ZENITHON AI

FASTER SIMULATIONS. BETTER DECISIONS.

Motivation

- Design iteration in fusion is slow
  - Many complicated physical systems with incomplete understanding of underlying physics
  - Most systems/components have interdependencies
  - Cost of high fidelity multiphysics simulations is prohibitively expensive for routine use in reactor design

- Meanwhile in last ~5 years, ML for physics has exploded
  - Significant advances in e.g. ML for CFD and weather modeling
  - 'Moonshot' promise of surrogate modeling is:
    - Training on experimental data yields models closer to ground truth
    - High fidelity, full physics models
    - Fully connected systems
    - Faster learning from experiments, reduced dependence on ROMS

Our team (hiring!)

ZENITHON AI

Leadership

Alex Higginbottom
CEO

Abetharan Antony
CTPO

Dev

David Voderholzer
Research Engineer

Hong Chul Nam
Research Engineer

Hrishikesh Viswanath
Research Engineer

Zara Ercan
Simulation Engineer

Vispi Karkaria
Research Engineer

Our advisors

Zongyi Li
Advisor
Inventor of Neural Operators
Incoming Asst Prof NYU

Currently @20hrs/month. Gradually increasing involvement

Dan Brunner
Advisor
Cofounder, CFS (largest fusion company)

Providing key PMF insights and connections

Lu Lu
Advisor
Inventor of DeepOnet, Significant contributor to PINNs
Asst Prof @ Yale

Providing research supervision and support

What we're currently working on

ZENITHON AI

Neutron transport

Monte Carlo / Neural Operator model

CFD/MHD

Transformer neural operator with Quadtree geometry construction at each timestep

Gyrokinetics

Function autoencoder with latent diffusion model

Neutronics

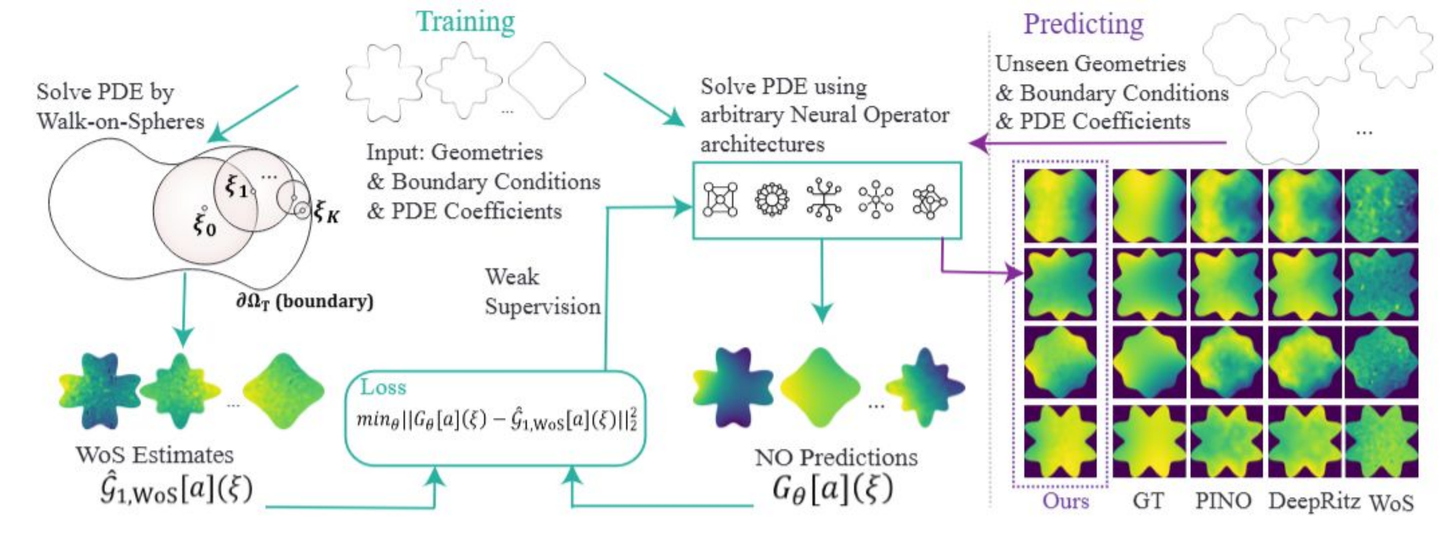

Training
Predicting
Solve PDE by
Walk-on-Spheres
Input: Geometries
& Boundary Conditions
& PDE Coefficients
Solve PDE using
arbitrary Neural Operator
architectures
Unseen Geometries
& Boundary Conditions
& PDE Coefficients
∂Ω_T (boundary)
Weak
Supervision
Loss
$min_\theta ||G_\theta[a](\xi) - \hat{G}_{1,\text{WoS}}[a](\xi)||_2^2$
WoS Estimates
$\hat{G}_{1,\text{WoS}}[a](\xi)$
NO Predictions
$G_\theta[a](\xi)$
Ours
GT
PINO
DeepRitz
WoS

CFD/MHD

First time-adaptive transformer neural operator

- Discretisation convergent model
- Incremental quadtree construction for AMR/transient features
- Transformer for superior generalisation across physics/geometries + scalability

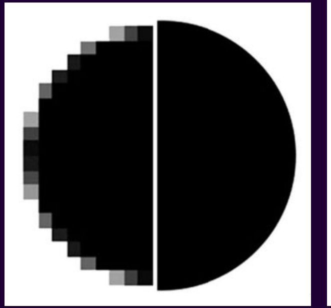

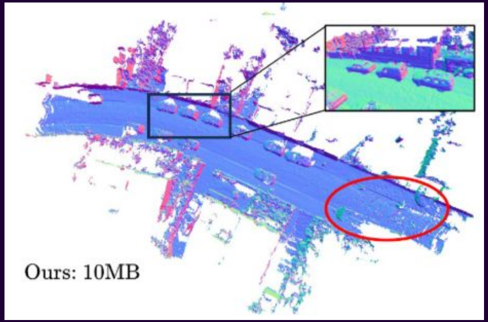

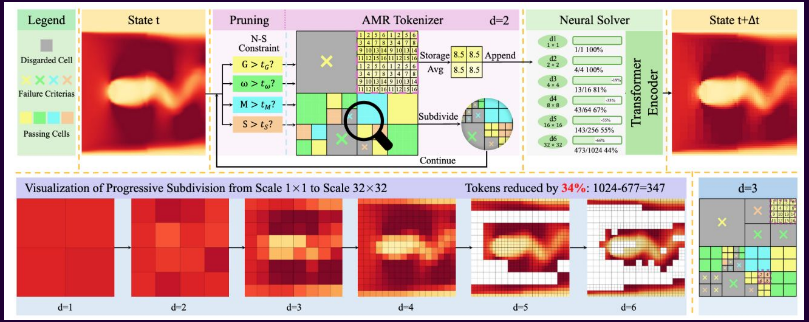

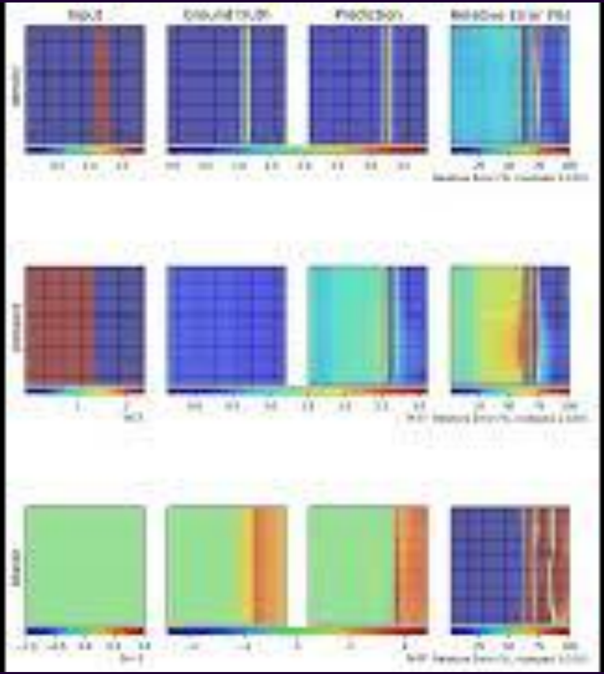

Figure: Z-Pinch simulation: Zenithon surrogate model (runtime < 8ms) vs. the full MHD simulation (runtime ~5 hours).

Physics Informed Function-Space Diffusion for Rapid Gyrokinetic Turbulence Modelling

A. Higginbottom[1], A. Antony[1], Z. Li[2], L. Lu[3]

[1]Zenithon AI

[2]MIT

[3]Yale

IAEA Workshop on Digital Engineering
December 10, 2025

https://www.zenithon.ai

Motivation

The Challenge

- **The Bottleneck:** High-fidelity Gyrokinetic (GK) solvers are too expensive for iterative profile optimization or real-time control ($\sim 10^5$ core-hours).
- **The Physics Need:** Accurate confinement prediction requires resolving electromagnetic turbulence across disjoint spatiotemporal scales.
- **Critical Outputs:** Beyond scalar fluxes, modern analysis requires full field data for potentials $(\phi, A_{||})$ to diagnose zonal flows and saturation mechanisms.

The Solution

A probabilistic diffusion framework bridging the speed of TGLF with the fidelity of nonlinear gyrokinetics.

Gyrokinetics: The High-Fidelity Standard

Schematic Gyrokinetic Equation

$$\frac{\partial h_a}{\partial t} - i(\omega_\theta + \omega_\xi + \omega_d)H_a + \frac{c}{\psi'}[f_{0a} + h_a, \Psi_a] = \sum_b C_{ab}^{\mathrm{GK}}$$

- **State Variables:** h_a (non-adiabatic dist. function); H_a (includes adiabatic response).
- **Linear Drifts:** Parallel streaming (ω_θ), particle trapping (ω_ξ), toroidal magnetic drifts (ω_d).
- **Nonlinearity:** Poisson bracket $[\cdot,\cdot]$ captures EM effects via generalized potential Ψ_a.
- **Collisions:** C_{ab}^{GK} is the linearized, gyro-averaged Fokker-Planck operator.

TGLF: The Reduced-Order Model (ROM)

Trapped Gyro-Landau Fluid (TGLF)

- Solves a linear gyrofluid eigenproblem:

$$-i\omega_k \vec{X}_k = \mathcal{L}(k_y; \mathbf{p})\, \vec{X}_k$$

- $\mathcal{L}$ depends on local parameters $\mathbf{p}$ (gradients $a/L_{T,n}$, geometry $q, \hat{s}$, β, collisions).
- **Quasilinear Fluxes:** Calculated by combining linear eigenmodes with analytic saturation rules (e.g., SAT2).

Limitation

TGLF is fast but relies on scalar approximations and cannot generate high-fidelity, nonlinear turbulence structures.

Function-Space Diffusion (FunDiff)

Concept

- An encoder-decoder pair $E, \mathcal{D}$ maps physical functions f to latent codes $z \in \mathbb{R}^D$.
- A latent diffusion model learns the density on $\mathbb{R}^D$.

Wasserstein Error Bound

$$W_1(P, \hat{P}) \leq \underbrace{\mathbb{E}\|f - \mathcal{D}(E(f))\|}_{\text{Reconstruction Error}} + \underbrace{\text{Lip}(\mathcal{D}) W_1(p, \hat{p})}_{\text{Generative Error}}$$

TurboTGLF Strategy:

- **Conditioning:** TGLF maps equilibrium $\mathbf{x}$ to state $r = R(\mathbf{x})$.
- **Generation:** Nonlinear GK defines a conditional distribution $P_r \equiv P(f, Q \mid r)$.
- **Goal:** Train a conditional latent model $\hat{p}_r(z) = p_\theta(z \mid r)$ to generate turbulence consistent with the TGLF state.

Model Architecture

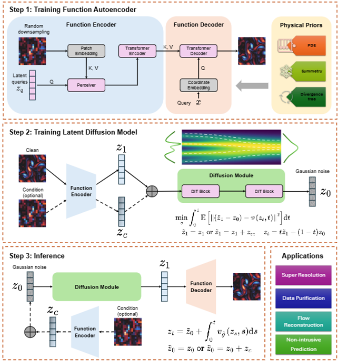

1. **Function Autoencoder (FAE):** Maps Fourier fields $U(k_x, k_y)$ to latent z. Enforces Hermitian symmetry and spectral constraints.
2. **Conditional Latent Diffusion:** A transformer models $p_\theta(z \mid r)$, where r includes TGLF growth rates, frequencies, and flux predictions.

Inference, Distillation, and UQ

Inference Pipeline

1. **Reverse-time ODE:** Integrate from Gaussian noise using the learned network + TGLF state r.
2. **Decoding:** Map latent z back to continuous Fourier fields $U(k_x, k_y)$ and flux Q.

Capabilities

- **Distillation (Future):** Student network mimics trajectories in few Euler steps for millisecond inference.
- **Uncertainty Quantification:** Drawing an ensemble $\{(U^{(m)}, Q^{(m)})\}$ for fixed r provides credible intervals for transport.

Benchmark: Canonical Kolmogorov Flow

Problem Setup:

- 2D Kolmogorov flow (Navier-Stokes) on $(0, 2\pi)^2$ with periodic forcing.
- **Constraint:** $\nabla \cdot u = 0$ (Incompressible).
- **Challenge:** Reconstruct full turbulence from sparse, noisy downsampled observations (simulating ROM conditioning).

Physics-Informed Decoder:

- Predicts stream function ψ rather than raw velocities.
- Velocity recovered as $(u, v) = (\partial_y \psi, -\partial_x \psi)$, guaranteeing divergence-free output by construction.

Results: Accuracy and Robustness

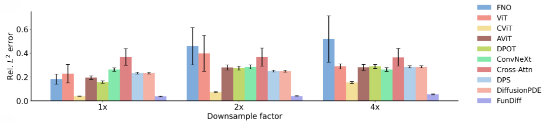

Relative ℓ_2 error vs. downsampling. FunDiff outperforms neural operators.

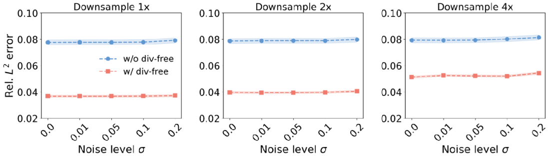

Physics-constrained model (Blue) beats data-only (Red), especially at high noise.

Results: Visual Reconstruction

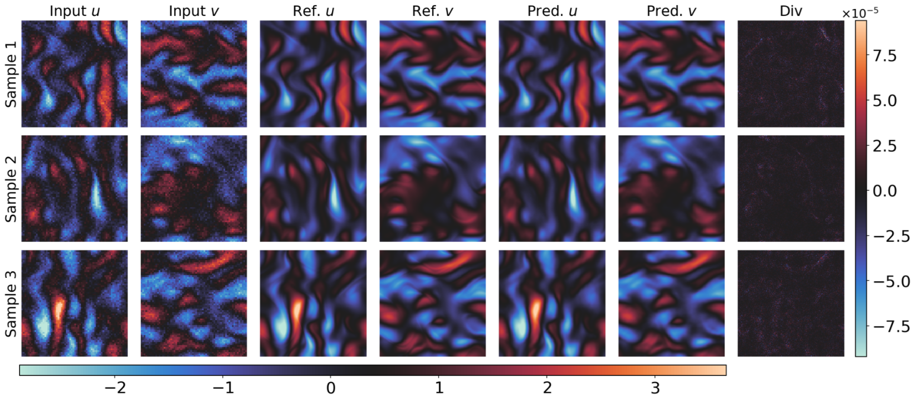

Observation: Generated fields remain numerically divergence-free while capturing complex turbulent structure, even under unconditional generation.

Contributions & Future Work

Key Contributions

- Rapid generative modelling framework for hybrid workflows (Qualikiz/TGLF).
- Continuous decoder for Fourier-space fields, robust to resolution changes.
- Physically constrained outputs (e.g., Divergence-free, Hermitian).
- Built-in ensemble uncertainty quantification.

Next Steps

- Train on larger, heterogeneous datasets.
- Ablation studies over different physics constraints.
- Test flow matching models for mean flow processing.
- Integrate with full-core transport solvers.
- Finalize distillation for real-time control.

References

1. S. Wang et al., “FunDiff: Diffusion models over function spaces for physics-informed generative modeling”, 2024.
2. Z. Zhou et al., “Simple and Fast Distillation of Diffusion Models”, NeurIPS 2024.
3. G. Staebler et al., “Quasilinear theory and modelling of gyrokinetic turbulent transport in tokamaks”, Nucl. Fusion 64, 103001 (2024).
4. J. Krommes, “The gyrokinetic description of microturbulence in magnetized plasmas”, PPPL report.

Thank You

alex@zenithon.ai