Journal of Applied Computational Intelligence

eISSN: 3136-6937

Vol. 4, No. 2, April, 2026

Rotationally Equivariant Fourier Neural Operator for Grain Orientation Evolution during Solidification

Zdravka Stamenković[1,*], Arsenija Mitrović[1], Gojka Vasiljev[1], Kraljevka Novčić[2]

[1] *Faculty of Materials Engineering, University of Belgrade, Belgrade, 11000, Serbia*

[2] *Faculty of Computer Science, University of Banja Luka, Banja Luka, 78000, Serbia*

Received 10 February 2026; Revised 20 June 2026; Accepted 25 June 2026; Published 30 June 2026

Abstract

Predicting grain orientation evolution during solidification requires a surrogate that distinguishes changes in the observation frame from changes in the physical microstructure. This paper develops a rotationally equivariant Fourier neural operator for two-dimensional solidification fields with fourfold crystal symmetry. Crystal orientation is represented by a symmetry-compatible complex field, while thermal gradients transform jointly with the spatial coordinates. Group-shared spectral layers are employed for long-range interaction and local interface correction, and a masked angular objective is used to avoid arbitrary orientation targets in the liquid region. The benchmark integrates orientation prediction, geometric consistency, and temporal evaluation within a common protocol. The benchmark comprises 1200 training trajectories, 150 validation trajectories, and 150 test trajectories, with rotation and thermal-shift evaluations. Across the five model configurations, the mean angular error of the proposed configuration is 2.8 degrees; that of the conventional spectral baseline is 6.8 degrees, and the late rollout error has been reduced from 10.8 to 4.7 degrees. As a result, the corresponding interface overlap rose to 0.934 and the normalized rotation inconsistency dropped to 0.014. The comparison jointly assesses angular fidelity, geometric consistency, interface reconstruction, and computational cost. The formulation provides a framework for evaluating symmetry-aware solidification surrogates through independent simulations and experiments. The prediction target describes planar grain selection and boundary motion, with crystallographic orientation retained within established grain interiors.

Keywords: Rotational Equivariance, Fourier Neural Operator, Grain Orientation Evolution, Crystal Symmetry, Solidification Microstructure

1. Introduction

Transient thermal fields and repeated solidification during wire arc additive manufacturing influence the resulting grain structure and crystal orientations [1]. Phase-field studies of directional solidification show that crystallographic alignment and the imposed temperature gradient influence competitive grain growth and grain boundary selection [2]. Coupled phase and orientation field models represent grain evolution during melting and solidification cycles in metal additive manufacturing [3]. Operator-learning surrogates have been developed to reduce the computational burden of repeated phase-field predictions of solidification morphology [4]. Reduced models must preserve the state variables that distinguish competing growth pathways. A surrogate that reproduces the solid fraction while assigning incorrect orientations can misrepresent texture, boundary misorientation, and the identity of surviving grains. Orientation field models couple structural order with crystallographic orientation. Grain selection under directional thermal forcing

* Corresponding author: stamen.zdr@tmf.bg.ac.rs

http://dx.doi.org/10.68406/jaci.2026.vol4iss2no14:196-211

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

involves interactions among neighboring interfaces. These considerations motivate a prediction problem that couples spatial morphology, crystallographic orientation, and thermal forcing. The present study investigates planar cross-sections with fourfold symmetry and aims to predict the evolution of the orientation distribution of solid material. The predicted evolution is mainly the movement of boundaries and the replacement of one grain by another, with crystal orientation preserved in the grain interior.

Learning a field evolution map provides a reusable surrogate across related process conditions. Azizzadenesheli et al. [5] described neural operators as a framework for learning mappings between functions defined on continuous domains. Spectral implementations are relevant to solidification because they can represent interactions over distances larger than the interface width. Spectral operator models learn frequency-space interactions, while multiscale prediction has motivated methods that address their bias toward low-frequency components [6]. A conventional spectral weight tensor can associate spatial patterns with the axes of the training grid. A consistent change of observation frame rotates the microstructure, thermal gradient, and boundary data together. Equivariance constraints have also been incorporated into learned dynamical operators through group-convolutional parameterizations [7]. The transformation must therefore cover the complete physical state. Rotation of the microstructure while keeping the thermal gradient constant will change the relative alignment of growth and thus result in a different experiment. A square grid supports exact quarter-turn permutations, whereas intermediate angles generally require interpolation; these cases are evaluated separately. Augmentation can expose the model to rotated configurations, while architectural constraints determine its transformation behavior. A consistent field transformation provides the basis for evaluating this distinction across observation frames.

Crystallographic equivalence introduces a separate representation issue. For a planar crystal with fourfold symmetry, orientations separated by a quarter turn describe equivalent crystallographic states, while the observation frame remains independently defined. Crystal symmetry has been incorporated into orientation-sensitive loss functions for reconstructing electron backscatter diffraction maps [8]. Encoding the raw angle as a regular scalar introduces a discontinuity at the boundary of the selected interval. Circular regression represents an angular response through its sine and cosine components, preserving its periodic structure during estimation [9]. Interface evaluation also requires a validity mask because crystallographic orientation is defined on solid material. Numerical treatment of diffuse-interface field requires careful interpretation of quantities in the vicinity of the transition zone. Angular and geometric errors should therefore be evaluated on spatially resolved regions. A small global error can conceal missed narrow grains or an inconsistent response to frame rotation. Evaluating temporal neural operators requires attention to error accumulation over successive autoregressive predictions [10]. Therefore, the research gap addressed here is a combined formulation and evaluation problem: how to introduce discrete spatial equivariance, maintain crystal symmetry in the target representation, and observe the independent effects of these on orientation prediction.

To connect the geometric structure of the prediction problem with the mechanisms of competitive grain growth, this study develops a rotationally equivariant Fourier neural operator that integrates a fourfold orientation representation, shared spectral interactions, and local reconstruction near moving solid-liquid interfaces. The formulation specifies the transformation of scalar fields, orientation channels and the thermal gradient prior to defining the prediction map. The formulation relates representation choice and spatial symmetry to the temporal stability of orientation prediction. Angular discrepancy, interface reconstruction, rotation consistency and process transfer are included in the numerical evaluation. Each diagnostic addresses a distinct prediction requirement and examines the effect of geometric constraints on local reconstruction quality. Section 2 reviews solidification surrogates, operator learning, and orientation-aware evaluation. Section 3 introduces the form and symmetry operation of the state. Section 4 describes optimization, data separation, and the reproducible validation protocol. Section 5 presents the dataset, numerical comparisons, ablations, and transfer analyses, relating the numerical observations to their physical interpretation. Section 6 identifies the limitations of the method and the extensions required for application across a broader range of materials and processing conditions. The contribution is the coordinated treatment of crystal-equivalent orientation targets, joint thermal-frame transformations, and interface-aware reconstruction within one solidification prediction task.

2. Related Work

2.1. Solidification models and learned microstructure evolution

Phase-field methods represent the motion of solidification interfaces. Quantitative phase-field studies show that anisotropy of the solid-liquid interfacial free energy affects the morphology selected during directional solidification [11]. This mechanism is directly related to competitive grain growth, as similarly located grains can respond differently to the same thermal forcing. The computational surrogate requires sufficient state information to distinguish these responses. Orientation-dependent selection requires crystallographic information, evolving solid masks, and the thermal field that specifies the driving direction. The predictive state therefore combines phase, orientation, and thermal fields.

Zhu et al. [12] developed a convolutional neural network model to predict the spatiotemporal evolution of grain morphologies using phase-field data. Predictive performance depends on the coverage of time, process conditions, and grain arrangements in the training distribution. Splitting consecutive frames from the same trajectory between training and testing can expose the model to nearly identical interfaces in both partitions. Trajectory-level separation is therefore necessary to evaluate generalization to new grain arrangements. Spatially distinct grain boundary networks can produce the same final orientation histogram, making spatial identity essential to solidification prediction. The predictor must retain the spatial distribution and connectivity of grain orientations over time, requiring trajectory-level validation and complementary morphological and angular metrics. A numerical reference model is employed to generate independent trajectories for validation of the learned evolution map.

2.2. Spectral operators and geometric constraints

Neural operators aim to find a reusable map of discretized fields and process related physical conditions with a single set of parameters. Neural-operator architectures include spectral, integral-kernel, and attention-based formulations with different mechanisms for exchanging information across a field [13]. Long-range field interactions are relevant when thermal forcing and neighboring grain competitors influence local growth. Spectral mode truncation removes fine spatial detail, and the periodic extension of a Fourier transform may introduce artificial communication at non-periodic boundaries. Local correction branches and padding can reduce these effects, with physical fidelity assessed against an independent reference. Resolution transfer therefore requires consistent representation of the physical interface at the target grid spacing.

Group-equivariant neural operators combine symmetry constraints with Fourier-based architectures to improve prediction under prescribed spatial transformations [14]. The cyclic group of quarter turns on a square lattice can be realized without resampling and is suitable for a controlled test of discrete rotation consistency. Continuous rotational equivariance requires a separate assessment beyond the discrete grid symmetry. Intermediate rotations can introduce interpolation error, altered boundary support, and spectral truncation effects. For directionally forced evolution, covariance requires the forcing and boundary metadata to transform jointly with the state. The output must retain the orientation information relevant to the physical problem. The constraint specifies the transformation of the output in the observation frame. This distinction is particularly relevant when anisotropic crystal kinetics interact with an imposed thermal gradient. The equivariant-operator literature supplies an architectural precedent, while the present study examines its adaptation to orientation-resolved solidification fields.

2.3. Orientation representations and validation gaps

Crystallographic orientation is represented through equivalence classes defined by crystal symmetry. Circular statistical methods treat angular observations on a periodic domain and require attention to bias arising from nonuniform angular sampling [15]. For the planar fourfold representation, multiplying the angle by the crystal symmetry order before forming a unit complex value preserves quarter-turn equivalence. The representation describes planar crystallographic orientation. In a three-dimensional extension, quaternion or rotation matrix representations would need to be introduced for an appropriate crystal symmetry group. Quaternion-based extensions require explicit treatment of crystal symmetry equivalence. Crystal equivalence identifies physically identical lattice orientations, whereas frame covariance specifies how the complete spatial state transforms under a change of coordinates.

Boundary-sensitive evaluation addresses localized geometric errors. Image-analysis validation frameworks recommend complementary overlap-based and boundary-based metrics when both region agreement and interface geometry are important [16]. Angular accuracy must be reported on physically meaningful solid regions together with the masking and rejection rules. Information-bottleneck uncertainty models for neural operators use confidence-aware latent representations and predictive output means and variances [17]. Their physical interpretation requires comparison with reference solutions and explicit consideration of correlations among pixels. Microstructure surrogates have therefore been evaluated over extended multistep prediction sequences to assess temporal stability under anisotropic grain growth [18]. The benchmark therefore evaluates one-step accuracy, late-rollout behavior, discrete rotation consistency, and transfer under altered thermal forcing separately. Separate reports keep the individual interpretations of the evaluation indicators.

3. Physical State and Equivariant Operator

3.1. State variables and crystal symmetry

The proposed state is defined on a square planar domain and contains a solid fraction, a normalized temperature, a vector thermal gradient, and a complex orientation field. The solid fraction distinguishes fully liquid material, diffuse interface material, and grain interiors. Temperature and solid fraction transform as spatial scalars, whereas the thermal gradient transforms as a vector under a change of frame. The orientation field stores a planar crystal axis modulo the

fourfold equivalence of the assumed crystal symmetry. Dataset assembly preserves this planar definition and requires a consistent projection of three-dimensional crystallographic measurements onto the stated orientation representation. Nucleation events and imposed boundary conditions are supplied consistently with the prediction interval. A joint rotation of all state and forcing fields changes the observation frame; changing the thermal-gradient direction relative to the grains changes the physical growth problem.

A symmetry compatible representation follows from the orientation encoding equation below. Here, z denotes the complex orientation field, θ is the crystallographic angle in radians, q is the crystal symmetry order, and i is the imaginary unit. The real and imaginary components provide the two orientation channels used in implementation. A binary validity mask identifies pixels whose solid fraction exceeds the reporting threshold. Liquid regions store a zero complex value, with angular interpretation restricted to valid solid pixels.

$$z = \exp(iq\theta), q = 4 \tag{1}$$

For a rotated observation frame, each spatial field is evaluated at the inverse rotated position. The orientation angle gains the frame rotation before symmetry reduction. This behavior is expressed through the rotation action equation, where R_α denotes the spatial rotation matrix and T_α is the corresponding transformation of the orientation field. At exact quarter turns, the phase factor equals unity because these rotations are crystal equivalent, while the orientation pattern still moves spatially. The thermal gradient components rotate jointly with the spatial fields. Evaluating this transformation separately from orientation perturbations preserves the distinction between a change of observation frame and a change in physical anisotropy.

$$[T_\alpha z](\mathbf{x}) = \exp(iq\alpha) z(R_\alpha^1 \mathbf{x}) \tag{2}$$

The state uses physical fields and boundary descriptors with explicitly defined transformation laws. Distances to physical boundaries can be supplied as scalar fields when boundary labels rotate jointly with the domain. Boundary driven temperature conditions are encoded through spatial boundary values that transform consistently with the observation frame. For the exact discrete tests, the square domain and its boundary data rotate together. Intermediate angle tests embed the domain in a larger padded square and restrict evaluation to the common interior support. These procedures distinguish interpolation free grid transformations from rotations affected by resampling and altered boundary support.

3.2. Group shared spectral blocks

The hidden representation contains four orientation indexed feature blocks associated with the cyclic group of quarter turns. A lifting map creates these blocks through shared operations applied to transformed inputs. Each spectral block uses a common underlying kernel across group related directions, together with the corresponding permutation of the orientation index. Complex Fourier coefficients satisfy the conjugate symmetry required for real hidden channels after the inverse transform. A radial mode cutoff maintains a rotation compatible spectral support. The local branch uses spatial kernels tied under the same group action, while nonlinearities operate identically across orientation indices. The retained spectral support must map onto itself under quarter-turn permutations; boundary handling and reconstruction require corresponding transformation checks.

Each group feature evolves according to the spectral update equation below. The indices g and h identify elements of the quarter turn group C_4, W_l denotes shared pointwise channel mixing, and σ is the activation function. The Fourier multiplier $K_l(g^{-1}h)$ includes the spatial frequency transformation associated with the output frame g; its frequency argument is suppressed for compactness. The summation couples the group indexed features through shared spectral interactions. Exact quarter turn transformations are implemented as grid index permutations.

$$v_{l+1}(g) = \sigma\left[W_l v_l(g) + \sum_{h \subset C_4} \mathcal{F}^{-1}(K_l(g^{-1}h)\mathcal{F}v_l(h))\right] \tag{3}$$

After the final block, a group compatible projection produces scalar phase and temperature channels together with the two orientation channels. Orientation outputs are normalized where their magnitude exceeds a numerical floor. Predictions below this floor are marked uncertain, and their rejection fraction is retained in evaluation. The phase output is clipped to its physical interval during rollout, with clipping frequency recorded as a stability diagnostic. These operations provide numerical safeguards and representation consistency. Validation against a physically justified reference establishes the fidelity of the learned dynamics. Coverage is reported with angular discrepancy so that exclusion of low-magnitude orientation vectors remains visible in the evaluation denominator.

Figure 1 presents the state transformation and architecture workflow, connecting scalar and vector inputs to group lifting, shared spectral blocks, local interface correction, and masked orientation reconstruction. The thermal gradient and microstructure rotate jointly along the input transformation path. A separate comparison path applies the inverse

frame transformation to the prediction. The local and spectral branches merge before the prediction heads, allowing interface detail and long-range grain competition to contribute to the same state update. The reconstruction stage retains the solid validity mask so that angular outputs remain associated with physically meaningful material regions.

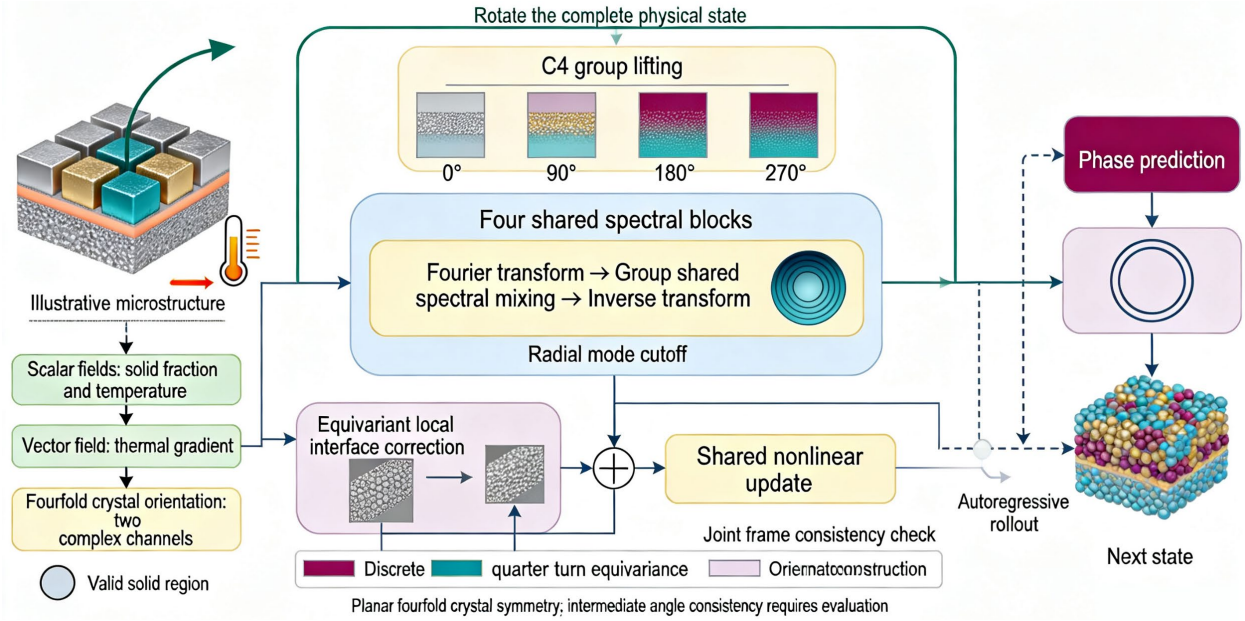


Figure 1. Joint frame transformation and rotationally equivariant Fourier operator architecture

3.3. Angular objectives and numerical scope

Angular discrepancy is evaluated on a common valid solid mask and reduced to the fundamental fourfold interval. The angular distance equation uses the wrapped phase difference between the predicted orientation $\hat{z}$ and reference orientation z, with the asterisk denoting complex conjugation. Dividing by the symmetry order returns the discrepancy on the physical angular scale. Consequently, orientations near opposite boundaries of the chosen interval are compared continuously through their crystal equivalence. Evaluation converts the resulting angle to degrees. Training uses a smooth chordal counterpart to limit sensitivity to the derivative discontinuity at the largest wrapped difference.

$$d_\theta(\hat{z}, z) = \frac{|\mathrm{Arg}(\hat{z}z^*)|}{q} \tag{4}$$

The common mask restricts the angular objective to valid solid material while introducing a coverage dependence that requires separate reporting. A predictor could obtain a smaller angular error by classifying difficult regions as liquid. Phase overlap, missed solid area, and orientation rejection fraction are therefore evaluated alongside angular discrepancy. Reference solid pixels omitted by the predicted mask contribute to the phase failure rate. Boundary regions and grain interiors are also assessed separately. This spatial conditioning helps identify whether a smaller average error reflects improved grain identification, more accurate interface placement, or a change in the evaluated solid area. Angular accuracy characterizes retained solid predictions, while missed reference-solid regions contribute separately to phase error and coverage loss.

The architecture enforces the discrete quarter turn action for the stated domain and field types. Continuous rotation consistency requires evaluation across interpolated observation frames. Solute conservation, energy evolution, and nucleation statistics require dedicated physical constraints and validation. Introducing a concentration field would require a conservation treatment and a transport equation consistent with the thermal and phase variables. The present formulation focuses on orientation and phase prediction, allowing geometric constraints to be examined within a defined state representation. Its scope also requires consistent treatment of boundary labels, interface masks, and crystallographic equivalence throughout data preparation, training, and autoregressive reconstruction. The architectural constraint applies to the quarter-turn grid group; consistency at intermediate angles is assessed numerically on the interpolated common support.

4. Training and Reproducible Evaluation Protocol

4.1. Data separation and optimization design

The solidification trajectories include metadata for initial grains, thermal forcing, mesh spacing, time step, and boundary conditions. Training, validation, and test partitions are assigned by trajectory identifier before temporal

windows are extracted. All augmented versions of a trajectory remain in the same partition. The dataset contains 1200 training trajectories, 150 validation trajectories, and 150 test trajectories, with 21 stored states per trajectory. Validation and test partitions support model selection and held-out assessment, respectively. An additional 120 trajectories support thermal-shift testing, while a further 120 trajectories support evaluation under altered nucleation density. Partition membership is fixed at the parent-trajectory level before temporal slicing and augmentation, and every derived window inherits that membership.

Optimization combines masked orientation error, phase reconstruction, interface gradient matching, and paired rotation consistency. These terms enter the objective equation through four nonnegative weights. The orientation term $\mathcal{L}_\theta$ evaluates crystal symmetry compatible discrepancy, $\mathcal{L}_\phi$ measures phase reconstruction, $\mathcal{L}_\mathrm{b}$ emphasizes interface gradients, and $\mathcal{L}_\mathrm{r}$ penalizes inconsistency under matched transformations. Supervised terms use reference masks fixed before training. The rotation term compares predictions obtained from consistently transformed versions of the same input. For an exact discrete group implementation, its residual primarily diagnoses numerical padding, reconstruction, or boundary handling effects.

$$\mathcal{L} = \lambda_\theta \mathcal{L}_\theta + \lambda_\phi \mathcal{L}_\phi + \lambda_\mathrm{b} \mathcal{L}_\mathrm{b} + \lambda_\mathrm{r} \mathcal{L}_\mathrm{r} \tag{5}$$

The optimizer configuration uses an initial learning rate of 0.001, a batch size of eight trajectory windows, and 200 epochs with cosine decay. Five independent seeds support assessment of initialization sensitivity. Validation selection uses a fixed weighted combination of angular and interface errors established before test evaluation. Gradient clipping at a norm of one limit unstable update. The architecture contains four spectral blocks, 24 retained radial modes, and 64 aggregate hidden channels distributed across the four group indices. Baselines receive parameter budgets within ten percent of the proposed configuration and an equal hyperparameter search allowance. Reporting both aggregate width and width per group makes channel allocation and computational comparisons explicit. The stated aggregate width of 64 corresponds to 16 hidden channels per group index, and parameter-budget comparisons refer to the aggregate model.

4.2. Rollout and transformation tests

Autoregressive evaluation advances the predicted phase and orientation while supplying the imposed thermal schedule from the test case. Each subsequent prediction consumes the preceding predicted state. Errors are sampled at five normalized times spanning the early transient and final stored state. A separate teacher forced evaluation diagnoses one step accuracy alongside the autoregressive assessment. Rotation testing begins with the three nonidentity quarter turns, for which spatial permutation is exact. Intermediate rotations from fifteen to seventy-five degrees are evaluated on shared cropped support using a common interpolation procedure for reference and predicted fields. Their residual therefore combines interpolation, architectural, and reconstruction contributions. Intermediate-angle residuals include resampling and common-support effects in addition to prediction error, whereas quarter-turn permutations provide the interpolation-free architectural diagnostic.

To quantify transformation consistency, the residual equation compares prediction after input rotation with rotation after prediction. The evolution map $\mathcal{G}$ acts on the complete state u, including the scalar fields, orientation channels, thermal gradient, and boundary descriptors. The norm uses standardized channels on the common evaluation support. A positive denominator floor ε stabilizes the ratio for states with small output magnitude. Exact grid rotations and interpolated rotations receive separate reporting labels because their numerical error floors differ. Multiple grains, directional gradients, and irregular interfaces exercise the complete transformation path.

$$\varepsilon_\mathrm{eq} = \frac{\| \mathcal{G}(\mathrm{T}_\alpha \mathrm{u}) - \mathrm{T}_\alpha \mathcal{G}(\mathrm{u}) \|_2}{\| \mathrm{T}_\alpha \mathcal{G}(\mathrm{u}) \|_2 + \varepsilon} \tag{6}$$

Figure 2 describes the evaluation workflow from trajectory preparation through partition assignment, temporal window extraction, optimization, fixed validation selection, and held out diagnostics. Separate branches distinguish exact grid rotations from interpolated rotations. Further branches route shifted thermal schedules and nucleation densities to process transfer evaluation. Confidence intervals are calculated by resampling complete trajectories within each seed, preserving spatial and temporal dependence. The same resampled trajectory identifiers are retained across competing configurations to support paired comparisons. Checkpoint errors, prediction coverage, and phase reconstruction are assessed together so that temporal instability can be localized to specific trajectories and interface regions. The resampling description defines the uncertainty-estimation procedure; numerical confidence intervals and the associated resampling settings remain necessary for statistical comparison.

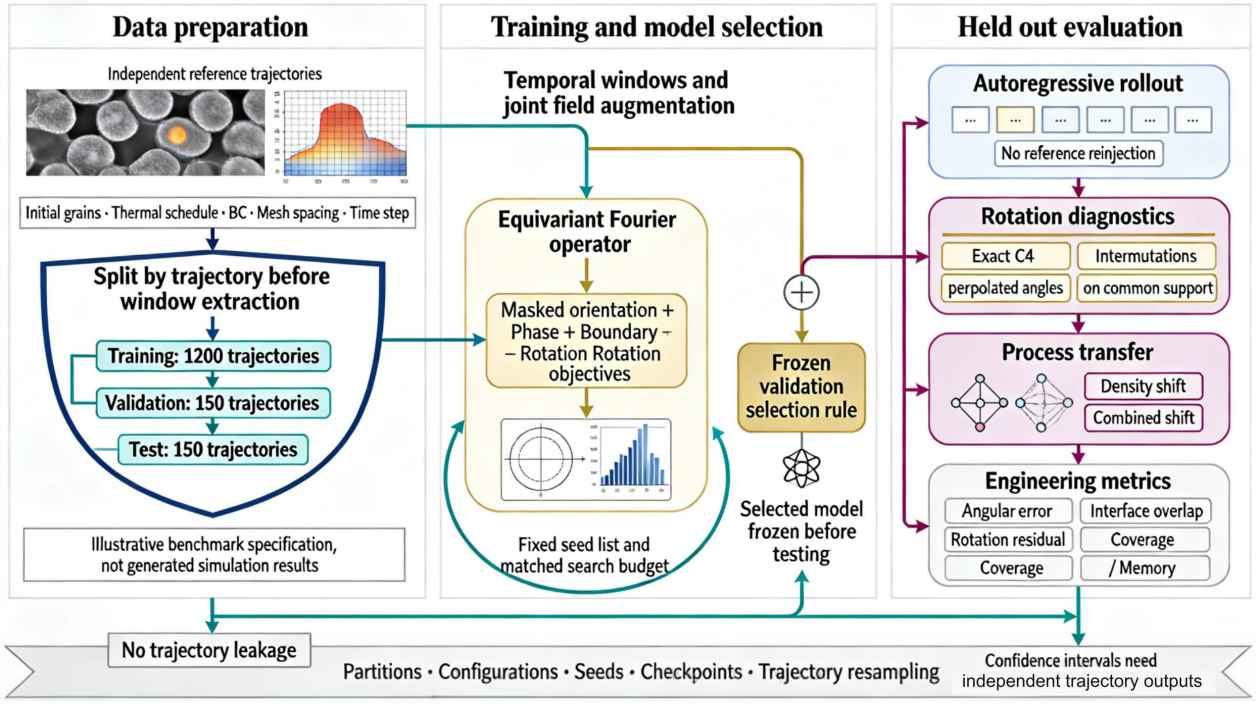


Figure 2. Training and evaluation workflow with trajectory level separation, fixed validation selection, autoregressive rollout, exact and interpolated rotation branches, and held out process shift tests

4.3. Reporting controls and reproducibility

The implementation records the configuration, partition manifest, random seeds, software versions, and checkpoints selected by the validation rule. The partition manifest contains trajectory identifiers and aggregate counts to support overlap detection. Every output retains its initial condition and thermal schedule identifiers. Prediction timing includes field transfer and reconstruction, while reference timing covers the same output interval with a consistent initialization boundary. Parameter count, peak memory, batch size, hardware, numerical precision, and mesh dimensions accompany the timing measurements. These details establish the computational context needed to relate reconstruction accuracy to the demands of each configuration. Independent reproduction additionally requires the numerical mask and normalization thresholds, loss coefficients, optimizer settings, seed values, and selected checkpoint identifiers.

Three controls examine the proposed modeling choices. Replacing the periodic orientation encoding with a raw scalar angle tests sensitivity to wrap boundaries. Removing group parameter sharing while retaining augmentation examines the contribution of the architectural constraint under comparable exposure to transformed inputs. Removing the interface branch tests local reconstruction at a fixed spectral bandwidth. Each control uses the same trajectory partitions, seed list, optimization budget, and validation rule. Checkpoint selection follows an identical procedure across configurations, preserving the distinction between architecture and model selection effects. Resolution transfer tests retain the physical domain and interface thickness while changing grid spacing, allowing mesh dependence to be assessed on a consistent physical basis. The encoding, group-sharing, and interface-branch controls respectively examine angular continuity, transformation consistency, and local boundary reconstruction.

The plotting package maintains agreement among plotted points, distributions, tables, and numerical comparisons. Each panel includes units or a normalized scale, and all panels use Times New Roman at a font size of 13. Figures are exported at a width of 800 pixels, with raster files carrying 300 dpi metadata and SVG files retaining editable vector elements. Axis ticks, numerical annotations, thresholds, and legends follow consistent conventions across related panels. Evaluation arrays are versioned together with configuration files and partition manifests. Revisions propagate to the manuscript, tables, and associated plots through the same data definitions, preserving traceability between numerical statements and their visual presentation.

5. Experimental Data and Analysis

5.1. Benchmark specification and comparative accuracy

The benchmark specifies 1200 training trajectories, 150 validation trajectories, and 150 test trajectories, with 21 stored states per trajectory on a nominal 128 by 128 grid. Separate thermal and density shift collections each contain 120 additional trajectories. Table 1 records the physical domain, cooling schedules, and nucleation densities used in the proposed evaluation. Angular accuracy, spatial overlap, and temporal sensitivity are examined using the same

configuration definitions and evaluation thresholds. Phase-field solidification models can exhibit unstable or distorted interface profiles under large driving forces, motivating explicit treatment of interface stability [19]. Accordingly, reference calculations should check interface stability and mesh convergence at each process condition before assembling learning targets. The five model configurations span convolutional recurrence, conventional spectral prediction, augmentation, group sharing, and the complete architecture. Their mean angular errors are 8.4, 6.8, 5.1, 3.6, and 2.8 degrees. Reference-solver convergence and surrogate-to-reference agreement are distinct checks; the latter inherits any unresolved discretization uncertainty in the targets.

Table 1. Benchmark allocation and processing conditions

Partition or setting	Specification
Training / validation / test	1200 / 150 / 150 trajectories
Thermal / density shift	120 / 120 additional trajectories
Stored states / nominal grid	21 / 128 by 128
Domain width	1 mm
Cooling rates	5, 10, 15, 20, 25 K/s
Shifted cooling rates	30 and 35 K/s
Nucleation densities	20, 40, 60, 80, 100 grains/mm2

Figure 3(a) compares the five mean angular errors against a five-degree reporting threshold. The conventional spectral configuration exceeds this threshold by 1.8 degrees, whereas the full configuration remains 2.2 degrees below it. The 4.0-degree difference corresponds to a 58.8 percent relative reduction within the comparison. The augmented configuration reaches 5.1 degrees, placing it close to the threshold and making its classification sensitive to trajectory variability. Figure 3(b) examines the associated distributions, with medians matching the means and interquartile intervals extending from approximately 4.78 to 8.82 degrees for the conventional operator and from 1.97 to 3.63 degrees for the full configuration. The interquartile width decreases from 4.04 to 1.66 degrees, indicating that the difference extends across the central portion of the error distribution. The lower center and narrower spread place more full configuration samples within the angular reporting target. The comparison should retain paired trajectory identifiers and inspect whether the apparent benefit is concentrated in a few difficult grain competitions. Large interiors can otherwise dominate sample counts and obscure errors associated with narrow grains near the advancing front. The five-degree guide is an evaluation threshold; its relation to a material acceptance requirement depends on the downstream engineering observable.

Figure 3(c) relates angular error to interface overlap, showing that the increase from 0.861 to 0.934 accompanies a decrease from 6.8 to 2.8 degrees. The overlap difference of 0.073 is smaller in numerical magnitude than the angular difference because the metrics have distinct scales and meanings. Overlap-based and boundary-based image metrics capture different error properties, and their interpretation depends on the evaluation task and target geometry [20]. The comparison therefore retains angular and geometric metrics as complementary measures. The group shared configuration occupies an intermediate position at 0.918 overlap and 3.6 degrees, suggesting a useful control for separating local reconstruction from the shared spectral representation. In this comparison, adding the interface branch changes overlap by 0.016 and angular error by 0.8 degrees. This pairing motivates examination of boundary displacement and label identity at the same physical locations. Such paired inspection would clarify whether local corrections preserve a narrow competitor or merely smooth a misplaced interface. The applicable interpretation depends on grain identity tracking through time, particularly where neighboring grains have similar symmetry reduced orientations. Overlap evaluates occupied-region geometry, whereas angular discrepancy evaluates crystallographic identity within the valid solid region; their numerical magnitudes have different units and scales.

Figure 3(d) resolves the cumulative distribution of angular errors. Across the evaluated samples, 28 percent of conventional spectral samples and all full configuration samples fall below five degrees. The augmented and group shared configurations yield corresponding fractions of 48 and 82 percent. These percentages expose threshold sensitivity that a mean alone would conceal. Bootstrap inference for dependent time series requires resampling procedures that account for temporal dependence [21]. Resampling complete trajectories would reveal whether the threshold fractions remain stable across independent initial grain arrangements. Each configuration is represented by 100 angular discrepancy samples. The threshold fractions depend on both the distribution center and its upper tail, making the comparison sensitive to the frequency of larger local errors. A useful physical evaluation should retain the largest errors, the affected grain areas, and the subsequent survival of those grains. A small population of highly misoriented cells near a competitive boundary may alter the predicted texture even when the global threshold fraction remains favorable. This case requires spatially resolved examination of the error tail and its temporal persistence. The cumulative fractions

describe the reported angular-discrepancy samples; trajectory-level reliability additionally requires aggregation that preserves spatial and temporal dependence.

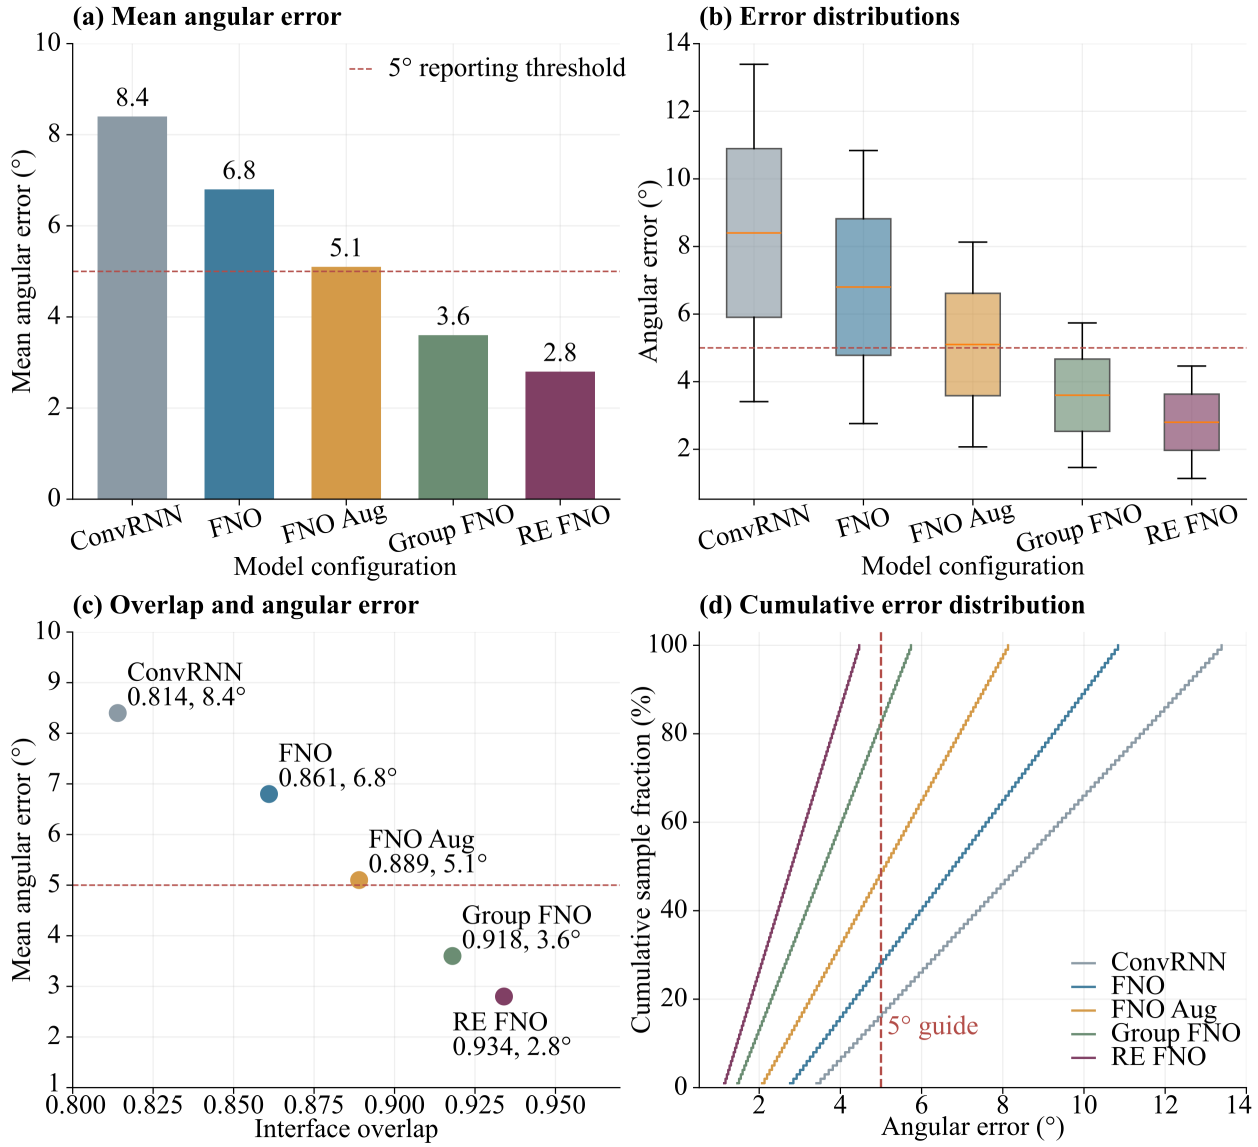


Figure 3. Comparative performance. (a) Mean angular errors. (b) Angular error distributions. (c) Interface overlap versus angular error. (d) Cumulative angular error distributions

Table 2 pairs these values with interface overlaps of 0.814, 0.861, 0.889, 0.918, and 0.934. The paired measurements distinguish orientation identification from geometric reconstruction. A displaced interface changes which grain occupies a location, whereas a mislabeled interior can retain the correct boundary geometry. These mechanisms warrant separate measurements because their implications for grain survival differ. The performance ordering is examined through paired angular and geometric metrics to identify the spatial origin of the differences.

Table 2. Comparative performance. Rotation residual uses interpolated rotations

Model	Error deg	Overlap	Residual	Latency ms
ConvRNN	8.4	0.814	0.174	8
FNO	6.8	0.861	0.126	12
FNO Aug	5.1	0.889	0.061	14
Group FNO	3.6	0.918	0.022	16
RE FNO	2.8	0.934	0.014	19

5.2. Rotation diagnostics rollout behavior and component sensitivity

The rotation analysis separates exact quarter turn permutations from intermediate angles that require interpolation. On regular grids, exact equivariance pertains to the discrete symmetry group of the grid, which must be distinguished from continuous spatial rotations [22]. Figure 4(a) presents seven sampled angles from zero to ninety degrees. The full configuration residual increases from 0.002 at the endpoints to 0.016 at forty-five degrees, remaining below the 0.02 reporting tolerance. The conventional sequence reaches 0.141 at forty-five degrees, giving a peak ratio of approximately 8.8 relative to the full configuration. Both profiles are symmetric around the midpoint. This shared shape provides a controlled setting for inspecting residual magnitude while holding angle dependence fixed. Figure 4(b) evaluates a separate source of sensitivity through five autoregressive checkpoints. Conventional spectral errors increase from 2.3 to 10.8 degrees, compared with 1.2 to 4.7 degrees for the full configuration. The early to late increases are 8.5 and 3.5 degrees, respectively. The need to represent evolving spatial structure has motivated latent-space operator models for complex time-dependent physical systems [23]. The separation therefore motivates checking whether boundary displacement grows gradually or appears abruptly after a grain identity change. Recording orientation magnitude and clipping frequency would help distinguish these behaviors. Endpoint residuals characterize the evaluated prediction pipeline, including preprocessing and reconstruction; an ideal layer-level identity should be checked separately from that aggregate diagnostic.

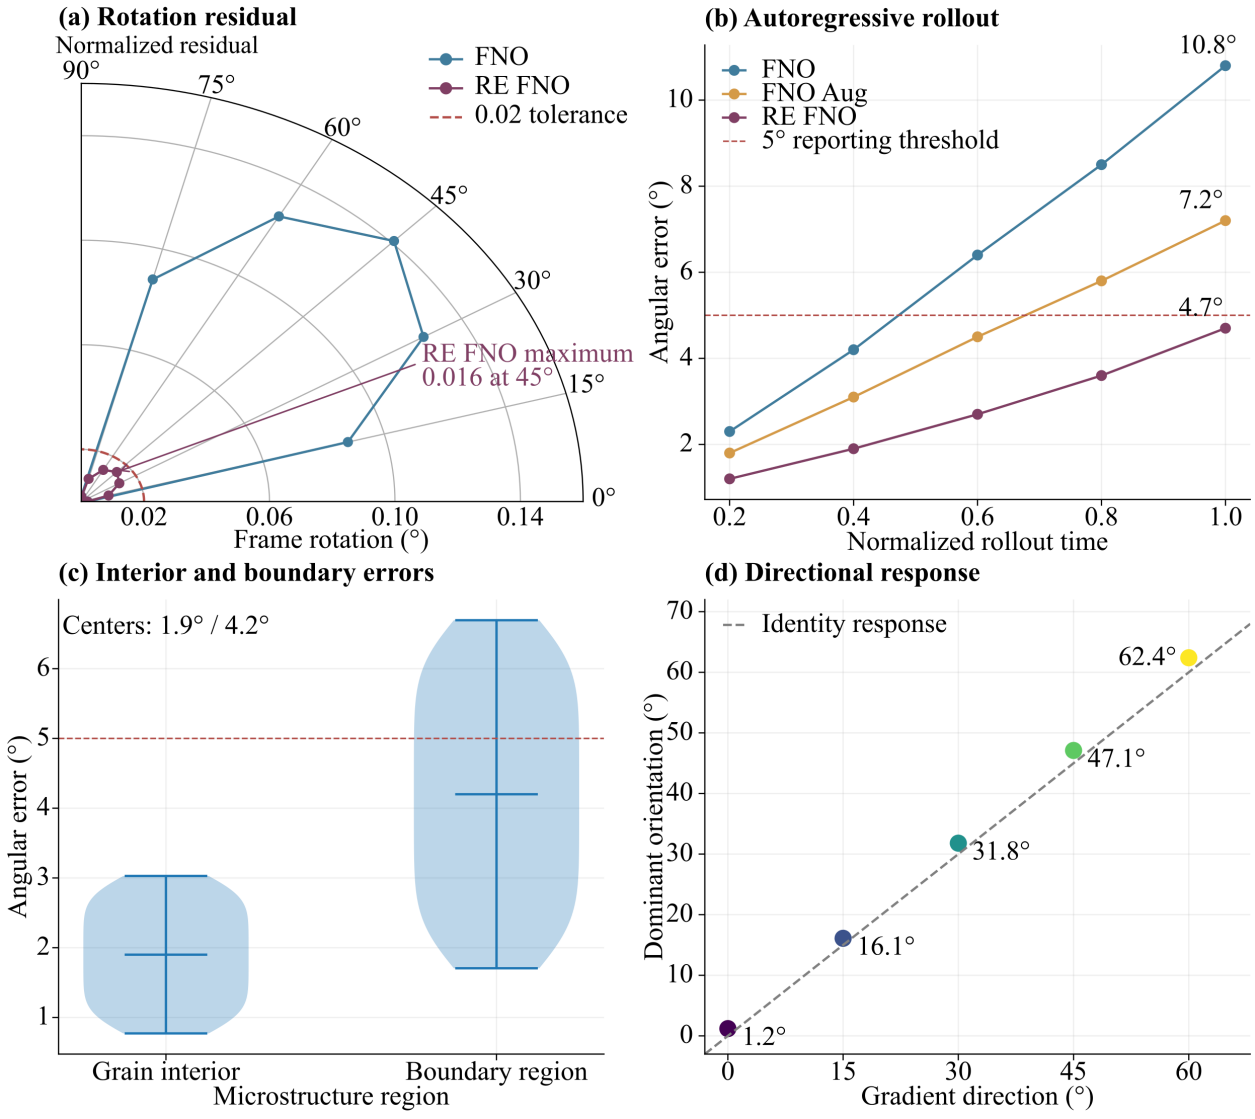


Figure 4. Rotation and rollout diagnostics. (a) Rotation residual profiles. (b) Rollout errors. (c) Interior and boundary distributions. (d) Directional response

Figure 4(c) localizes the full configuration error to grain interiors and boundaries, with distribution centers of 1.9 and 4.2 degrees. Their ratio of approximately 2.2 identifies a substantial interface sensitivity within the comparison. The boundary errors extend beyond the five-degree guide while the interior distribution remains below it. This contrast indicates which locations would merit closer inspection in a physical validation. Boundary-guided segmentation networks explicitly incorporate boundary information to improve delineation of target structures [24]. Their

effectiveness should be assessed together with phase coverage because changing the solid mask also changes which angular targets are evaluated. Figure 4(d) examines directional response at imposed gradient directions of zero, fifteen, thirty, forty-five, and sixty degrees. The dominant orientations are 1.2, 16.1, 31.8, 47.1, and 62.4 degrees, giving positive offsets of 1.2, 1.1, 1.8, 2.1, and 2.4 degrees. The increasing offset at larger directions motivates a check for interpolation or reconstruction bias after consistent angular unwrapping. The dominant direction reflects the changing distribution of occupied grains; its evolution should be interpreted through grain selection and boundary motion. The dominant orientation records the distribution of surviving grain regions under thermal forcing and therefore represents selection behavior rather than a lattice rotation rate.

Figure 5(a) compares five component configurations with angular errors of 5.1, 4.7, 3.6, 3.3, and 2.8 degrees. Removing group sharing produces the largest increase, at 2.3 degrees above the complete model. Replacing the periodic orientation encoding adds 1.9 degrees, while removing the local branch or boundary loss adds 0.8 or 0.5 degrees. These differences identify candidate mechanisms for investigation, with representation continuity and geometric sharing receiving distinct controls. Automated hyperparameter exploration has been introduced for neural operators to support systematic investigation of architecture and training choices [25]. Each variant should consequently be trained with the same trajectory partitions and seed list. Figure 5(b) examines thermal sensitivity across cooling rates of 5, 10, 15, 20, and 25 kelvin per second. The full configuration errors rise from 2.2 to 4.1 degrees, while conventional errors rise from 5.4 to 9.7 degrees. The endpoint increases are 1.9 and 4.3 degrees, respectively. The full configuration remains below the five-degree guide throughout this interval, whereas the conventional configuration exceeds it at every sampled rate. Intermediate observations reveal a progressively steeper increase in both sequences toward stronger cooling, motivating finer sampling near the upper end of the proposed operating range.

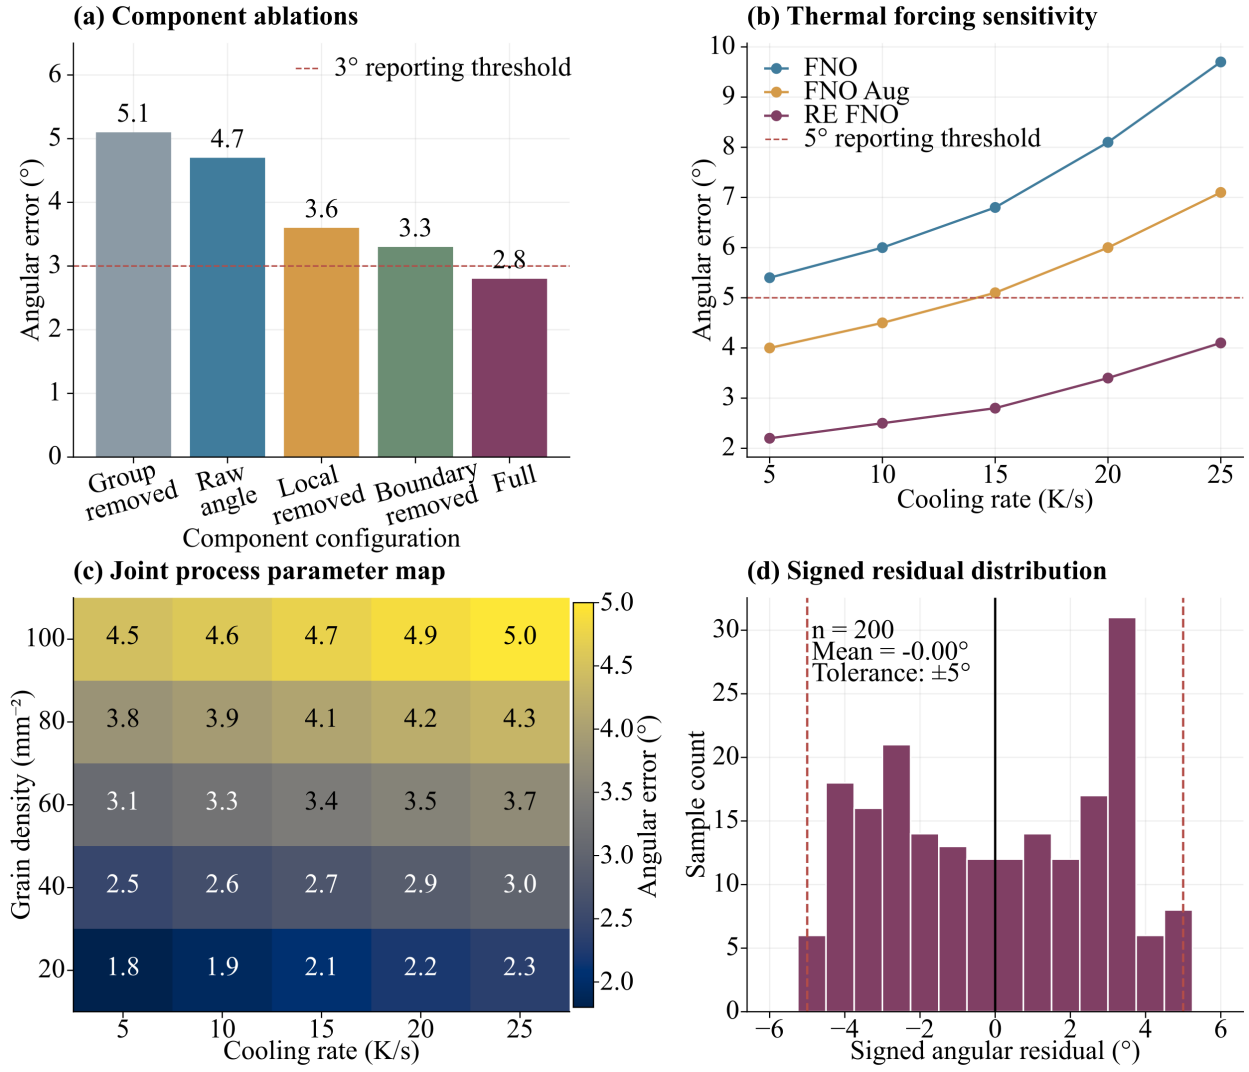


Figure 5. Component and process sensitivity. (a) Component ablations. (b) Cooling rate sensitivity. (c) Cooling rate and nucleation density map. (d) Signed residual histogram

Figure 5(c) extends this sensitivity analysis to twenty-five combinations of cooling rate and nucleation density. The error ranges from 1.8 degrees at the lowest rate and density to 5.0 degrees at the highest combination. Along the lowest

density row, increasing cooling rate changes the error from 1.8 to approximately 2.3 degrees; along the highest density row, the corresponding change is approximately 4.5 to 5.0 degrees. The change across density levels exceeds the change across cooling rates within individual rows. This pattern concentrates attention on the relation between competition density and boundary reconstruction under a fixed spatial resolution. Figure 5(d) examines 200 signed residuals with a mean of zero and guides at plus and minus five degrees. The distribution contains both signs, allowing bias and spread to be considered separately. A vanishing signed mean can coexist with appreciable absolute errors when positive and negative deviations balance. Direction conditioned residual histograms would help determine whether this balance persists across gradient orientations or masks opposing directional biases. At the high rate and high-density corner, reference mesh refinement and time step sensitivity should be checked alongside the learned residuals to identify the contribution of numerical target uncertainty. The parameter map describes the evaluated rate-density combinations; interpolation between those combinations requires additional validation of error behavior and reference resolution.

5.3. Training data availability process transfer and computational cost

The training fraction analysis uses nested allocations of 10, 25, 50, 75, and 100 percent of the trajectory collection. Figure 6(a) reports full configuration errors of 5.9, 4.5, 3.6, 3.1, and 2.8 degrees. The corresponding conventional spectral values are 10.2, 8.8, 7.7, 7.2, and 6.8 degrees. Increasing the allocation from ten to twenty five percent changes the full configuration error by 1.4 degrees; increasing it from seventy-five to one hundred percent changes the error by 0.3 degrees. The decreasing marginal response concentrates attention on coverage of rare grain competition patterns as the training allocation increases. Standardized evaluation pipelines have been developed to examine how predictive performance depends on training sample size [26]. The proposed nested manifests preserve that consistency while allowing added trajectories to be traced to their process conditions and initial grain arrangements. The full configuration crosses the five-degree guide at the twenty five percent allocation, whereas the conventional configuration remains above it at the complete allocation. The interpretation of such a crossing would depend on trajectory level variability and on whether each subset adequately represents difficult boundary interactions. Dataset size alone leaves the diversity of grain competition patterns unspecified. A smaller marginal gain with added trajectories can reflect limited new pattern diversity, model capacity, or optimization constraints, so the curve alone leaves the saturation mechanism unresolved.

Figure 6(b) compares nominal, rotated, thermal shift, density shift, and combined shift regimes, with central errors of 2.8, 3.0, 4.4, 4.8, and 6.2 degrees for the full configuration. The nominal to rotated increase is 0.2 degrees, while the nominal to combined shift increase is 3.4 degrees. This contrast separates sensitivity to an observation frame change from sensitivity to altered physical conditions. Neural-operator risk analyses explicitly distinguish learning under the training distribution from prediction under shifted input distributions [27]. Here the thermal and density shifts alter the process specification, whereas the jointly rotated inputs preserve the intended physical configuration. The thermal and density centers lie below five degrees but their upper tails extend beyond that guide. The combined regime has a central value above it, concentrating attention on prediction coverage and fallback behavior. Approximately 34 percent of combined shift samples remain below five degrees, compared with complete coverage under nominal conditions. Process metadata should be retained so that the affected trajectories could be inspected for changes in front velocity, competition duration, and surviving grain count. Joint frame rotations preserve the specified physical configuration, while thermal-magnitude and nucleation-density shifts change the evolution problem and require separate transfer interpretation.

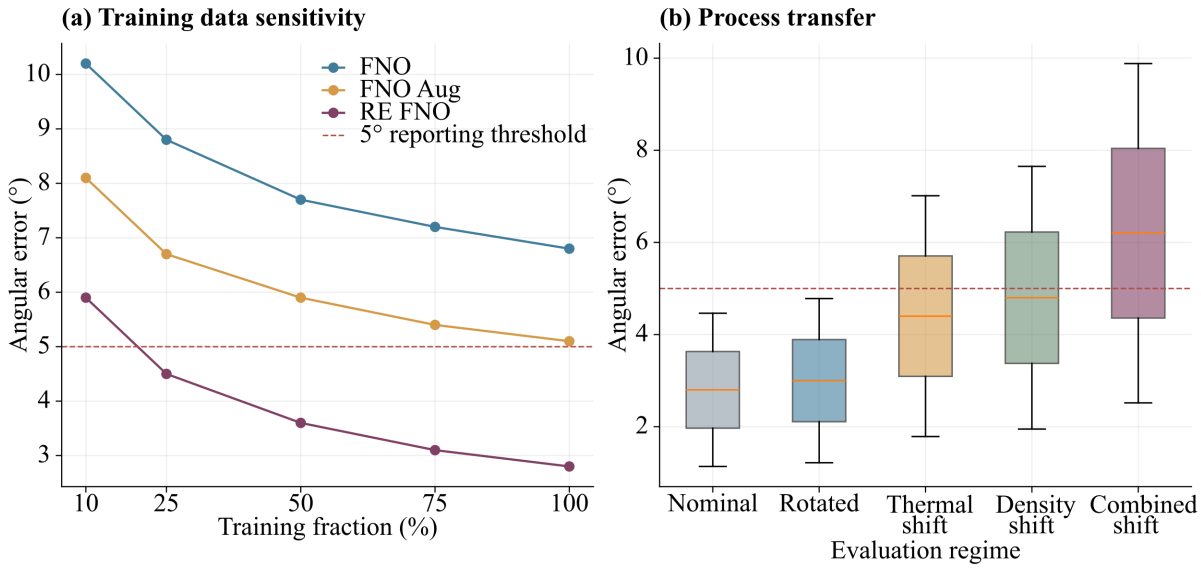


Figure 6. Training availability and transfer. (a) Nested learning curves. (b) Nominal and shifted regime distributions

Figure 7 places the five configurations on an angular accuracy and latency plane, with latencies of 8, 12, 14, 16, and 19 milliseconds. Marker areas encode parameter counts of 1.8, 2.1, 2.1, 2.4, and 2.6 million. The complete configuration trades an additional seven milliseconds relative to the conventional spectral configuration for a 4.0-degree reduction in error. Relative to the group shared variant, the corresponding changes are three milliseconds and 0.8 degrees. Both configurations are in the shaded two-to-five-degree reporting band; therefore, a practical choice will be based on the allowed angular error and required response speed. Inference benchmarking shows that throughput and efficiency depend on accelerator scale, architecture, memory, host processors, and workload scenario [28]. These operating conditions influence whether the latency differences persist under practical workloads. The accuracy and latency coordinates reveal the computational cost of local reconstruction and group-indexed feature processing. Parameter sharing reduces parameter storage, while additional group operations can increase execution cost. Therefore, peak activation memory and transform workspace should be considered in the assessment of a deployable implementation. Deployment conclusions require matched hardware, precision, batching, and timing scope; the reported latency coordinates should be interpreted within those measurement conditions.

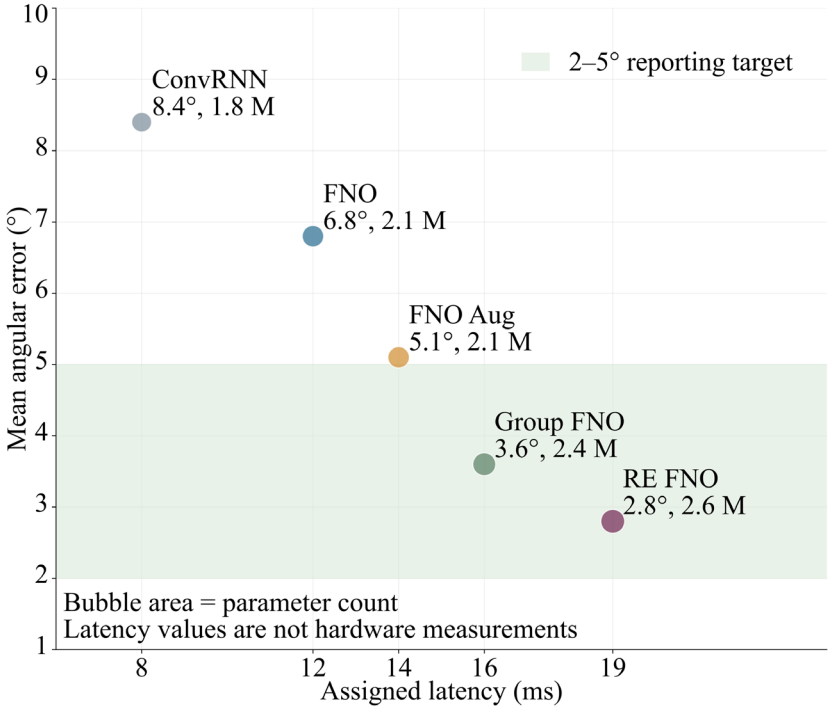


Figure 7. Accuracy and computational cost

Process transfer and computational cost are related to the handling of uncertain predictions. A fallback policy that uses a reference solver for difficult trajectories adds a regime-dependent cost to the surrogate call. The combined shift distribution provides a suitable environment for the study of this policy because its upper tail is substantially larger than that of the nominal distribution. The evaluation should record the number of rejected states, the duration of each numerical correction, and the length of the subsequent stable rollout. Correction at isolated checkpoints can reduce trajectory restart costs, with its effectiveness depending on consistent state transfer between the solver and surrogate. State transfer requires consistent conventions for the orientation field, solid fraction, and thermal metadata. Latency can also be measured with batch sizes that are representative of the process exploration and advance multiple trajectories simultaneously. The ranking based on single-trajectory latency may change when memory limits constrain batching. The existing five parameter counts and latency points provide a starting set of configurations for these measurements, and high-density and combined-shift cases are used to generate demanding conditions for tracking memory use, rejection frequency, and time spent in reconstruction. Solver fallback is an application strategy whose effective cost depends on rejection frequency, state-transfer overhead, and numerical correction duration.

6. Conclusion

A rotationally equivariant Fourier neural operator was developed in this paper to address the problem of planar grain orientation evolution during solidification, and, as shown by benchmark experiments, it improved angular accuracy, interface reconstruction and frame-change consistency.

1. Orientation accuracy and interface reconstruction. The mean angular error of the proposed configuration was 2.8°, and that of the conventional Fourier neural operator was 6.8°; thus, a reduction of 58.8% was achieved. Interface overlap

rose from 0.861 to 0.934. The paired improvements indicate complementary contributions of orientation representation and local interface correction to grain identification and boundary reconstruction.

2. Rotation consistency and temporal performance. The normalized rotation residual decreased from 0.126 to 0.014, while the final rollout error decreased from 10.8° to 4.7°. Across cooling rates from 5 to 25 K/s, the proposed configuration maintained angular errors between 2.2° and 4.1°, whereas the conventional operator ranged from 5.4° to 9.7°. The comparison identifies improved tolerance to observation frame changes alongside reduced error accumulation during temporal prediction.

3. Applicable conditions and remaining limitations. Four-fold crystal orientation encoding and discrete quarter-turn equivariance are used in the formulation. The present formulation covers planar microstructures and evaluates intermediate rotations through interpolation-sensitive diagnostics. Under combined thermal and nucleation-density shifts, the central angular error reached 6.2°, indicating sensitivity to training coverage under altered physical conditions. The combined-shift error identifies a practical validation priority: coverage of coupled thermal and nucleation conditions beyond nominal training settings.

7. Declarations

7.1 Author Contributions

Zdravka Stamenković contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Arsenija Mitrović contributes to validation, analysis, investigation, data collection, draft preparation. Gojka Vasiljev contributes to conceptualization, validation, analysis. Kraljevka Novčić contributes to validation, investigation, data collection. All authors have read and agreed with the manuscript before its submission and publication.

7.2 Funding and Acknowledgments

This research received no specific financial support from any funding agency.

7.3 Institutional Review Board Statement

Not applicable.

7.4 Informed Consent Statement

Not applicable.

7.5 Date Availability

Not applicable.

7.6 Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

7.7 Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammarly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

8. References

[1] Herbeaux, A., Aboleinein, H., Villani, A., Maurice, C., Bergheau, J.-M., & Klöcker, H. (2024). Combining phase field modeling and deep learning for accurate modeling of grain structure in solidification. Additive Manufacturing, 81, Article 103994. https://doi.org/10.1016/j.addma.2024.103994

[2] Boukellal, A. K., Sarebanzadeh, M., Orozco-Caballero, A., Sket, F., LLorca, J., & Tourret, D. (2024). Grain growth competition and formation of grain boundaries during solidification of hcp alloys. Acta Materialia, 269, Article 119830. https://doi.org/10.1016/j.actamat.2024.119830

[3] Secer, R. O., Hossein Nedjad, S., & Yildiz, M. (2024). Extending a phase field model of polycrystalline solidification for simulating grain evolution in metal additive manufacturing. Computational Materials Science, 240, Article 112997.

https://doi.org/10.1016/j.commatsci.2024.112997

[4] Ciesielski, D., Li, Y., Hu, S., King, E., Corbey, J., & Stinis, P. (2025). Deep operator network surrogate for phase-field modeling of metal grain growth during solidification. Computational Materials Science, 246, Article 113417. https://doi.org/10.1016/j.commatsci.2024.113417

[5] Azizzadenesheli, K., Kovachki, N., Li, Z., Liu-Schiaffini, M., Kossaifi, J., & Anandkumar, A. (2024). Neural operators for accelerating scientific simulations and design. Nature Reviews Physics, 6(5), 320–328. https://doi.org/10.1038/s42254-024-00712-5

[6] Liu, X., Xu, B., Cao, S., & Zhang, L. (2024). Mitigating spectral bias for the multiscale operator learning. Journal of Computational Physics, 506, Article 112944. https://doi.org/10.1016/j.jcp.2024.112944

[7] Harder, H., Nüske, F., Philipp, F. M., Schaller, M., Worthmann, K., & Peitz, S. (2025). Group-convolutional extended dynamic mode decomposition. Physica D: Nonlinear Phenomena, 480, Article 134725. https://doi.org/10.1016/j.physd.2025.134725

[8] Jangid, D. K., Brodnik, N. R., Echlin, M. P., Gudavalli, C., Levenson, C., Pollock, T. M., Daly, S. H., & Manjunath, B. S. (2024). Q-RBSA: High-resolution 3D EBSD map generation using an efficient quaternion transformer network. npj Computational Materials, 10(1), Article 27. https://doi.org/10.1038/s41524-024-01209-6

[9] Meilán-Vila, A., Crujeiras, R. M., & Francisco-Fernández, M. (2024). Nonparametric estimation for a functional-circular regression model. Statistical Papers, 65(2), 945–974. https://doi.org/10.1007/s00362-023-01420-5

[10] Diab, W., & Al Kobaisi, M. (2025). Temporal neural operator for modeling time-dependent physical phenomena. Scientific Reports, 15(1), Article 32791. https://doi.org/10.1038/s41598-025-16922-5

[11] Kim, G., Takaki, T., Shibuta, Y., Ko, H., & Ohno, M. (2024). Morphological diversity in directionally-solidified microstructures with varying anisotropy of solid-liquid interfacial free energy. Journal of Materials Research and Technology, 30, 4044–4052. https://doi.org/10.1016/j.jmrt.2024.04.126

[12] Zhu, C., Liu, S., Gao, Z., Wang, L., & Miao, J. (2024). Spatiotemporal evolution of grain microstructure: A CNN perspective. Materials Today Communications, 40, Article 110005. https://doi.org/10.1016/j.mtcomm.2024.110005

[13] Liu, S., Yu, Y., Zhang, T., Liu, H., Liu, X., & Meng, D. (2025). Architectures, variants, and performance of neural operators: A comparative review. Neurocomputing, 648, Article 130518. https://doi.org/10.1016/j.neucom.2025.130518

[14] Erzmann, D., & Dittmer, S. (2024). Equivariant neural operators for gradient-consistent topology optimization. Journal of Computational Design and Engineering, 11(3), 91–100. https://doi.org/10.1093/jcde/qwae039

[15] Alsammani, A., Stacey, W. C., & Gliske, S. V. (2024). Estimation of circular statistics in the presence of measurement bias. IEEE Journal of Biomedical and Health Informatics, 28(2), 1089–1100. https://doi.org/10.1109/jbhi.2023.3334684

[16] Maier-Hein, L., Reinke, A., Godau, P., Tizabi, M. D., Buettner, F., Christodoulou, E., Glocker, B., Isensee, F., Kleesiek, J., Kozubek, M., Reyes, M., Riegler, M. A., Wiesenfarth, M., Kavur, A. E., Sudre, C. H., Baumgartner, M., Eisenmann, M., Heckmann-Nötzel, D., Rädsch, T., … Jäger, P. F. (2024). Metrics Reloaded: Recommendations for image analysis validation. Nature Methods, 21(2), 195–212. https://doi.org/10.1038/s41592-023-02151-z

[17] Guo, L., Wu, H., Wang, Y., Zhou, W., & Zhou, T. (2024). IB-UQ: Information bottleneck based uncertainty quantification for neural function regression and neural operator learning. Journal of Computational Physics, 510, Article 113089. https://doi.org/10.1016/j.jcp.2024.113089

[18] Zhang, K., Li, Y., Wang, Z., Song, Y., Du, Y., Liu, L., & Zhang, Y. (2026). Accelerating phase-field prediction of grain growth in strongly anisotropic systems using a CNN-transformer surrogate model. Computational Materials Science, 274, Article 115020. https://doi.org/10.1016/j.commatsci.2026.115020

[19] Uddagiri, M., Tegeler, M., & Steinbach, I. (2024). Interface stabilization and propagation in phase field models of solidification: Resolving the issue of large driving forces. Modelling and Simulation in Materials Science and Engineering, 32(6), Article 065034. https://doi.org/10.1088/1361-651x/ad691d

[20] Reinke, A., Tizabi, M. D., Baumgartner, M., Eisenmann, M., Heckmann-Nötzel, D., Kavur, A. E., Rädsch, T., Sudre, C. H., Acion, L., Antonelli, M., Arbel, T., Bakas, S., Benis, A., Buettner, F., Cardoso, M. J., Cheplygina, V., Chen, J., Christodoulou, E., Cimini, B. A., … Maier-Hein, L. (2024). Understanding metric-related pitfalls in image analysis validation. Nature Methods, 21(2), 182–194. https://doi.org/10.1038/s41592-023-02150-0

[21] Arteche, J. (2025). Frequency domain local bootstrap in short and long memory time series. Econometric Reviews, 44(2), 163–191. https://doi.org/10.1080/07474938.2024.2401719

[22] Gregory, W. G., Hogg, D. W., Blum-Smith, B., Arias, M. T., Wong, K. W. K., & Villar, S. (2025). Equivariant geometric convolutions for dynamical systems on vector and tensor images. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 383(2298), Article 20240247. https://doi.org/10.1098/rsta.2024.0247

[23] Kontolati, K., Goswami, S., Karniadakis, G. E., & Shields, M. D. (2024). Learning nonlinear operators in latent spaces for real-time predictions of complex dynamics in physical systems. Nature Communications, 15(1), Article 5101. https://doi.org/10.1038/s41467-024-49411-w

[24] Xu, R., Xu, C., Li, Z., Zheng, T., Yu, W., & Yang, C. (2024). Boundary guidance network for medical image segmentation. Scientific Reports, 14(1), Article 17345. https://doi.org/10.1038/s41598-024-67554-0

[25] Ghiotto, M. (2026). HyperNOs: Automated and parallel library for neural operators research. Bollettino dell'Unione Matematica Italiana, 19(3), 709–743. https://doi.org/10.1007/s40574-025-00516-0

[26] Barberis, A., Aerts, H. J. W. L., & Buffa, F. M. (2024). Robustness and reproducibility for AI learning in biomedical sciences: RENOIR. Scientific Reports, 14(1), Article 1933. https://doi.org/10.1038/s41598-024-51381-4

[27] Lara Benitez, J. A., Furuya, T., Faucher, F., Kratsios, A., Tricoche, X., & de Hoop, M. V. (2024). Out-of-distributional risk bounds for neural operators with applications to the Helmholtz equation. Journal of Computational Physics, 513, Article 113168. https://doi.org/10.1016/j.jcp.2024.113168

[28] Castaño, J., Bustillo, J., Franch, X., & Martínez-Fernández, S. (2026). Data-driven multi-objective optimization of ML inference hardware configurations for energy, performance and cost. Sustainable Computing: Informatics and Systems, 50, Article 101347. https://doi.org/10.1016/j.suscom.2026.101347