

# Research Proposal

## CC-FNO-E: Conservation-Constrained Fourier Neural Operators for Regime-Generalizable Fluid Dynamics

**Team:** AI Research Collaborative (Group)

**Date:** May 24, 2026

**Project:**

CC-FNO-E: Conservation-Constrained Fourier Neural Operators for Regime-Generalizable Fluid Dynamics

**Target Compute:** 4x AMD Instinct MI300X GPUs (single node), approx. 1 week

### 1. PROBLEM STATEMENT

Neural operators have emerged as highly efficient surrogates for partial differential equation (PDE) solvers. However, their accuracy degrades sharply when they are evaluated outside the training distribution of physical parameters. In fluid dynamics, achieving out-of-distribution (OOD) generalization across different Reynolds number regimes in 2D incompressible Navier-Stokes equations remains a core challenge.

Existing standard Fourier Neural Operators (FNOS) fail to respect fundamental physics, leading to severe error accumulation and instability when transitioning between laminar and turbulent flows. This project tests whether enforcing hard conservation constraints and physical consistency can yield near-regime-invariant accuracy and stable autoregressive rollouts without requiring target-regime training labels.

### 2. PROPOSED APPROACH

We are building CC-FNO-E and its variant CC-FNO-E-PF, which introduce hard conservation constraints into the Fourier Neural Operator architecture:

- Leray Projection Layer:** Enforces divergence-free velocity fields at every spectral layer of the network.
- Dissipation-Aware Scaling Constraint:** Bounds kinetic energy at the output to act as an amplitude regularizer.
- Pushforward Training (CC-FNO-E-PF):** Combines conservation constraints with multi-step pushforward training to drastically close the accuracy-stability Pareto gap during long rollouts.

This stack is selected over a standard FNO because the energy constraint provides excellent long-term stability, while the combination of both constraints yields the best overall per-step accuracy.

### 3. DATASET

The datasets consist of numerical simulations of 2D incompressible Navier-Stokes equations across varying fluid regimes:

Regime Category	Corpus / Configuration	Size / Properties	Role
In-Distribution (Train)	2D Navier-Stokes Simulation Data (Re $\in$ {500, 1000})	~50 GB	Used strictly for training the base operators.
Out-of-Distribution (Eval)	Laminar OOD Evaluation (Re = 100)	Tabular / Mesh Tensor	Downstream evaluation of zero-shot transfer accuracy.

Out-of-Distribution (Eval)	Turbulent OOD Evaluation ( $Re = 2000$ )	Tabular / Mesh Tensor	Downstream evaluation of zero-shot stability.
----------------------------	--	-----------------------	---

## Preprocessing

- Grids are downsampled to a uniform spectral mesh size (e.g.,  $64 \times 64$  or  $128 \times 128$ ).
- Data is stored as memory-mapped tensors for optimized data-parallel loading.
- Total on-disk footprint is estimated around 120 GB, well within standard storage limits.

## 4. EVALUATION METRICS

Experiment Type	Primary Metric	Target / Interpretation
OOD Regime Accuracy	Single-step relative $L^2$ error across all Reynolds numbers	Measures generalization across unseen fluid regimes
Long-term Rollout Stability	Percentage drift over 100-step rollouts	Enforces non-divergent physical bounds over time
Autoregressive Horizon	Maximum rollout horizon achieved while maintaining stability	Targeting 90-95 steps before error breaks physics bounds
Statistical Significance	Wilcoxon signed-rank test ( $p = 0.031$ )	Evaluated over 5 random seeds to guarantee consistency

## 5. COMPUTE REQUIREMENTS

Estimates are based on initial single-GPU benchmarks running standard FNO blocks with custom projection layers, scaled up to a multi-GPU AMD node configuration.

Phase	Hardware	Estimated Wall-Clock	Notes / Details
Phase 1: CC-FNO-E Base Training	4x AMD Instinct MI300X	~2 days	Training across $Re \in \{500, 1000\}$ with mixed-precision
Phase 2: Pushforward Fine-tuning (CC-FNO-E-PF)	4x AMD Instinct MI300X	~2 days	Multi-step lookahead backpropagation to close the stability gap
Phase 3: OOD Sweeps & Ablations	2x AMD Instinct MI300X	~2 days	Sweeping 10 model variants over 5 random seeds across 4 Reynolds numbers

### Dependencies & Memory Allocation

- PyTorch + ROCm Toolkit compatibility framework.
- torchharmonic or custom FFT-based spectral layers.
- Custom Leray projection and dissipation scaling modules implemented in pure PyTorch for high portability.

## 6. EXPECTED DELIVERABLES

1. **Model Checkpoints:** Converged weights for both CC-FNO-E and CC-FNO-E-PF variants.

2. **Benchmark Headline Table:** Relative  $L^2$  error comparison showing a reduction from the standard FNO error (0.039–0.105) down to the target conservation-constrained error (0.013–0.019).
3. **Ablation Analysis:** Clear data separating the distinct impacts of the energy constraint (amplitude regularization) versus the Leray projection.
4. **Conference Draft:** A fully compiled manuscript formatted for the International Conference on Machine Learning (ICML).

## 7. RISK & MITIGATION

Risk	Likelihood	Mitigation
Pushforward training instability	Medium	Fall back to incremental horizon scheduling (start with 2-step lookahead, scale up to 5+ steps dynamically).
High computational overhead from spectral Leray projections	Low-Medium	Optimize via inline execution; if needed, compute projections every N blocks rather than at every single layer.
Energy constraint acts strictly as a dampener rather than preserving true forcing dynamics	Medium	Calibrate the dissipation-aware scaling factor using a small validation sweep to ensure fluid structures remain sharp.