Continuous Meteorological Field Reconstruction using Neural Transformer Operators

Sumanth Kumar Boya, Azhar Gafoor, Deepak N. Subramani

Indian Institute of Science Bengaluru · {ksumanth, azharctp, deepakns}@iisc.ac.in

INTRODUCTION

In meteorological and oceanographic sciences, critical physical variables are available only as **sparse, irregular, or incomplete observations** due to sensor limitations, extreme environments, and large spatio-temporal scales.

Existing neural operators and implicit neural representation (INR) networks **struggle to maintain reconstruction accuracy under highly variable sparsity** and require costly retraining for each new configuration.

We introduce the **Implicit Neural Transformer Operator (INTO)**, a context-aware neural operator that reconstructs PDE fields via data assimilation from limited observations, achieving robust generalization across sparsity levels **without task- or sparsity-specific retraining**.

•Sparsity-Agnostic •No Retraining •Cross-Attention Operators •Data Assimilation

THE CHALLENGE

Reconstruction of continuous physical fields from sparse observations is an ill-posed inverse problem. Classical methods (3D/4D-Var, ensemble Kalman filters) depend heavily on accurate PDE models and grow computationally expensive with system dimensions.

Deep learning methods such as INRs offer expressive continuous representations but are designed for fixed sensor layouts, they generalize poorly when sparsity levels or sensor configurations change, demanding separate training for every new regime.

The key insight: treat sparse observations as key–value tokens and dense query coordinates as queries in a cross-attention framework, enabling the model to naturally adapt to any sensor configuration at inference time.

•INRs: one model per config •INTO: one model for all

MODEL ARCHITECTURE

INTO consists of iterative kernel integral operator units that embed domain coordinates into sparse observation-aware representation vectors through cross-attention mechanisms.

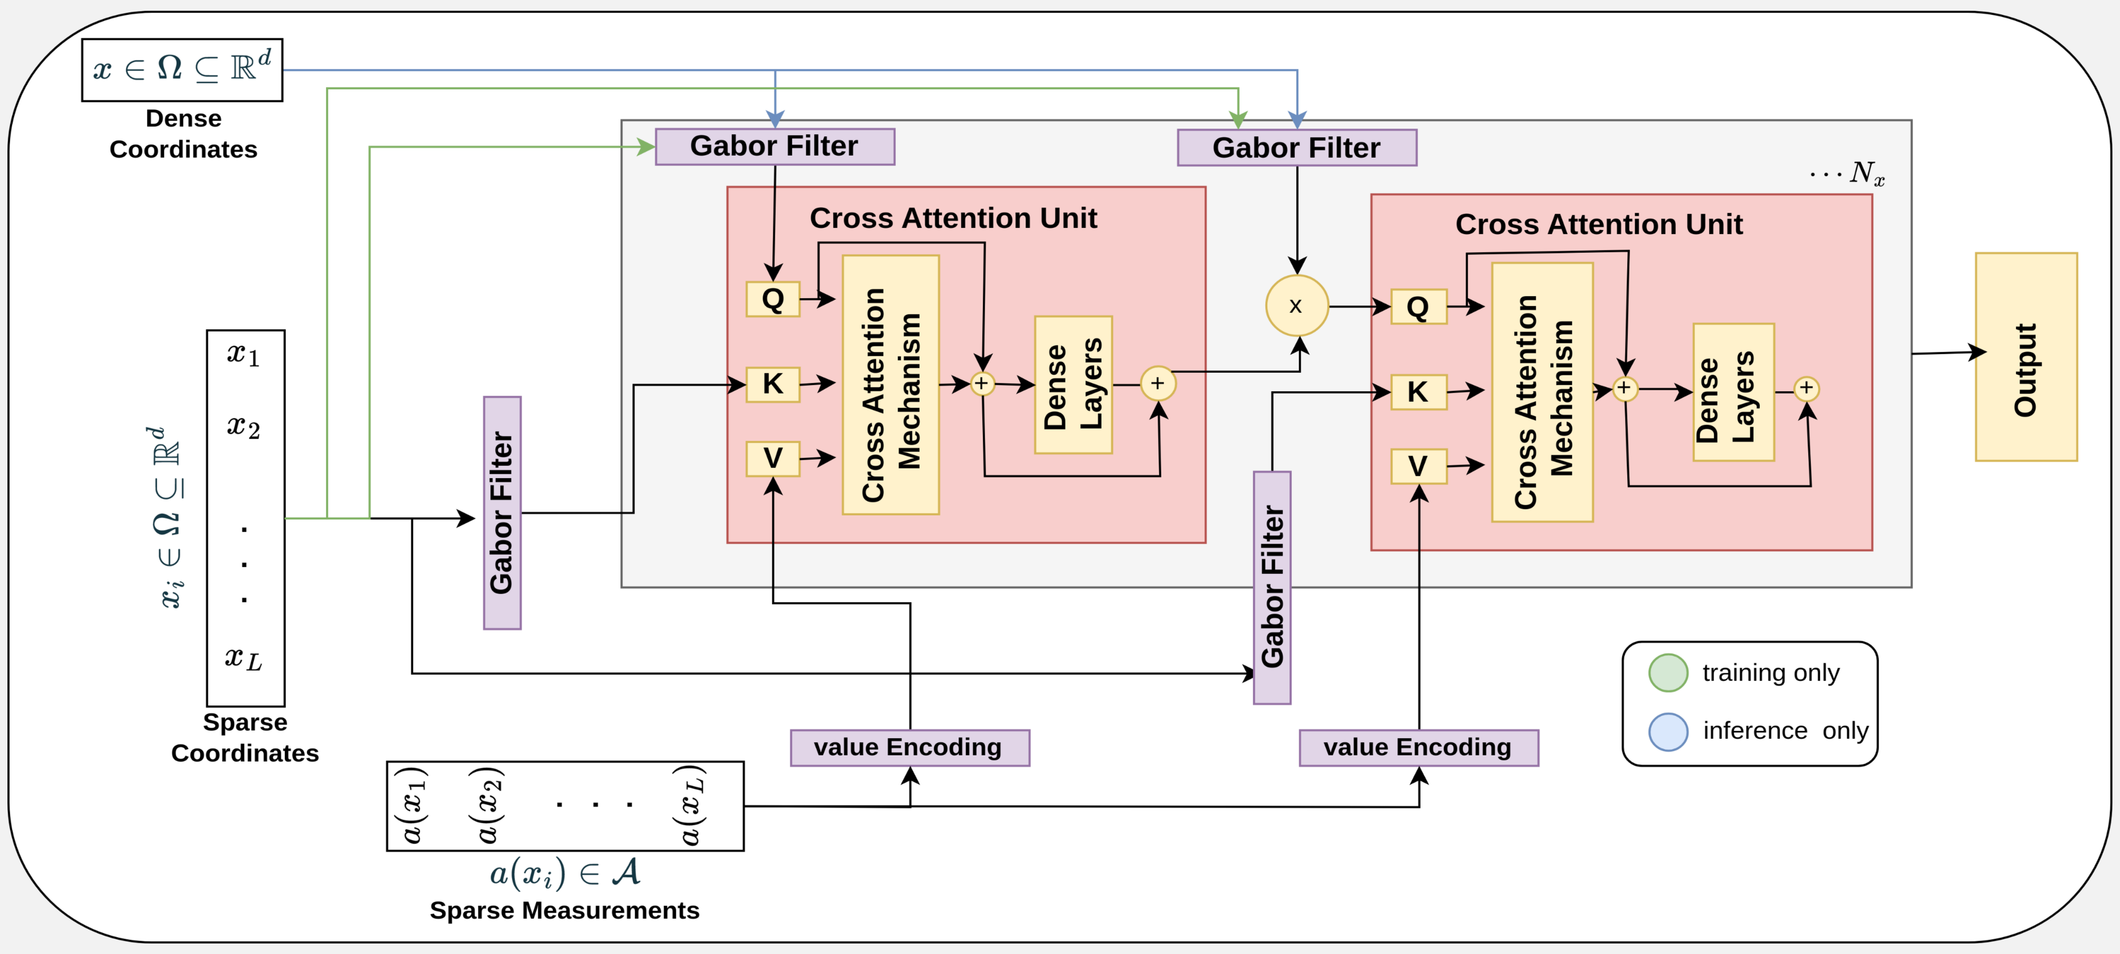

GEOPHYSICAL BENCHMARKS

Four large-scale real-world datasets spanning climate, ocean, and biogeochemical variables.

GST — Global Surface Temperature
CESM2 climate model · 1024 snapshots · 192×288 grid · sparsity 5%, 25%, 50%

SST — Sea Surface Temperature
GHRSST MUR Level-4 · 360 daily frames · 901×1001 · sparsity 0.1–0.5%

SST — Sea Surface Temperature
Copernicus DUACS altimetry · 365 daily frames · 180×360 · sparsity 10–50%

CHL — Chlorophyll Concentration
NEMO-PISCES biogeochemistry · 579 frames · 170×360 · sparsity 10–50%

QUANTITATIVE PERFORMANCE

40–55% MSE reduction on GST & SST vs. best INR baseline

25–45% MSE reduction on SSH & CHL vs. best INR baseline

36/48 Experiments where INTO ranks best overall

4× Models trained (vs. 48 for INR baselines)

Model	5%	25%	50%
SIREN	13.63	12.44	12.21
FFN-P	7.60	6.51	6.56
MMGN	3.79	2.91	2.82
MFN	15.85	15.18	15.11
INTO (per-config)	2.955	1.847	1.727
INTO* (single model)	**2.184**	**1.788**	**1.750**

INTO* = single model trained once on T4 at 5% sparsity, deployed across all 12 GST configurations - consistently best or second-best.

QUALITATIVE EVALUATION

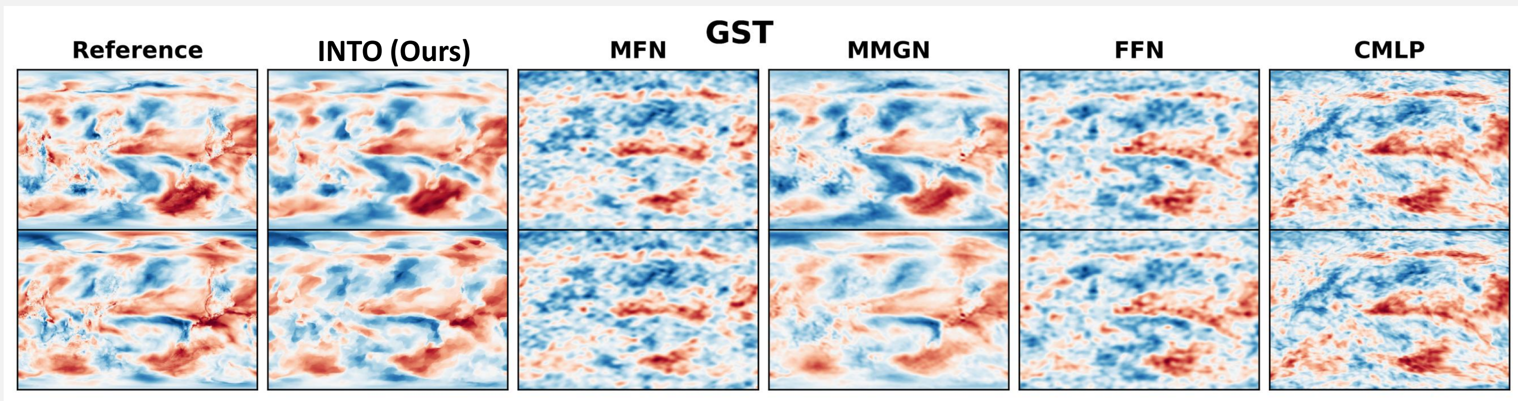

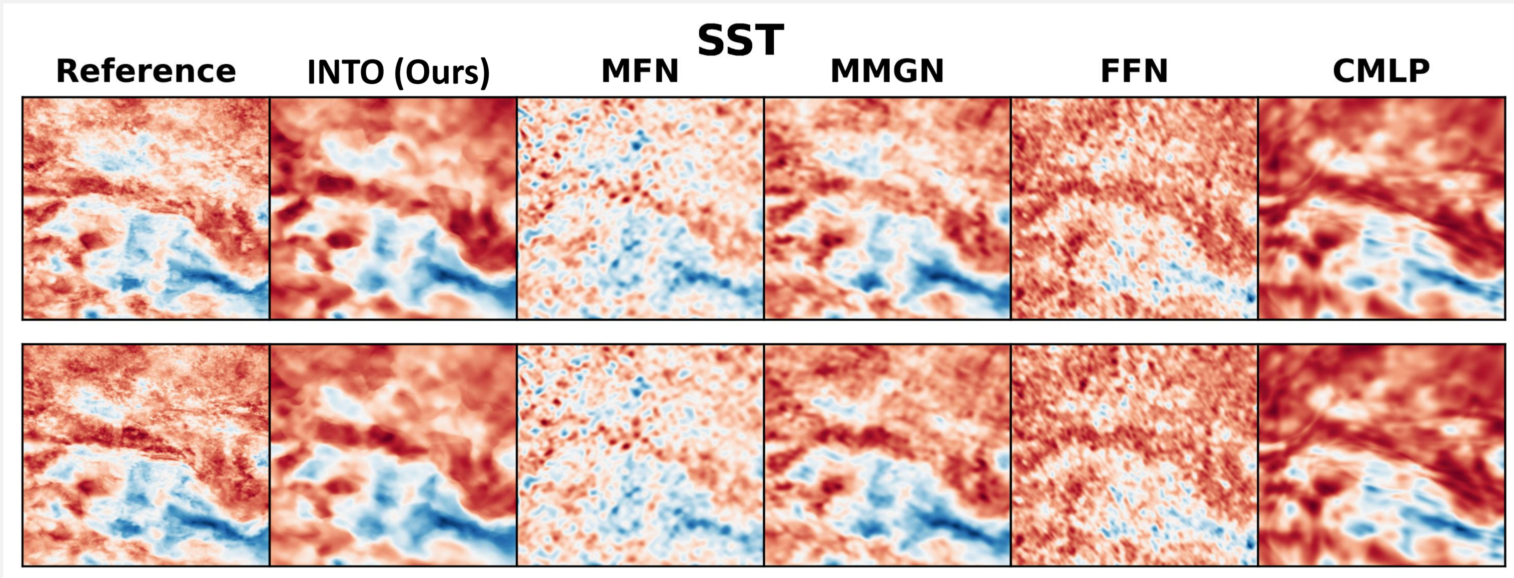

INTO VS. PRIOR APPROACHES

Aspect	Traditional INRs	Neural Operators	INTO (Ours)
Generalization	Retrain per config	Operator level	Single model all configs
Training data	Sparse samples	Dense data required	Sparse samples only
Inference input	Dense coords	Sparse observations	Both (sparse K/V + dense Q)
Core mechanism	Coord encoding + MLP	Architecture-dependent	Cross-attn kernel integral
Sparsity-agnostic	No	Partial	Yes

KEY CONTRIBUTIONS

- We introduce INTO, the first architecture integrating neural operators within an INR framework for continuous PDE field reconstruction from sparse data.
- Sparsity-agnostic generalization: a single trained model handles all observation tasks and sparsity levels, reducing the training cost by 12x.
- Superior reconstruction quality: 25-55% MSE reduction over state-of-the-art INRs across GST, SST, SSH, and CHL datasets, preserving eddies, gradients, and fine-scale spatial structures.
- Real-world applicability: scalable, unified framework for PDE-constrained data assimilation, suitable for real-time field reconstruction in large-scale physical science applications.