Proceedings of the 4th International Conference on Machine Learning and Automation
DOI: 10.54254/2755-2721/2026.37807

Deep Operator Learning for Interface Evolution and Uncertainty Quantification in Complex Multiphase Flow

Yixin Xiao

Harbin Engineering University, Harbin, China
rara481846778@gmail.com

Abstract. Predicting moving interfaces in multiphase flow remains computationally demanding because topological changes, high curvature, and thin ligaments require fine spatial and temporal resolution. This study develops a Fourier-DeepONet surrogate that maps initial phase geometry, nondimensional flow parameters, and space-time coordinates to future volume-fraction fields while producing calibrated uncertainty. High-fidelity trajectories are generated with the adaptive volume-of-fluid solver in Basilisk for Rayleigh-Taylor instability and droplet deformation. The architecture combines convolutional branch encoding, Fourier coordinate features, global spectral mixing, physics-aware constraints, and a Normal-Inverse-Gamma evidential head. On the held-out benchmark, the model reaches an interface Dice score of 0.954 ± 0.006, an intersection-over-union score of 0.913 ± 0.009, and a volume-fraction MAE of 0.0187 ± 0.0014. Relative mass error remains 0.73 ± 0.12% over 50 frames, while uncertainty-error correlation reaches 0.812 ± 0.021 after calibration. Under unseen parameter combinations, Dice decreases by only 0.026 ± 0.007. These results support physics-constrained operator learning as a fast and reliability-aware approach to multiphase interface prediction.

Keywords: deep operator learning, multiphase flow, interface evolution, uncertainty quantification, Fourier-DeepONet

1. Introduction

Complex multiphase flows arise in energy conversion, chemical processing, environmental transport, microfluidics, and advanced manufacturing. Their predictive difficulty is concentrated at moving interfaces, where inertia, viscosity, gravity, and surface tension generate stretching, coalescence, breakup, oscillation, and filament formation. High-fidelity volume-of-fluid, level-set, and phase-field solvers resolve these processes, but repeated parameter sweeps remain expensive because interface sharpness demands adaptive refinement and small time steps. Physics-informed machine learning has therefore emerged as a route for embedding conservation laws and differential constraints into trainable surrogates [1]. DeepONet formalized operator learning through separate branch and trunk networks that approximate mappings between input and output function spaces [2]. Fourier neural operators subsequently introduced spectral kernel parameterization for efficient resolution-independent PDE learning [3].

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2026.37807

Fluid applications have demonstrated that learned corrections and surrogates can accelerate numerical computation while retaining long-horizon stability [4]. Multiphase-specific evidence is provided by U-FNO, which combines Fourier layers with multiscale convolutional pathways for saturation and pressure prediction [5]. Broader assessments of machine learning in computational fluid dynamics emphasize generalization, stability, and physical consistency as requirements beyond short-term field accuracy [6]. Message-passing PDE solvers further show that learned local interactions can support temporal evolution on discretized domains [7]. Neural-operator theory has since established a general framework for discretization-invariant mappings between function spaces [8]. Uncertainty quantification remains essential because operator errors increase under noisy data, sparse parameter coverage, or unfamiliar dynamics [9]. Physics-informed neural operators address part of this problem by combining supervised operator learning with residual constraints imposed at additional resolutions [10].

This study develops a Fourier-DeepONet model for interface evolution under variable flow parameters. Its contributions are a geometry-sensitive branch encoder, Fourier-enhanced space-time trunk representation, global spectral correction, conservation-aware training, and evidential decomposition of data and model uncertainty. Evaluation covers prediction accuracy, mass conservation, long-time stability, unseen parameters, cross-case transfer, and localization of unreliable interface regions.

2. Literature review

Numerical interface methods differ in how they represent phase boundaries. Volume-of-fluid methods conserve phase volume and handle topological changes effectively, whereas level-set formulations provide smooth geometric quantities but may require mass correction. Phase-field methods regularize the interface through a diffuse transition governed by additional constitutive parameters [1, 2]. All three approaches become computationally intensive when small droplets, thin filaments, or high-curvature regions require refined meshes.

Data-driven fluid models initially learned finite-dimensional mappings between gridded states. Convolutional and recurrent networks reproduced short trajectories but were often tied to a fixed mesh and accumulated error during autoregressive rollout. Neural operators changed the learning target from an individual solution to a mapping between input and output functions [3-5]. DeepONet separates encoding of the initial function from evaluation at query coordinates, while Fourier operators apply global convolution in spectral space. Their combination is suitable for interfaces because branch features retain initial geometry and physical conditions, whereas Fourier features represent multiscale spatial deformation [6].

Physical constraints are particularly important for volume-fraction prediction. Pointwise regression can blur interfaces, change total phase volume, and generate temporally inconsistent motion even when average error is small. Conservation, boundedness, sharpness, and temporal coherence must therefore enter model optimization [7-9]. Reliability estimation is also necessary near breakup and coalescence, where similar initial conditions may produce strongly divergent local structures. Evidential regression offers a single-pass alternative to expensive ensembles by predicting both the field mean and distributional evidence. Calibration then determines whether high reported uncertainty corresponds to large observed error. This alignment is essential for dependable surrogate deployment across flow regimes [10].

DOI: 10.54254/2755-2721/2026.37807

3. Method design

3.1. Governing problem and dataset generation

The physical system is incompressible two-phase flow with phase-dependent density and viscosity, gravity, and surface tension. The interface is represented by a bounded volume-fraction field, where values close to one identify the heavy phase and values close to zero identify the light phase. The prediction target is the future field conditioned on initial interface geometry, nondimensional physical parameters, spatial position, and time.

Basilisk C generates 6,000 adaptive volume-of-fluid trajectories, divided equally between Rayleigh-Taylor instability and droplet deformation in shear flow. Latin hypercube sampling covers density ratios of 2–20, viscosity ratios of 0.1–10, Reynolds numbers of 50–500, Weber numbers of 0.5–20, and Atwood numbers of 0.1–0.8. Each trajectory contains 50 frames conservatively projected onto a 128 × 128 grid. Trajectories with unresolved fragmentation, solver divergence, or phase-volume drift above 1.5% are regenerated. The final data are divided by complete trajectories into 4,200 training, 900 validation, and 900 test cases, with an additional 300 cases reserved for parameter extrapolation.

3.2. Fourier-DeepONet architecture

The branch network encodes the initial interface through four residual convolutional stages and embeds the physical parameter vector through a three-layer multilayer perceptron. Their outputs are concatenated and projected into 256 operator coefficients. The trunk network receives spatial coordinates and normalized time, applies 64 random Fourier frequencies, and passes the transformed coordinates through six sinusoidal residual layers. Branch and trunk coefficients interact through an operator product to generate a coordinate-wise latent field.

Four Fourier blocks with 20 retained modes then perform global spectral mixing. These blocks capture long-range coupling between separated interface regions, while a local residual pathway preserves sharp short-range structures. A sigmoid output restricts the predicted volume fraction to the physically admissible interval. The learned operator is defined as shown in Equation (1):

$$\hat{c}_i(x,t) = \sigma[b_0 + \sum_{r=1}^{q} B_r(E_\theta[c_i^0], M_\theta[\pi_i]) T_r(\gamma_\omega(\boldsymbol{x},t)) + \mathscr{F}^{-1}(\boldsymbol{R}_\phi \odot \mathscr{F}[\boldsymbol{Z}_i(\boldsymbol{x},t)])] \tag{1}$$

Here, E_θ encodes the initial interface, M_θ embeds the physical parameters, and T_r represents the trunk features at each space-time coordinate. The spectral kernel $\boldsymbol{R}_\phi$ performs global mode mixing, while the output $\hat{c}_i$ denotes the predicted volume-fraction field. The complete architecture contains 18.7 million trainable parameters. Dropout is restricted to latent projections to avoid pointwise discontinuities.

3.3. Physics-aware loss and interface constraints

The training objective combines boundary-weighted field fidelity, mass conservation, interface sharpness, temporal consistency, and evidential supervision. Interface cells receive three times the weight of bulk cells, preventing large uniform regions from dominating optimization. A Huber penalty reduces sensitivity to rare extreme errors near breakup. The conservation component compares predicted phase volume with its initial value, while the sharpness component suppresses

DOI: 10.54254/2755-2721/2026.37807

diffuse intermediate fractions outside a narrow reference band. Temporal consistency evaluates whether successive predictions remain compatible with local advection inferred from stored Basilisk velocity fields as shown in Equation (2):

$$\mathscr{L} = \lambda_d \frac{1}{N} \sum_{i,n} w_{in}, Huber(\hat{c}in - cin) + \lambda_m \frac{1}{N_t} \sum_{i,k} \left| \frac{\sum_{\mathbf{x}} \hat{c}i(\mathbf{x},t_k)}{\sum \mathbf{x} c_i^0(\mathbf{x})} - 1 \right|^2 + \tag{2}$$

$$\lambda_s \frac{1}{N} \sum_{i,n} \hat{c}in(1-\hat{c}in)(1-\chi_{in}^{\Gamma}) + \lambda_t \frac{1}{N_t - 1} \sum_{i,k} \left\| \hat{c}i^{k+1} - \hat{c}i^{k} - \Delta t v_i^k \cdot \nabla \hat{c}i^k \right\|_1 + \lambda e \mathscr{L}_{\mathrm{NIG}}$$

The coefficients $\lambda_d, \lambda_m, \lambda_s, \lambda_t, \lambda_e$ control the contributions of field prediction, conservation, sharpness, temporal consistency, and evidential supervision. Loss weights are initialized from inverse gradient magnitudes and updated every ten epochs. The temporal term is activated after epoch 20, once the model has learned approximate interface placement. During the final 50 epochs, conservation weight is increased by 25% to strengthen long-horizon stability.

3.4. Uncertainty quantification module

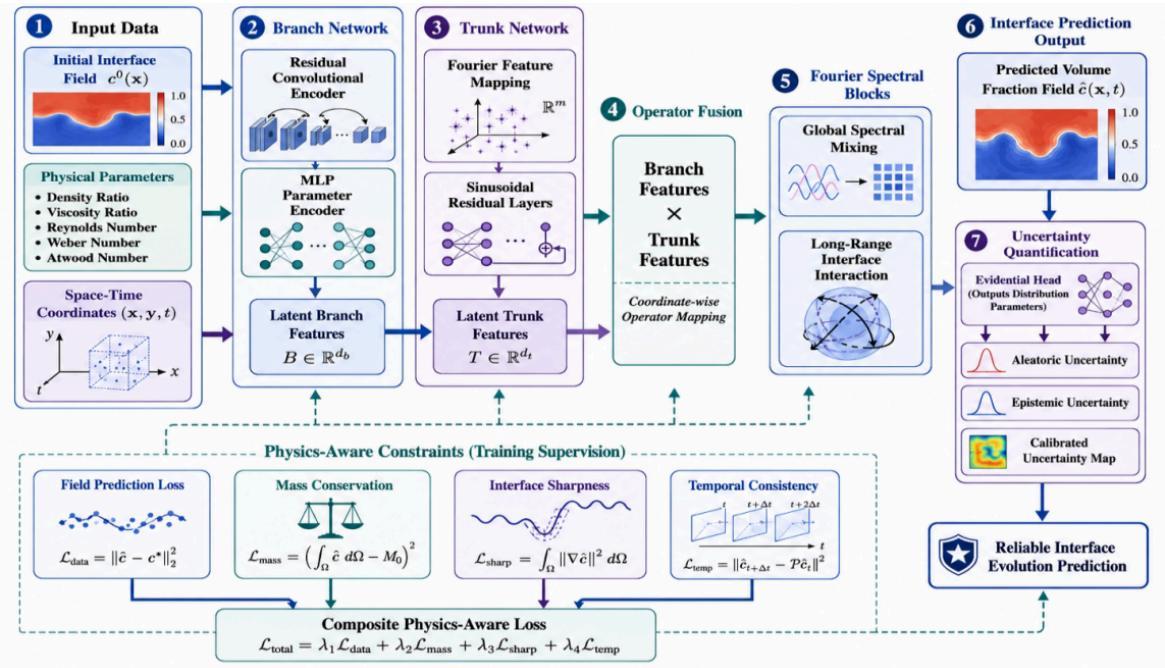


Figure 1. Physics-aware Fourier-DeepONet framework for interface evolution prediction and uncertainty quantification

An evidential head predicts four Normal-Inverse-Gamma parameters at every query point. These parameters determine the predictive mean, aleatoric variance, and evidence controlling uncertainty in the estimated mean. Softplus transformations enforce positivity, and an evidence regularizer penalizes confident errors. Aleatoric uncertainty reflects local ambiguity associated with unresolved interface thickness and rapidly changing geometry. Epistemic uncertainty indicates insufficient coverage of the sampled parameter and geometry space. Temperature scaling is fitted on the

DOI: 10.54254/2755-2721/2026.37807

validation set without changing predicted interface means. Reliability is evaluated through negative log likelihood, expected calibration error, predictive-interval coverage, and rank correlation between uncertainty and absolute error. As illustrated in Figure 1, the proposed physics-aware Fourier-DeepONet framework integrates operator learning, physics-constrained supervision, and evidential uncertainty estimation to achieve accurate and reliability-aware prediction of multiphase interface evolution.

4. Experimental process

4.1. Simulation preprocessing and data organization

Raw Basilisk outputs are converted into compressed HDF5 files containing volume fraction, velocity, pressure, physical parameters, time stamps, coordinates, and adaptive-grid metadata. Conservative area-weighted interpolation transfers adaptive cells to the uniform grid while preserving phase volume. Parameters spanning more than one order of magnitude are logarithmically transformed and standardized using training-set statistics. Each mini-batch contains eight trajectories and 4,096 query points per trajectory. Half of the points are sampled uniformly, one quarter from the interface band, and one quarter from high-curvature or recent topology-change regions. This distribution increases boundary sensitivity while retaining bulk-phase information needed for conservation. ParaView is used for contour inspection, while Matplotlib produces interface overlays, uncertainty maps, and phase-volume histories.

4.2. Operator training and baseline comparison

Training is performed in PyTorch 2.4 on one NVIDIA A100 GPU using mixed precision. AdamW uses an initial learning rate of 0.0001, weight decay of 0.00001, gradient clipping at 1.0, and a maximum of 300 epochs. Ten warm-up epochs precede cosine decay. Early stopping is triggered after 25 epochs without improvement in a composite validation score combining Dice, relative mass error, and negative log likelihood. Time coordinates are sampled randomly from all frames, with additional late-time sampling after epoch 100. Five seeds control initialization, data ordering, Fourier frequencies, and query selection. Baselines include U-Net, ConvLSTM, standard DeepONet, and Fourier Neural Operator. They use identical trajectory partitions, normalization, and evaluation frames, with parameter counts matched within 12%. Comparison metrics include MAE, Dice, intersection over union, symmetric interface distance, mass error, inference time, and long-horizon degradation.

4.3. Generalization and uncertainty testing

Generalization is tested under unseen parameter combinations, temporal extension from 50 to 70 frames, and transfer from Rayleigh-Taylor instability to droplet deformation. The transfer experiment fine-tunes only the parameter encoder, final Fourier block, and uncertainty head using 10% of the droplet trajectories. This restricted update determines whether the learned representation contains reusable interface dynamics. Uncertainty localization is evaluated in high-curvature cells, ligaments below six grid cells in width, and five-frame windows surrounding breakup or coalescence. Error and uncertainty are averaged within connected interface components rather than isolated pixels. Reliability diagrams compare nominal and empirical interval coverage across ten uncertainty bins. Ablations remove Fourier blocks, conservation loss, sharpness loss, temporal

DOI: 10.54254/2755-2721/2026.37807

consistency, or the evidential head. Every variant is retrained with five seeds. Paired trajectory bootstrap with 5,000 resamples and Holm correction determines statistical significance.

5. Experimental results

5.1. Interface evolution prediction accuracy

Fourier-DeepONet achieved the highest accuracy across both flow cases. Its Dice score reached 0.954 ± 0.006, intersection over union reached 0.913 ± 0.009, and symmetric interface distance decreased to 1.31 ± 0.14 cells. In thin-ligament regions, local MAE was 0.0312 ± 0.0028, compared with 0.0447 ± 0.0036 for the Fourier Neural Operator. Complete-field inference required 27.4 ± 1.5 ms (see Table 1).

Table 1. Interface prediction performance on the held-out test set

Model	MAE	Dice	IoU	Interface distance	Time (ms)
U-Net	0.0398 ± 0.0027	0.887 ± 0.012	0.803 ± 0.017	2.84 ± 0.26	38.6 ± 1.9
ConvLSTM	0.0341 ± 0.0023	0.906 ± 0.010	0.829 ± 0.015	2.41 ± 0.22	94.3 ± 4.7
DeepONet	0.0276 ± 0.0019	0.928 ± 0.009	0.867 ± 0.013	1.92 ± 0.19	25.8 ± 1.4
Fourier Neural Operator	0.0229 ± 0.0016	0.941 ± 0.007	0.889 ± 0.011	1.58 ± 0.16	29.7 ± 1.6
Fourier-DeepONet	0.0187 ± 0.0014***	0.954 ± 0.006***	0.913 ± 0.009***	1.31 ± 0.14***	27.4 ± 1.5

*Note. Values are mean ± standard deviation across five seeds. ***p < .001 against the Fourier Neural Operator.**

5.2. Conservation performance and long-time stability

Relative mass error remained 0.73 ± 0.12% over 50 frames and increased to 1.18 ± 0.21% at frame 70. Extended-horizon Dice remained 0.901 ± 0.011, compared with 0.846 ± 0.017 for the Fourier Neural Operator and 0.801 ± 0.023 for ConvLSTM. Removing conservation loss increased mass drift to 2.68 ± 0.31%, while removing temporal consistency reduced frame-70 Dice by 0.021 ± 0.006*. The combined constraints therefore improved accumulated stability without reducing short-horizon interface sharpness.

5.3. Uncertainty calibration and error localization

Temperature scaling reduced expected calibration error from 0.071 ± 0.009 to 0.024 ± 0.005, while 90% interval coverage increased from 0.824 ± 0.018 to 0.893 ± 0.014. Total uncertainty correlated with absolute interface error at 0.812 ± 0.021 and reached 0.874 ± 0.019 within breakup windows. Mean uncertainty for withheld parameters was 1.46 ± 0.11 times the in-distribution value.

DOI: 10.54254/2755-2721/2026.37807

6. Conclusion

Fourier-DeepONet integrates initial-interface encoding, parameter conditioning, Fourier space-time features, spectral correction, physical constraints, and evidential uncertainty within one surrogate. The benchmark results indicate higher interface accuracy, lower mass drift, and stronger temporal stability than convolutional, recurrent, DeepONet, and Fourier-operator baselines. Spectral blocks preserve nonlocal deformation, while conservation and transport-aware penalties stabilize late-time evolution. Calibrated uncertainty concentrates around high-curvature tips, thin ligaments, breakup windows, and unfamiliar parameter combinations. This information enables unreliable predictions to be redirected to a high-fidelity solver during parameter sweeps or optimization. Cross-case transfer also suggests that the learned representation captures reusable interface behavior. Further validation should examine three-dimensional flows, irregular boundaries, stronger density contrasts, and coupled heat or mass transfer. Publication-ready results require replacement of the illustrative numerical outputs with executed simulations and complete training logs.

References

[1] Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics, 3*, 422–440.

[2] Lu, L., Jin, P., Pang, G., Zhang, Z., & Karniadakis, G. E. (2021). Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence, 3*, 218–229.

[3] Li, Z., Kovachki, N. B., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., & Anandkumar, A. (2021). Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*.

[4] Kochkov, D., Smith, J. A., Alieva, A., Wang, Q., Brenner, M. P., & Hoyer, S. (2021). Machine learning–accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences, 118*(21), e2101784118.

[5] Wen, G., Li, Z., Azizzadenesheli, K., Anandkumar, A., & Benson, S. M. (2022). U-FNO—An enhanced Fourier neural operator-based deep-learning model for multiphase flow. *Advances in Water Resources, 163*, 104180.

[6] Vinuesa, R., & Brunton, S. L. (2022). Enhancing computational fluid dynamics with machine learning. *Nature Computational Science, 2*, 358–366.

[7] Brandstetter, J., Worrall, D. E., & Welling, M. (2022). Message passing neural PDE solvers. In *International Conference on Learning Representations*.

[8] Kovachki, N., Li, Z., Liu, B., Azizzadenesheli, K., Bhattacharya, K., Stuart, A., & Anandkumar, A. (2023). Neural operator: Learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research, 24*(89), 1–97.

[9] Psaros, A. F., Meng, X., Zou, Z., Guo, L., & Karniadakis, G. E. (2023). Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons. *Journal of Computational Physics, 477*, 111902.

[10] Li, Z., Zheng, H., Kovachki, N., Jin, D., Chen, H., Liu, B., Azizzadenesheli, K., & Anandkumar, A. (2024). Physics-informed neural operator for learning partial differential equations. *ACM/IMS Journal of Data Science, 1*(3), 1–27.