

Lawrence Berkeley National Laboratory

LBL Publications

Title

Spatio-temporal Fourier Transformer (StFT) for Long-term Dynamics Prediction

Permalink

<https://escholarship.org/uc/item/2f58c012>

Authors

Long, Da

Zhe, Shandian

Williams, Samuel

et al.

Publication Date

2025-03-14

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Spatio-temporal Fourier Transformer (StFT) for Long-term Dynamics Prediction

Da Long^{*1} Shandian Zhe¹ Samuel Williams² Leonid Oliker² Zhe Bai^{*2}

Abstract

Simulating the long-term dynamics of multi-scale and multi-physics systems poses a significant challenge in understanding complex phenomena across science and engineering. The complexity arises from the intricate interactions between scales and the interplay of diverse physical processes. Neural operators have emerged as promising models for predicting such dynamics due to their flexibility and computational efficiency. However, they often fail to effectively capture multi-scale interactions or quantify the uncertainties inherent in the predictions. These limitations lead to rapid error accumulation, particularly in long-term forecasting of systems characterized by complex and coupled dynamics. To address these challenges, we propose a spatio-temporal Fourier transformer (StFT), in which each transformer block is designed to learn dynamics at a specific scale. By leveraging a structured hierarchy of StFT blocks, the model explicitly captures dynamics across both macro- and micro-spatial scales. Furthermore, a generative residual correction mechanism is integrated to estimate and mitigate predictive uncertainties, enhancing both the accuracy and reliability of long-term forecasts. Evaluations conducted on three benchmark datasets (plasma, fluid, and atmospheric dynamics) demonstrate the advantages of our approach over state-of-the-art ML methods.

in fields such as climate modeling, plasma science, fluid dynamics, and beyond. Traditional approaches rely heavily on numerical solvers, which discretize the domain and iteratively solve PDEs using methods like finite difference, finite element, and spectral methods (Tadmor, 2012). While effective in many scenarios, these techniques face significant limitations when applied to complex and large-scale simulations. They require substantial computational resources and exhibit poor scalability with increasing problem size, rendering them impractical for high-dimensional, large-scale, or long-term physics systems due to excessive computational costs and memory demands.

Recent advances in deep learning have revolutionized the field of PDE modeling by introducing data-driven methodologies that significantly accelerate computations for science while maintaining high accuracy. Supported by universal approximation theorem (Chen & Chen, 1995), neural operators that learn the mapping between two function spaces have demonstrated great success in simulating various PDE systems across multiple scientific disciplines without retraining for new conditions (Li et al., 2020; Lu et al., 2021). With the success of transformer in natural language processing and computer vision (Vaswani, 2017; Dosovitskiy, 2020), transformer-based neural operators can process multiple input functions while enabling arbitrary querying of output function locations, offering enhanced flexibility in handling complex functional mappings (Hao et al., 2023; Li et al., 2022). A series of neural operators have been developed to address complex scientific problems, including applications in weather forecasting, turbulent fluid dynamics, and boiling phenomena (Pathak et al., 2022; Li et al., 2023b; Bi et al., 2023; Hassan et al., 2023; Lin et al., 2021).

Despite success of these methods, accurately predicting the long-term evolution of these systems presents unique challenges in terms of both computational efficiency and long-term stability. The inherent multi-scale nature and multi-physics complexity of such systems necessitate methodologies that can efficiently represent and integrate dynamics across disparate spatial and temporal scales while simultaneously capturing the complex interactions between distinct physical processes, such as the influence of micro-scale turbulence on macro-scale flow (Peters, 2009; Natrajan et al., 2007). For example, large-scale atmospheric pressure sys-

1. Introduction

Predicting long-term spatio-temporal dynamics governed by partial differential equations (PDEs) is a cornerstone of scientific and engineering research, with broad applications

¹School of Computing, University of Utah, Salt Lake City, UT, United States ²Applied Mathematics and Computational Research Division, Lawrence Berkeley National Lab, Berkeley, CA, United States. Correspondence to: Zhe Bai <zhebai@lbl.gov>, Da Long <da.long@utah.edu>.

tems, such as high-pressure ridges and low-pressure troughs, play a crucial role in shaping local weather patterns. Inaccurate representation of these systems can cause significant errors in forecasting rainfall, wind speed, and temperature (Wang et al., 2006; Barlow et al., 2019). In magnetically confined plasmas, the magnetohydrodynamic (MHD) instabilities caused by current or pressure gradients can limit burning plasma performance and threaten fusion device integrity (von Goeler et al., 1974; Graves et al., 2012; Seo et al., 2024). Although neural operators present advantages over traditional approaches, they still encounter challenges associated with computational complexity and the demands for scientific fidelity, especially when the underlying physics involve rapid changes or high-frequency components. These challenges are further intensified in high-resolution simulations of multi-scale scenarios. Recent efforts to address these limitations include decomposing learnable parameters (Kossaifi et al., 2023), and employing linear attention to reduce computational time (Hao et al., 2023; Cao, 2021). Furthermore, integrating uncertainty quantification (UQ) into modeling frameworks is essential for assessing the confidence and reliability of predictions, particularly in complex systems such as in fluids and weather forecasting (Cheung et al., 2011; Scher & Messori, 2018). However, most existing neural operators do not incorporate UQ, which is particularly important for modeling long-term dynamics, where even small errors can propagate across scales and result in significant inaccuracies.

On the other hand, existing approaches for predicting spatio-temporal dynamics can be classified into two primary categories. The first category comprises models that directly forecast future states at a predetermined number of time steps based on historical observed states (Wang et al., 2024; Kontolati et al., 2024). The second category includes models that utilize an autoregressive manner, which addresses challenges of scaling and fitting complexities as a continuous-time emulator (Pathak et al., 2022; Rühling Cachay et al., 2024; Lippe et al., 2023; McCabe et al., 2023). Nevertheless, prediction errors that arise in short-term accumulate over time, resulting in unstable and inaccurate long-term forecasts. Previous work has sought to alleviate these issues through techniques such as the pushforward trick, invariance preservation, and iterative refinement (McCabe et al., 2023; Lippe et al., 2023; Brandstetter et al., 2022). However, developing robust, multi-scale modeling frameworks for long-term dynamic predictions of these systems remains crucial for improving prediction accuracy and understanding the complex interplay between scales.

In this work, we introduce a spatio-temporal Fourier transformer StFT for the forecasting of multi-scale and multi-physics systems. Each StFT block captures the underlying dynamics of the physical processes at one specific scale. Through a hierarchical composition of StFT blocks across

multiple scales, augmented by a generative residual correction block, our model is capable of learning the intricate interactions both within the same scales and across different scales, while also providing meaningful uncertainty quantification. The cascading StFT blocks enable our model to predict high-resolution dynamics across a spectrum of varying scales for correlated physical states. By integrating StFT within an auto-regressive framework, our method achieves superior accuracy in both short-term and long-term predictions compared to existing state-of-the-art autoregressive baselines.

Our contributions can be summarized as follows:

- We propose Spatio-temporal Fourier transformer (StFT), a novel ML model that learns underlying dynamics across various scales while incorporating interactions among diverse physical processes.
- We propose StFT-F, which incorporates a residual correction mechanism to refine the forecasting of StFT and capture intrinsic uncertainties.
- In each StFT block, we propose an overlapping tokenizer and a detokenizer that share boundaries between adjacent regions, enhancing the smoothness of vision representation.
- We demonstrate the effectiveness of StFT in an autoregressive framework on a diverse set of applications. Evaluating performance across variables in each complex system, StFT outperforms the best baseline, achieving over 600%, 20%, and 50% higher accuracy in the plasma, fluid, and atmospheric dynamics datasets, respectively. Compared with StFT, StFT-F has an average of 5 percent improvement while simultaneously providing meaningful uncertainty quantification.

2. Related Work

Neural Operators. There have been many popular neural operators and their variants including Fourier neural operators (Li et al., 2020; Gupta et al., 2021; Tran et al., 2021; Cao et al., 2023; Li et al., 2023a; Rahman et al., 2022), DeepONet (Lu et al., 2021; Wang et al., 2021; Jin et al., 2022; Wang et al., 2022; Kontolati et al., 2023; Prasthofer et al., 2022), transformer based operators (Hao et al., 2023; Cao, 2021; Li et al., 2022), image-to-image inspired operators (Gupta & Brandstetter, 2022; Rühling Cachay et al., 2024; Long et al., 2024). U-Net, a fundamentally hierarchical structure model, has inspired several neural operators (Rahman et al., 2022; Liu et al., 2022b; Gupta & Brandstetter, 2022), which allows to solve multi-scale PDEs by hierarchically aggregating of feature representations progressing from fine to coarse levels. Recent work in computer vision (Liu et al., 2021; Fan et al., 2021; Zhang et al.,

2022) have introduced methods for extracting multi-scale features through hierarchical architectures. However, these hierarchical models do not explicit forecast the multi-scale structures of physical processes. In contrast, our method begins with a coarse approximation that captures large-scale, low-frequency phenomena, and incrementally refines the representation over layers to resolve finer details. These structured decompositions allow for error diagnosis, enhancing interpretability of model performance across different scales, and enable targeted improvements. In contrast, U-Net like hierarchical models lack explicit refinement mechanisms, which limits the transparency and interpretability of their multi-scale representations.

Generative Models

Recent years, generative models, especially diffusion models have demonstrated huge success in various domains, including vision, audio, robotics, and medical field (Ho et al., 2020; Song et al., 2020; Tian et al., 2024; Kong et al., 2020; Wolleb et al., 2022). Video generation techniques are closely related to spatio-temporal dynamics prediction, as demonstrated in (Ho et al., 2022; Voleti et al., 2022; Singer et al., 2022). Recent work emulates the forward and inverse transformations of diffusion models to learn the spatio-temporal dynamics (Rühling Cachay et al., 2024). As an alternative approach in generative modeling, flow matching has been introduced to enable fast sampling and has since been applied to video generation (Lipman et al., 2022; Liu et al., 2022a; Albergo & Vanden-Eijnden, 2022; Polyak et al., 2024; Esser et al.). While video generation typically explores a range of creative and diverse possibilities from text or image prompts, forecasting spatial-temporal dynamics driven by PDEs necessitates more than mere statistical resemblance - it requires each prediction is firmly grounded in the underlying physical laws. To achieve accurate forecasting while capturing the inherent stochasticity of physical processes, our work incorporates a flow matching block following the proceeding StFT blocks. This approach allows our model to precisely recover the true dynamics while simultaneously offering meaningful uncertainty quantification.

3. Method

We consider an autoregressive formulation for long-term multi-scale spatiotemporal physical processes. We define a vector \tilde{u}_t representing the historical snapshots of the multi-physics variables at timestamps from $t - k + 1$ to t , in a total of k snapshots. Specifically, $\tilde{u}_t = [u_t, u_{t-1}, \dots, u_{t-k+1}]$. We formulate the probabilistic one-step forward neural operator StFT-F as

$$u_t = \mathcal{F}_{\theta_d}(\tilde{u}_{t-1}) + r_t, \quad (1)$$

$$r_t \sim \mathcal{P}_{\theta_g}(r|\tilde{u}_{t-1}, \mathcal{F}_{\theta_d}(\tilde{u}_{t-1})), \quad (2)$$

where \mathcal{F}_{θ_d} denote the StFT operator, a deterministic forecasting parameterized by θ_d , and \mathcal{P}_{θ_g} is the generative flow matching block parameterized by θ_g for refining the forecasting of StFT and providing uncertainty quantification. \mathcal{F}_{θ_d} represents the deterministic evolution of the system that encapsulates the predictive dynamics. The term $\mathcal{P}_{\theta_g}(r|\tilde{u}_{t-1}, \mathcal{F}_{\theta_d}(\tilde{u}_{t-1}))$ captures the probabilistic nature of the residual of the model. It represents a distribution conditioned on the current state \tilde{u}_{t-1} , and the deterministic prediction, modeling the uncertainty or variations that are not captured by the deterministic component. The residual r_t calculates deviations from the deterministic prediction, and its distribution allows the model to account for noise or inherent stochasticity in the physical processes. Therefore, by sampling residual r_t , our model learns stochastic trajectories from data. Besides providing predictive uncertainties caused by inherent noise, these stochastic trajectories can help study the long-term behavior, stabilities, and bifurcations in stochastic systems.

Figure 1 presents the overview of the StFT’s architecture. In the following subsections, we introduce the overlapping tokenizer, detokenizer, StFT block, and the residual refinement mechanism based on flow matching.

3.1. StFT Block

Figure 2 illustrates the architecture of the StFT block. Assuming the collected data samples \tilde{u}_t are mapped to a tensor of $W \times H \times C$, where W, H are spatial dimensions, and C is for the number of physical variables. Each StFT block is designed to handle a specific scale; therefore, by employing a specific patch size, we partition the input at a corresponding level of granularity. To ensure that the first block captures the coarsest features or the largest scale, we set the patch size $P_w^1 \times P_h^1$ to a large value, allowing it to model broad spatial structures effectively. Overlaps between patches are not considered in this section. As a result, $\frac{W}{P_w^1} \times \frac{H}{P_h^1} \times C$ patches or tokens are fed into the variable mixer block, where tokens corresponding to different physical processes but sharing the same spatial and temporal locations are mixed into a single token. Following this, two transformation paths are performed: one in the spatio-temporal domain, and the other in the frequency domain. In the spatio-temporal path, the tokens first pass through a spatio-temporal embedder, after which the embeddings are processed by multiple standard transformer layers. In the frequency path, the tokens are first processed by a 2D/3D Fourier transform, where only low-frequency components are retained. These low-frequency components are then passed through a frequency embedder to obtain frequency embeddings. Subsequently, these frequency embeddings are fed to the standard transformer layers for further processing. Finally, an inverse 2D/3D Fourier transform is applied to map the frequency tokens back to the spatio-temporal domain. Once the tokens from both paths

are obtained, a linear layer is used to project them back to the original space, where each token has a shape of $P_w^1 \times P_h^1$. Next, a detokenizer yields the first block prediction u_{t+1}^1 . Note that the number of tokens remains unchanged within the transformer block. Each patch represents a significant portion of the historical snapshots, encapsulating macroscopic structural features. This coarse-level partitioning reduces the complexity of modeling fine-grained details. By maintaining a lower granularity, the model prioritizes structural coherence over extraneous details, enabling it to focus on capturing and predicting global relationships between regions more effectively.

3.2. A Hierarchy of StFT Blocks

In the subsequent blocks, we shift our focus to smaller scales with fine details. Consequently, we concatenate the prediction u_{t+1}^1 with the input \tilde{u}_t , and consider this combined input as the input for the next StFT block. We further subdivide each patch from the previous level into smaller patches. The smaller patches allows for less information to be aggregated within a single patch, thereby minimizing the risk of losing local variations and enhancing the richness and informativeness of the fine-scale representation. As a result, the subdivision increases the resolution or the total number of patches to the next StFT block. By leveraging the finer granularity of the patches to focus on smaller regions, it allows the model to better localize features and capture their details. In addition, conditioning on the coarser prediction allows the models to iteratively refine its estimates, beginning with a broad global summary of the corresponding regions. This iterative refinement process *mirrors* how humans process images, first focusing on global structures and then honing in on finer details. Through the repeated subdivision of the patches, the model progressively refines its predictions across multiple scales.

3.3. Residual Refinement Based on Flow Matching

Finally, the model refines its predictions through a rectified flow block, which belongs to the family of flow matching models (Liu, 2022; Liu et al., 2022a). Flow matching is formulated as an ordinary differential equation on time $t \in [0, 1]$,

$$\frac{d}{dt}\psi_t(x) = \nu_t(\psi_t(x)), \quad (3)$$

where the learnable velocity field ν_t directs the transformation of each sample X_0 from a source distribution p_0 , typically a Gaussian distribution, toward the target distribution $X_1 \sim p_1$, with p_1 representing the data distribution. If we prescribe the velocity field ν_t such that it guides every sample along a straight-line trajectory from X_0 to X_1 , it is referred to as a rectified flow.. In this case, X_t represents the linear interpolation across the entire timespan between

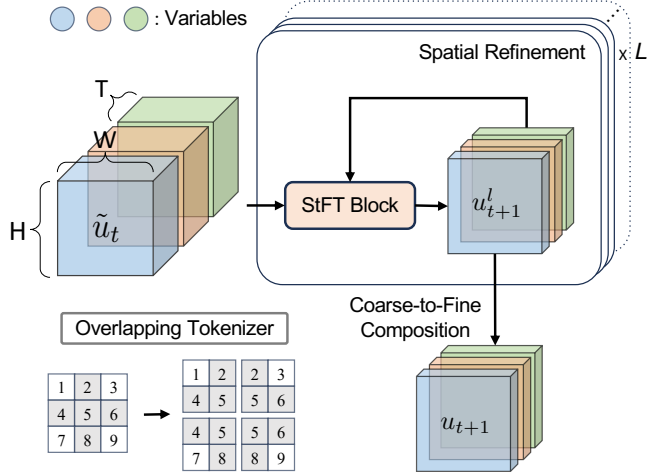


Figure 1. Overview of the proposed StFT model. Top right: the model predicts u_{t+1} using the past k snapshots \tilde{u}_t , employing L spatial refinements from coarse to fine scales through the proposed spatiotemporal Fourier transformer blocks. Bottom left: an illustration of the overlapping tokenizer, where the patch size is 2×2 , and the overlapping number is 1×1 .

X_0 and X_1 , which can be expressed as

$$X_t = tX_1 + (1 - t)X_0. \quad (4)$$

We employ a parameterized \mathcal{M}_{θ_g} to approximate ν_t , leading to the following learning objective:

$$\mathcal{L}(\theta_g) = \mathbb{E}_{t, X_0, X_1} \|\mathcal{M}_{\theta_g}(X_t) - (X_1 - X_0)\|^2. \quad (5)$$

In our model, the block takes the prediction u_{t+1}^L from the final block of StFT and the past snapshots \tilde{u}_t as conditioning inputs. Its objective is to generate the distribution of residuals $y - \sum_j u_{t+1}^j$, where y is the ground truth for the solution at $t + 1$. Our training loss then becomes:

$$\mathbb{E}_{X_0 \sim \mathcal{N}(0, \mathbf{I}), t \sim (0, 1)} [(\mathcal{M}_{\theta_g}(\tilde{u}_T, \mathcal{F}_{\theta_d}(\tilde{u}_T), (1 - t) * X_0 + t * (y - \sum_j u_{T+1}^j)) - (y - \sum_j u_{T+1}^j - X_0))^2], \quad (6)$$

where we use T to denote the timestamp in \tilde{u}_T and u_{T+1} .

3.4. Overlapping Tokenizer and Detokenizer

Next, we introduce the overlapping tokenizer and the overlapping detokenizer in our model. The overlapping tokenizer is illustrated in bottom left of Figure 1. For any adjacent two patches, they share a boundary. For instance, a 3×3 input generates four 2×2 patches, with gray boxes representing the shared boundaries. The overlapping detokenizer reconstructs these shared boundaries by computing the mean over the overlapping regions of any two adjacent

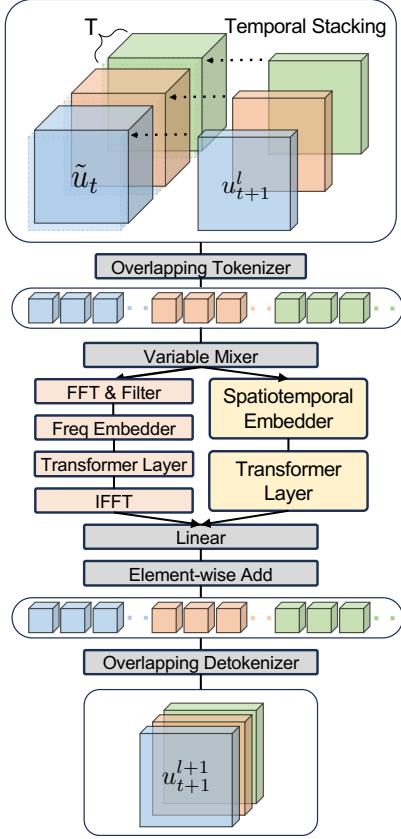


Figure 2. Illustration of the proposed spatiotemporal Fourier transformer block for StFT model: first, the past snapshots \tilde{u}_t and a coarser prediction from the previous layer are temporally stacked. The stacked discretized function is then passed to the overlapping tokenizer to generate tokens for each variable. Next, tokens corresponding to different variables in the same spatial and temporal locations are mixed through a variable mixer. Following this, two paths of transformation are applied on the tokens. After passing through the overlapping detokenizer, a finer prediction for times-tamp $t + 1$ is obtained.

patches. This strategy effectively mitigates discontinuity issues at the patch boundaries, which is particularly important when the target function that the model aims to predict is continuous and smooth. In addition, it improves feature learning and extraction by incorporating shared boundary features into the patch embeddings. This not only preserves local continuity, but also enables the model to capture spatial relationships between patches more effectively, particularly during the self-attention mechanism.

4. Experiments

4.1. Datasets

In this section, we consider three spatio-temporal multi-physics systems arising from time-dependent PDEs of

Algorithm 1 StFT

Initialize: For each StFT block $l \in [1, L]$, patch sizes $(p_h, p_w)_l$, truncation frequencies $(m_h, m_w)_l$, overlapping numbers $(o_h, o_w)_l$, u_{t+1}^0 as None, and output u_{t+1} as zeros.

Input: $\tilde{u}_t \leftarrow (u_t, u_{t-1}, \dots, u_{t-k+1})$

Output: u_{t+1}

$i \leftarrow i$ -th token

$v \leftarrow v$ -th variable in V

for $l = 1$ **to** L **do**

$x_t \leftarrow \text{TemporalStacking}(\tilde{u}_t, u_{t+1}^{l-1})$

$\{x_{i,t}\}_v \leftarrow \text{OLTTokenizer}(x_t, p_h, p_w, o_h, o_w)$

Frequency Path

$\bar{x}_{i,t,1} \leftarrow \text{VariableMixer}^1(\{x_{i,t}\}_{v \in V})$

$f_{i,t} \leftarrow \text{FreqEmbedder}(\text{FFTFilter}(\bar{x}_{i,t,1}, m_h, m_w))$

$f_{i,t} \leftarrow \text{TransformerBlock}^1(f_{i,t})$

$\{x_{i,t+1,1}\}_v \leftarrow \text{Linear}^1(\text{IFFT}(f_{i,t}, m_h, m_w))$

Spatiotemporal Path

$\bar{x}_{i,t,2} \leftarrow \text{VariableMixer}^2(\{x_{i,t}\}_{v \in V})$

$e_{i,t} \leftarrow \text{StEmbedder}(\bar{x}_{i,t,2})$

$e_{i,t} \leftarrow \text{TransformerBlock}^2(e_{i,t})$

$\{x_{i,t+1,2}\}_v \leftarrow \text{Linear}^2(e_{i,t})$

Merge

$\{x_{i,t+1}\}_v \leftarrow \{x_{i,t+1,1}\}_v + \{x_{i,t+1,2}\}_v$

$u_{t+1}^l \leftarrow \text{OLDetokenizer}(\{x_{i,t+1}\}_v, p_h, p_w, o_h, o_w)$

$u_{t+1} \leftarrow u_{t+1} + u_{t+1}^l$

end for

return u_{t+1}

a variety of complexities, including a high-dimensional plasma dynamics system based on reconstructed equilibrium of DIII-D experimental discharges, a 2D incompressible Navier-Stokes equation in velocity-pressure form within a squared boundaries driven by an external force, and a viscous shallow-water equation modeling the dynamics of large-scale atmospheric flows on a spherical domain. The problem setup and data generation are detailed in Appendix A.

4.2. Experimental Setup and Baselines

Our goal is to simulate long-time trajectories given a few initial observations. This task is particularly challenging due to the multiple correlated variables present in the Navier-Stokes and plasma magnetohydrodynamics (MHD), with the test trajectories consisting of snapshots that vary from 71 to 244. We employ an autoregressive framework for all the methods. During training, each model utilizes five historical snapshots to predict the subsequent snapshot. For testing, given the initial five snapshots of a trajectory, all models forecast the entire trajectories in an autoregressive manner, iteratively generating future states based on the previously forecasted outputs from the trained model.

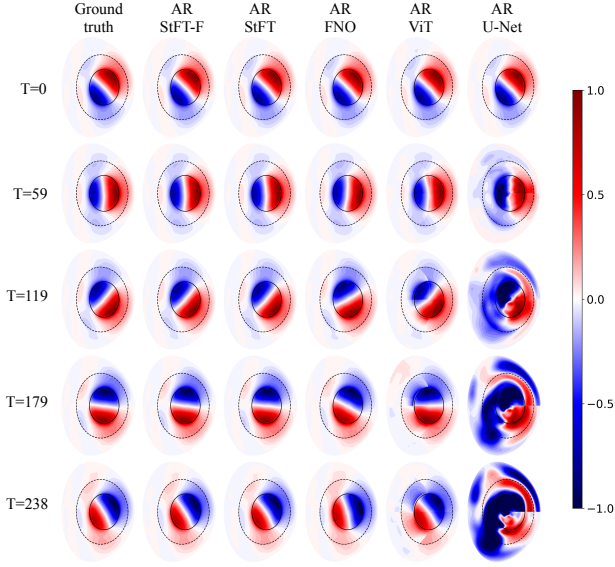


Figure 3. Temporal evolution of perturbed electrostatic potential $\delta\phi$ contours predicted by different models: StFT-F, StFT, FNO, ViT and U-Net. Significant phase differences between the predictions of the models appear after $T = 59$, where StFT and StFT-F perform stable across the forecasting window.

We evaluate these datasets using the following well-known and state-of-the-art methods for comparison: autoregressive Fourier Neural Operator (AR-FNO), autoregressive U-Net (AR-UNet), and autoregressive vision transformer (AR-ViT). For FNO, we use the authors’ open source implementation. For U-Net, we employ the implementation of the modified U-Net as evaluated in the recent BubbleML work (Hassan et al., 2023), where the modified U-Net demonstrated superior performance across most of their proposed benchmarks.

We divide the trajectories of each dataset into training, test, and validation sets. For each method, we identify the tunable hyperparameters, specify a range for each hyperparameter, and conduct a grid search across all the hyperparameters. For our methods, we implement both StFT, the deterministic component of our model, and StFT-F, our model with a generative residual refinement block. First, we evaluate the forecasting performance of all methods by calculating the mean L_2 relative error over the test trajectories. For StFT-F, in order to obtain the mean prediction for all the test trajectories, we generate 50 prediction samples at each autoregressive step, and then inject their mean into the next step. Additionally, we assess the uncertainty quantification capability of our method by sampling 100 trajectories for each test case. For each trajectory, at each autoregressive step, we generate a single prediction sample, which is subsequently fed into the next autoregressive step to iteratively forecast the full sequence. We ensure that all models are




Figure 4. Results of autoregressive prediction in L_2 relative error across the timespan: (top) perturbed parallel vector potential δA_{\parallel} in plasma MHD, and (bottom) magnitude of velocity in shallow-water equation. The shaded region indicates the uncertainty distribution of 3σ in the relative error of StFT-F.

thoroughly trained by running sufficient epochs. We used the AdamW optimizer with a learning rate of $1e^{-4}$ to train those models on an A100 GPU with a total memory capacity of 40 GB. A comprehensive list of all the hyperparameters along with their respective ranges used for grid search is provided in the Appendix C.

4.3. Main Results

Table 1 shows the forecasting performance of all the models on the three applications. Further details regarding the the problem setup and dataset can be found in Appendix A. We present several test trajectories and visualize their uncertainties across the forecasting time horizon estimated by StFT-F, as illustrated in Figure 6.

StFT performs significantly better than all other baselines across all physical processes in the three applications. In the plasma MHD dataset, the test trajectory has a total of six coupled physics variables and 244 snapshots. On average, StFT achieves a reduction in error by a factor of six compared to the best baseline, AR-ViT. Although AR-FNO maintains a high resolution in its long-term prediction, it fails to capture the correct dynamics of mode evolution, leading to out-of-phase predictions as shown in Figure 3.

We examine the error growth by plotting the L_2 relative errors over time, as illustrated in Figure 4 for several rep-

Table 1. Relative L_2 error over three spatiotemporal prediction systems, including linear plasma dynamics, Navier-Stokes equations, and shallow-water equations for multi-scale, multi-physics study. AR-StFT refers to the deterministic model’s results. AR-StFT-F denotes the outcomes of the deterministic model with a refinement of flow matching layer for uncertainty quantification, where at each autoregressive step, 50 sampled are generated to get a mean prediction. The best and second best prediction results for each case are marked in blue and orange.

Dataset	Variable(s)	AR-StFT	AR-StFT-F	AR-FNO	AR-ViT	AR-UNet
Plasma MHD	$\delta\phi$	2.80e-2	2.24e-2	2.28e-1	1.73e-1	1.02e0
	$\delta A_{ }$	2.45e-2	2.30e-2	2.30e-1	3.24e-1	8.13e-1
	$\delta B_{ }$	3.05e-2	2.66e-2	2.33e-1	1.95e-1	7.79e-1
	δn_e	2.84e-2	2.45e-2	2.33e-1	2.08e-1	1.01e0
	δn_i	3.28e-2	2.93e-2	2.33e-1	2.18e-1	1.04e0
	δu_e	3.99e-2	3.73e-2	3.18e-1	2.99e-1	6.96e-1
Navier-Stokes	u	3.38e-2	3.30e-2	4.46e-2	5.09e-2	6.16e-2
	v	3.60e-2	3.17e-2	4.57e-2	4.60e-2	6.15e-2
	p	5.16e-2	4.44e-2	5.90e-2	7.03e-2	7.84e-2
Shallow-Water	\mathbf{V}	6.25e-2	6.53e-2	9.53e-2	1.33e-1	2.02e-1

representative variables. In the bottom figure for the shallow-water dataset, AR-FNO first appears to slightly better than all other methods during the short term from timestamp 0 to 0.2, and StFT-F shows superior performance soon after. For the plasma dataset, StFT-F demonstrates dominance starting from $t = 0.2$ with a stable performance, while the errors of all baseline methods begin to propagate from that point onward, resulting in a rapid decline compared to StFT-F. Notably, StFT-F exhibits long-term stability relative to the other methods. We also compare the error over time comparing StFT and StFT-F in Figure 5. StFT-F consistently shows improved performance throughout the latter half of temporal domain for the Navier-Stokes dataset. For the plasma dataset, StFT-F begins to prevail from $t = 0.6$. These results indicate that our method achieves superior long-term stability and accuracy among all other methods.

For the shallow-water and Navier-Stokes datasets, on average, StFT reduces the errors by an average of 50% and 25%, respectively. In Navier-Stokes and plasma datasets, StFT-F not only surpasses StFT but also offers the additional capability of uncertainty quantification, achieving error reductions of 10% in both cases. In the shallow-water dataset, we observe a slight increase in error with StFT-F. More experiment and visualization and results are included in Appendix D.

4.4. Uncertainty Quantification

Figure 6 presents the distribution of the empirical standard deviation along with the mean prediction. The generation of test example predictions follows a stochastic recursive process. We generate $S = 100$ sample trajectories for each test example, where at each time step t , the trajectory sample $u_t^{(s)}$ for $s = 1, \dots, S$ is computed as: $u_t^{(s)} = \mathcal{F}_{\theta_d}(\tilde{u}_{t-1}^{(s)}) + r_t^{(s)}$, where the residual term is sampled from: $r_t^{(s)} \sim P_{\theta_g}(r|\tilde{u}_{t-1}^{(s)}, \mathcal{F}_{\theta_d}(\tilde{u}_{t-1}^{(s)}))$. With




Figure 5. Results of comparing StFT and StFT-F for the autoregressive prediction in L_2 relative error across different timestamps. (top) velocity component u in Navier-Stokes equation, and (bottom) magnitude of velocity in shallow-water equation. The shaded region represents the standard deviation distribution of the relative error of StFT-F. In the bottom figure, the uncertainty is negligible that it is not visually discernible. StFT-F demonstrates better performance in the latter stages of the forecasting time horizon compared to StFT.




Figure 6. Evaluation of the forecasting results for three applications: the ground truth, StFT-F prediction, residual of StFT-F prediction, and the associated uncertainty across the time horizon, presented across the initial state (left panel) and final state (right panel) for multiple variables: (a-b) plasma: perturbed electrostatic potential ϕ and electron number density n_e , (c-d) Navier-Stokes: velocity component u and pressure p , (e) shallow-water: velocity field \mathbf{V} .

the generated S trajectory samples, we compute the empirical standard deviation as: $\bar{u}_t = \frac{1}{S} \sum_{s=1}^S u_t^{(s)}$, and $\sigma_t = \sqrt{\frac{1}{S} \sum_{s=1}^S (u_t^{(s)} - \bar{u}_t)^2}$. As observed, regions with large errors correspond to those exhibiting significant uncertainties predicted by StFT-F. Additionally, it is evident that uncertainties increase with time. This aligns with our expectation, as errors accumulate during the autoregressive forecasting process. The StFT-F model provides aleatoric uncertainty quantification for studying and analyzing the long-term behaviors of complex physical systems.

4.5. Contribution of Each StFT Block

To assess the contribution of each StFT block corresponding to a specific scale in fitting the training data, we quantify the weight of each block using the following measure:

$$\mathcal{W}_i = \frac{\|\mathbf{y}_i\|_2}{\|\mathbf{y}\|_2}, \quad (7)$$

where \mathbf{y}_i represents the prediction from the i -th StFT block, and \mathbf{y} denotes the ground truth. We normalize the contributions, and present the contribution of each StFT block for each dataset in Figure 7. A greater contributing factor from the fine-scale layer in StFT is observed in the Navier-Stokes and shallow-water equations, attributed to




Figure 7. The contribution of each StFT block across the three datasets. From bottom to top, the patch size decreases. For this plasma data, the first two levels are sufficient to capture the multi-scale structures, whereas for Navier-Stokes and shallow-water equations, the finest scale contributes more significantly.

the sharper changes and smaller scale structures inherent in the dynamics of higher-order nonlinearities. Ablation experiments regarding the impact of multi-scale StFT blocks and frequency path is elaborated in Appendix B.

5. Conclusion

In this paper, we propose a spatio-temporal Fourier transformer (StFT) for multi-scale and multi-physics long-term

dynamics forecasting. Specifically, each StFT block is tailored to address a particular spatial scale, and through a hierarchical composition of multiple StFT blocks spanning different scales, StFT learns the interplay between multiple scales and interactions between multiple physical processes, resulting in stable and accurate long-term dynamics forecasting in an autoregressive manner. Furthermore, we propose and demonstrate the use of a generative residual correction mechanism, which enables meaningful quantification of uncertainties in the predictive model. Our approach demonstrates superior performance compared to several state-of-the-art baselines in the domains of plasma, fluid, and atmospheric dynamics.

Impact Statement

Turbulence remains one of the great unsolved problems in physics. Yet turbulence, whether driven by gravity, heating, or magnetic fields, manifests in physical phenomena spanning multi-scale fluid flow to plasma fusion reactors to planetary atmospheres to convective layers in stars to the flow of interstellar gas in stellar nurseries that span trillions of kilometers. Where mathematics failed, data-driven machine learning models provide a pathway to understanding turbulence and gaining insights into not only the origins and ultimate fate of stars and planetary ecosystems, but also opens pathways to nearly infinite sources of clean energy. Our StFT model presented in this paper is a stepping stone towards obviating the exponentially increasing computational costs of simulating geophysical systems (an inherently chaotic dynamical system requiring vast ensembles of high-resolution small time step simulations) as well as realization of digital twins to aid in the design and operational control of tokamak-based fusion power plants.

Acknowledgement

ZB, DL, SW and LO gratefully acknowledge support from the U.S. Department of Energy, Office of Science, SciDAC/Advanced Scientific Computing Research under Award Number DE-AC02-05CH11231. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract Number DE-AC02-05CH11231.

References

Albergo, M. S. and Vanden-Eijnden, E. Building normalizing flows with stochastic interpolants. *arXiv preprint arXiv:2209.15571*, 2022.

Barlow, M., Gutowski, W. J., Gyakum, J. R., Katz, R. W.,

Lim, Y.-K., Schumacher, R. S., Wehner, M. F., Agel, L., Bosilovich, M., Collow, A., et al. North american extreme precipitation events and related large-scale meteorological patterns: a review of statistical methods, dynamics, modeling, and trends. *Climate Dynamics*, 53:6835–6875, 2019.

Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., and Tian, Q. Accurate medium-range global weather forecasting with 3d neural networks. *Nature*, 619(7970):533–538, 2023.

Brandstetter, J., Worrall, D., and Welling, M. Message passing neural pde solvers. *arXiv preprint arXiv:2202.03376*, 2022.

Burns, K. J., Vasil, G. M., Oishi, J. S., Lecoanet, D., and Brown, B. P. Dedalus: A flexible framework for numerical simulations with spectral methods. *Physical Review Research*, 2(2):023068, 2020.

Cao, Q., Goswami, S., and Karniadakis, G. E. Lno: Laplace neural operator for solving differential equations. *arXiv preprint arXiv:2303.10528*, 2023.

Cao, S. Choose a transformer: Fourier or galerkin. *Advances in neural information processing systems*, 34: 24924–24940, 2021.

Chen, T. and Chen, H. Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems. *IEEE transactions on neural networks*, 6(4):911–917, 1995.

Cheung, S. H., Oliver, T. A., Prudencio, E. E., Prudhomme, S., and Moser, R. D. Bayesian uncertainty analysis with applications to turbulence modeling. *Reliability Engineering & System Safety*, 96(9):1137–1149, 2011.

Dosovitskiy, A. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al. Scaling rectified flow transformers for high-resolution image synthesis, 2024. [URL https://arxiv.org/abs/2403.03206](https://arxiv.org/abs/2403.03206), 1.

Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., and Feichtenhofer, C. Multiscale vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 6824–6835, 2021.

Graves, J., Chapman, I., Coda, S., Lennholm, M., Albergante, M., and Jucker, M. Control of magnetohydrodynamic stability by phase space engineering of energetic ions in tokamak plasmas. *Nature communications*, 3(1): 624, 2012.

- Gupta, G., Xiao, X., and Bogdan, P. Multiwavelet-based operator learning for differential equations. *Advances in neural information processing systems*, 34:24048–24062, 2021.
- Gupta, J. K. and Brandstetter, J. Towards multi-spatiotemporal-scale generalized pde modeling. *arXiv preprint arXiv:2209.15616*, 2022.
- Hao, Z., Wang, Z., Su, H., Ying, C., Dong, Y., Liu, S., Cheng, Z., Song, J., and Zhu, J. Gnot: A general neural operator transformer for operator learning. In *International Conference on Machine Learning*, pp. 12556–12569. PMLR, 2023.
- Hassan, S. M. S., Feeney, A., Dhruv, A., Kim, J., Suh, Y., Ryu, J., Won, Y., and Chandramowlishwaran, A. Bubbleml: a multi-physics dataset and benchmarks for machine learning. *arXiv preprint arXiv:2307.14623*, 2023.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., and Fleet, D. J. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022.
- Jin, P., Meng, S., and Lu, L. Mionet: Learning multiple-input operators via tensor product. *SIAM Journal on Scientific Computing*, 44(6):A3490–A3514, 2022.
- Kong, Z., Ping, W., Huang, J., Zhao, K., and Catanzaro, B. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020.
- Kontolati, K., Goswami, S., Karniadakis, G. E., and Shields, M. D. Learning in latent spaces improves the predictive accuracy of deep neural operators. *arXiv preprint arXiv:2304.07599*, 2023.
- Kontolati, K., Goswami, S., Em Karniadakis, G., and Shields, M. D. Learning nonlinear operators in latent spaces for real-time predictions of complex dynamics in physical systems. *Nature Communications*, 15(1):5101, 2024.
- Kossaifi, J., Kovachki, N., Azizzadenesheli, K., and Anandkumar, A. Multi-grid tensorized fourier neural operator for high-resolution pdes. *arXiv preprint arXiv:2310.00120*, 2023.
- Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.
- Li, Z., Meidani, K., and Farimani, A. B. Transformer for partial differential equations’ operator learning. *arXiv preprint arXiv:2205.13671*, 2022.
- Li, Z., Huang, D. Z., Liu, B., and Anandkumar, A. Fourier neural operator with learned deformations for pdes on general geometries. *Journal of Machine Learning Research*, 24(388):1–26, 2023a.
- Li, Z., Peng, W., Yuan, Z., and Wang, J. Long-term predictions of turbulence by implicit u-net enhanced fourier neural operator. *Physics of Fluids*, 35(7), 2023b.
- Lin, C., Li, Z., Lu, L., Cai, S., Maxey, M., and Karniadakis, G. E. Operator learning for predicting multiscale bubble growth dynamics. *The Journal of Chemical Physics*, 154(10), 2021.
- Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- Lippe, P., Veeling, B., Perdikaris, P., Turner, R., and Brandstetter, J. Pde-refiner: Achieving accurate long rollouts with neural pde solvers. *Advances in Neural Information Processing Systems*, 36:67398–67433, 2023.
- Liu, Q. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint arXiv:2209.14577*, 2022.
- Liu, X., Gong, C., and Liu, Q. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022a.
- Liu, X., Xu, B., and Zhang, L. Ht-net: Hierarchical transformer based operator learning model for multiscale pdes. 2022b.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 10012–10022, 2021.
- Long, D., Xu, Z., Yang, G., Narayan, A., and Zhe, S. Arbitrarily-conditioned multi-functional diffusion for multi-physics emulation. *arXiv preprint arXiv:2410.13794*, 2024.
- Lu, L., Jin, P., Pang, G., Zhang, Z., and Karniadakis, G. E. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.
- McCabe, M., Harrington, P., Subramanian, S., and Brown, J. Towards stability of autoregressive neural operators. *arXiv preprint arXiv:2306.10619*, 2023.

- Natrajan, V. K., Yamaguchi, E., and Christensen, K. T. Statistical and structural similarities between micro-and macroscale wall turbulence. *Microfluidics and Nanofluidics*, 3:89–100, 2007.
- Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., Kurth, T., Hall, D., Li, Z., Azizzadenesheli, K., et al. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. *arXiv preprint arXiv:2202.11214*, 2022.
- Peebles, W. and Xie, S. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.
- Peters, N. Multiscale combustion and turbulence. *Proceedings of the Combustion Institute*, 32(1):1–25, 2009.
- Polyak, A., Zohar, A., Brown, A., Tjandra, A., Sinha, A., Lee, A., Vyas, A., Shi, B., Ma, C.-Y., Chuang, C.-Y., et al. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024.
- Prasthofer, M., De Ryck, T., and Mishra, S. Variable-input deep operator networks. *arXiv preprint arXiv:2205.11404*, 2022.
- Rahman, M. A., Ross, Z. E., and Azizzadenesheli, K. U-no: U-shaped neural operators. *arXiv preprint arXiv:2204.11127*, 2022.
- Rühling Cachay, S., Zhao, B., Joren, H., and Yu, R. Dyffusion: A dynamics-informed diffusion model for spatiotemporal forecasting. *Advances in Neural Information Processing Systems*, 36, 2024.
- Scher, S. and Messori, G. Predicting weather forecast uncertainty with machine learning. *Quarterly Journal of the Royal Meteorological Society*, 144(717):2830–2841, 2018.
- Seo, J., Kim, S., Jalalvand, A., Conlin, R., Rothstein, A., Abbate, J., Erickson, K., Wai, J., Shousha, R., and Kolenmen, E. Avoiding fusion plasma tearing instability with deep reinforcement learning. *Nature*, 626(8000):746–751, 2024.
- Singer, U., Polyak, A., Hayes, T., Yin, X., An, J., Zhang, S., Hu, Q., Yang, H., Ashual, O., Gafni, O., et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- Tadmor, E. A review of numerical methods for nonlinear partial differential equations. *Bulletin of the American Mathematical Society*, 49(4):507–554, 2012.
- Tian, S., Zheng, M., and Liang, X. Transfusion: A practical and effective transformer-based diffusion model for 3d human motion prediction. *IEEE Robotics and Automation Letters*, 2024.
- Tran, A., Mathews, A., Xie, L., and Ong, C. S. Factorized fourier neural operators. *arXiv preprint arXiv:2111.13802*, 2021.
- Vaswani, A. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017.
- Voleti, V., Jolicœur-Martineau, A., and Pal, C. Mcvdmasked conditional video diffusion for prediction, generation, and interpolation. *Advances in neural information processing systems*, 35:23371–23385, 2022.
- von Goeler, S., Stodiek, W., and Sauthoff, N. Studies of internal disruptions and $m=1$ oscillations in tokamak discharges with soft-x-ray techniques. *Physical Review Letters*, 33(20):1201, 1974.
- Wang, B., Hoskins, B., and Wang, B. Large-scale atmospheric dynamics. *The Asian Monsoon*, pp. 357–415, 2006.
- Wang, S., Wang, H., and Perdikaris, P. Learning the solution operator of parametric partial differential equations with physics-informed deep operators. *Science advances*, 7(40):eabi8605, 2021.
- Wang, S., Wang, H., and Perdikaris, P. Improved architectures and training algorithms for deep operator networks. *Journal of Scientific Computing*, 92(2):35, 2022.
- Wang, S., Seidman, J. H., Sankaran, S., Wang, H., Pappas, G. J., and Perdikaris, P. Bridging operator learning and conditioned neural fields: A unifying perspective. *arXiv preprint arXiv:2405.13998*, 2024.
- Wolleb, J., Bieder, F., Sandkühler, R., and Cattin, P. C. Diffusion models for medical anomaly detection. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 35–45. Springer, 2022.
- Zhang, Z., Zhang, H., Zhao, L., Chen, T., Arik, S. Ö., and Pfister, T. Nested hierarchical transformer: Towards accurate, data-efficient and interpretable visual understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 3417–3425, 2022.

A. Problem Setup and Datasets

A.1. Plasma magnetohydrodynamic (MHD) equations

We consider magnetohydrodynamic (MHD) equations that characterize the plasma instabilities in fusion tokamaks. The coupled multi-physics system includes the continuity equation solving charge density δn , Poisson's equation solving $\delta\phi$, the Ampere's law to solving δu_{\parallel} , the Faraday's law with the assumption $E_{\parallel} = 0$ to solving δA_{\parallel} , and the perpendicular force balance equation to solving δB_{\parallel} . The first continuity equation for gyrocenter charge density is expressed as,

$$\begin{aligned} & \frac{\partial \delta n}{\partial t} + \mathbf{B}_0 \cdot \nabla \left(\frac{n_0 \delta u_{\parallel}}{B_0} \right) - n_0 \mathbf{v}_* \cdot \frac{\nabla B_0}{B_0} + \delta \mathbf{B}_{\perp} \cdot \nabla \left(\frac{n_0 u_{\parallel 0}}{B_0} \right) \\ & - \frac{\nabla \times \mathbf{B}_0}{e B_0^2} \cdot \left(\nabla \delta P_{\parallel} + \frac{(\delta P_{\perp} - \delta P_{\parallel}) \nabla B_0}{B_0} \right) \\ & + \nabla \cdot \left(\frac{\delta P_{\parallel} \mathbf{b}_0 \nabla \times \mathbf{b}_0 \cdot \mathbf{b}_0}{e B_0} \right) - \frac{\mathbf{b}_0 \times \nabla \delta B_{\parallel}}{e} \cdot \nabla \left(\frac{P_0}{B_0^2} \right) \\ & - \frac{\nabla \times \mathbf{b}_0 \cdot \nabla \delta B_{\parallel}}{e B_0^2} P_0 = 0, \end{aligned} \quad (8)$$

where n is the density, B is the magnetic field, u_{\parallel} is the parallel flow velocity, and P is the pressure. The perturbed quantities are denoted by δ with the equilibrium states including temperature, density, magnetic field and the flux surface from the reconstruction of DIII-D experiments. Here, $\delta n = \delta n_e + q_i \delta n_i / q_e$ stands for the difference of ion and electron density, and $\delta u_{\parallel} = \delta u_{\parallel e} + q_i \delta u_{\parallel i} / q_e$ denotes the difference of ion and electron flow. We have $\mathbf{v}_* = \mathbf{b}_0 \times \nabla (\delta P_{\parallel} + \delta P_{\perp}) / (n_0 m_e \Omega_e)$, where m_e is the electron mass, and $\Omega_e = e B_0 / m_e$ is the electron cyclotron frequency. The perturbed electron parallel flow δu_{\parallel} can be solved from Ampere's law,

$$\delta u_{\parallel} = \frac{1}{\mu_0 e n_0} \nabla_{\perp}^2 \delta A_{\parallel}, \quad (9)$$

where μ_0 is the permeability of vacuum. δA_{\parallel} is the perturbed vector potential. In the single fluid model, $E_{\parallel} = 0$ is assumed. Then δA_{\parallel} can be solved from

$$\frac{\partial A_{\parallel}}{\partial t} = \mathbf{b}_0 \cdot \nabla \phi, \quad (10)$$

and the electrostatic potential ϕ can be solved from gyrokinetic Poisson's equation (the quasi-neutrality condition)

$$\frac{c^2}{v_A^2} \nabla_{\perp}^2 \phi = \frac{e \delta n}{\epsilon_0}, \quad (11)$$

where c is the speed of light, v_A is the Alfvén velocity, and ϵ_0 is the dielectric constant of vacuum. The parallel magnetic perturbation δB is given by the perpendicular force balance,

$$\frac{\delta B_{\parallel}}{B_0} = -\frac{\beta_e}{2} \frac{\delta P_{\perp}}{P_0} = -\frac{\beta_e}{2} \frac{\partial P_0}{\partial \psi_0} \frac{\delta \psi}{P_0}. \quad (12)$$

The perturbed pressure in the fluid limit can be calculated by

$$\begin{aligned} \delta P_{\perp} &= \frac{\partial P_0}{\partial \psi_0} \delta \psi - 2 \frac{\delta B_{\parallel}}{B_0} P_0, \\ \delta P_{\parallel} &= \frac{\partial P_0}{\partial \psi_0} \delta \psi - \frac{\delta B_{\parallel}}{B_0} P_0. \end{aligned} \quad (13)$$

In these equations, ψ_0 and $\delta \psi$ is the equilibrium and perturbed magnetic flux, and the evolution of $\delta \psi$ is solved from

$$\frac{\partial \delta \psi}{\partial t} = -\frac{\partial \phi}{\partial \alpha_0}, \quad (14)$$

where α_0 is from the Clebsch representation of \mathbf{B} field, and $\mathbf{B}_0 = \nabla \psi_0 \times \nabla \alpha_0$. We run a linear gyrokinetic simulation with a $100 \times 250 \times 24$ mesh in radial, poloidal and parallel directions. The time step is set to $\Delta t = 0.005 R_0 / C_s = 1.483 \times 10^{-8} \text{s}$. We keep both $n = 0, 1$ modes, generate a trajectory of 128,000 time steps, and collect the data every 100 snapshots. We focus on emulating the dynamics of electrostatic potential $\delta \phi$, parallel vector potential δA_{\parallel} , electron number density δn_e , ion number density δn_i , and electron velocity δu_e in their trajectories.




Figure 8. Temporal evolution of normalized perturbed parallel vector potential δA_{\parallel} and perturbed electron density δn_e contours predicted by different models: StFT-F, StFT, FNO, ViT and U-Net. Significant phase differences between the predictions of the models appear after $T = 59$, where StFT and StFT-F perform stable across the forecasting time horizon.

A.2. 2D incompressible Navier-Stokes equations

We consider the 2D incompressible Navier-Stokes (NS) equation on a rectangular domain $(x, y) \in [0, 1]^2$,

$$\begin{aligned} \frac{\partial u}{\partial t} + \frac{\partial p}{\partial x} &= -u \frac{\partial u}{\partial x} - v \frac{\partial u}{\partial y} + \frac{1}{\text{Re}} \nabla^2 u + f(x, y), \\ \frac{\partial v}{\partial t} + \frac{\partial p}{\partial y} &= -u \frac{\partial v}{\partial x} - v \frac{\partial v}{\partial y} + \frac{1}{\text{Re}} \nabla^2 v + f(x, y), \\ \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} &= 0, \end{aligned} \quad (15)$$

where u and v represent the velocity components in the x and y directions, and p represents the pressure. $f(x, y)$ is the source term, and we set it to $e^{-100((x-0.5)^2 + (y-0.5)^2)}$. The Reynolds number is set to 1000. We run a finite difference solver to compute the solutions on a 50×50 spatial grid, with the temporal domain discretized into a total of 101 timestamps over $T \in [0, 20]$. We generate a total of 100 trajectories by sampling the four boundary conditions uniformly from $(0.1, 0.6)$.

A.3. Spherical shallow-water equations

We consider the viscous shallow-water equations modeling the dynamics of large-scale atmospheric flows:

$$\begin{aligned} \frac{D\mathbf{V}}{Dt} &= -f\mathbf{k} \times \mathbf{V} - g\nabla h + \nu \nabla^2 \mathbf{V}, \\ \frac{Dh}{Dt} &= -h\nabla \cdot \mathbf{V} + \nu \nabla^2 h, \quad x \in \Omega, \quad t \in [0, 1], \end{aligned} \quad (16)$$

where \mathbf{V} is the velocity vector tangential to the spherical surface, \mathbf{k} is the unit vector normal to the surface, h is the thickness of the fluid layer, $f = 2\Xi \sin \phi$ is the Coriolis parameter (Ξ being the Earth's angular velocity), g is the gravitational acceleration, and ν is the diffusion coefficient. The equations are defined over a spherical domain $\Omega = (\lambda, \phi)$, with longitude λ and latitude ϕ .

As an initial condition, a zonal flow typical of a mid-latitude tropospheric jet is defined for the velocity component u as a

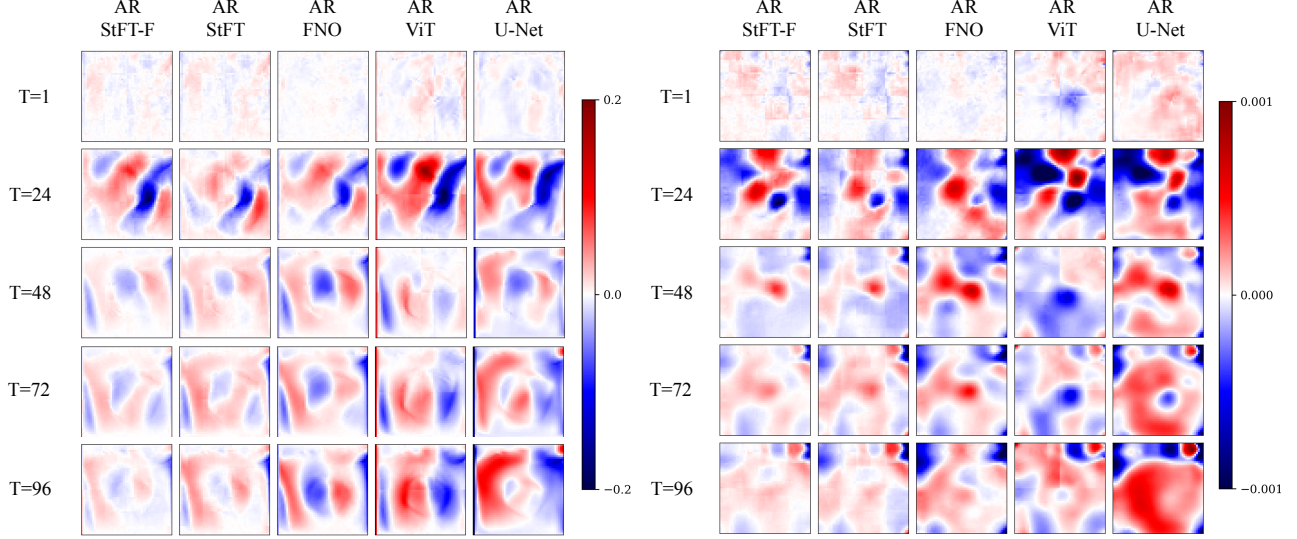


Figure 9. 2D incompressible Navier-Stokes equation: pointwise error of the predicted evolution of velocity component u and pressure p contours across different models: StFT-F, StFT, FNO, ViT and U-Net. For long-term predictions, StFT and StFT-F demonstrate lower residuals compared to other models.

function of latitude ϕ :

$$u(\phi, t = 0) = \begin{cases} 0, & \phi \leq \phi_0, \\ \frac{u_{\max}}{n} \exp \left[\frac{1}{(\phi - \phi_0)(\phi - \phi_1)} \right], & \phi_0 < \phi < \phi_1, \\ 0, & \phi \geq \phi_1, \end{cases}$$

where u_{\max} is the maximum zonal velocity, ϕ_0 and ϕ_1 represent the southern and northern boundaries of the jet in radians, and $n = \exp[-4/(\phi_1 - \phi_0)^2]$ normalizes u_{\max} at the midpoint of the jet. To induce barotropic instability, a localized Gaussian perturbation is added to the height field, expressed as:

$$h'(\lambda, \phi, t = 0) = \hat{h} \cos(\phi) \exp \left[- \left(\frac{\lambda}{\alpha} \right)^2 \right] \exp \left[- \left(\frac{\phi_2 - \phi}{\beta} \right)^2 \right],$$

where $-\pi < \lambda < \pi$, and parameters \hat{h} , ϕ_2 , α , and β control the shape and location of the perturbation. The parameters α and β are sampled from uniform distributions $\alpha \sim U[0.1, 0.5]$ and $\beta \sim U[0.03, 0.2]$. We ran the solver from Dedalus (Burns et al., 2020) on a 256×256 spherical grid, and the temporal dimension is discretized into 72 timestamps. We collect a total of 200 trajectories by sampling α and β .

B. Ablation Study

To evaluate the effectiveness of the hierarchical structure and the frequency path in StFT, we conduct an ablation study of our model. First, we only keep one layer of StFT while removing the hierarchical structure. Second, we keep the hierarchical structure, and remove the frequency path in each hierarchical layer.

Table 2 shows the L_2 relative errors averaged over all the variables. Note that “mono-scale” denotes a single scale of StFT, spanning from the coarsest to the finest level, and “multi-scale” represents the multiple layer setting in the hierarchical structure without frequency path. We find that “multi + \mathcal{F} ” outperforms both ablation cases. These results demonstrate the effectiveness of the hierarchical composition of StFT blocks and the frequency path in StFT block. With the hierarchical composition, for Plasma MHD, the error drops to 0.0307 from 0.0805, and for shallow-water equations (SWE), the error drops to 0.0625 from 0.0975. We observe that the fine-level layer setting in the SWE achieves the best performance among the single layer results, and the multi-layer settings further decrease the prediction error. The frequency path in StFT also




Figure 10. Spherical shallow-water equations: pointwise error of the temporal evolution of velocity field predicted by all the autoregressive models: StFT-F, StFT, FNO, ViT and U-Net. The prediction error exhibits a temporal growth trend, with our model StFT and StFT-F consistently demonstrate lower residuals over the forecasting time horizon.

plays a crucial role, where in Plasma MHD, the error drops to 0.0307 from 0.0404, and in the SWE dataset, the error drops to 0.0625 from 0.0956.

	MHD	SWE
mono-scale + \mathcal{F}	0.0805/0.105	2.5729/0.101/0.0975
multi-scale	0.0404	0.0956
multi-scale + \mathcal{F}	0.0307	0.0625

Table 2. Ablation study results. We run models with the combinations of StFT blocks and the frequency path. “Mono-scale” denotes a single scale of StFT block from the coarsest to the finest level, and “multi-scale” represents the multiple layer setting in the hierarchical structure. \mathcal{F} stands for the frequency path in StFT blocks.

C. Experimental Details

Training/validation/test data sets. For the plasma MHD data, we split the trajectory of 1221 snapshots into a training set (the first 927 snapshots), a validation set (the middle 50 snapshots), and a test set (the last 244 snapshots). For the Navier-Stokes dataset, we have a total of 100 trajectories (101 snapshots for each trajectory), and split them into 80 trajectories for training, 10 for validation, and 10 for testing. For the shallow-water dataset, we have a total of 200 trajectories (72 snapshots for each trajectory), and split them into 160 trajectories for training, 20 for validation, and 20 for testing.

Generative residual correction block. We follow a two-step training protocol in training StFT -F: first, we train StFT thoroughly, and then we train the generative residual correction block. We employ a rectified flow to learn distributions of the residuals given the prediction of the last layer of StFT and the past snapshots. We implement a similar structure to the Diffusion Transformer (DiT) as the backbone model (Peebles & Xie, 2023). In each DiT block, we apply adaptive layer normalization before a self-attention layer and an MLP layer. We use adaLN-Zero for time conditioning. For the historical snapshots \tilde{u}_t and the prediction $\mathcal{F}_{\theta_d}(\tilde{u}_t)$ from the last layer of StFT, these conditions are incorporated as extra input tokens.

Hyperparameters. For StFT on the plasma dynamics, 3D FFT is used to encode the spatio-temporal inputs in the frequency path. We use the patch size of 128 for the first StFT block and 64 for the second StFT block. The overlapping size is set to 1. The hidden dimension is set to 128. The depth for each StFT block is set to 6. We keep the lowest 8 frequencies for each spatial dimension. For the rectified flow block, the depth is set to 8, and the hidden dimension is set to 128. For the Navier-Stokes equation, StFT uses a patch size of 25 for the coarse block (the first block), 13 for the middle block, and 8 for the last block. The overlapping size is set to 0, and the frequency path is not used. For each block in the hierarchical structure, the depth is set to 8, and we use a hidden dimension of 512. In the rectified flow block, we use a depth of 4 and set the hidden dimension to 128. For the shallow-water equation, three levels of StFT blocks are employed, and their patch sizes are set to 128, 64, and 32 respectively. For each block, the depth is set to 6, and the hidden dimension is set to 512. We use the 2D FFT to encode the spatio-temporal inputs, and the lowest 8 frequencies are kept for each spatial dimension. The overlapping size is set to 1. For the rectified flow model, we use a depth of 8 and a hidden dimension of 128.

Baseline. We run enough epochs to make sure all methods are thoroughly trained. For AR-FNO, the number of modes are selected through a search in [16, 20, 24], the number of layers are searched in [4, 5], and the hidden dimension is set to 256. For AR-ViT, we vary the hidden dimension from [256, 512], and the patch size from [16, 32, 64]. For AR-UNet, the hidden dimension of bottleneck embeddings are searched in [64, 128, 256, 512].

Training. We run enough epochs to make sure all methods are thoroughly trained. For the generative residual correction block, we run 200 epochs after StFT is trained in the first step. We notice that when training the rectified flow on the residuals, it is crucial to normalize them, and we normalize the residuals by dividing them by their standard deviations.

D. More Visualization Results

Figure 8 illustrates the ground truth and predicted temporal evolution of normalized perturbed parallel vector potential δA_{\parallel} and perturbed electron density δn_e in plasam MHD using AR-StFT, AR-StFT -F, AR-FNO, AR-ViT and AR-UNet methods. StFT and StFT-F perform stable across the forecasting time horizon. Figure 9 and 10 show the pointwise error of all the models compared to the ground truth data in the Navier-Stokes and shallow-water equations, where StFT and StFT-F demonstrate lower residuals compared to other baseline models.