Sandia National Laboratories
SAND2026-26072PE

GEOMETRY-AWARE NEURAL OPERATORS FOR FAST AND ACCURATE NON-PARAMETRIC DESIGN

Interlab PI: Michael Penwarden[1] | Yucheng Fu[2] | Pratanu Roy[3]

Team: Tian Yu Yen[1] | Zachary Morrow[1] | Rishi Roy[1] | Wenqian Chen[2] | Zhi-Feng Wei[2] | Mamud Lal[2] | Maruti Mudunuru[2] | Stephen Castonguay[3] | Kevin Chung[3] | Sumanta Roy[3] | Guang Lin[4] | Ujin Jeon[4]

[1] Sandia National Laboratories | [2] Pacific Northwest National Laboratory | [3] Lawrence Livermore National Laboratory | [4] Purdue University | 2026 LDRD Annual Meeting

ArGEnT [1] separates geometric context from solution queries, enabling accurate predictions on irregular domains, flexible query discretizations, and previously unseen geometries.

1 ArGEnT decouples geometry and queries

Cross-attention ArGEnT represents geometry with an independent point cloud used for keys and values, while arbitrary solution locations construct the queries in the transformer-based framework.

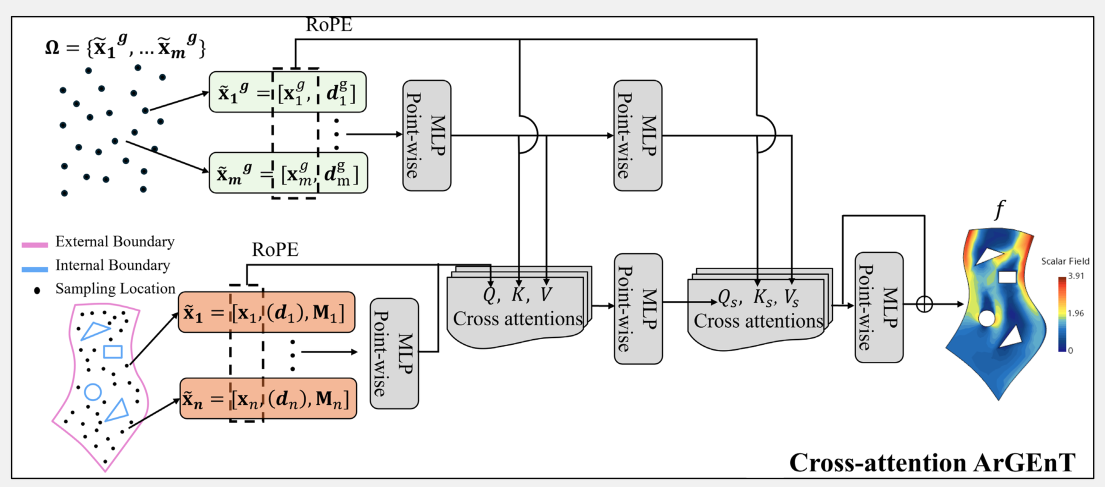

Geometry sampling is decoupled from prediction locations. The same compact geometry representation can condition different query sets without resampling.

ARGENT AS A DEEPONET TRUNK

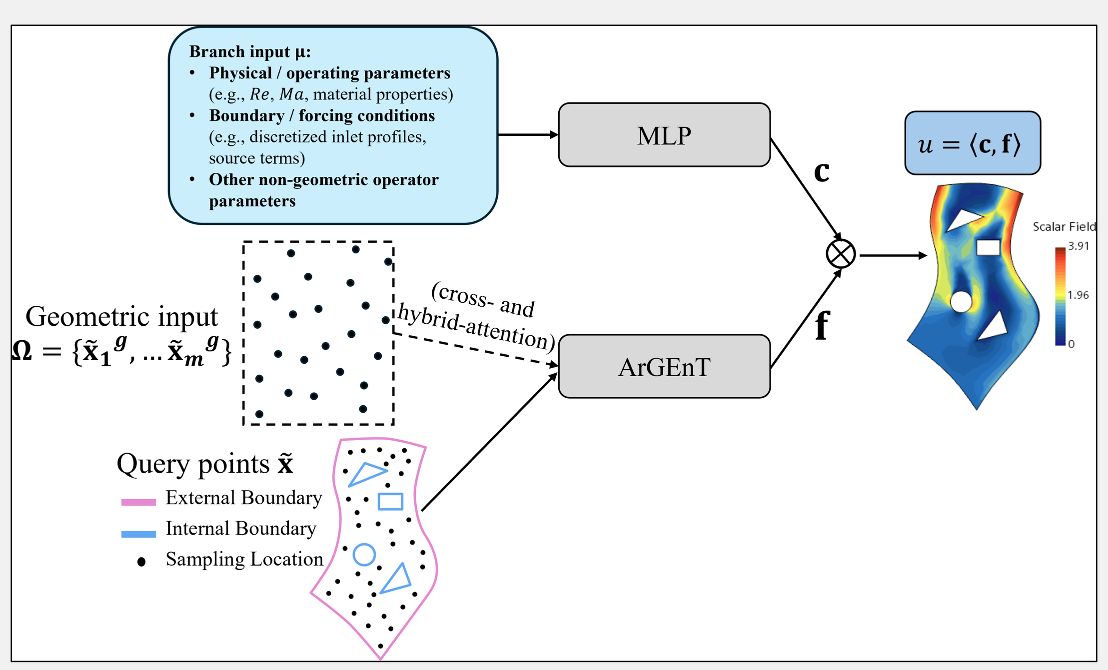

The DeepONet branch can encode information beyond geometry, including boundary and loading conditions, material properties, forcing functions, and operating parameters for flexible geometry-aware neural operator learning.

CROSS-ATTENTION DECOUPLES TWO POINT CLOUDS

Geometry representation: Point clouds with signed-distance function (SDF) labels describe complex domains without requiring a mesh as model input.

Flexible evaluation: Solution fields are predicted at arbitrary locations and resolutions.

Transformer cross-attention: Each solution query selectively draws information from the independently encoded geometry.

Reusable geometry context: the same encoded geometry can serve different query discretizations without resampling.

2 3D surrogates on non-parametric brackets

The DeepJEB-derived dataset contains non-parametric 3D geometries with large topological variation and 120,000-380,000 finite-element nodes per example.

TRAIN AND TEST GEOMETRIES

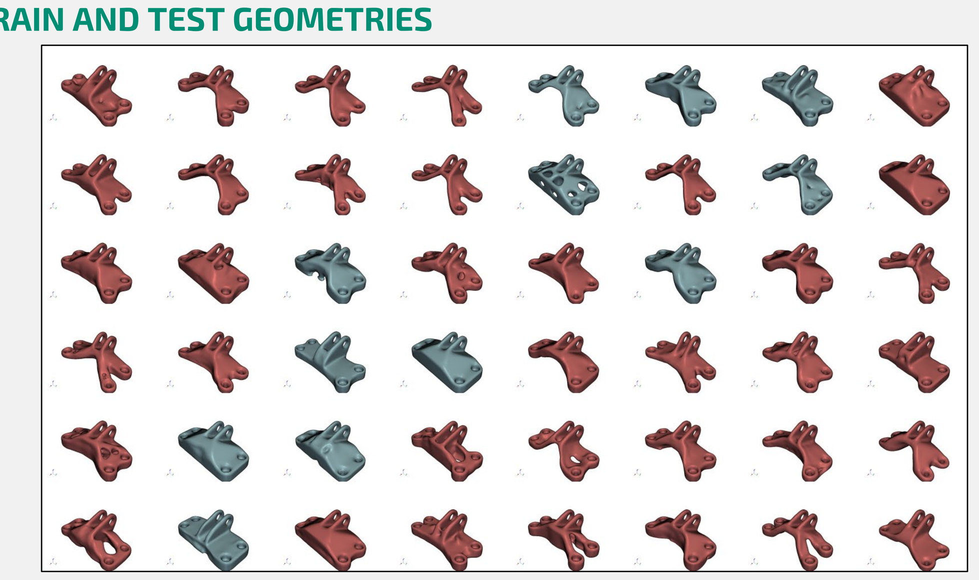

Training geometries are red; held-out test geometries are blue. Both sets contain pronounced changes in topology, thickness, openings, and load-path structure.

VERTICAL-LOAD MEAN ABSOLUTE ERROR - LOWER IS BETTER

Model	ux (mm)	uy (mm)	uz (mm)	σvm (MPa)
PointNet	0.008	0.003	0.013	11.666
DeepONet	0.026	0.007	0.048	21.842
Point-DeepONet	0.007	0.003	0.012	10.541
Transolver	0.0085	0.0022	0.0154	7.4013
Transolver+	0.0084	0.0024	0.0118	7.3257
Cross-attn. ArGEnT	0.0039	0.0018	0.0076	7.9454
Cross-attn. ArGEnT (scaled)	**0.0023**	**0.0013**	**0.0042**	**6.1865**

PREDICTION AND TRUTH

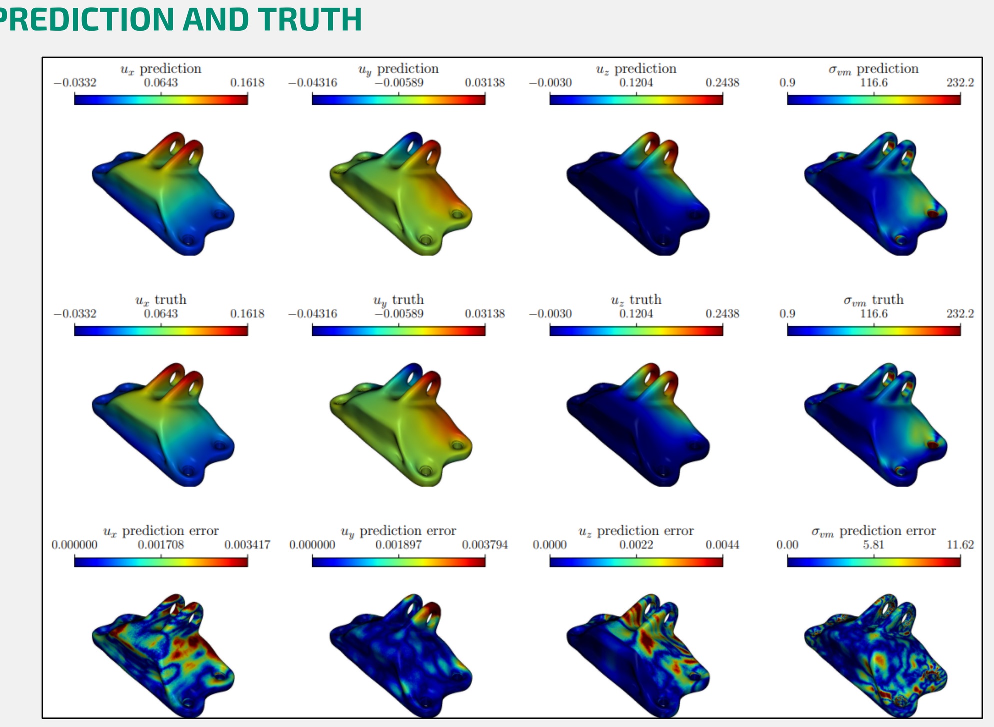

The scaled cross-attention model gives the lowest vertical-load error in every reported field while resolving displacement and von Mises stress on the original 3D geometries.

3 GOLD-FLOW extends ArGEnT to design

GOLD-FLOW - Geometry-aware Operator Learning for Design in Flow Energy Components - continues the LDRD through DOE Genesis Mission Phase I.

OUT-OF-DISTRIBUTION FLOW PREDICTIONS SPELL GENESIS

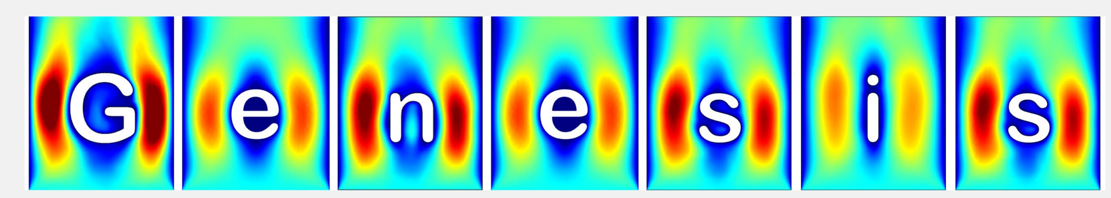

Trained on 1,000 flow-over-cylinder cases, ArGEnT predicts flow over unseen, non-parametric letter geometries, demonstrating transfer beyond the training geometry family.

CLOSED-LOOP NON-PARAMETRIC DESIGN

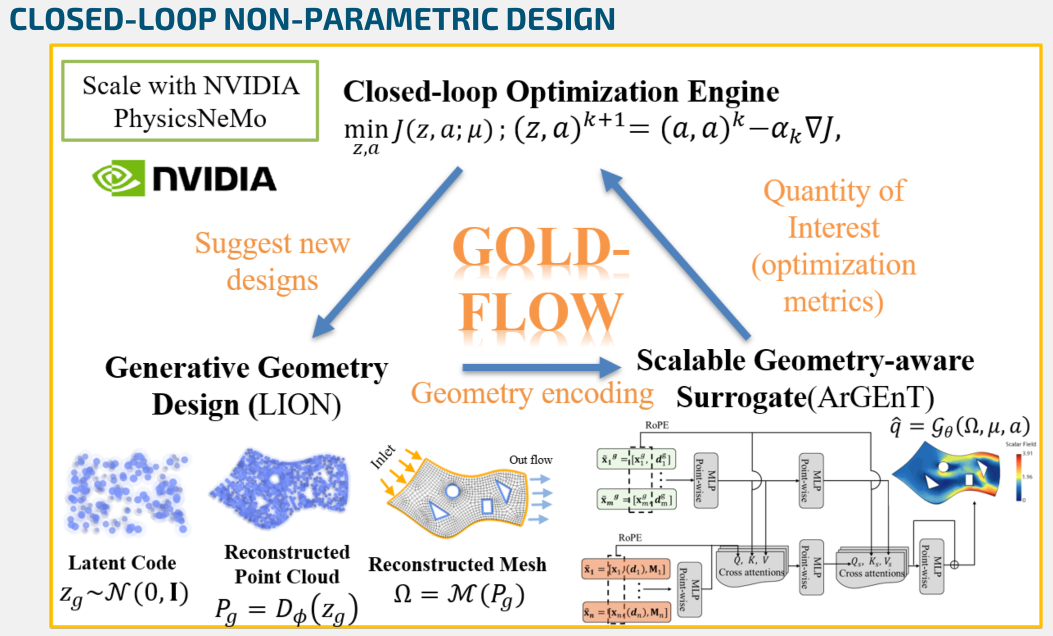

FLOW COMPONENT TARGETS

REDOX FLOW BATTERY

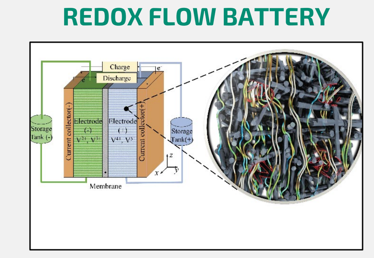

TURBINE NOZZLE

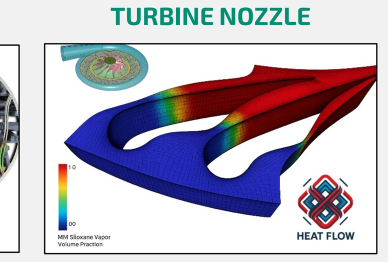

CROSS-FLOW TURBINE

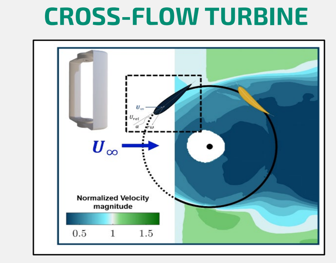

The same generative-surrogate loop supports electrochemical flow fields, turbine nozzles, and cross-flow rotors without reducing the designs to a common parametric geometry family.

TARGET: >100x FASTER DESIGN CYCLES

Phase I target relative to conventional simulation-driven design workflows

FOLLOW-ON WORKFLOW AND SCALE

- **Surrogate-driven design:** Latent generative models propose non-parametric geometries; ArGEnT predicts the fields needed to evaluate physics-based quantities of interest.
- **Deployment and scale:** Separable generator, surrogate, and optimizer modules support NVIDIA PhysicsNeMo deployment and distributed screening of large design spaces.
- **Geometry-general workflow:** The operator evaluates point clouds directly, allowing the optimization loop to compare designs with different topology and discretization.

4 Related work strengthens the workflow

Companion papers, preprints, and drafts extend the project across geometry encoding, uncertainty, efficient training, and hard enforcement of interface physics.

PNNL - GEONET FOR HETEROGENEOUS POROUS MEDIA

Topology-aware porous-media surrogates: A convolutional Leaf Network encodes solid-void topology separately from material fields and query coordinates, preserving geometry-induced pressure transitions that coordinate-only models smooth. About 5,000 PFLOTRAN simulations (about 30 TB) yield NRMSE ≈ 0.03, MAPE ≈ 1.5%, and $R^2 \approx 0.99$ on unseen geometries.

SANDIA - RELIABLE AND SURFACE-AWARE OPERATORS

Epistemic UQ diagnostics for operators: Ensembles separate uncertainty associated with optimization, input discretization, and sampled training functions, exposing failure modes that a single aggregate uncertainty estimate can hide, studied across a comprehensive set of neural operators.

Curvature-aware surface operators: Normals, curvature invariants, and shape-operator features replace off-surface SDF ghost points for fields predicted directly on surfaces, improving accuracy and train/test costs for ArGEnT on shell mechanics problems.

Physics-informed shallow networks: Shallow learnable polynomial bases reduce repeated differentiation depth while targeting optimization stiffness in high-frequency and high-order PDEs. The construction retains analytic derivatives for physics-informed residuals.

PURDUE – SANDIA UNIVERSITY PARTNER PLUS-UP WORK

Precision-coupled replica exchange: FP32 cold-chain refinement is paired with mixed-precision FP16 hot-chain exploration and FP32 exchange-energy evaluation, reducing peak GPU memory for training ArGEnT on large scale problems .

Bayesian multifidelity active learning: Multifidelity Laplace Neural Operator uses reSGLD uncertainty to select informative high-fidelity trajectories. It improves data efficiency on Lorenz, Duffing, and composite-beam dynamics, including frequency mismatch between fidelity levels.

LLNL - INTERFACE AND DISCONTINUITY OPERATORS

Hard-constrained interface PINNs: Windowing and buffer ansatzes embed interface continuity and flux balance into the solution representation instead of balancing penalty terms for interface problems.

φ-DeepONet for discontinuities: Multiple branches represent discontinuous inputs, while a domain-decomposition embedding in the trunk captures interface-driven output discontinuities under physics- and interface-informed training.

[1] W. Chen, Z.-F. Wei, Y. Fu, M. Penwarden, P. Roy, and P. Stinis, "ArGEnT: Arbitrary Geometry-encoded Transformer for Operator Learning," arXiv:2602.11626, 2026. (Under review in the Journal of Neural Networks)

Michael Penwarden - Abstract

Evaluating complex engineering designs requires costly simulations that limit how many alternatives can be explored. This interlaboratory LDRD between Sandia, PNNL, and LLNL develops neural operators for rapid physics predictions across diverse geometries. We present ArGEnT, a geometry-aware surrogate model, and demonstrate its capabilities on three-dimensional structural components and flow around unseen geometries. Related work investigates epistemic uncertainties, mixed precision training, and physics-informed learning. Direct follow-on work includes GOLD-FLOW, a DOE Genesis Mission Phase I project that combines generative AI with fast surrogate predictions for non-parametric design optimization of energy-system flow components.