	96.0%	0.959	37.7	90.0%	0.909
SrvfRegNet+FuncNN	7.3	95.7%	0.957	136.0	89.0%	0.908	14.8	93.5%	0.933	6.5	94.7%	0.947	37.7	90.0%	0.889
SrvfRegNet+ADAFNN	7.3	94.6%	0.949	136.0	73.4%	0.753	14.8	89.0%	0.894	6.5	94.3%	0.943	37.7	85.0%	0.857
SrvfRegNet+TSLANet	7.3	**96.4%**	0.961	136.0	89.3%	0.884	14.8	95.5%	0.955	6.5	**96.3%**	0.960	37.7	**95.0%**	**0.952**

- DeepFRC achieves best or second-best ATV, accuracy, and F_1 on all datasets.
- Outperforms TTN and sequential SrvfRegNet+classifier pipelines.
- Comparable to TSLANet but with explicit alignment.

Alignment Visualization

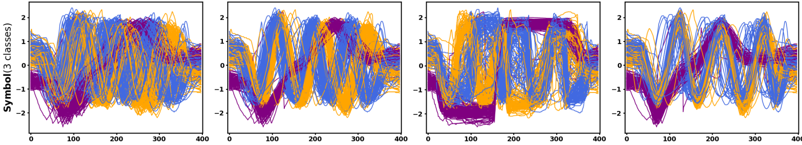

(a) Raw Data (b) Alignment by DeepFRC (c) Alignment by TTN (d) Alignment by SrvfRegNet

- DeepFRC produces smooth, class-separated alignments.
- TTN distorts trajectories; SrvfRegNet loses inter-class separation.
- Interpretability: reveals canonical shapes (e.g., biomechanical events).

* Corresponding e-mail: zengpch@shanghaitech.edu.cn IMS, ShanghaiTech University

Ablation Study

Table 2: Ablation study: contributions of three components in DeepFRC.

Model	Wave ATV	Wave ACC	Wave F_1-score	Yoga ATV	Yoga ACC	Yoga F_1-score	Symbol (2 classes) ATV	Symbol (2 classes) ACC	Symbol (2 classes) F_1-score	Symbol (3 classes) ATV	Symbol (3 classes) ACC	Symbol (3 classes) F_1-score	MotionSense ATV	MotionSense ACC	MotionSense F_1-score
DeepFRC	**5.6**	**96.4%**	**0.965**	**16.2**	**89.8%**	**0.909**	**4.8**	**96.0%**	**0.959**	**3.2**	**96.3%**	**0.963**	**25.0**	**95.0%**	**0.952**
DeepFRC w/o N.D.O.	–	94.4%	0.946	–	83.1%	0.846	–	91.0%	0.905	–	94.7%	0.947	–	90.0%	0.909
DeepFRC w/o S.R.	5.8	95.3%	0.955	17.7	89.2%	0.903	5.3	94.5%	0.945	3.3	93.3%	0.933	28.8	90.0%	0.900
DeepFRC w/o C.N.	7.3	–	–	136.0	–	–	14.8	–	–	6.5	–	–	37.7	–	–

Table 3: P-values from paired t-tests (10 runs) comparing full DeepFRC against ablated variants. Bold indicates significance ($p < 0.05$).

Hypothesis Tests (p-value)	Wave ATV	Wave ACC	Wave F_1-score	Yoga ATV	Yoga ACC	Yoga F_1-score	Symbol (2 classes) ATV	Symbol (2 classes) ACC	Symbol (2 classes) F_1-score	Symbol (3 classes) ATV	Symbol (3 classes) ACC	Symbol (3 classes) F_1-score	MotionSense ATV	MotionSense ACC	MotionSense F_1-score
Full vs. w/o N.D.O.	–	**0.0045**	**0.0051**	–	**0.0023**	**0.0033**	–	**0.0005**	**0.0006**	–	**0.0000**	**0.0000**	–	**0.0000**	**0.0000**
Full vs. w/o S.R.	**0.0453**	**0.0081**	**0.0080**	**0.0370**	**0.0349**	**0.0449**	**0.0382**	**0.0150**	**0.0179**	**0.0056**	**0.0001**	**0.0003**	**0.0025**	**0.0001**	**0.0000**
Full vs. w/o C.N.	**0.0000**	–	–	**0.0000**	–	–	**0.0000**	–	–	**0.0000**	–	–	**0.0000**	–	–

- Removing any component (Neural Deformation Operator, Spectral Representation, Classifier) degrades performance.
- p-values < 0.05 confirm statistical significance.

* Corresponding e-mail: zengpch@shanghaitech.edu.cn IMS, ShanghaiTech University

Robustness Analyses

- **Noise**: DeepFRC maintains high correlation ($\rho > 0.96$) up to $\sigma = 0.20$ (synthetic).
- **Missing data**: $5 - 10\%$ missing points imputed via Fourier splines; performance nearly unchanged.
- **Scalability**: Linear complexity $\mathcal{O}(Nn)$; works on $100\times$ larger Symbol dataset.
- **Sparse data**: Robust classification even with few samples/time points; registration requires $n \geq 100$.
- **Non-diffeomorphic warpings**: Graceful degradation; accuracy $> 90\%$ even with severe violations.

* Corresponding e-mail: zengpch@shanghaitech.edu.cn IMS, ShanghaiTech University

* Corresponding e-mail: zengpch@shanghaitech.edu.cn
IMS, ShanghaiTech University

Conclusion and Future Work

- **DeepFRC**: First end-to-end deep learning framework for joint functional registration and classification.
- Theoretical guarantees (approximation and generalization).
- Strong empirical results on multiple datasets; robust to noise, missing data, and scale.
- Limitations: Assumes smooth, diffeomorphic warpings; sensitive to extreme label noise.
- **Future work**: Adaptive architectures, robust losses for non-smooth functions and noisy labels.

* Corresponding e-mail: zengpch@shanghaitech.edu.cn IMS, ShanghaiTech University