

000 001 002 003 004 005 P3D: HIGHLY SCALABLE 3D NEURAL SURROGATES 006 FOR PHYSICS SIMULATIONS WITH GLOBAL CONTEXT 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

ABSTRACT

We present a scalable framework for learning deterministic and probabilistic neural surrogates for high-resolution 3D physics simulations. We introduce *P3D*, a hybrid CNN-Transformer backbone architecture targeted for 3D physics simulations, which significantly outperforms existing architectures in terms of speed and accuracy. Our proposed network can be pretrained on small patches of the simulation domain, which can be fused to obtain a global solution, optionally guided via a scalable sequence-to-sequence model to include long-range dependencies. This setup allows for training large-scale models with reduced memory and compute requirements for high-resolution datasets. We evaluate our backbone architecture against a large set of baseline methods with the objective to simultaneously learn 14 different types of **PDE dynamics** in 3D. We demonstrate how to scale our model to high-resolution isotropic turbulence with spatial resolutions of up to 512^3 . Finally, we show the versatility of our architecture by training it as a diffusion model to produce probabilistic samples of highly turbulent 3D channel flows across varying Reynolds numbers, accurately capturing the underlying flow statistics.

1 INTRODUCTION

Training neural networks on high-resolution data substantially increases the required GPU memory and compute costs. Scaling models and their input dimensions typically requires substantial engineering effort, posing a major barrier to the widespread and cost-effective adoption of machine learning across application domains. Scientific machine learning and engineering are especially affected due to the multi-scale nature of relevant phenomena whose modeling often requires specialized and highly computationally demanding numerical solutions. In this paper, we focus on learning surrogate models for simulations focusing on fluid dynamics that have downstream applications in fields such as aerospace (Arranz et al., 2024), climate science (Bodnar et al., 2024), energy systems (Degrave et al., 2022), and biomedical engineering (Morris et al., 2016). Machine learning models inherently compete with existing solvers, which are often employed to create the reference targets for learning. As such, they need to either significantly outperform the corresponding solvers while maintaining an acceptable level of accuracy (Kochkov et al., 2021; Pestourie et al., 2023), or yield solutions where traditional solvers fall short, for example working with noisy (Franz et al., 2023) or only partial input data (Shu et al., 2023), or by providing uncertainty estimates (Jacobsen et al., 2023). A large fraction of papers in this area address learning problems for either low-dimensional or comparatively smooth data in 2D. In this paper, we deliberately focus on high-resolution 3D phenomena, covering a wide variety of different types of PDEs.

We propose *P3D*: a hybrid CNN-Transformer backbone built on PDE-Transformer (Holzschuh et al., 2025) that combines fast processing of local features via convolutions and windowed attention mechanisms for learning generalizable token representations. P3D makes important extensions for scaling to very high resolutions in 3D: *crop-based pre-training* combined with a designated module for processing *global dependencies*. This design achieves superior scalability and accuracy compared to existing 3D baselines. We evaluate the P3D backbone architecture through an extensive comparison with existing architectures for modeling physics simulations in 3D, training on 14 different types of dynamics simultaneously. We then employ P3D as a surrogate model for isotropic turbulence at resolution 512^3 and demonstrate how the P3D architecture pretrained on crops of 128^3 can be scaled to the entire domain, while achieving high accuracy with temporally stable autoregressive rollouts.

054 A significant difficulty when modeling large systems is how to aggregate and distribute information
 055 globally across the network. We propose to link the bottleneck layers of P3D with a sequence-to-
 056 sequence model, called *context model*, for an efficient global processing and information aggregation
 057 utilizing highly optimized self-attention mechanisms in LLM layers, which is combined with different
 058 scalable and memory-aware training and inference strategies. Additionally, we propose a direct
 059 mechanism to let aggregated global information flow back to the decoder modules via adaptive
 060 instance normalization layers. In our final experiment, we train P3D as a diffusion model to learn
 061 the distribution of velocity and pressure fields of a turbulent channel flow on a non-equidistant grid.
 062 This setup requires access to global information like the relative position to the walls, and cannot be
 063 addressed by learning local representations alone. We verify that velocity profiles from the generated
 064 samples of P3D closely match the ground truth, demonstrating that high distributional accuracy
 065 can be obtained even when the solution fields are constructed from smaller regions, which are only
 066 coordinated through the information flow from the context model. To summarize our contributions:
 067

- 068 • We introduce *P3D*, a hybrid CNN-Transformer architecture for autoregressive prediction of
 069 *high-resolution physics simulations in 3D*, combining convolutions for fast learning of local
 070 features and windowed self-attention for deep representation learning.
- 071 • We demonstrate the efficiency and versatility of P3D in three experiments: (1) jointly
 072 learning multiple types of simulations (2) scaling P3D pretrained on crops of 128^3 to a
 073 high-resolution simulation of isotropic turbulence at 512^3 , and, (3) generating probabilistic
 074 samples from P3D trained via flow matching for the velocity and pressure fields of a turbulent
 075 channel flow, closely matching the ground truth flow statistics.
- 076 • We propose mechanisms for efficient global information processing, including linking
 077 bottleneck layers with a sequence-to-sequence *global context model* and injecting global
 078 information into adaptive instance normalization layers via region tokens.
- 079 • We evaluate *different setups for finetuning*, which enable a more fine-grained control of
 080 precomputation and gradient backpropagation through encoder and decoder blocks to reduce
 081 VRAM and compute requirements.

082 2 RELATED WORK

083 **PDEs and machine learning** Machine learning has sparked much resonance in accelerating and
 084 improving numerical PDE solvers as well as fully replacing them. Approaches that are combined
 085 with existing PDE solvers can replace components of the solver (Bar-Sinai et al., 2019), learn closure
 086 models (Duraisamy et al., 2019; Sirignano & MacArt, 2023) or learn corrections to a fast solver
 087 on a coarse grid (Um et al., 2021; Kochkov et al., 2021; Dresdner et al., 2023). Other directions
 088 target problems that are more difficult to address with numerical solvers, such as inverse problems
 089 (Raissi et al., 2019; Bruna et al., 2022; Holzschuh et al., 2023), or uncertainty quantification (Xiao &
 090 Cinnella, 2019; Liu & Thurey, 2024). Leveraging diffusion models for autoregressive prediction
 091 and inverse problems for PDEs has been explored by several works in the past (Lippe et al., 2023;
 092 Kohl et al., 2024; Shu et al., 2023; Shehata et al., 2025), albeit limited to data in 2D.

093 **Pretraining and 3D neural surrogates** Subramanian et al. (2023); Yang et al. (2023); McCabe
 094 et al. (2023) pretrain foundation models for PDEs on multiple PDE dynamics simultaneously. Such
 095 models can be used either for zero-shot-predictions or for finetuning when adapting to new dynamics,
 096 which allows for improved results with fewer training samples. Previous works have demonstrated
 097 learning PDEs in 3D, e.g., learning large-eddy simulations (Peng et al., 2023; Li et al., 2023a; 2024a;
 098 Jiang et al., 2025), and elastic wave propagation (Lehmann et al., 2024). Smoke buoyancy in 3D was
 099 targeted as a test case by Li et al. (2023b) Notably, most previous work targets resolutions of up to
 100 64^3 , an exception being probabilistic experiments at a resolution of 128^3 (Molinaro et al., 2024), and
 101 surrogate training with up to $128 \times 128 \times 256$ (Ohana et al., 2024). This motivates our contributions
 102 for scalable architectures, as surrogate models for truly high-resolution 3D physics simulations are of
 103 paramount interest in different scientific areas.

104 **Scalable transformer architectures** Transformers have become one of the dominant backbone
 105 architectures in deep learning due to their high computational efficiency and their ability to model
 106 long-range causal relationships (Vaswani et al., 2017; Devlin et al., 2018). Transformers have also
 107 become a popular competitor to CNNs in vision and understanding tasks (Dosovitskiy et al., 2020;
 108 Rodrigo et al., 2024), scale to large images (Gupta et al., 2024), and have recently been adopted
 109 for learning surrogate models for physics simulations (McCabe et al., 2023; Wu et al., 2024; Alkin

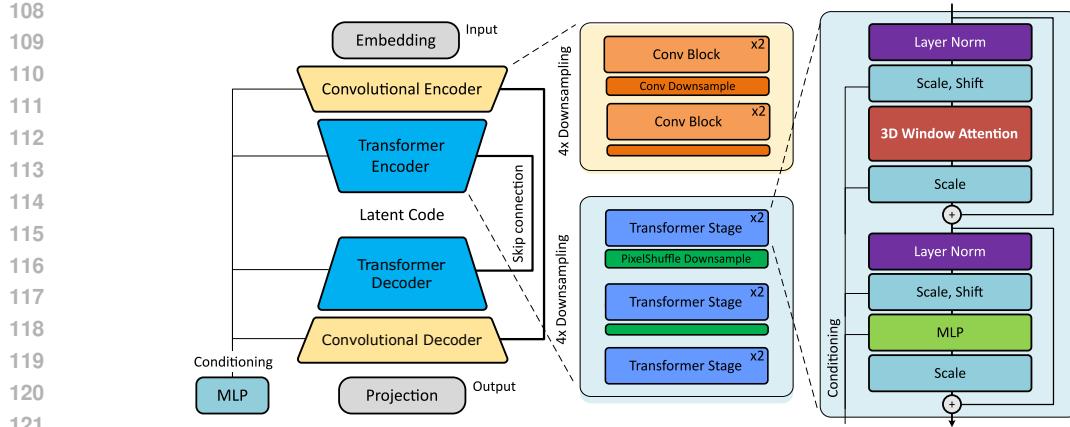


Figure 1: Overview of P3D. Convolutional blocks for local feature processing are combined with transformers for deep representation learning, yielding a U-shaped multi-scale architecture. The transformer backbone combines windowed attention and conditioning via adaptive instance normalization, which are modified and optimized for 3D.

et al., 2024). A major computational difficulty is the quadratic complexity of the global self-attention mechanism. Liu et al. (2021) address this limitation by restricting the computation of the attention operation to a local window and Ho et al. (2019) by computing the attention only across the data axes.

3 METHOD

Problem formulation Let Ω denote a spatial domain with n physical quantities $u(x, t) : \Omega \times [0, T] \rightarrow \mathbb{R}^n$ that are discretized in time and space and described by the temporal sequence $[\mathbf{u}_0, \mathbf{u}_{\Delta t}, \dots, \mathbf{u}_T]$. We consider all additional information about the sequence such as the type of PDE or hyperparameters of the simulator to be encoded in an m -dimensional conditioning vector $\mathbf{c} \in \mathbb{R}^m$. We assume the availability of many such sequences as training data, representing the temporal evolution of different types of PDEs with varying initial conditions or coefficients. We denote our proposed network architecture by \mathcal{M}_Θ with weights Θ .

We address two main tasks in this paper. The first is *autoregressive prediction*: For a given sequence of P preceding states $[\mathbf{u}_{t-P\Delta t}, \dots, \mathbf{u}_{t-\Delta t}]$, denoted by \mathbf{u}_{in} , our target is to predict the next state $\mathbf{u}_t := \mathbf{u}_{\text{out}}$. The second task is to train a *probabilistic sampler* to draw samples from a distribution of states representing solutions for a PDE as specified by the parameter vector \mathbf{c} . In this case $\mathbf{u}_{\text{in}} = \emptyset$.

3.1 BACKBONE ARCHITECTURE: P3D TRANSFORMER

The key components of our proposed hybrid CNN-Transformer architecture P3D are the hierarchical U-shape structure with the hybrid encoder and decoder-pair based on convolutional and transformer blocks. A visual overview is given in Figure 1. In the following, we highlight the main components of the architecture and explain how they support its central goal to enable the efficient training and inference for high-resolution 3D simulations.

Hybrid encoder/decoder We utilize convolutional en- and decoders with skip connections. Fully transformer-based architectures that work in the pixel space for 2D data and images like ViTs rely on a patchification operation to transform patches of size $p \times p$ into tokens. A corresponding approach in 3D would transform patches of size p^3 into a single token, significantly increasing the amount of information encoded in each token. To balance both the number of tokens for the transformer as well as the information density of each token, we learn local features via the convolutional encoder to obtain an optimized compressed representation.

Attention and positional encoding The self-attention operation used by transformers has quadratic complexity in the number of tokens. For 3D data, this becomes a major computational issue as the number of tokens grows cubically with the spatial discretization, leading to computational blow-up as the domain size increases. The central building block of our transformer encoder is the windowed multi-head self-attention (Liu et al., 2021), which only computes self-attention within a local

162 windows. For computing the attention scores between tokens, we use the log-spaced relative positions
 163 of tokens inside the same window.
 164

165 The architecture of the P3D combines Swin transformers (Liu et al., 2021), diffusion transformers
 166 (Peebles & Xie, 2023, DiT) and UNets (Ronneberger et al., 2015) into a 3D variant, which can be
 167 seen as an extension of PDE-Transformer (Holzschuh et al., 2025) in 3D. The main alteration is
 168 the replacement of the patchification with large convolutional en- and decoders. We also removed
 169 the shifting of windows during the computation of self-attention. The convolutional en- and de-
 170 coder follow the design of modern UNet blocks, using adaptive instance normalization and group
 171 normalization. We give a detailed description of the architecture in Section A.4.
 172

3.2 CONTEXT NETWORK

173 The P3D backbone architecture does not include any absolute positional embeddings as well as no
 174 operations aggregating and distributing information globally. Thus, it has to rely on learning local
 175 features and dynamics within its perceptual field. This promotes translation-equivariance, which is
 176 an important inductive bias for surrogate modeling of PDEs. Our design choice to not include other
 177 physical inductive biases is intentional to keep the P3D backbone architecture flexible. For learning
 178 large-scale simulations, we pretrain P3D on smaller crops of the simulation domain and then scale the
 179 pretrained network to larger inputs. However, global information and long-range dependencies often
 180 play a crucial role to obtain correct solutions, which is also a major limitation of PDE-Transformer.
 181 To address this shortcoming, we introduce the context network.

182 **Token embeddings** The bottleneck layer of the
 183 transformer encoder consist of tokens, which are em-
 184 bedded into latent tokens via a linear layer. The P3D
 185 encoder compresses a crop of size 16^3 into a single
 186 latent token. Then, a frequency-based positional em-
 187 bedding vector is added to each latent token, similar
 188 to Dosovitskiy et al. (2020), which constitutes an
 189 absolute positional encoding of latent tokens. In ad-
 190 dition, we partition the domain into *regions* and we
 191 match the size of regions with the size of the domain
 192 crops P3D was pretrained on. For each region, we
 193 include a corresponding so-called *region token* in the
 194 sequence of latent tokens, similar to the classification
 195 token in ViTs. Each region token is initialized via a
 196 learnable embedding layer and we add a frequency-
 197 based positional embedding vector. The purpose of
 198 region tokens is to serve as a more direct feedback
 199 mechanism to the decoder, which we describe in the
 200 next paragraph. Our implementation processes the
 201 token sequence via $n = 6$ transformer layers. Figure
 202 2 provides an overview of this setup. In principle,
 203 any efficient sequence model can be used. After the
 204 sequence of region and latent tokens is processed, the
 205 latent tokens are added to the input of the decoder
 206 via a skip connection. See Section A.5 for a full
 207 description.
 208

209 **Region tokens** The region tokens are retained and
 210 used as a more direct mechanism to let information
 211 flow through the decoder network. Region tokens
 212 are initialized as a learnable embedding vector with
 213 frequency-based positional encoding and are pro-
 214 cessed in the sequence model together with the latent
 215 tokens. Each region token corresponds to a crop that
 216 is processed independently of other crops by the
 217 encoder and decoder blocks. We use scale and shift
 218 operations to condition the decoder block of each re-
 219 gion on the region token. Within each decoder block,

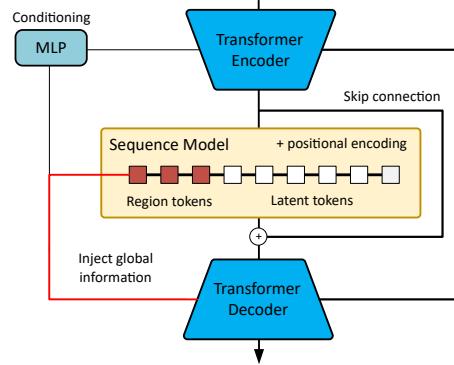


Figure 2: Global context via a sequence model. The bottleneck layers are connected to the sequence model, which embeds the bottleneck representation as latent tokens. Region tokens are used to inject global information directly into the decoder.

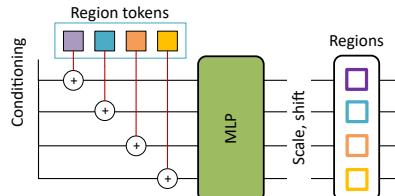


Figure 3: Conditioning via the region tokens. The input domain is partitioned into regions/crops, each of which has a corresponding messenger token. The regions are modulated individually via learnable scale and shifts based on the region tokens.

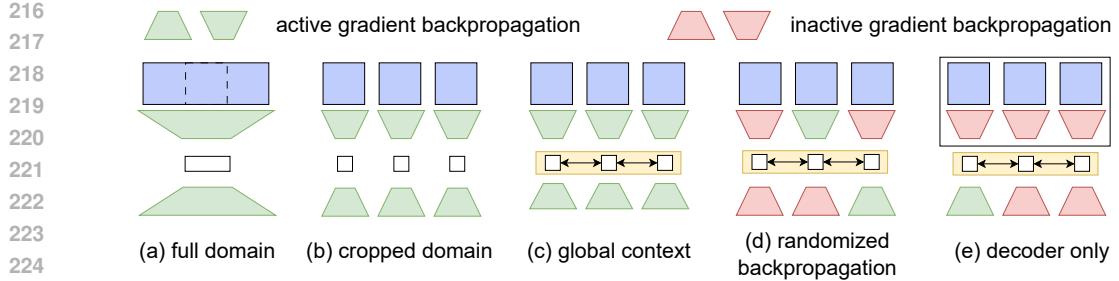


Figure 4: Different training and inference setups. (a) shows training on the full domain and (b) on domain crops. (c) includes the context network for global information processing, which can also be trained by randomly disabling gradient backpropagation for a percentage of the encoders and decoders, see (d). In (e) the latent codes from a pretrained encoder can be precomputed and only the context network and decoder are trained.

for each adaptive instance normalization layer with scale or shift operations, we transform the region token via a linear layer to a region embedding vector which is added to the embedding vector of the conditioning c . Each region gets modulated differently based on the region token. This is visualized in Figure 3.

3.3 SCALING OUTPUT DOMAINS

We consider different setups for training and inference, see Figure 4, which include training on the full domain (a), on crops (b) and different training variants in combination with the context network (c) to (e). While it is preferable to train on small crops due to compute requirements, for inference, we generally want to process the full domain. We consider two strategies: (1) we scale to the full domain via the translation equivariance of the P3D architecture, i.e., we combine the domain crops and process them as a single input, and, (2) we encode and decode each crop of the full domain independently and combine the network outputs. We tag a model that is trained on crops of resolution x^3 and which internally runs inference on resolution y^3 by $\langle x|y \rangle$. For example, a network trained on crops of size 64^3 that is scaled via strategy (1) to resolution 128^3 is tagged $\langle 64|128 \rangle$, while the same network scaled via strategy (2) is tagged $\langle 64|64 \rangle$. For strategy (2), if we use the context network for communication between the latent codes, we use the tag $\langle Xx|Xy \rangle$.

4 EXPERIMENTS

We evaluate P3D as well as different scaling and finetuning setups in three experiments. P3D has 3 different configurations: S , B and L that determine the embedding dimension d (32, 64 and 128 respectively) of the first layer. The corresponding models have 11M to 180M parameters. We denote the configurations with P3D- S for our model with the S config, changed accordingly for B and L .

4.1 JOINTLY LEARNING MULTIPLE PDEs

Our dataset for this task comprises 14 different types of [PDE dynamics](#) including Burger's equation, Kuramoto-Shivashinsky, Gray-Scott, Swift-Hohenberg and many others. The dataset is based on APEBench (Koehler et al., 2024), and a full description of each PDE with visualizations can be found in Section B.2 in the appendix. For all PDEs except Gray-Scott, we consider 60 different simulations with varying initial conditions and PDE-specific parameters such as viscosity, domain extent or diffusivity. For Gray-Scott, we include 10 simulations for each of its hyperparameters. Each spatially periodic simulation contains 20 snapshots discretized at resolution 320^3 for Gray-Scott and 384^3 for all other PDEs. We evaluate and benchmark models on random crops of the simulation domain of size 128^3 , 64^3 and 32^3 using supervised training, see Section A.1. Even with full information about simulation hyperparameters and the type of PDE, the behavior is not fully deterministic as quantities beyond the cropped regions influence the solution inside it. Simulations have different numbers of channels and we zero-pad data with fewer channels than the number of maximum channels $N_C = 3$.

Training on cropped data Cropped data has an implicit time-dependent boundary condition, which is not known by the model. This can be seen as an extension and more difficult variant of the

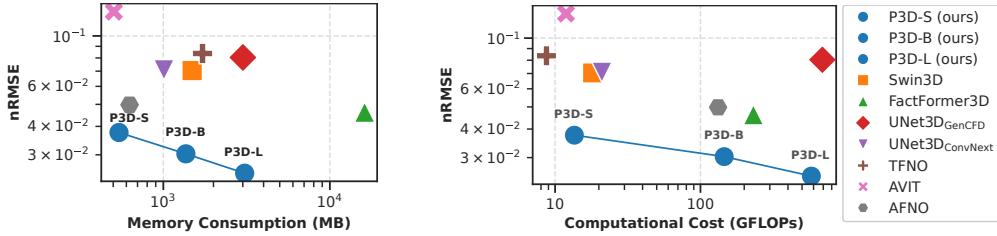


Figure 5: Comparison of model accuracy vs. (left) memory usage during backpropagation and (right) computational costs for inference for jointly learning different types of PDEs with crops of size 64^3 for P3D and baselines.

multi-physics training (Subramanian et al., 2023; Yang et al., 2023; McCabe et al., 2023), where in addition to not knowing the PDE or simulation hyperparameters, the model has to estimate the boundary conditions in a data-driven manner. Mathematically, the model \mathcal{M}_Θ is trained to regress

$$\operatorname{argmin}_{\Theta} \mathbb{E}_{(s_t, s_{t+\Delta t}, \mathbf{c}) \in \mathcal{D}_{\text{train}}} [\mathbb{E}_{(s_t^{\text{crop}}, s_{t+\Delta t}^{\text{crop}})} [\|\mathcal{M}_\Theta(s_t^{\text{crop}}, \mathbf{c}) - s_{t+\Delta t}^{\text{crop}}\|_2^2]], \quad (1)$$

where we sample $(s_t, s_{t+\Delta t}, \mathbf{c}) \in \mathcal{D}_{\text{train}}$ from the training dataset and apply a random cropping to obtain $(s_t^{\text{crop}}, s_{t+\Delta t}^{\text{crop}})$. The input \mathbf{u}_{in} corresponds to s_t^{crop} and \mathbf{u}_{out} to $s_{t+\Delta t}^{\text{crop}}$. The mapping $s_t^{\text{crop}} \mapsto s_{t+\Delta t}^{\text{crop}}$ is not deterministic since the boundary conditions are not prescribed. The model \mathcal{M}_Θ has to learn a prediction that minimizes the prediction error w.r.t. all possible simulation states that are outside the cropped domain, i.e., the optimal prediction s^* for $s_{t+\Delta t}^{\text{crop}}$ minimizes

$$s^* = \operatorname{argmin}_s \mathbb{E}_{(\hat{s}_t, \hat{s}_{t+\Delta t}) \sim \mathcal{D}_{\text{train}}} [\|\text{crop}(\hat{s}_{t+\Delta t}) - s\|_2^2 \mid \text{crop}(\hat{s}_t) = s_t^{\text{crop}}], \quad (2)$$

where $\text{crop}(\cdot)$ is the crop operation used for s_t^{crop} . The performance depends on how well the model is able to extrapolate the dynamics outside the cropped input for a short prediction horizon.

nRMSE evaluation We evaluate the nRMSE, see Section A.3, and consider a wide range of sota baseline architectures. Specifically, we include Swin3D, our own implementation of the SwinV2 architecture (Liu et al., 2021) extended to 3D, AViT, an axial vision transformer (McCabe et al., 2023), Adaptive FNOs (Guibas et al., 2021, AFNO), Tucker-Factorized FNOs (Kossaifi et al., 2023, TFNO), and FactFormer (Li et al., 2024b). Additionally, we consider two different convolutional UNet architectures, UNet_{ConvNeXt} as used in Ohana et al. (2024) and UNet_{GenCFD} (Molinaro et al., 2024). We train all models for 1000 epochs on four H100 GPUs. We use a fixed learning rate of $2.0 \cdot 10^{-4}$ for all models with the AdamW optimizer with weight decay 10^{-15} and batch size 256 in bf16-mixed precision. See Table 1 for the results and Tables B3 to B5 for detailed evaluations of each PDE type. As we train on bigger crop sizes, the observed simulation domains becomes larger, increasing the amount of information that becomes available to the network. At the same time, the relative volume of the boundary becomes smaller, thus decreasing its relative weight and causing an according decrease in the nRMSE. P3D performs best across all crop sizes. Performance significantly improves when scaling the model size from S to L .

Table 1: Comparison of the nRMSE ($\times 10^{-2}$) on the test dataset (averaged over all PDEs).

Method	Crop size		
	32^3	64^3	128^3
epochs = 1000			
TFNO	8.46	8.37	-
FactFormer	6.24	4.62	-
UNet _{ConvNeXt}	8.59	7.09	-
UNet _{GenCFD}	7.61	8.04	8.27
AViT	20.9	25.0	15.1
Swin3D	7.92	7.04	5.03
AFNO	9.95	4.98	4.79
P3D-S	6.27	3.76	3.33
P3D-B	4.69	3.03	2.52
P3D-L	4.13	2.49	2.08

Memory and compute For scaling an architecture to high-resolution 3D simulations, the memory requirements as well as the inference speed are essential. Transformer architectures have been shown to achieve improved accuracy as the number of parameters and floating point operations increase, therefore comparing different architectures needs to take both factors into account. In Figure 5, we

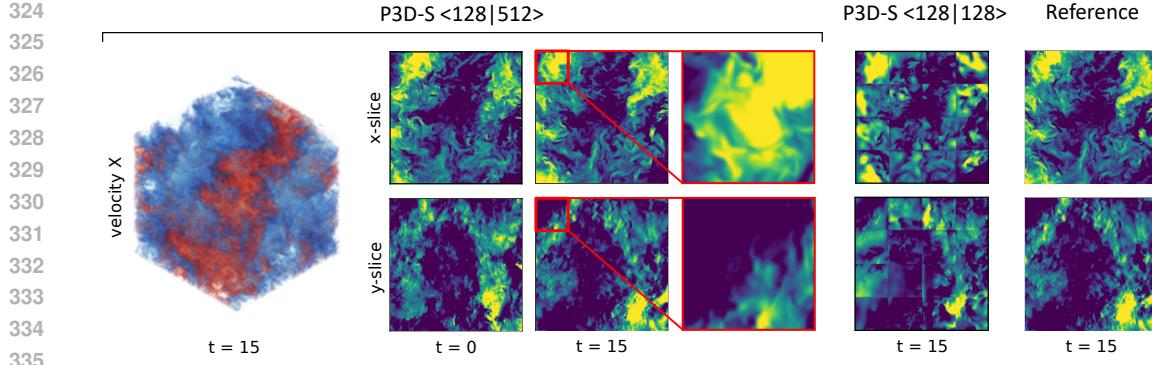


Figure 6: Forcible isotropic turbulence. Prediction on the test set at resolution 512^3 with an autoregressive rollout of 15 steps. The model is pretrained on patches of size 128^3 , without finetuning on 512^3 . P3D-S <128|512> successfully scales to the full domain, whereas for P3D-S <128|128> discontinuities at the borders of crop regions become apparent.

compare the nRMSE at patch size 64^3 against (1) the computational cost for inference measured in GFLOPs, and (2) the VRAM consumption in MB for a backward pass with batch size 1. The P3D networks provide the best tradeoff between accuracy and computational cost/memory requirements. See Table A2 for additional information.

Ablation on network design We empirically verify our network design and consider: (1) P3D without the transformer backbone, which purely relies on the convolutional encoder and decoder, denoted *P3D-conv*, (2) a patch-based P3D with linear tokenizer for patch size $p = 4$ and no convolutional encoder and decoder, denoted *P3D-patch*. This baseline (*) extends the mixed channel version of *PDE-Transformer to 3D*. (3) disabling PDE type conditioning by setting $\mathbf{c} = \mathbf{0}$, denoted *P3D-no-c*, (4) P3D with window shifting enabled, *P3D-shift*. Additionally, we evaluate different values of the window size $w = 2$, $w = 4$ (default) and $w = 8$. P3D achieves a relative improvement of 16.0% (32^3), 26.7% (64^3) and 37.5% (128^3) over 3D PDE-Transformer(*), see Table 2. Shifting and larger window sizes give no significant improvements here.

4.2 ISOTROPIC TURBULENCE

The goal of the this experiment is to scale P3D to a high-resolution simulation involving complex dynamics. For this, we consider forced isotopic turbulence simulated via direct numerical simulation (DNS) at resolution 1024^3 provided by the John Hopkins Turbulence Database (Perlman et al., 2007).

The dataset is cropped from the original resolution to 512^3 with a total of 500 snapshots, which are saved after reaching a statistical stationary state and comprises channels for the velocity X/Y/Z and pressure. The data is split into test and training sets, where the first 420 snapshots are used for training and the last 80 snapshots for testing.

RMSE and enstrophy error for crops 128^3
 We first evaluate the performance of P3D compared to baselines from the previous experiment trained on domain crops of size 128^3 . We consider an evaluation of autoregressive rollouts

Table 2: Validation MSE ($\times 10^{-3}$, time-weighted EMA with $\lambda = 0.99$).

Method	Crop size		
	32^3	64^3	128^3
epochs = 400			
P3D-S-conv	8.33	5.40	3.25
P3D-S-patch*	6.48	3.78	2.16
P3D-S-no-c	5.69	2.94	1.41
P3D-S-shift	5.41	2.84	1.37
P3D-S	5.44	2.77	1.35
P3D-S $w = 2$	5.68	2.96	1.49
P3D-S $w = 4$	5.44	2.77	1.35
P3D-S $w = 8$	5.44	2.90	1.32

Table 3: RMSE ($\times 10^{-2}$) and the enstrophy spectrum error ($\times 10^2$) at different autoregressive rollout steps on the test set with crop size 128^3 .

Method	RMSE		Enstrophy	
	1	15	1	15
UNet _{GenCFD}	5.48	67.7	4.2	140
Swin3D	3.22	24.3	4.0	156
AViT	9.45	37.7	26.7	112
AFNO	3.69	29.8	7.6	190
P3D-S	2.81	28.2	2.15	31.9
P3D-B	2.04	31.5	0.72	19.2

378 from 1 to 15 steps. We evaluate the RMSE and a spectral error based on the enstrophy spectrum,
 379 which we compute based on the vorticity that is derived from the velocity fields of the data, see
 380 Section C.3. We train all models for 4000 epochs using the same setup as in Section 4.1, but reduce
 381 the batch size to 32. P3D performs best and achieves the lowest RMSE for few autoregressive rollout
 382 steps. It consistently has by far the best spectral error, see Table 3 and Table C1 for an extended
 383 evaluation. For many autoregressive rollout steps, the RMSE becomes less informative, as the
 384 prediction starts to deviate from the reference due to the uncertainty from the boundary of the crop
 385 regions. See Figures C6 and C7 for visualizations.

386 **Scaling P3D to 512^3** We scale P3D pre-
 387 trained on crops of size 128^3 to the full do-
 388 main 512^3 . Training P3D on cropped data re-
 389quires less than a day on four A100 GPUs. See
 390 Section C.4 for details. We compare the per-
 391 formance of P3D when scaling the network to
 392 larger domain sizes using the scaling strategies
 393 introduced in Section 3.3. P3D $\langle 128|128 \rangle$ pro-
 394 cesses blocks independently using the original
 395 training resolution. Thus the domain 512^3 is
 396 split into 64 blocks of size 128^3 that are processed independent of each other. P3D $\langle 128|512 \rangle$
 397 processes the full domain, leveraging the translation-equivariance of the architecture. See Table 4
 398 for an evaluation of the RMSE. Increasing the domain size during inference consistently gives
 399 improvements. With increasing domain size, the relative volume of the boundary shrinks, thus
 400 the uncertainty of turbulent motions is reduced and networks are able to provide more accurate
 401 predictions. Similar to previous results, larger networks improve performance as well. Note that
 402 the RMSE is different between Tables 3 and 4, since we do not consider longer rollouts in Table 4
 403 which affects the distribution of samples in the test set. In this experiment, due to the isotropic and
 404 homogeneous nature of the simulation, we achieve highly accurate results without finetuning or
 405 global information.

406 **Long auto-regressive rollouts**

407 To evaluate accuracy over pro-
 408 longed periods, we evaluated ad-
 409 dditional tests comparing the full-
 410 scale an individual crop models
 411 for 50 rollout steps. We per-
 412 formed inference with the scaled-
 413 up model P3D-S $\langle 128|512 \rangle$ on
 414 the full domain, while for eval-
 415 uation, we computed the spectral er-
 416 rror and RMSE on the 128^3 crops
 417 to ensure the metrics are directly
 418 comparable to those in Table 3.

419 P3D-S $\langle 128|512 \rangle$ maintains an excellent enstrophy spectrum error below 10×10^2 for the full
 420 period, whereas P3D-S $\langle 128|128 \rangle$, which processes individual regions, achieves this for only 10 steps.
 421 We note that the 512^3 domain itself is cropped from an even larger 1024^3 resolution, introducing
 422 boundary uncertainties that complicate auto-regressive unrolling, even for the scaled-up model. Not
 423 being able to reliably compare small scale structures, the RMSE values are still slightly lower for
 424 the scaled-up $\langle 128|512 \rangle$ model. This P3D-S version outperforms the mean-field prediction for up to
 425 40 to 50 steps. In particular, the enstrophy comparison highlights the advantages of applying P3D
 426 models pre-trained on 128^3 crops to the full 512^3 domain for auto-regressive rollouts.

427 **4.3 TURBULENT CHANNEL FLOW**

428 For the last experiment we train P3D as a generative model for learning to sample from a turbulent
 429 channel flow simulation with a periodic channel with no-slip boundaries at $\pm y$ and is driven by a
 430 dynamic forcing to prevent a loss of energy. The seemingly simple geometry represents a well-studied
 431 and relevant scenario that is highly challenging as it requires fine spatial and temporal discretizations

Table 4: RMSE ($\times 10^{-2}$) on the test set for P3D trained on crop size 128^3 and evaluated at 512^3 for different scaling strategies.

Method	Scaling strategy		
	$\langle 128 128 \rangle$	$\langle 128 256 \rangle$	$\langle 128 512 \rangle$
P3D-S	1.90	1.68	1.60
P3D-B	1.38	1.24	-

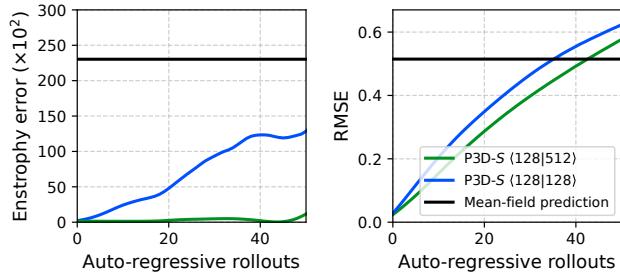


Figure 7: Enstrophy spectrum error ($\times 10^2$) and RMSE over auto-
 regressive rollout steps at resolution 512^3 .

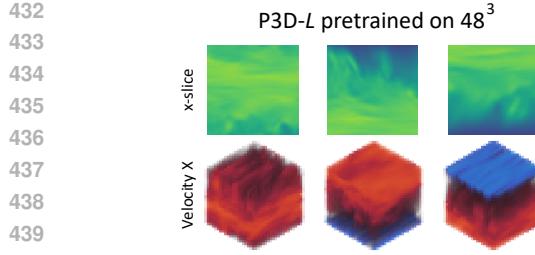


Figure 8: Turbulent channel flow. We pretrain PDE- L on crops of size 48^3 . For P3D- L $<48|192>$ scaled to the full domain $96^2 \times 192$, the relative positions of crops to the wall is critical. Without absolute positional encodings and global information passing, the generated samples are incorrect. P3D- L $<X48|X48>$ finetuned with the context network resolves this issue.

with correspondingly long simulation times to provide converged turbulence statistics (Hoyas & Jiménez, 2008).

Due to the very costly initial transient phase of these simulations, it is especially attractive to phrase the turbulent channel flow (TCF) problem as a probabilistic learning problem, where states from the relevant equilibrium phase can be sampled directly, i.e. without resolving the initial warm-up phase. We generate a dataset comprising 20 simulations with Reynolds numbers Re within the interval $[400, 800]$ spaced equidistantly, see Section D.1. The computational grid comprises $96 \times 96 \times 192$ spatially adaptive cells with a finer discretization near the wall. The data contains channels for the velocity in X/Y/Z direction as well as pressure.

P3D- L pretrained on 48^3 We train P3D- L as a diffusion model following Section A.2 on crops of size 48^3 . Pretraining on crops requires significantly less compute and VRAM and the network converges much faster. For P3D- L with scaling strategy $<48|192>$, the network is not aware of the position of the wall relative to the position of the individual regions, causing the global structure of the generated flow to be incorrect, see Figure 8. For a quantitative statistical evaluation including training on the full domain and baselines, see Table 5.

Ablation on finetuning and context We evaluate finetuning P3D pretrained on the cropped domain 48^3 using the context model, see Table 5. Finetuning corresponds to P3D- L $<X48|X48>$ with different training setups, see Figure 4. *finetune* and *finetune w/o region tokens* follows the setup of Figure 4c, which backpropagates gradients through all crops and network modules. While this requires more VRAM than training on the full domain, only few epochs to achieve good results are necessary. For *finetune, decoder only*, we only backpropagate gradients through the decoder and context network, and also only backpropagate through 10% randomly selected decoder blocks, which corresponds to the setup in Figure 4e. We do not precompute the bottleneck representations from the frozen encoders, but this could be done to further reduce the VRAM and compute requirements. This setup

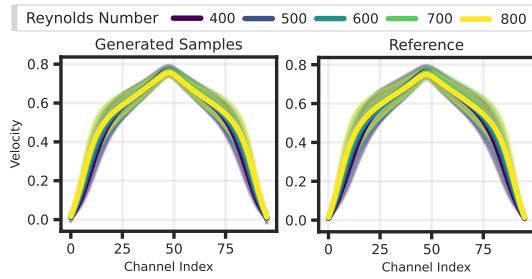


Figure 9: Flow statistics of samples produced by P3D- L conditioned with varying Reynolds number (left) and the time-resolved DNS (right).

Table 5: L2 error for the mean μ ($\times 10^{-5}$) and variance σ^2 ($\times 10^{-5}$) of the velocity profiles along the channel flow direction, see Section D.2 for details. For finetuning, we report the mean and std. dev. over 5 different seeds.

Model	Mean L2	Variance L2	VRAM	epochs
UNet _{GenCFD} full domain	132.38	17.66	17.4 GB	400
AFNO full domain	28.73	1849.3	3.4 GB	400
P3D- L full domain	3.02	13.20	14.9 GB	400
P3D- L $<48 192>$	5862.81	233.77	2.8 GB	2000
finetune w/o region tokens	4541 ± 495	2026 ± 267	15.8 GB	20
finetune	23.6 ± 21.4	40.4 ± 49.4	15.8 GB	20
finetune, decoder only	941 ± 484	1170 ± 392	6.0 GB	20
finetune, decoder only	97.7 ± 102.6	131 ± 149	6.0 GB	100
finetune, decoder only	16.8 ± 5.0	24.1 ± 17.2	6.0 GB	500

486 achieves a significant reduction in VRAM, but requires more training epochs. The generated sample
 487 for P3D-*L* <X48IX48>, see Figure 8, qualitatively shows the correct global structure and content of
 488 individual regions. See also Figure 9 and Figures D6 to D9 for comparisons of the velocity profiles.
 489 While finetuning matches the statistics well, there are still visible discontinuities between generated
 490 crop regions, which leaves further room for improvement.

491 **Benchmarking speedups** We compare the runtime per
 492 sample (s) and the mean L2 error of the different methods,
 493 see Figure 10. We also compare against running the
 494 simulation at half resolution $48^2 \times 96$. The warmup phase of
 495 the DNS takes on average 2 hours and 24 minutes, which
 496 cannot be skipped. After the warumup, we simulate a total
 497 time of $t = 20$. To calculate statistics, we take samples
 498 with $\Delta t = 1$ to avoid high autocorrelation. The reference
 499 simulation at full resolution has a runtime of 454s per sample.
 500 For the generative models, we consider 20 samples for
 501 each Reynolds number and 100 inference steps. Timings
 502 are obtained on a L40 GPU. P3D-*L* and AFNO represent
 503 the Pareto frontier for the generative models regarding the
 504 mean L2 error.

5 LIMITATIONS

505 P3D relies on grids to achieve high computational efficiency. Extending P3D to arbitrary geometries
 506 is not trivial, and requires careful modification of each architectural layer. Also, training P3D on
 507 datasets featuring diverse data modalities, resolutions, and dimensionalities remains an unaddressed
 508 task. The finetuning process using the context network decodes each region independently, which
 509 can cause discontinuities at the boundary of each region. Further strategies to prevent and mitigate
 510 discontinuities are an interesting direction for future work.

6 CONCLUSION

511 We have presented *P3D*: an efficient hybrid CNN-Transformer architecture for learning surrogates
 512 for high-resolution 3D physics simulations. We demonstrated the strong advantages of P3D over a
 513 comprehensive list of baselines for simultaneously learning different types of **PDE dynamics**, showing
 514 improved accuracy and stable training while at the same time being faster and more memory efficient
 515 than the strongest competitors. We scaled P3D to a high-resolution isotropic turbulence simulation
 516 by pretraining on smaller crops from the domain, and demonstrated its capabilities as a probabilistic
 517 generative model. The P3D model accurately predicts distributions of high-resolution velocity and
 518 pressure fields for a turbulent channel flow with varying Reynolds numbers, demonstrating how to
 519 include global information and coordinated pretrained networks via a global context model. Our
 520 architecture **enables the scaling of neural surrogate** models to very high resolutions, unlocking their
 521 potential to deliver real-world impact across scientific domains.

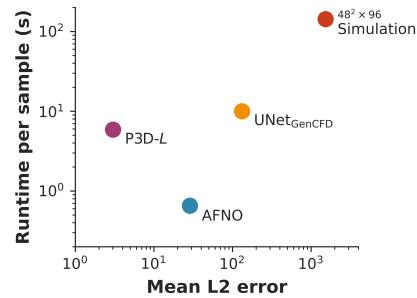
522 REFERENCES

523 Benedikt Alkin, Andreas Fürst, Simon Schmid, Lukas Gruber, Markus Holzleitner, and Johannes
 524 Brandstetter. Universal physics transformers: A framework for efficiently scaling neural operators.
 525 *Advances in Neural Information Processing Systems*, 37:25152–25194, 2024.

526 Gonzalo Arranz, Yuenong Ling, Sam Costa, Konrad Goc, and Adrian Lozano-Duran. Building-block
 527 flow model for computational fluids. *arXiv:2403.09000*, 2024.

528 Yohai Bar-Sinai, Stephan Hoyer, Jason Hickey, and Michael P. Brenner. Learning data driven
 529 discretizations for partial differential equations. *Proceedings of the National Academy of Sciences*,
 530 116(31):15344–15349, 2019. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1814058116.

531 Cristian Bodnar, Wessel P Bruinsma, Ana Lucic, Megan Stanley, Johannes Brandstetter, Patrick
 532 Garvan, Maik Riechert, Jonathan Weyn, Haiyu Dong, Anna Vaughan, et al. Aurora: A foundation
 533 model of the atmosphere. *arXiv preprint arXiv:2405.13063*, 2024.



534 Figure 10: Runtime per sample (s) vs.
 535 Mean L2 error, see Table 5, w.r.t. to
 536 reference simulation at $96^2 \times 192$.

540 Joan Bruna, Benjamin Peherstorfer, and Eric Vanden-Eijnden. Neural galerkin scheme with active
 541 learning for high-dimensional evolution equations, 2022.

542 Yifan Chen, Mark Goldstein, Mengjian Hua, Michael S Albergo, Nicholas M Boffi, and Eric
 543 Vanden-Eijnden. Probabilistic forecasting with stochastic interpolants and föllmer processes. In
 544 *Proceedings of the 41st International Conference on Machine Learning*, pp. 6728–6756, 2024.

545 Jonas Degrave, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco Carpanese,
 546 Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de Las Casas, et al. Magnetic control of
 547 tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897):414–419, 2022.

548 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
 549 bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

550 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
 551 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
 552 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint
 553 arXiv:2010.11929*, 2020.

554 Gideon Dresdner, Dmitrii Kochkov, Peter Christian Norgaard, Leonardo Zepeda-Nunez, Jamie
 555 Smith, Michael Brenner, and Stephan Hoyer. Learning to correct spectral methods for simulating
 556 turbulent flows. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL
 557 <https://openreview.net/forum?id=wNBARGxoJn>.

558 Karthik Duraisamy, Gianluca Iaccarino, and Heng Xiao. Turbulence modeling in
 559 the age of data. *Annual Review of Fluid Mechanics*, 51(1):357–377, jan 2019.
 560 doi: 10.1146/annurev-fluid-010518-040547. URL <https://doi.org/10.1146/annurev-fluid-010518-040547>.

561 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
 562 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English,
 563 and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis.
 564 In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria,
 565 July 21-27, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=FPnUhsQJ5B>.

566 Erik Franz, Barbara Solenthaler, and Nils Thuerey. Learning to estimate single-view volumetric flow
 567 motions without 3d supervision. In *International Conference on Learning Representations*, 2023.

568 John Guibas, Morteza Mardani, Zongyi Li, Andrew Tao, Anima Anandkumar, and Bryan Catanzaro.
 569 Adaptive fourier neural operators: Efficient token mixers for transformers. *CoRR*, abs/2111.13587,
 570 2021. URL <https://arxiv.org/abs/2111.13587>.

571 Ritwik Gupta, Shufan Li, Tyler Zhu, Jitendra Malik, Trevor Darrell, and Karttikeya Mangalam. xt:
 572 Nested tokenization for larger context in large images. In *Forty-first International Conference on
 573 Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. URL
 574 <https://openreview.net/forum?id=wDDprThYeT>.

575 Insu Han, Rajesh Jayaram, Amin Karbasi, Vahab Mirrokni, David P Woodruff, and Amir Zandieh.
 576 Hyperattention: Long-context attention in near-linear time. *arXiv:2310.05869*, 2023.

577 Jonathan Ho, Nal Kalchbrenner, Dirk Weissenborn, and Tim Salimans. Axial attention in multi-
 578 dimensional transformers, 2019.

579 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models.
 580 In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances
 581 in Neural Information Processing Systems*, volume 33, pp. 6840–6851. Curran Asso-
 582 ciates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf>.

583 Benjamin Holzschuh, Simona Vegetti, and Nils Thuerey. Solving inverse physics problems with score
 584 matching. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey
 585 Levine (eds.), *Advances in Neural Information Processing Systems 36: Annual Conference on
 586 Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December
 587 10 - 16, 2023*, 2023.

594 Benjamin Holzschuh, Qiang Liu, Georg Kohl, and Nils Thuerey. Pde-transformer: Efficient and
 595 versatile transformers for physics simulations. 2025.

596

597 Sergio Hoyas and Javier Jiménez. Reynolds number effects on the Reynolds-stress budgets in turbulent
 598 channels. *Physics of Fluids*, 20(10):101511, 10 2008. ISSN 1070-6631. doi: 10.1063/1.3005862.

599

600 Christian Jacobsen, Yilin Zhuang, and Karthik Duraisamy. Cocogen: Physically-consistent and
 601 conditioned score-based generative models for forward and inverse problems. *arXiv preprint*
 602 *arXiv:2312.10527*, 2023.

603

604 Yuchi Jiang, Zhijie Li, Yunpeng Wang, Huiyu Yang, and Jianchun Wang. An implicit adaptive
 605 fourier neural operator for long-term predictions of three-dimensional turbulence. *arXiv preprint*
 606 *arXiv:2501.12740*, 2025.

607

608 Dmitrii Kochkov, Jamie A. Smith, Ayya Alieva, Qing Wang, Michael P. Brenner, and Stephan
 609 Hoyer. Machine learning-accelerated computational fluid dynamics. *Proceedings of the National
 610 Academy of Sciences*, 118(21):e2101784118, 2021. doi: 10.1073/pnas.2101784118. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2101784118>.

611

612 Felix Koehler, Simon Niedermayr, Rüdiger Westermann, and Nils Thuerey. Apebench: A benchmark
 613 for autoregressive neural emulators of pdes. In *Advances in Neural Information Processing Systems*
 614 *37 (Datasets and Benchmarks Track)*, 2024. URL <https://openreview.net/forum?id=iWc0qE116u>.

615

616 Georg Kohl, Li-Wei Chen, and Nils Thuerey. Benchmarking autoregressive conditional diffusion
 617 models for turbulent flow simulation, 2024. URL <https://arxiv.org/abs/2309.01745>.

618

619 Jean Kossaifi, Nikola Kovachki, Kamyar Azizzadenesheli, and Anima Anandkumar. Multi-grid
 620 tensorized fourier neural operator for high-resolution pdes. *arXiv preprint arXiv:2310.00120*,
 621 2023.

622

623 Fanny Lehmann, Filippo Gatti, Michaël Bertin, and Didier Clouteau. 3d elastic wave propagation
 624 with a factorized fourier neural operator (f-fno). *Computer Methods in Applied Mechanics and
 625 Engineering*, 420:116718, 2024.

626

627 Zhijie Li, Wenhui Peng, Zelong Yuan, and Jianchun Wang. Long-term predictions of turbulence by
 628 implicit u-net enhanced fourier neural operator. *Physics of Fluids*, 35(7), 2023a.

629

630 Zhijie Li, Tianyuan Liu, Wenhui Peng, Zelong Yuan, and Jianchun Wang. A transformer-based neural
 631 operator for large-eddy simulation of turbulence. *Physics of Fluids*, 36(6), 2024a.

632

633 Zijie Li, Dule Shu, and Amir Barati Farimani. Scalable transformer for pde surrogate modeling.
 634 *Advances in Neural Information Processing Systems*, 36:28010–28039, 2023b.

635

636 Zijie Li, Kazem Meidani, and Amir Barati Farimani. Scalable transformer for pde surrogate modeling.
 637 *Advances in Neural Information Processing Systems*, 36, 2024b.

638

639 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
 640 for generative modeling, 2023. URL <https://arxiv.org/abs/2210.02747>.

641

642 Phillip Lippe, Bas Veeling, Paris Perdikaris, Richard Turner, and Johannes Brandstetter. Pde-
 643 refiner: Achieving accurate long rollouts with neural pde solvers. *Advances in Neural Information
 644 Processing Systems*, 36:67398–67433, 2023.

645

646 Qiang Liu and Nils Thuerey. Uncertainty-aware surrogate models for airfoil flow simulations
 647 with denoising diffusion probabilistic models. *AIAA Journal*, 62(8):2912–2933, 2024. doi:
 648 10.2514/1.J063440. URL <https://doi.org/10.2514/1.J063440>.

649

650 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate
 651 and transfer data with rectified flow. In *The Eleventh International Conference on Learning
 652 Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL
 653 <https://openreview.net/forum?id=XVjTT1nw5z>.

648 Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo.
 649 Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the*
 650 *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.

651 Michael McCabe, Bruno Régaldo-Saint Blancard, Liam Parker, Ruben Ohana, Miles Cranmer,
 652 Alberto Bietti, Michael Eickenberg, Siavash Golkar, Geraud Krawezik, Francois Lanusse, Mariel
 653 Pettee, Tiberiu Tesileanu, Kyunghyun Cho, and Shirley Ho. Multiple physics pretraining for
 654 physical surrogate models. In *NeurIPS 2023 AI for Science Workshop*, 2023. URL <https://openreview.net/forum?id=M121mQKuxa>.

655 Roberto Molinaro, Samuel Lanthaler, Bogdan Raonic, Tobias Rohner, Victor Armegioiu, Zhong Yi
 656 Wan, Fei Sha, Siddhartha Mishra, and Leonardo Zepeda-Núñez. Generative AI for fast and accurate
 657 statistical computation of fluids. *CoRR*, abs/2409.18359, 2024. doi: 10.48550/ARXIV.2409.18359.
 658 URL <https://doi.org/10.48550/arXiv.2409.18359>.

659 Paul D Morris, Andrew Narracott, Hendrik von Tengg-Kobligk, Daniel Alejandro Silva Soto, Sarah
 660 Hsiao, Angela Lungu, Paul Evans, Neil W Bressloff, Patricia V Lawford, D Rodney Hose, et al.
 661 Computational fluid dynamics modelling in cardiovascular medicine. *Heart*, 102(1):18–28, 2016.

662 Ruben Ohana, Michael McCabe, Lucas Thibaut Meyer, Rudy Morel, Fruzsina Julia Agocs, Miguel
 663 Beneitez, Marsha Berger, Blakesley Burkhart, Stuart B. Dalziel, Drummond Buschman Fielding,
 664 Daniel Fortunato, Jared A. Goldberg, Keiya Hirashima, Yan-Fei Jiang, Rich Kerswell, Surya-
 665 narayana Maddu, Jonah M. Miller, Payel Mukhopadhyay, Stefan S. Nixon, Jeff Shen, Romain
 666 Watteaux, Bruno Régaldo-Saint Blancard, François Rozet, Liam Holden Parker, Miles Cranmer,
 667 and Shirley Ho. The well: a large-scale collection of diverse physics simulations for machine
 668 learning. In *Advances in Neural Information Processing Systems 37 (Datasets and Benchmarks*
 669 *Track*), 2024. URL <https://openreview.net/forum?id=00Sx577BT3>.

670 John E. Pearson. Complex patterns in a simple system. *Science*, 261(5118):189–192, 1993. doi:
 671 10.1126/science.261.5118.189. URL <https://www.science.org/doi/abs/10.1126/science.261.5118.189>.

672 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *2023 IEEE/CVF*
 673 *International Conference on Computer Vision (ICCV)*, pp. 4172–4182, 2023. doi: 10.1109/ICCV51070.2023.00387.

674 Wenhui Peng, Zelong Yuan, Zhijie Li, and Jianchun Wang. Linear attention coupled fourier neural
 675 operator for simulation of three-dimensional turbulence. *Physics of Fluids*, 35(1), 2023.

676 Eric A. Perlman, Randal C. Burns, Yi Li, and Charles Meneveau. Data exploration of turbulence
 677 simulations using a database cluster. In *Proceedings of the ACM/IEEE Conference on High*
 678 *Performance Networking and Computing, SC 2007*, pp. 23, 2007. doi: 10.1145/1362622.1362654.
 679 URL <https://doi.org/10.1145/1362622.1362654>.

680 Raphaël Pestourie, Youssef Mroueh, Chris Rackauckas, Payel Das, and Steven G Johnson. Physics-
 681 enhanced deep surrogates for partial differential equations. *Nature Machine Intelligence*, 5(12):
 682 1458–1465, 2023.

683 Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A
 684 deep learning framework for solving forward and inverse problems involving nonlinear partial
 685 differential equations. *Journal of Computational physics*, 378:686–707, 2019.

686 Marcos Rodrigo, Carlos Cuevas, and Narciso García. Comprehensive comparison between vision
 687 transformers and convolutional neural networks for face recognition tasks. *Scientific reports*, 14
 688 (1):21392, 2024.

689 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical
 690 image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI*
 691 *2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III*
 692 *18*, pp. 234–241. Springer, 2015.

693 Youssef Shehata, Benjamin Holzschuh, and Nils Thuerey. Improved sampling of diffusion models in
 694 fluid dynamics with tweedie’s formula. In *The Thirteenth International Conference on Learning*
 695 *Representations*, 2025.

702 Dule Shu, Zijie Li, and Amir Barati Farimani. A physics-informed diffusion model for high-fidelity
 703 flow field reconstruction. *Journal of Computational Physics*, 478:111972, 2023.

704

705 Justin Sirignano and Jonathan F. MacArt. Deep learning closure models for large-eddy simulation of
 706 flows around bluff bodies. *Journal of Fluid Mechanics*, 966, jul 2023. doi: 10.1017/jfm.2023.446.
 707 URL <https://doi.org/10.1017/jfm.2023.446>.

708

709 Shashank Subramanian, Peter Harrington, Kurt Keutzer, Wahid Bhimji, Dmitriy Morozov, Michael
 710 Mahoney, and Amir Gholami. Towards foundation models for scientific machine learning: Charac-
 711 terizing scaling and transfer behavior, 2023.

712

713 Kiwon Um, Robert Brand, Yun, Fei, Philipp Holl, and Nils Thuerey. Solver-in-the-loop: Learning
 714 from differentiable physics to interact with iterative pde-solvers, 2021.

715

716 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
 717 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing*
 718 *systems*, 30, 2017.

719

720 Haixu Wu, Huakun Luo, Haowen Wang, Jianmin Wang, and Mingsheng Long. Transolver: A fast
 721 transformer solver for pdes on general geometries. *arXiv preprint arXiv:2402.02366*, 2024.

722

723 Heng Xiao and Paola Cinnella. Quantification of model uncertainty in rans simulations: A review.
 724 *Progress in Aerospace Sciences*, 108:1–31, 2019.

725

726 Liu Yang, Siting Liu, Tingwei Meng, and Stanley J Osher. In-context operator learning for differential
 727 equation problems. *arXiv preprint arXiv:2304.07993*, 2023.

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 A TRAINING METHODOLOGY AND NETWORK ARCHITECTURES
757758 A.1 SUPERVISED TRAINING
759760 For tasks that have a deterministic solution, such as training a surrogate model for a numerical solver,
761 the P3D can be trained in a supervised manner using mean squared error (MSE) loss, enabling fast,
762 single-step inference. In this case the network is directly optimized with the MSE
763

764
$$\mathcal{L}_S = \mathbb{E} [||\mathcal{M}_\Theta(\mathbf{u}_{\text{in}}, \mathbf{c}) - \mathbf{u}_{\text{out}}||_2^2]. \quad (3)$$

765
766

767 A.2 PROBABILISTIC TRAINING
768769 To allow for sampling from the full posterior distribution rather than approximating an averaged
770 outcome, diffusion training is preferable for probabilistic solutions. For such cases, we employ the
771 flow matching formulation of diffusion models (Lipman et al., 2023; Liu et al., 2023; Ho et al.,
772 2020). Given the input \mathbf{u}_{in} and a conditioning vector \mathbf{c} , samples \mathbf{x}_0 drawn from a noise distribution
773 $p_0 = \mathcal{N}(0, I)$ are transformed into samples \mathbf{x}_1 from the posterior distribution p_1 by solving an
774 ordinary differential equation (ODE) of the form $d\mathbf{x}_t = v(\mathbf{x}_t, t); dt$. Then the model \mathcal{M}_Θ learns
775 the velocity field v by regressing a vector field that defines a probabilistic trajectory from p_0 to p_1 .
776 Samples along this trajectory are produced through the forward process
777

778
$$\mathbf{x}_t = t \mathbf{u}_{\text{out}} + [1 - (1 - \sigma_{\text{min}})t] \epsilon \quad (4)$$

779 for $t \in [0, 1]$ with $\epsilon \sim \mathcal{N}(0, I)$ and a time threshold of $\sigma_{\text{min}} = 10^{-4}$. The velocity v can be regressed
780 by training via

781
$$\mathcal{L}_{\text{FM}} = \mathbb{E} [||\mathcal{M}_\Theta(\mathbf{u}_{\text{in}}, \mathbf{x}_t, \mathbf{c}, t) - \mathbf{u}_{\text{out}} + (1 - \sigma_{\text{min}})\epsilon||_2^2]. \quad (5)$$

782

783 After training, samples can be drawn from the posterior conditioned on \mathbf{u}_{in} and \mathbf{c} , by sampling
784 $\mathbf{x}_0 \sim \mathcal{N}(0, I)$ and integrating the corresponding ODE $d\mathbf{x}_t = \mathcal{M}(\mathbf{u}_{\text{in}}, \mathbf{x}_t, \mathbf{c}, t) dt$ over the time
785 interval $t = 0$ to $t = 1$. We typically employ explicit Euler steps with a suitable, chosen step size Δt .
786787 A.3 NORMALIZED RMSE
788789 The normalized RMSE (nRMSE) is defined as
790

791
$$\text{nRMSE} = \frac{1}{M} \sum_{i=1}^M \sqrt{\frac{\text{MSE}(\hat{\mathbf{u}}_{\text{out}}, \mathbf{u}_{\text{out}})}{\text{MSE}(\mathbf{0}, \mathbf{u}_{\text{out}})}}, \quad (6)$$

792

793 where $\hat{\mathbf{u}}_{\text{out}}$ is the network prediction and M corresponds to the number of trajectories for a specific
794 PDE in the test dataset. In our evaluation, we calculate the nRMSE for each PDE dataset and report
795 the nRMSE averaged over all PDE datasets.
796797 We can also define the nRMSE for autoregressive rollouts over the entire simulation trajectory. We
798 define the nRMSE $_t$ at time t by comparing the predicted state $\hat{\mathbf{u}}_t^S$ for a given system S at time t with
799 the reference \mathbf{u}_t^S . We average over all systems S for each PDE dataset.
800801 A.4 P3D ARCHITECTURE
802803 We provide additional details on the backbone architecture of P3D below.
804805 **Embedding of time, class labels and physical parameters** We combine all three types of
806 conditionings within a combined embedding layer. Time for flow matching/diffusion and physical
807 parameters are implemented via timestep embeddings. Class labels are implemented via label em-
808 beddings. The embedding vectors of all three types are added and used as the joint embedding. The
809 embedding dimension for each in the convolutional encoder/decoder is 64. In our experiments, class
labels are used in Section 4.1. Physical parameters and timestep embeddings are used in Section 4.3.

810 **Convolutional encoder** The convolution encoder first embeds the input using a Conv3D layer
 811 (kernel size 3, padding 1) with filters that correspond to the embedding dimensions of the configuration.
 812 This is followed by downsampling layers implemented via Conv3D layers (kernel size 3, padding 1,
 813 stride 2). Intermediate states before each downsampling operations are saved for residual connections.
 814 Encoder blocks and consecutive downsampling are applied twice. For each layer, the corresponding
 815 number of filters is shown in Table A1. Encoder blocks are repeated twice. Each encoder block
 816 consists of GroupNormalization layers, followed by GELU activations, Conv3D layer (kernel size
 817 3, padding 1), GroupNormalization, modulation via shift and scale operations depending on the
 818 conditioning, GELU and an additional Conv3D layer (kernel size 3, padding 1). The input and output
 819 of each encoder block are connected via skip connections. The shift and scale vectors are learned via
 820 linear layers from the embedding vectors of the convolutional encoder/decoder.

821 **Convolutional decoder** The design of the convolutional decoder mirrors the convolutional encoder
 822 in a U-shape architecture with residual connections. Upsampling layers are implemented via a
 823 combination of Conv3D layers to increase the number of filters and PixelShuffle3D layers. For an
 824 input number of channels C_{in} and a target number of channels in the upsampled output C_{out} , the
 825 Conv3D operation first expands the number of channels $C_{\text{in}} \times H \times W \times D \rightarrow 8C_{\text{out}} \times H \times W \times D$
 826 and PixelShuffle3D spatially rearranges the pixels $8C_{\text{out}} \times H \times W \times D \rightarrow C_{\text{out}} \times 2H \times 2W \times 2D$.
 827

828 A.5 CONTEXT MODEL

830 The context module is a transformer-based architecture for processing 3D volumetric data in a
 831 sequence-like fashion. The model's design and operation can be broken down into the following key
 832 components:

833 **Latent tokens** The model accepts a 5D tensor of shape (B, C, H, W, D) with batch size B , token
 834 embedding dimension C , height H , width W , and depth D , corresponding to the token representations
 835 of the transformer encoder block arranged on a 3D spatial grid according to their positions. This
 836 input is first reshaped into a 2D sequence of tokens with shape $(B, (H \times W \times D), C)$ by flattening
 837 the spatial dimensions (H, W, D) into a single sequence length. This token sequence is then passed
 838 through an input projection layer to map it to the model's internal hidden size $d_C = 512$. The new
 839 tokens are called *latent tokens*.
 840

841 **Region tokens** The model utilizes learnable region tokens, which are a set of `nn.Parameter`
 842 tensors. The number of tokens is defined by the number of regions. The region tokens are initialized
 843 from a normal distribution with a standard deviation of 10^{-6} . Region tokens are appended to the
 844 input sequence and are used to obtain embeddings for each region, which are injected into the adaptive
 845 instance normalization layers of the decoder modules.
 846

847 **Positional encoding** Positional information is incorporated using 3D sine-cosine positional em-
 848 beddings. The embeddings are added to the input token sequence after the initial projection and
 849 concatenation with region tokens. The positional embedding for a given position (w, h, d) and
 850 dimension i is calculated by concatenating three separate 1D sine-cosine embeddings for each spatial
 851 axis. The 1D positional embedding at position p and dimension i are calculated via:

$$\text{PE}(p, 2i) = \sin\left(\frac{p}{T^{2i/d_P}}\right) \quad (7)$$

$$\text{PE}(p, 2i + 1) = \cos\left(\frac{p}{T^{2i/d_P}}\right), \quad (8)$$

852 where $d_P = d_C/3$ is the embedding dimension per axis, and $T = 10$ is a temperature hyperparameter.
 853 Both latent tokens and region tokens have their own positional encoding.
 854

855 **Core architectural blocks** The sequence of tokens is processed by a stack of $n_L = 6$ `transformer`
 856 blocks. Each block consists of:

- **RMSNorm**: A root mean square normalization layer that normalizes the input to a unit RMS, followed by a learned scaling factor. It is a more computationally efficient alternative to traditional layer normalization.

864
 865
 866
 867
 868
 869
 870
 871
 872
 873
 874
 875
 876
 877
 878
 879
 880
 881
 882
 883
 884
 885
 886
 887
 888
 889
 890
 891
 892
 893
 894
 895
 896
 897
 898
 899
 900
 901
 902
 903
 904
 905
 906
 907
 908
 909
 910
 911
 912
 913
 914
 915
 916
 917
 918

- **Self-Attention**: A multi-head self-attention module that computes attention weights to model relationships between tokens in the input sequence. This block receives the input query tokens and calculates query, key, and value vectors. **There is no causal masking**. We use hyper attention (Han et al., 2023), this is a specialized attention variant designed for efficient long-sequence processing.
- **MLP**: A gated multi-layer perceptron with a GELU activation function.

Output The processed latent tokens are projected back to the original token size C via a linear layer and reshaped to match the input of the context network. The resulting tensor is added to the input of the transformer decoder, forming a skip connection. For each adaptive instance normalization module in the decoder, the regions tokens are embedded via a unique linear layer. The region embeddings are added to the input of the adaptive instance normalization layer of their corresponding region, resulting in a different modulation for each region.

Table A1: Different configurations S , B and L of P3D. Table shows the total number of weights, the number of filters within the convolutional encoder/decoder and the number of groups for GroupNormalization layers.

Configuration	Number of parameters	Embedding dimensions	Number of groups
S	11.2M	[32, 32, 64]	16
B	46.2M	[64, 128, 128]	32
L	181M	[128, 256, 256]	32

A.6 BASELINE ARCHITECTURES

Performance comparison of baseline architectures Below, we summarize the numbers of the different architectures used in the experiments for an input of size 128^3 with a single channel, see Table A2.

Table A2: Performance comparison of architectures.

Model	Params	GFLOPS	Memory	Throughput
Swin3D	50.3M	289.6	2.9GB	9.86it/s
FactFormer	5.0M	-	20.4GB	0.82it/s
UNet _{GenCFD}	100.0M	11039.6	4.8GB	2.03it/s
UNet _{ConvNext}	9.2M	335.6	2.4GB	11.9it/s
TFNO	75.6M	139.4	5.1GB	4.7it/s
AViT	60.0M	143.0	0.3GB	188.3it/s
AFNO	64.1M	2116.4	0.4GB	31.7it/s
P3D- <i>S</i>	11.2M	238.7	0.9GB	35.2it/s
P3D- <i>B</i>	46.1M	2330.6	2.1GB	10.0it/s
P3D- <i>L</i>	181.2M	9276.2	4.6GB	3.7it/s

Regarding the number of parameters, GFLOPs, memory and throughput, the three configurations S , B and L of P3D are well positioned compared to the baselines we chose. Timings were obtained on a RTX A5000 GPU with batch size 1. Importantly, across all tasks, P3D outperforms the baselines in terms of accuracy. We fixed the training setup for all architectures, using learning rate/optimizer/batch size that are common for training large transformer and UNet models (Esser et al., 2024). Due to the number and size of the different models, we do not perform any hyperparameter tuning for individual models.

Swin3D Our implementation of extends the SwinV2 (Liu et al., 2021) transformer in 3D. The code is based `Swinv2PreTrainedModel` in the Python `transformers` package, with pixel shuffling in 3D for upsampling. For the B (default) configuration of Swin3D, we use a token embedding dimension of 96 with patch size $p = 2$, four stages with two layers each, window size $w = 4$ and 12 attention heads for each stage.

918 **FactFormer** We use the official code repository of FactFormer <https://github.com/BaratiLab/FactFormer>. We adapt the FactFormer3D module used for the 3D Navier-Stokes experiment (Li et al., 2023b). Hyperparameters are taken from the configuration file in the repository: `dim=128, depth=4, dim_head=64, heads=6`.

919

920

921

922

923 **UNet_{GenCFD}** We use the official code repository for UNet_{GenCFD} <https://github.com/camlab-ethz/GenCFD>. We remove any modifications of the core UNet structure specific to the diffusion training setup in Molinaro et al. (2024). Hyperparameters are taken from the repository: `num_channels=[128, 256, 256], downsample_ratio=[2, 2, 2], num_blocks=4, padding_method='circular', use_position_encoding=True and num_heads=8`.

924

925

926

927

928

929 **UNet_{ConvNeXt}** We use the implementation of UNet_{ConvNeXt} with ConvNeXt blocks provided by the GitHub repository https://github.com/PolymathicAI/the_well. This architecture was used as a benchmark in Ohana et al. (2024). Hyperparameters are kept the same: `spatial_resolution=[128, 128, 128], stages=4, blocks_per_stage=1, blocks_per_neck=1, init_features=32`.

930

931

932

933

934

935 **TFNO** We use the implementation of TFNO (Kossaifi et al., 2023) from the Python neuralop package <https://github.com/neuraloperator/neuraloperator>. Hyperparameters are: `spatial_resolution=[128, 128, 128], modes_1=16, modes_2=16, modes_3=16, hidden_channels=64`.

936

937

938

939

940 **AViT** We use the implementation of axial vision transformers (Ho et al., 2019) provided by the GitHub repository https://github.com/PolymathicAI/the_well. This architecture was used as a benchmark in Ohana et al. (2024). Hyperparameters for the *B* configuration are kept the same: `spatial_resolution=[128, 128, 128], hidden_dim=768, num_heads=12, processor_blocks=8`.

941

942

943

944

945

946 **AFNO** We use the implementation of AFNO (Guibas et al., 2021) provided by the GitHub repository https://github.com/PolymathicAI/the_well. This architecture was used as a benchmark in Ohana et al. (2024). Hyperparameters for the *B* configuration are kept the same: `spatial_resolution=[128, 128, 128], hidden_dim=768, n_blocks=12, cmlp_diagonal_blocks=8 and patch_size=8`.

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972 **B EXPERIMENT 1: JOINTLY LEARNING PDEs**
973974 The datasets used in this experiment were carefully selected for a diverse list of partial differential
975 equations (PDEs) with a focus on high spatial resolution. This is combined with variations of the
976 physical parameters of the PDEs and different initial conditions, creating diverse dynamics across the
977 different simulations for each type of PDE.
978979 The dataset encompasses linear, reaction-diffusion, and nonlinear PDEs. We utilized the
980 *Exponax* solver, as detailed by Koehler et al. (2024) in the APEBench benchmark. The solver
981 utilizes Exponential Time Differencing Runge-Kutta (ETDRK) methods. We intentionally opted
982 not to use the APEBench dataset directly from the original authors. This decision was driven by
983 our goal to create datasets with enhanced resolution and greater diversity in the underlying physical
984 behaviors, rather than relying solely on variations in initial conditions as done in APEBench. It is a
985 characteristic of the ETDRK methods that they operate within the Fourier domain. Consequently,
986 their application is limited to scenarios with periodic domains and cannot accommodate complex
987 boundary conditions.
988989 Table B1: Summary of datasets produced for the joint PDE learning task, covering linear, reaction-
990 diffusion, and nonlinear PDEs. The table details the dimensions of each dataset: number of simula-
991 tions (s), time steps (t), fields/channels (f), and spatial dimensions (x, y, z). Beyond the explicitly
992 varied quantities listed for each dataset, the initial conditions for every simulation (s) are also distinct.
993

Dataset	s	t	f	x	y	z	Varied Quantities across s	Test Set
hyp	60	30	1	384	384	384	hyper-diffusivity	$s \in [50, 60[$
fisher	60	30	1	384	384	384	diffusivity, reactivity	$s \in [50, 60[$
sh	60	30	1	384	384	384	reactivity, critical number	$s \in [50, 60[$
gs-alpha	10	30	2	320	320	320	initial conditions only	separate: $s=3, t=100$
gs-beta	10	30	2	320	320	320	initial conditions only	separate: $s=3, t=100$
gs-gamma	10	30	2	320	320	320	initial conditions only	separate: $s=3, t=100$
gs-delta	10	30	2	320	320	320	initial conditions only	$s \in [8, 10[$
gs-epsilon	10	30	2	320	320	320	initial conditions only	separate: $s=3, t=100$
gs-theta	10	30	2	320	320	320	initial conditions only	$s \in [8, 10[$
gs-iota	10	30	2	320	320	320	initial conditions only	$s \in [8, 10[$
gs-kappa	10	30	2	320	320	320	initial conditions only	$s \in [8, 10[$
burgers	60	30	3	384	384	384	viscosity	$s \in [50, 60[$
kdv	60	30	3	384	384	384	domain extent, viscosity	$s \in [50, 60[$
ks	60	30	1	384	384	384	domain extent	separate: $s=5, t=200$

1008 **B.1 DATA GENERATION SETUP**
10091010 A key aspect of the simulations, in addition to parameter variation, is the use of randomized initial
1011 conditions. The standard approach for constructing these conditions involves randomly selecting one
1012 of three initialization methods, each providing a unique spectral energy distribution. The first method,
1013 a random truncated Fourier series initializer, involves summing multiple Fourier series up to a cutoff
1014 frequency, chosen as a uniformly random integer between 2 and 10 (exclusive of 11). The second,
1015 the Gaussian random field initializer, produces a power-law spectrum in Fourier space, where energy
1016 diminishes polynomially with the wavenumber; its exponent is uniformly randomly selected from
1017 [2.3,3.6[. The third method, the diffused noise initializer, generates a tensor of values from normally
1018 distributed white noise, subsequently diffusing it. This results in a spectrum that decays exponentially
1019 quadratically, with an intensity rate uniformly random from [0.00005,0.01[. After generation, all
1020 initializers ensure the initial conditions' values are normalized to a maximum absolute value of one.
1021 For vector quantities, the randomly chosen initializer is applied independently to each component.
10221023 **B.2 PDE TYPES**
10241025 We make use of Exponential Time Differencing Runge-Kutta (ETDRK) methods to efficiently
1026 simulate different PDEs via Exponax. While the chosen linear PDEs are simple and analytically

1026
 1027
 1028
 1029
 1030
 1031
 1032
 1033
 1034
 1035
 1036
 1037
 1038
 1039
 1040
 1041
 1042
 1043
 1044
 1045
 1046
 1047
 1048
 1049
 1050
 1051
 1052
 1053
 1054
 1055
 1056
 1057
 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079

solvable, the underlying dynamics are essential for more complicated PDEs. These linear PDEs can be understood as representing a scalar attribute, such as density. Unless stated otherwise, sampling from intervals is consistently performed using a uniform random distribution. In the following, we describe a range of different three-dimensional PDE problems that are employed in our experiments ranging from linear, over reaction-diffusion, to non-linear PDEs. The class of non-linear PDEs is particularly challenging, as these cases more closely resemble real-world problems. [Table B2 gives an overview of the different PDE equations and dynamics.](#)

Hyper-Diffusion (hyp) behaves similarly to diffusion, where density dissipates inside a periodic domain due to the effects of hyper-diffusion. Unlike diffusion, hyper-diffusion does not treat all wavelengths equally, analogous to the relation between dispersion and advection. Higher frequency components are damped even more aggressively compared to normal diffusion, leading to visually stronger blur effect of the density field over time.

- Dimensionality: $s = 60$, $t = 30$, $f = 1$, $x = 384$, $y = 384$, $z = 384$
- Initial Conditions: random truncated Fourier / Gaussian random field / diffused noise
- Boundary Conditions: periodic
- Time Step of Stored Data: 0.01
- Spatial Domain Size of Simulation: $[0, 1] \times [0, 1]$
- Fields: density
- Varied Parameters: hyper-diffusivity $\zeta \in [0.00005, 0.0005[$
- Validation Set: random 15% split of all sequences from $s \in [0, 50[$
- Test Set: all sequences from $s \in [50, 60[$

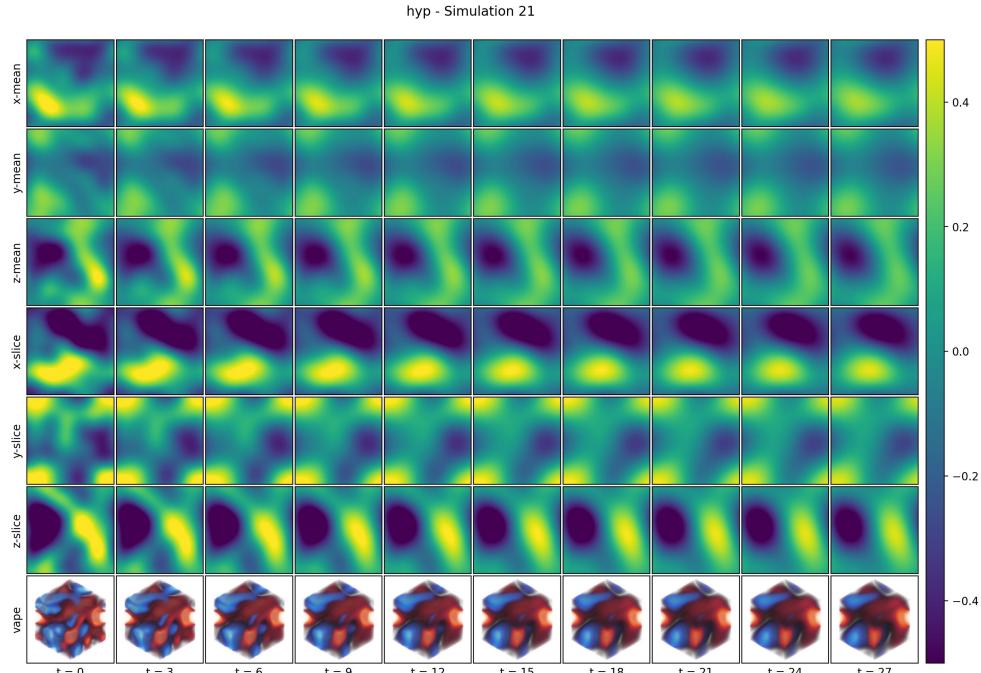


Figure B1: Random example simulation from hyp.

1080
 1081
 1082
 1083
 1084
 1085
 1086
 1087
 1088
 1089
 1090
 1091
 1092
 1093
 1094
 1095
 1096
 1097
 1098
 1099
 1100
 1101
 1102
 1103
 1104
 1105
 1106
 1107
 1108
 1109
 1110
 1111
 1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133

Table B2: Overview of the PDE dynamics. Table adapted from Koehler et al. (2024).

Name	Dynamics	PDE	Coefficients
Hyper-Diffusion	hyp	$\frac{\partial u}{\partial t} = -\zeta((\nabla \odot \nabla) \cdot (\nabla \odot \nabla))u$	ζ varied
Fisher-KPP	fisher	$\frac{\partial u}{\partial t} = \nu(\nabla \cdot \nabla)u + ru(1-u)$	ν varied r varied
Swift-Hohenberg	sh	$\frac{\partial u}{\partial t} = ru - (k + \nabla)^2 u + u^2 - u^3$	r varied k varied
Gray-Scott	gs-alpha gs-beta gs-delta gs-iot-a gs-theta gs-theta gs-epsilon gs-kappa	$\frac{u_0}{\partial t} = \nu_0 \nabla \cdot \nabla u_0 - u_0 u_1^2 + f(1 - u_0)$ $\frac{u_1}{\partial t} = \nu_1 \nabla \cdot \nabla u_1 + u_0 u_1^2 - f(f + k)u_1$	$\nu_0 = 0.002$ $\nu_1 = 0.001$ f varied k varied
Burgers	burgers	$\frac{\partial \mathbf{u}}{\partial t} = -\frac{1}{2} \nabla \cdot (\mathbf{u} \otimes \mathbf{u}) + \nu(\nabla \cdot \nabla) \mathbf{u}$	ν varied
Korteweg-de-Vries	kdv	$\frac{\partial \mathbf{u}}{\partial t} = -\frac{1}{2} b \nabla \cdot (\mathbf{u} \otimes \mathbf{u}) - a \vec{1} \cdot (\nabla \odot \nabla \odot (\nabla \mathbf{u})) + \nu(\nabla \odot \nabla) \mathbf{u} - \zeta((\nabla \odot \nabla) \cdot (\nabla \odot \nabla)) \mathbf{u}$	$b = -6.0$ $a = 1.0$ ν varied $\zeta = 0.01$
Kuramoto-Sivashinsky	ks	$\frac{\partial \mathbf{u}}{\partial t} = -b \frac{1}{2} \nabla u _2^2 + \nu(\nabla \odot \nabla)u - \zeta((\nabla \odot \nabla) \cdot (\nabla \odot \nabla))u$	$b = 1.0$ $\nu = 1.0$ $\zeta = 1.0$

1134 **Fisher-KPP (fisher)** is a foundational reaction-diffusion PDE. These systems are used to model
 1135 biological or geological processes, often resulting in pattern formation. This equation details how the
 1136 concentration of a substance varies over time and space, influenced by a reaction process governed by
 1137 a reactivity parameter, and its dispersal through diffusion, which is defined by a diffusivity parameter.
 1138 Its applications extend to various domains, including wave propagation, population dynamics, ecology,
 1139 and plasma physics. Figure B2 shows example visualizations from `fisher`.

1140

- Dimensionality: $s = 60, t = 30, f = 1, x = 384, y = 384, z = 384$
- Initial Conditions: random truncated Fourier / Gaussian random field / diffused noise (with clamping to $[0, 1]$)
- Boundary Conditions: periodic
- Time Step of Stored Data: 0.005
- Spatial Domain Size of Simulation: $[0, 1] \times [0, 1]$
- Fields: concentration
- Varied Parameters: diffusivity $\nu \in [0.0001, 0.02]$ and reactivity $r \in [5, 15]$
- Validation Set: random 15% split of all sequences from $s \in [0, 50[$
- Test Set: all sequences from $s \in [50, 60[$

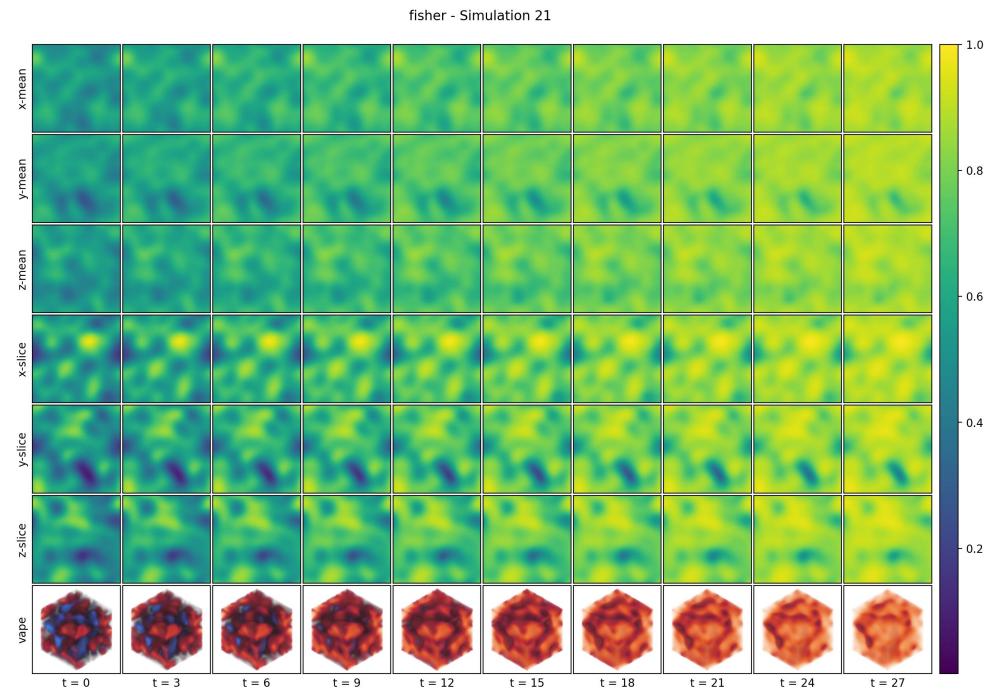


Figure B2: Random example simulation from `fisher`.

1188 **Swift-Hohenberg (sh)** is known for depicting various pattern formation processes. This equation
 1189 can be applied to illustrate the structure of wrinkles in curved elastic bilayer materials. A prime
 1190 example is the formation of human fingerprints, where tensions between skin layers generate their
 1191 unique wrinkling. Figure B3 shows example visualizations from `sh`.

1192

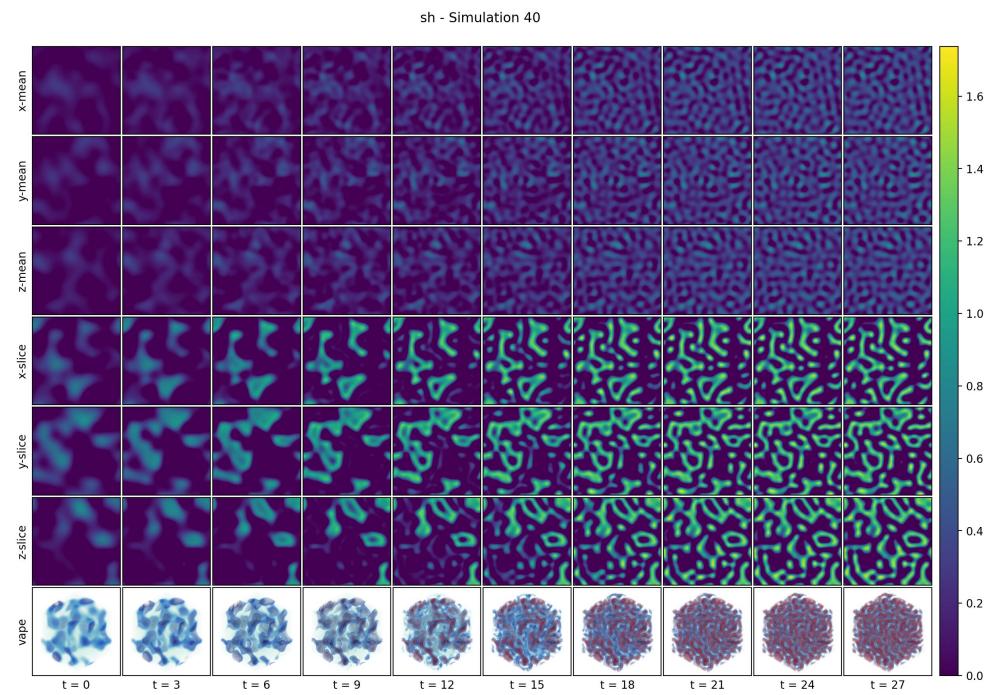
- 1193 • Dimensionality: $s = 60, t = 30, f = 1, x = 384, y = 384, z = 384$
- 1194 • Initial Conditions: random truncated Fourier / Gaussian random field / diffused noise
- 1195 • Boundary Conditions: periodic
- 1196 • Time Step of Stored Data: 0.5 (with 5 substeps for the simulation)
- 1197 • Spatial Domain Size of Simulation: $[0, 20\pi] \times [0, 20\pi]$
- 1198 • Fields: concentration
- 1199 • Varied Parameters: reactivity $r \in [0.4, 1[$ and critical number $k \in [0.8, 1.2[$
- 1200 • Validation Set: random 15% split of all sequences from $s \in [0, 50[$
- 1201 • Test Set: all sequences from $s \in [50, 60[$

1202

1203

1204

1205



1227

1228

Figure B3: Random example simulation from `sh`.

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

1242 **Gray-Scott (gs)** illustrates the dynamic interplay of two reacting and diffusing chemical
 1243 substances. Substance s_a with concentration c_a is depleted through reaction but resupplied based
 1244 on a defined feed rate. Meanwhile, substance s_b , the reaction's product with concentration c_b , is
 1245 eliminated from the domain at a given kill rate. The balance between these two rates profoundly
 1246 influences the simulation outcomes, leading to diverse stable or evolving patterns. We simulate
 1247 several cases: four with temporally steady configurations, which result in a state that does not
 1248 substantially change anymore (gs-delta, gs-theta, gs-iota, and gs-kappa), and four
 1249 temporally unsteady configurations, which continuously evolve over time (gs-alpha, gs-beta,
 1250 gs-gamma, and gs-epsilon). For the unsteady case, separate test sets with longer temporal
 1251 rollouts are created. Figure B4 shows example visualizations from the steady configurations, and
 1252 Figure B5 from the unsteady configurations and corresponding test sets. For further details, see
 1253 (Pearson, 1993).

1254 For all simulations, the diffusivity of the substances is fixed to $d_a = 0.00002$ and $d_b = 0.00001$. In
 1255 addition to that, datasets are initialized with a Gaussian blob initializer. The initializer creates four
 1256 Gaussian blobs at random positions and variances in the center 60% (20% for gs-kappa) of the
 1257 domain, where the initialization of c_a is the complement of c_b , i.e. $c_a = 1 - c_b$.

1258 **Steady Configurations (gs-delta, gs-theta, gs-iota, and gs-kappa):**

- Dimensionality: $s = 10, t = 30, f = 2, x = 320, y = 320, z = 320$ (per configuration)
- Initial Conditions: random Gaussian blobs
- Boundary Conditions: periodic
- Time Step of Simulation: 1.0 (all configurations)
- Time Step of Stored Data:
 - gs-delta: 130.0
 - gs-theta: 200.0
 - gs-iota: 240.0
 - gs-kappa: 300.0
- Number of Warmup Steps (discarded, in time step of data storage):
 - gs-delta: 0
 - gs-theta: 0
 - gs-iota: 0
 - gs-kappa: 15
- Spatial Domain Size of Simulation: $[0, 2.5] \times [0, 2.5]$
- Fields: concentration c_a , concentration c_b
- Varied Parameters: feed rate and kill rate determined by configuration (i.e., initial conditions
 only within configuration)
 - gs-delta: feed rate: 0.028, kill rate: 0.056
 - gs-theta: feed rate: 0.040, kill rate: 0.060
 - gs-iota: feed rate: 0.050, kill rate: 0.0605
 - gs-kappa: feed rate: 0.052, kill rate: 0.063
- Validation Set: random 15% split of all sequences from $s \in [0, 8[$
- Test Set: all sequences from $s \in [8, 10[$

1286 **Unsteady Configurations (gs-alpha, gs-beta, gs-gamma, and gs-epsilon):**

- Dimensionality: $s = 10, t = 30, f = 2, x = 320, y = 320, z = 320$ (per configuration)
- Initial Conditions: random Gaussian blobs
- Boundary Conditions: periodic
- Time Step of Simulation: 1.0 (all configurations)
- Time Step of Stored Data:
 - gs-alpha: 30.0
 - gs-beta: 30.0
 - gs-gamma: 75.0

```

1296     - gs-epsilon: 15.0
1297 • Number of Warmup Steps (discarded, in time step of data storage):
1298     - gs-alpha: 75
1299     - gs-beta: 50
1300     - gs-gamma: 70
1301     - gs-epsilon: 300
1302
1303 • Spatial Domain Size of Simulation:  $[0, 2.5] \times [0, 2.5]$ 
1304 • Fields: concentration  $c_a$ , concentration  $c_b$ 
1305 • Varied Parameters: feed rate and kill rate determined by configuration (i.e., initial conditions
1306 only within configuration)
1307     - gs-alpha: feed rate: 0.008, kill rate: 0.046
1308     - gs-beta: feed rate: 0.020, kill rate: 0.046
1309     - gs-gamma: feed rate: 0.024, kill rate: 0.056
1310     - gs-epsilon: feed rate: 0.020, kill rate: 0.056
1311
1312 • Validation Set: random 15% split of all sequences from  $s \in [0, 10[$ 
1313 • Test Set: separate simulations with  $s = 30, t = 100, f = 2, x = 320, y = 320, z = 320$ 
1314 (per configuration)
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

```

1350
1351
1352
1353
1354

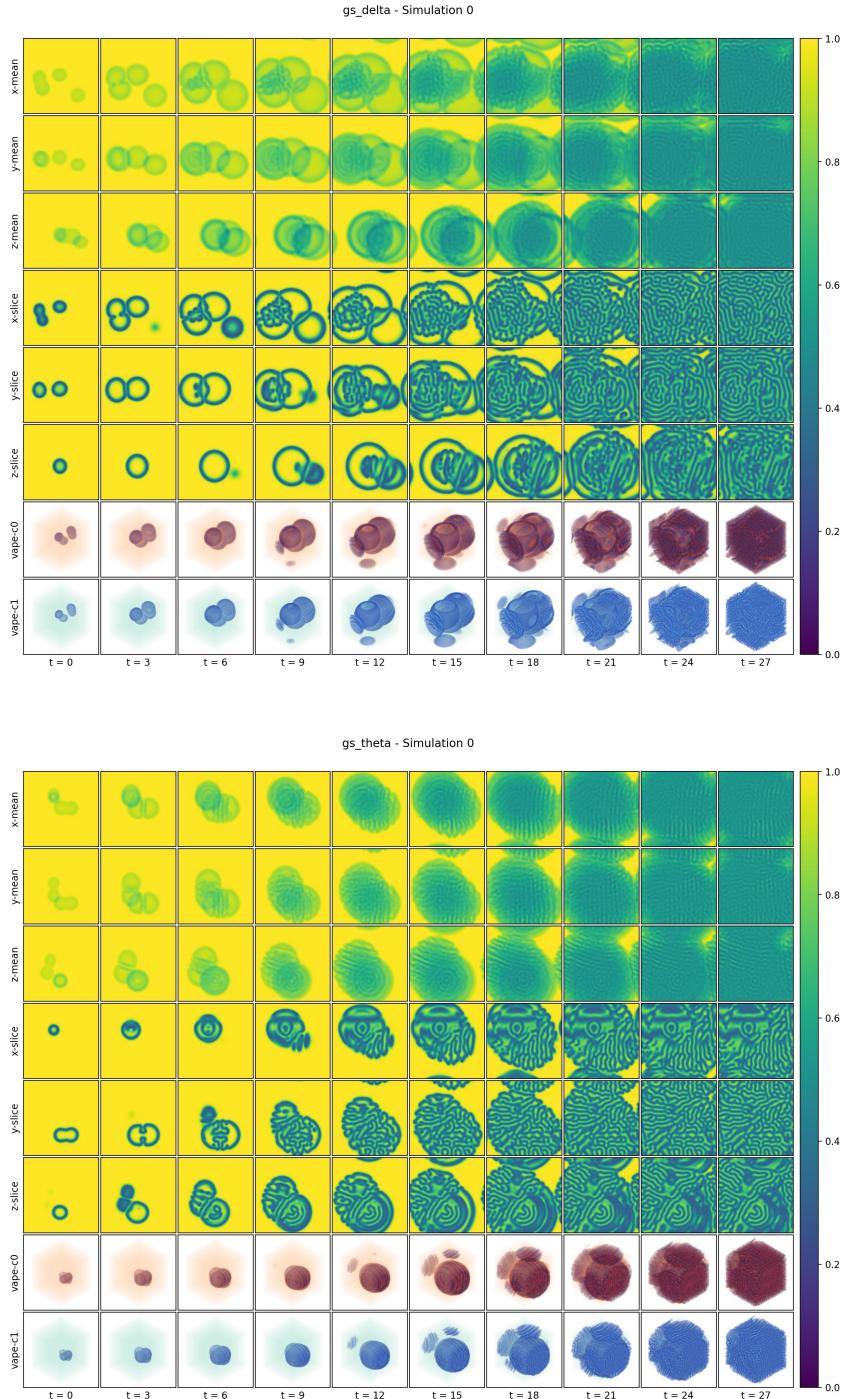


Figure B4: Random example simulations from steady configurations of the Gray-Scott model of a reaction-diffusion system: `gs-delta`, `gs-theta`.

1401

1404
1405
1406
1407
1408

Figure B5: Random example simulations from unsteady configurations of the Gray-Scott model of a reaction-diffusion system: $gs\text{-}\alpha$, $gs\text{-}\beta$.

Burgers (burgers) bears resemblance to an advection-diffusion problem. Instead of modeling the transport of a scalar density, this equation describes how a flow field itself evolves due to the combined effects of advection and diffusion. This process can result in the formation of abrupt discontinuities, often referred to as shock waves, which present a significant challenge for accurate simulation. Burgers' equation also finds utility in fields such as nonlinear acoustics and the modeling of traffic flow. Figure B6 shows example visualizations from `burgers`.

- Dimensionality: $s = 60$, $t = 30$, $f = 3$, $x = 384$, $y = 384$, $z = 384$
- Initial Conditions: random truncated Fourier / Gaussian random field / diffused noise
- Boundary Conditions: periodic
- Time Step of Stored Data: 0.01 (with 50 substeps for the simulation)
- Spatial Domain Size of Simulation: $[0, 1] \times [0, 1]$
- Fields: velocity (x, y)
- Varied Parameters: viscosity $\nu \in [0.001, 0.005[$
- Validation Set: random 15% split of all sequences from $s \in [0, 50[$
- Test Set: all sequences from $s \in [50, 60[$

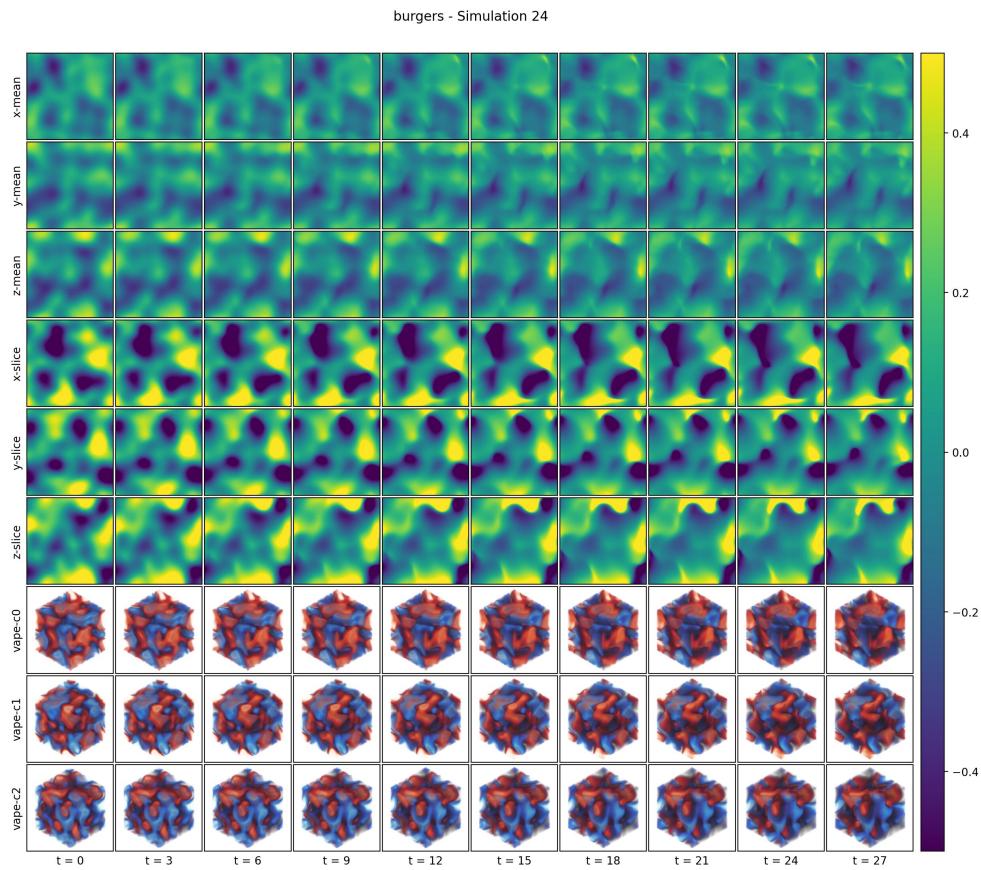


Figure B6: Random example simulation from burgers.

1512
 1513 **Korteweg-de-Vries (kdv)** presents simulations of the Korteweg-de-Vries equation within a pe-
 1514 riodic domain. This equation models dispersive, non-dissipative wave propagation and is a classic
 1515 example of an integrable PDE. It poses a challenge because energy is transferred to high spatial
 1516 frequencies, resulting in distinct, moving soliton waves that maintain their shape and propagation
 1517 speed. Throughout these simulations, the convection coefficient remains constant at -6 , and the
 1518 dispersivity coefficient is consistently 1 . Figure B7 shows example visualizations from kdv.
 1519

1520

- 1521 • Dimensionality: $s = 60, t = 30, f = 2, x = 384, y = 384, z = 384$
- 1522 • Initial Conditions: random truncated Fourier / Gaussian random field / diffused noise
- 1523 • Boundary Conditions: periodic
- 1524 • Time Step of Stored Data: 0.05 (with 10 substeps for the simulation)
- 1525 • Spatial Domain Size of Simulation: varied per simulation
- 1526 • Fields: velocity (x, y)
- 1527 • Varied Parameters: domain extent $\in [30, 120[$ identically for x, y, z , i.e. a square domain,
 1528 and viscosity $\nu \in [0.1, 0.25[$
- 1529 • Validation Set: random 15% split of all sequences from $s \in [0, 50[$
- 1530 • Test Set: all sequences from $s \in [50, 60[$

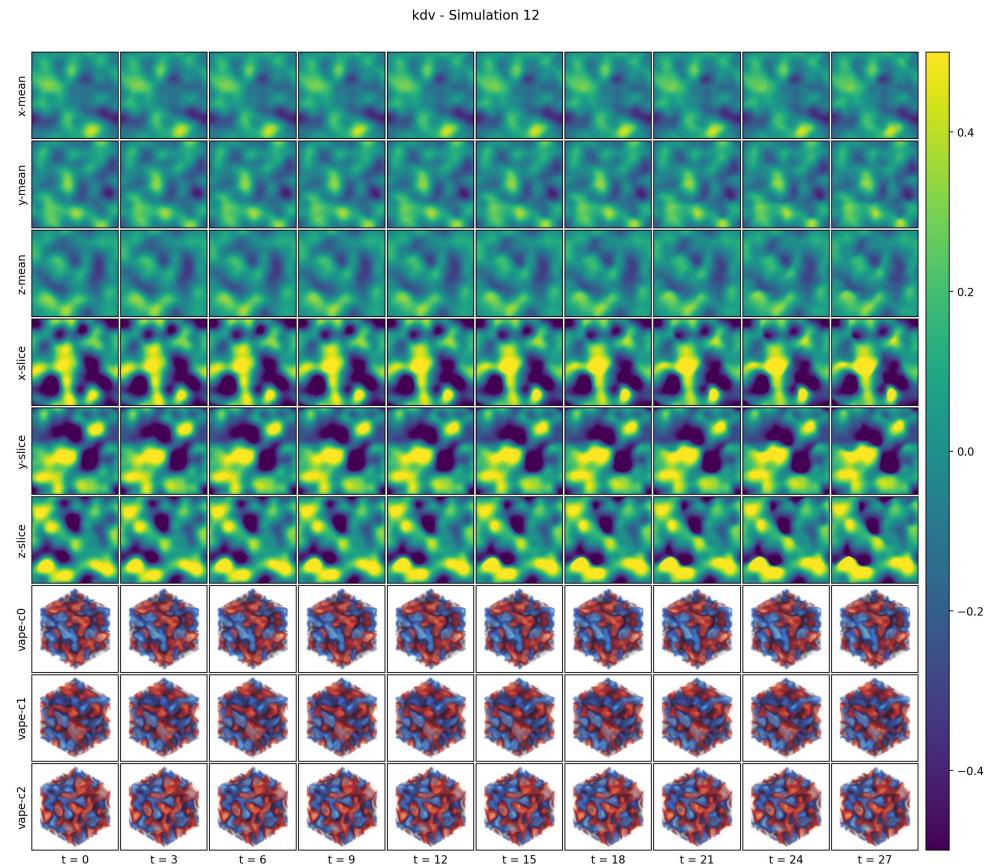


Figure B7: Random example simulation from kdv.

1566 **Kuramoto-Sivashinsky (ks)** models thermo-diffusive flame instabilities in combustion and also
 1567 finds use in reaction-diffusion systems on a periodic domain. It's notable for its chaotic behavior,
 1568 where even slightly different initial conditions can lead to wildly divergent temporal trajectories over
 1569 time. The initial transient phase of the simulations is not included in the dataset. Figure B8 shows
 1570 example visualizations from **ks**.

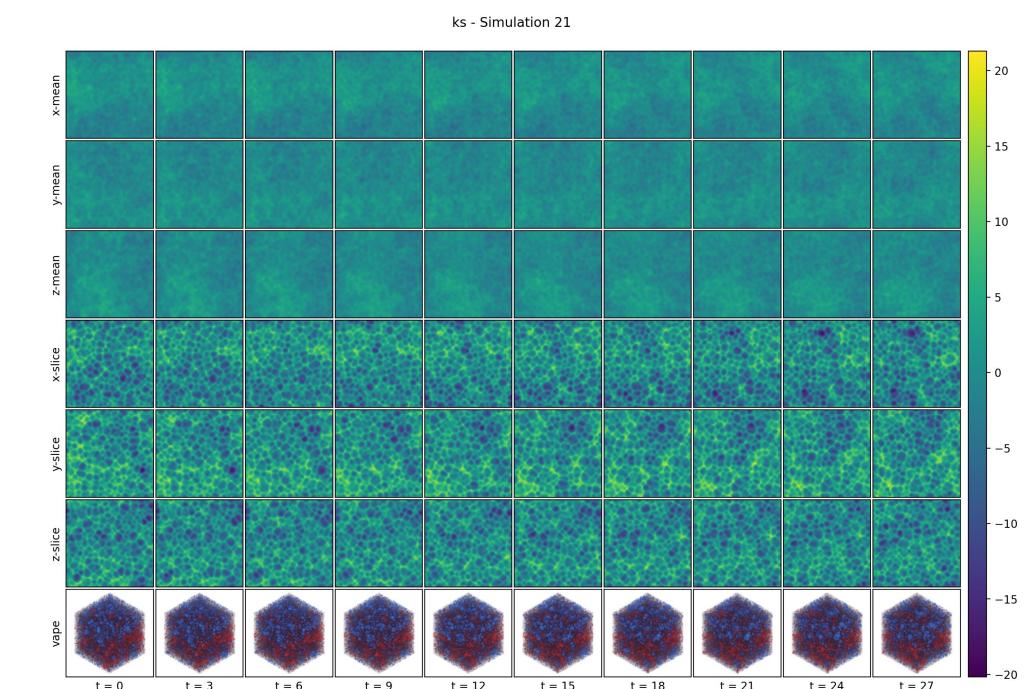
1571

- Dimensionality: $s = 60$, $t = 30$, $f = 1$, $x = 384$, $y = 384$, $z = 384$
- Initial Conditions: random truncated Fourier / Gaussian random field / diffused noise
- Boundary Conditions: periodic
- Time Step of Stored Data: 0.2 (with 2 substeps for the simulation)
- Number of Warmup Steps (discarded, in time step of data storage): 200
- Spatial Domain Size of Simulation: varied per simulation
- Fields: density
- Varied Parameters: domain extent $\in [10, 130[$ identically for x, y , i.e. a square domain
- Validation Set: random 15% split of all sequences from $s \in [0, 600[$
- Test Set: separate simulations with $s = 50$, $t = 200$, $f = 1$, $x = 384$, $y = 384$, $z = 384$

1583

1584

1585

1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619Figure B8: Random example simulation from **ks**.

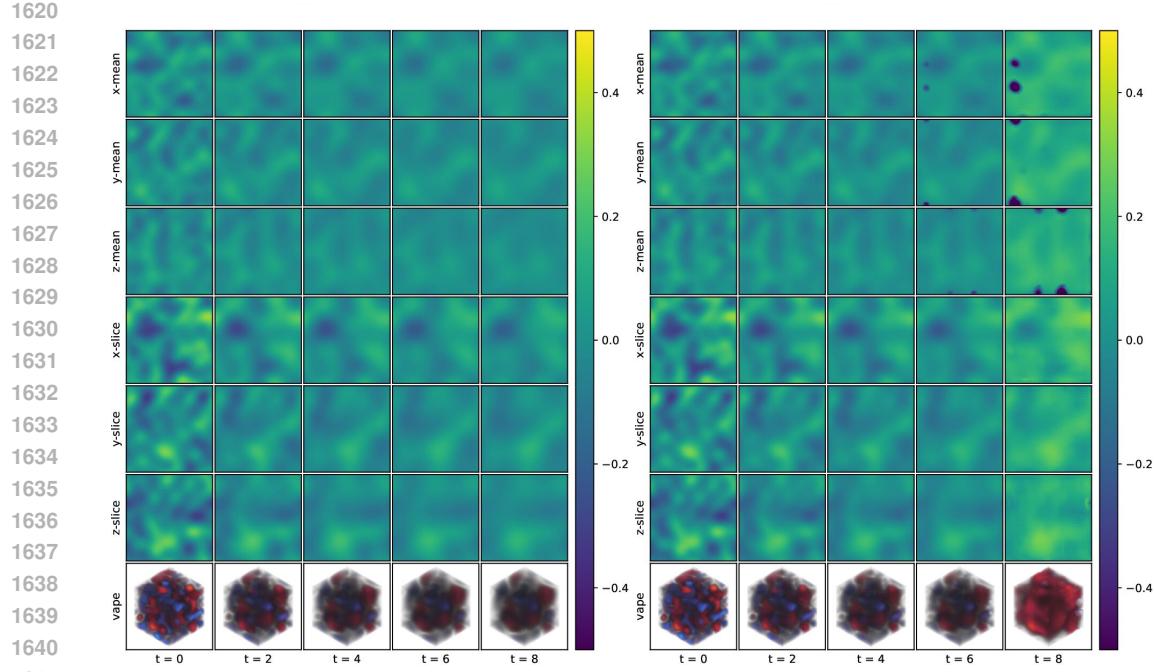


Figure B9: hyp. Reference (left) and autoregressive prediction for $t = 8$ steps with P3D- $L < 128 | 384 >$ (right) on the test set at resolution 384^3 .

B.3 VISUALIZATIONS OF PREDICTIONS

Below, we visualize several example predictions from the P3D- L network trained on crop size 128^3 on the test datasets. See Figures B9 to B17. During inference, we apply the network to larger crops of domain size 320^3 for the Gray-Scott PDEs and 384^3 for all other PDEs. The resolution is significantly higher than what the network was originally trained on. In addition, since the full domain is periodic, we change the padding mode of convolutional layers within the convolutional encoder and decoder from "zeros" to "circular". We consider autoregressive rollouts of up to $t = 8$ steps.

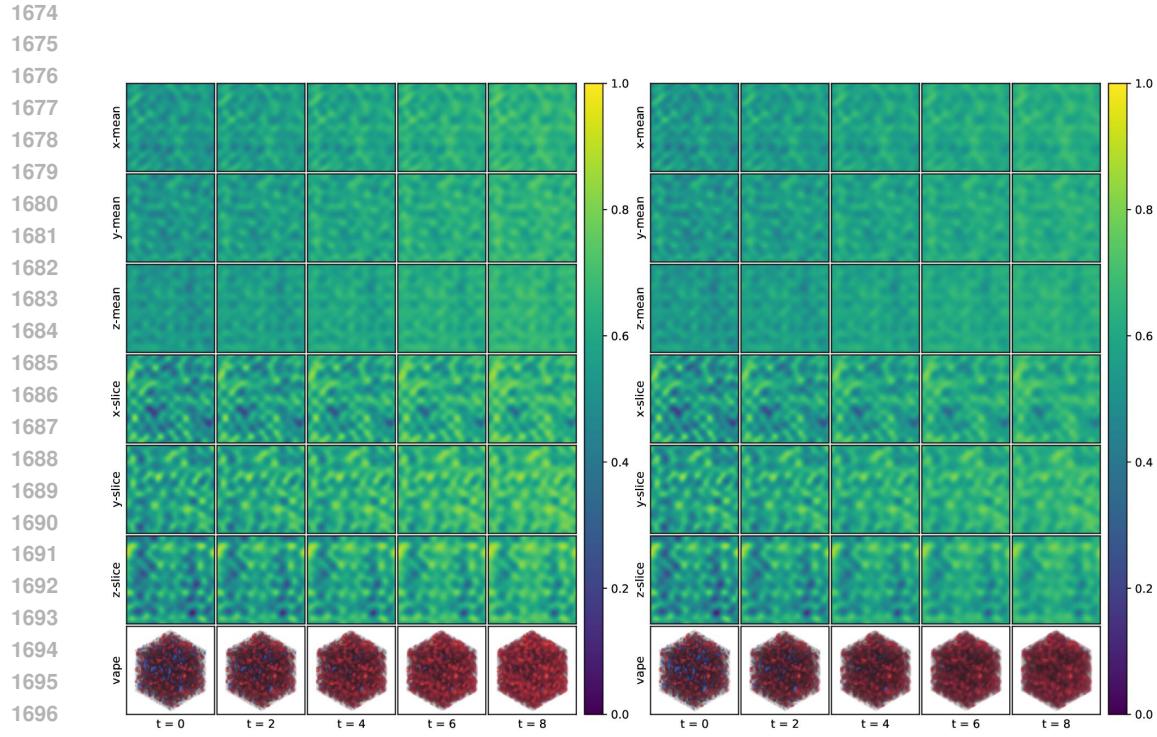


Figure B10: *fisher*. Reference (left) and autoregressive prediction for $t = 8$ steps with P3D- $L < 128|384 >$ (right) on the test set at resolution 384^3 .

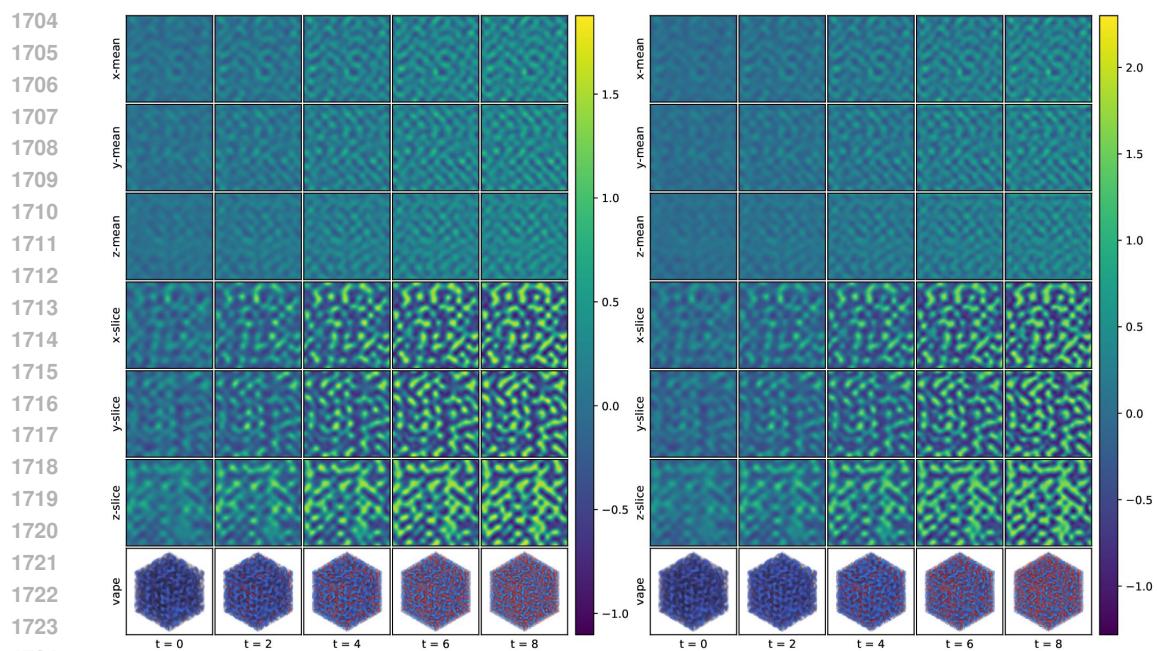


Figure B11: *sh*. Reference (left) and autoregressive prediction for $t = 8$ steps with P3D- $L < 128|384 >$ (right) on the test set at resolution 384^3 .

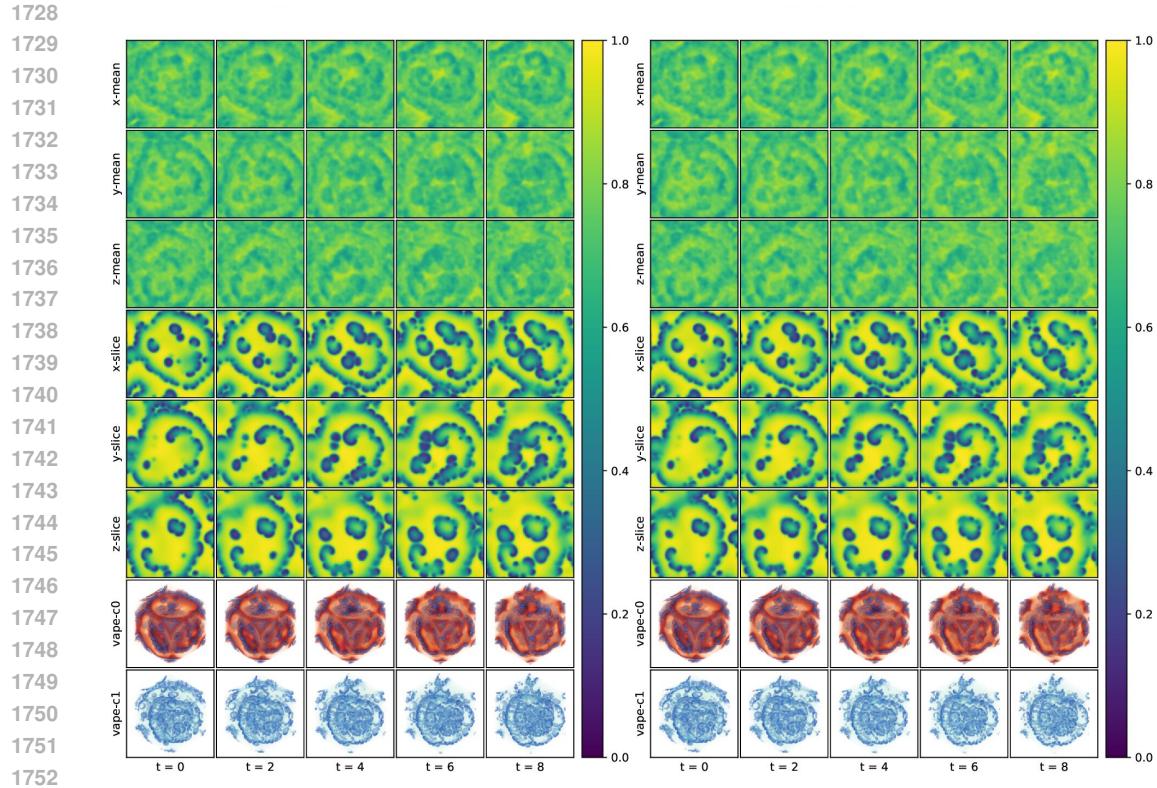


Figure B12: $gs\text{-}\alpha$. Reference (left) and autoregressive prediction for $t = 8$ steps with $P3D\text{-}L <128|320>$ (right) on the test set at resolution 320^3 .

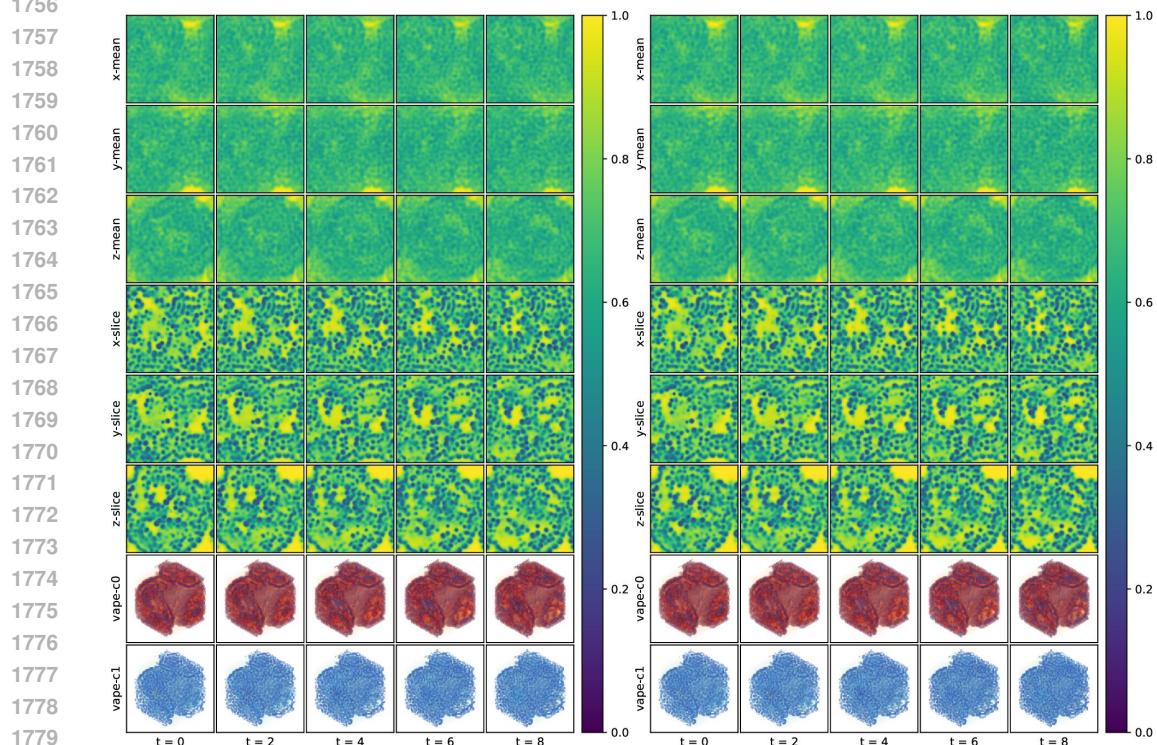


Figure B13: $gs\text{-}\epsilon$. Reference (left) and autoregressive prediction for $t = 8$ steps with $P3D\text{-}L <128|320>$ (right) on the test set at resolution 320^3 .

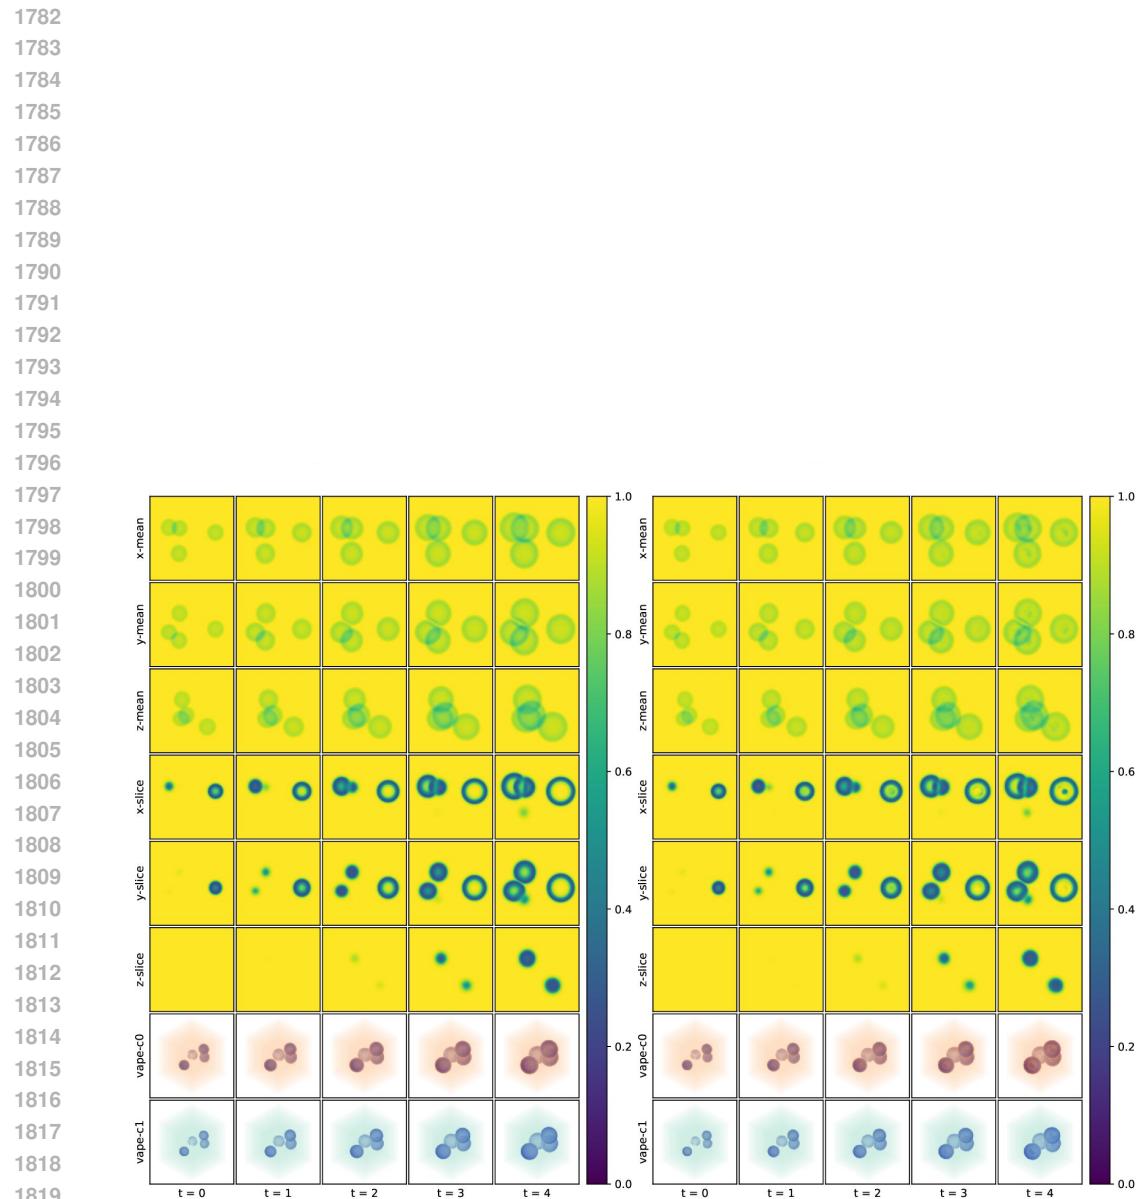


Figure B14: $gs\text{-}\delta\alpha$. Reference (left) and autoregressive prediction for $t = 8$ steps with $P3D-L <128|320>$ (right) on the test set at resolution 320^3 .

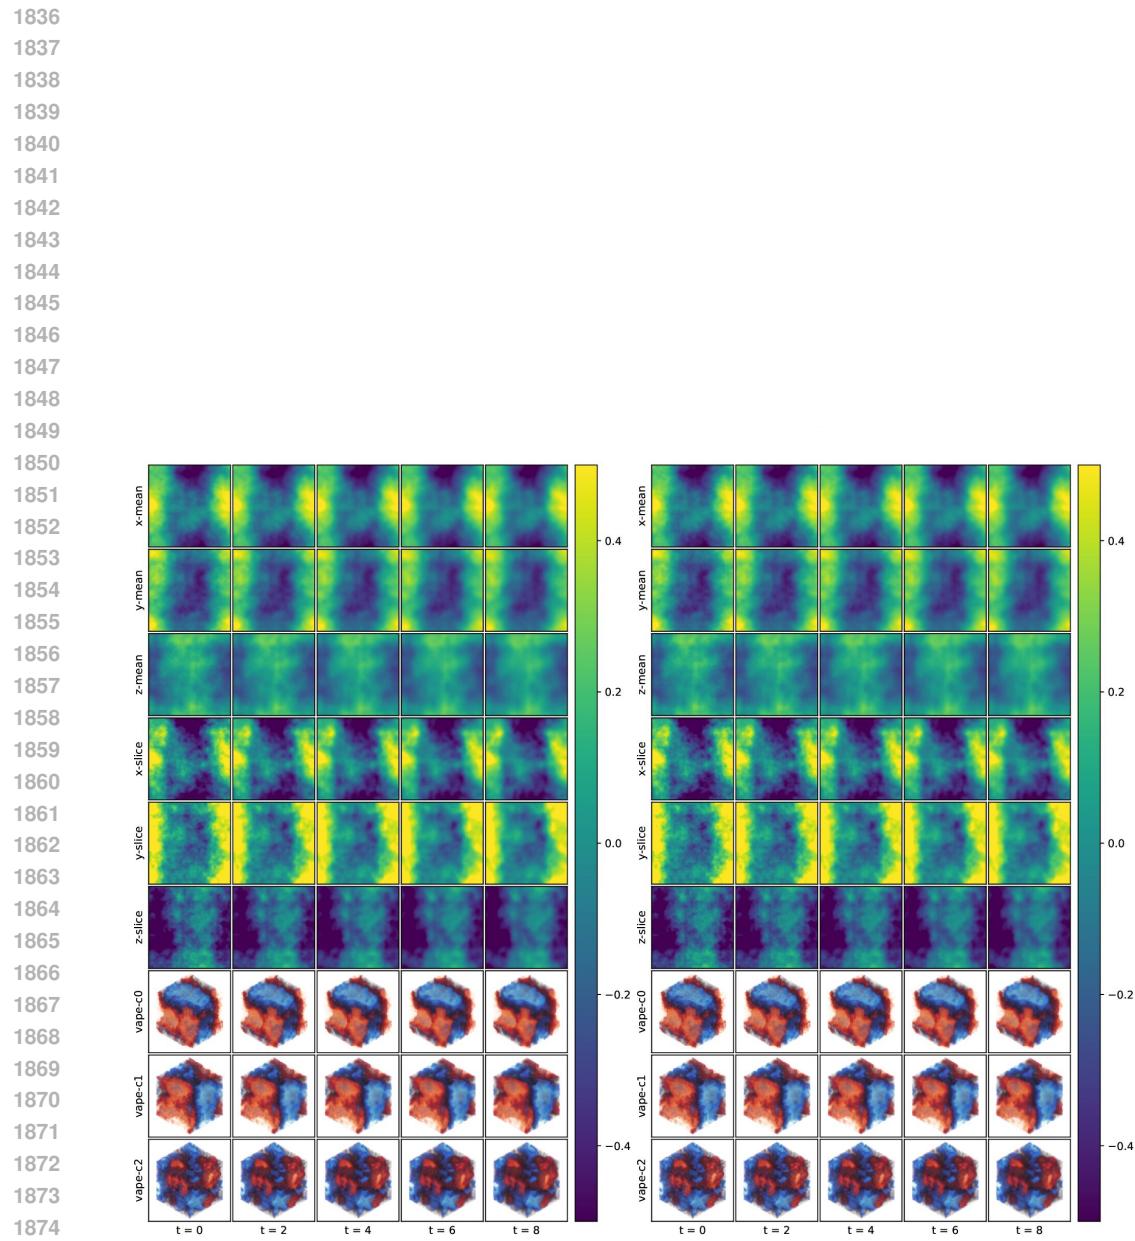


Figure B15: *burgers*. Reference (left) and autoregressive prediction for $t = 8$ steps with P3D- $L <128|384>$ (right) on the test set at resolution 384^3 .

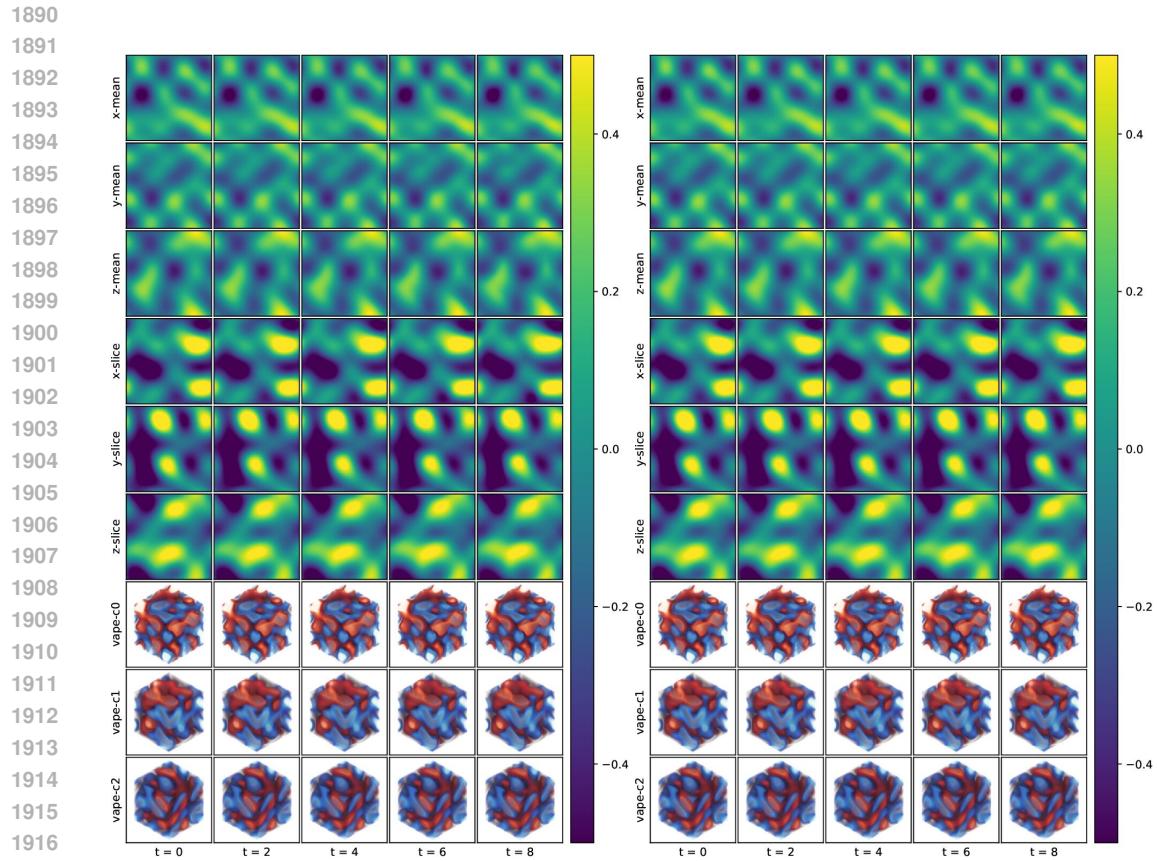


Figure B16: kdv . Reference (left) and autoregressive prediction for $t = 8$ steps with $P3D-L <128|384>$ (right) on the test set at resolution 384^3 .

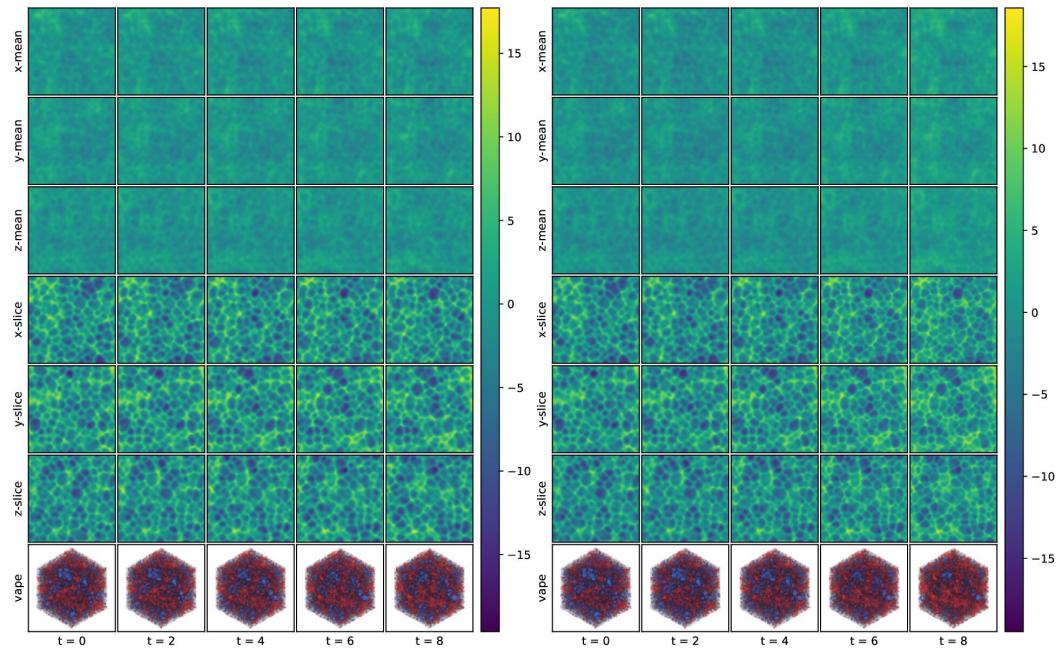


Figure B17: ks . Reference (left) and autoregressive prediction for $t = 8$ steps with $P3D-L <128|384>$ (right) on the test set at resolution 384^3 .

Table B3: Normalized RMSE ($\times 10^{-3}$) for crop size 128.

Model Name	Fisher	GS α	GS β	GS δ	GS ϵ	GS γ	GS ι	GS κ	GS θ	Hyp	KDV	KS	SH	Burgers	Average
P3D- <i>S</i>	8.3	15.7	29.3	14.1	13.4	47.4	12.9	16.0	6.3	73.4	92.0	18.9	91.4	27.2	33.3
P3D- <i>B</i>	7.1	9.0	14.9	8.1	7.7	19.9	9.1	10.6	4.7	75.3	75.3	19.0	75.6	16.8	25.2
P3D- <i>L</i>	6.9	5.7	8.9	5.7	5.4	16.0	6.0	7.1	2.3	69.1	65.8	18.4	62.7	11.1	20.8
AFNO	13.0	20.7	30.4	18.8	23.3	43.3	17.1	20.0	6.1	146.5	81.5	32.4	183.5	34.6	47.9
AVIT	124.6	49.7	60.8	51.7	62.7	101.0	47.6	48.1	56.7	619.2	191.2	65.1	470.7	177.0	151.9
Swin3D	11.7	24.3	30.7	22.0	23.0	50.0	22.8	17.5	5.2	153.6	90.1	27.0	190.8	36.7	50.4
UNetGenCFD	17.3	35.8	19.5	6.4	24.9	50.7	6.4	6.4	6.4	214.2	107.7	43.2	573.1	45.8	82.7

Table B4: Normalized RMSE ($\times 10^{-3}$) for crop size 64.

Model Name	Fisher	GS α	GS β	GS δ	GS ϵ	GS γ	GS ι	GS κ	GS θ	Hyp	KDV	KS	SH	Burgers	Average
P3D- <i>S</i>	10.4	21.0	25.4	5.2	18.1	43.3	4.6	4.6	4.6	92.2	100.4	50.5	111.8	34.4	37.6
P3D- <i>B</i>	8.3	11.7	10.9	4.2	10.1	26.1	3.8	3.8	3.8	83.6	92.2	42.1	98.7	26.2	30.4
P3D- <i>L</i>	6.8	7.3	10.1	7.1	6.6	17.5	8.8	8.3	1.4	69.6	73.7	28.5	88.1	15.0	24.9
AFNO	19.7	20.0	22.7	3.1	21.0	38.0	2.3	2.3	2.3	189.4	90.8	39.4	210.7	36.3	49.9
AVIT	175.6	81.9	53.6	19.5	88.3	106.0	19.0	19.0	19.0	2004.7	206.6	92.7	450.8	213.9	253.6
Swin3D	20.7	67.4	56.9	5.0	59.7	98.4	4.5	4.5	4.5	155.6	115.9	61.2	285.3	46.2	70.4
FactFormer	19.06	18.01	13.94	3.75	16.10	30.76	3.75	3.75	3.75	182.78	95.13	31.53	184.20	40.48	46.2
UNetGenCFD	36.6	35.1	72.7	14.0	42.5	65.5	13.7	13.6	13.7	272.4	128.4	37.0	324.4	56.2	80.4
UNetConvNeXt	18.6	40.3	49.8	6.8	36.2	66.8	6.5	6.5	6.5	190.8	108.8	42.1	364.5	49.0	70.9
TFNO	23.2	111.6	56.2	6.3	49.5	96.7	6.0	6.0	6.0	175.5	113.8	222.1	263.1	36.2	83.7

Table B5: Normalized RMSE ($\times 10^{-3}$) for crop size 32.

Model Name	Fisher	GS α	GS β	GS δ	GS ϵ	GS γ	GS ι	GS κ	GS θ	Hyp	KDV	KS	SH	Burgers	Average
P3D- <i>S</i>	9.2	49.7	56.3	6.3	33.7	57.2	6.3	6.3	6.3	146.5	127.3	158.4	176.2	38.8	62.7
P3D- <i>B</i>	8.3	24.5	14.8	2.5	15.9	40.1	2.5	2.5	2.5	131.1	126.4	117.5	131.7	37.9	46.9
P3D- <i>L</i>	7.9	18.5	11.4	1.9	14.2	32.4	1.9	1.9	1.9	109.3	110.8	110.9	129.4	27.3	41.4
AFNO	21.7	43.8	22.0	5.8	33.4	65.9	5.8	5.8	5.8	188.7	118.3	74.0	619.5	43.1	89.5
AVIT	237.1	126.7	50.7	13.6	107.6	122.7	13.6	13.6	13.6	1089.4	256.7	177.6	476.0	230.1	209.2
Swin3D	16.6	154.5	38.6	3.8	56.2	92.6	3.8	3.8	3.8	189.1	143.6	83.4	263.5	55.7	79.2
FactFormer	20.1	43.1	23.1	5.8	32.0	43.9	5.8	5.8	5.8	222.9	105.0	54.4	263.7	43.3	62.5
UNetGenCFD	18.5	38.6	22.1	3.5	20.9	53.6	3.5	3.5	3.6	196.3	105.8	37.0	517.2	41.9	76.1
UNetConvNeXt	20.3	111.6	38.3	7.4	49.2	89.9	7.4	7.4	7.4	166.3	128.5	77.4	442.5	48.9	85.9
TFNO	21.9	176.8	41.2	5.6	46.4	84.4	5.6	5.6	5.6	175.3	121.8	204.1	252.4	38.3	84.7

1998 C EXPERIMENT 2: ISOTROPIC TURBULENCE

1999
2000 For the second experiment, we make use of the Johns Hopkins Turbulence Database (*JHTDB*). It
2001 contains data from various direct numerical simulations of homogeneous and wall-bounded turbulent
2002 flows (Perlman et al., 2007). The simulations are stored with space-time history and allow for arbitrary
2003 spatiotemporal query points.

2004
2005 **Isotropic Turbulence (iso-turb)** a direct numerical simulation of the Navier-Stokes equations
2006 at Reynolds number around 433 simulated on a 1024^3 periodic grid via a pseudo-spectral parallel
2007 code. It contains homogeneous isotropic turbulence, i.e., an idealized version of realistic turbulence
2008 with statistical properties that are invariant to translations and rotations of the coordinate axes. The
2009 following overview summarizes key characteristics of the dataset (for further details see Perlman
2010 et al., 2007):

- 2011 • Dimensionality: $s = 1, t = 500, f = 4, x = 512, y = 512, z = 512$
- 2012 • Boundary conditions: periodic
- 2013 • Time step of stored data: 0.002
- 2014 • Spatial domain size: $[0, 2\pi] \times [0, 2\pi] \times [0, 2\pi]$
- 2015 • Fields: velocity X/Y/Z, pressure
- 2016 • Validation set: random 15% split of all timesteps from $t \in [0, 420]$
- 2017 • Test set: all sequences from $t \in [420, 500]$

2018 C.1 EXTENDED EVALUATION

2019 See Table C1 for an extended evaluation expanding upon Table 3. The B config corresponds to the
2020 previous size of the baselines models for Swin3D, AViT and AFNO. The S config decreases the
2021 hidden dimensionality of the model architecture, which we change from 768 (B) to 384 (S) for AViT
2022 and AFNO. For Swin3D, we decrease it from 96 (B) to 48 (S). For Swin3D and AViT, the hidden
2023 dimensionality corresponds to the token embedding dimension. We also trained an L config for the
2024 AViT model with hidden dimensionality 1536. All baseline models were trained for 4000 epochs.
2025 Additionally, we kept training P3D- S and P3D- B for a total of 20000 epochs. The evaluation shows
2026 that the RMSE and spectral error keep improving, showing no indication of overfitting.

2027 We also include a benchmark comparison regarding the number of parameters, GFLOPs, VRAM and
2028 throughput for inference of P3D and the different baseline architectures in Table C2.

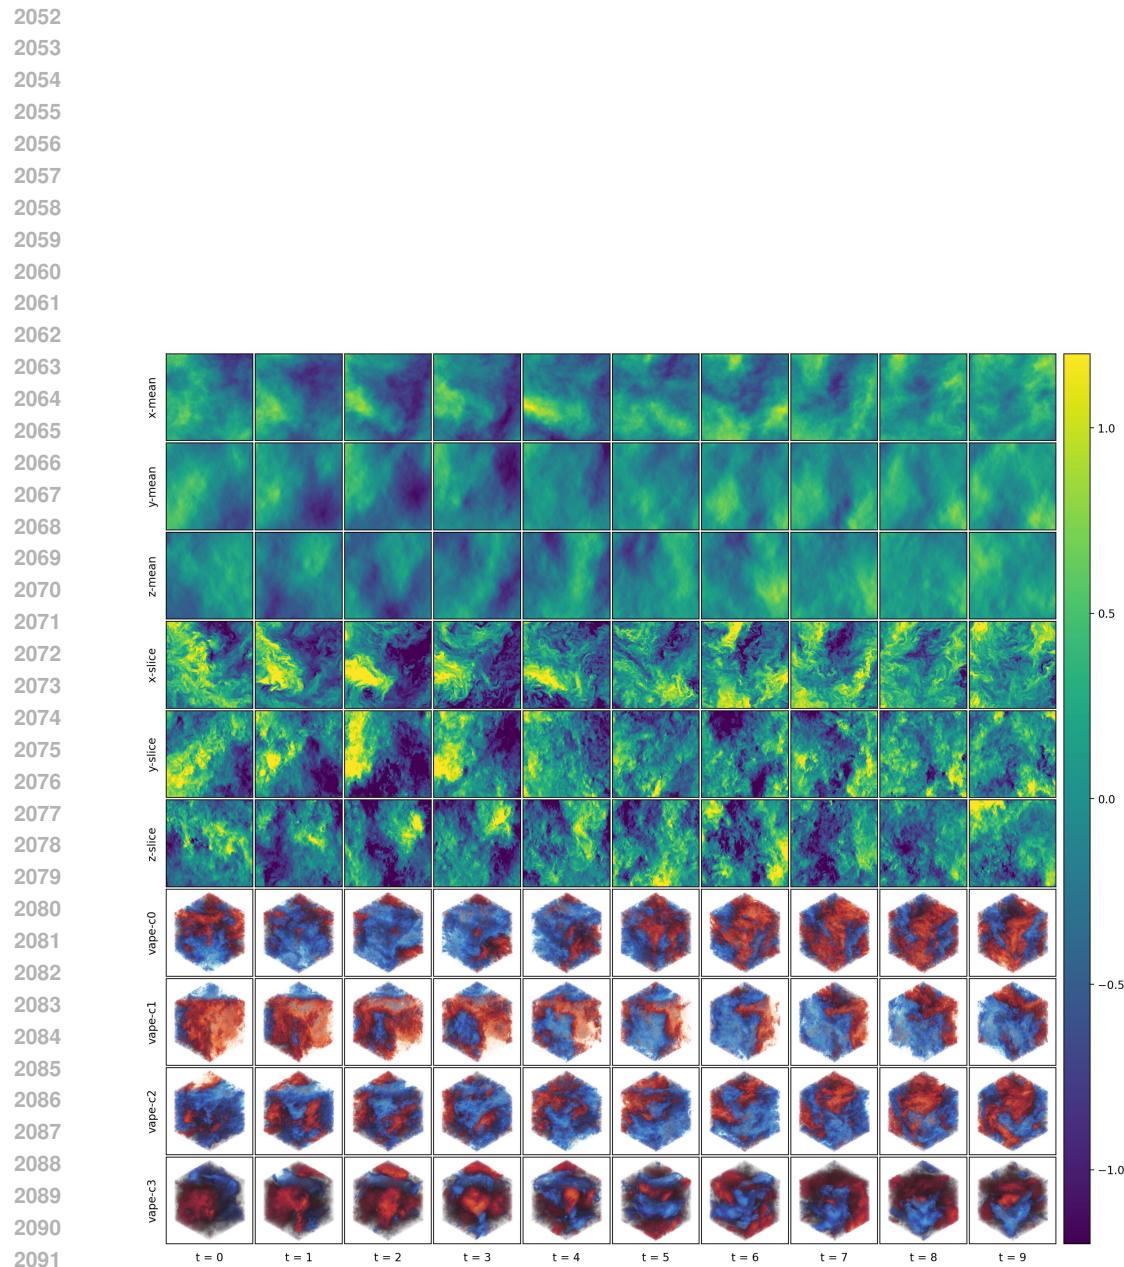


Figure C1: Isotropic Turbulence. Training dataset visualization at resolution 512^3 showing the velocity X/Y/Z and pressure from $t = 0$ until $t = 420$. The first six rows are 2D mean projections and slices of the velocity in x -direction. The 3D renderings show velocity X/Y/Z and pressure.

Table C1: Performance comparison on the test set with crop size 128^3 for RMSE ($\times 10^{-2}$) and the L2 enstrophy graph error ($\times 10^2$) at different autoregressive rollout steps.

Model	RMSE				L2 Enstrophy				epochs
	1	5	10	15	1	5	10	15	
UNet _{GenCFD}	5.48	25.42	48.60	67.72	4.25	14.5	22.7	140	4000
Swin3D- <i>S</i>	4.04	11.53	18.66	24.30	6.28	59.7	120	158	4000
Swin3D- <i>B</i>	3.22	10.45	18.15	24.36	4.06	52.1	112	156	4000
AViT- <i>S</i>	10.09	21.47	32.60	40.36	31.60	44.1	99.2	144	4000
AViT- <i>B</i>	9.45	19.57	30.00	37.77	26.70	49.1	84.6	112	4000
AViT- <i>L</i>	9.18	18.80	29.76	39.20	24.20	54.6	96.8	129	4000
AFNO- <i>S</i>	4.25	14.80	23.58	30.40	8.30	99.9	171	196	4000
AFNO- <i>B</i>	3.69	13.33	23.52	29.80	7.69	88.7	158	190	4000
P3D- <i>S</i>	2.81	9.87	20.50	28.25	2.15	8.23	21.6	31.9	4000
	2.17	8.99	19.40	27.40	1.29	6.68	16.0	25.3	20000
P3D- <i>B</i>	2.04	8.79	20.23	31.52	0.72	1.39	3.38	19.2	4000
	1.54	8.11	21.09	44.92	0.21	0.71	3.49	14.7	20000

Table C2: Architecture benchmark comparison on crop size 128^3 measured on a H100 GPU using CUDA 12.8. VRAM is measured at inference with batch size 1. Throughput with batch size 16.

Model	Params	GFLOPs	VRAM	Throughput
Swin3D- <i>S</i>	18.9M	79.6	2.2G	51.2it/s
Swin3D- <i>B</i>	50.3M	289.6	2.9G	34.1it/s
AViT- <i>S</i>	15.1M	48.8	0.16G	859it/s
AViT- <i>B</i>	60.0M	143.0	0.3G	346it/s
AViT- <i>L</i>	239.2M	467.0	1.1G	128it/s
AFNO- <i>S</i>	17.0M	942.2	0.2G	167it/s
AFNO- <i>B</i>	64.1M	2116.4	0.4G	58.6it/s
P3D- <i>S</i>	11.2M	238.7	0.9G	88.9it/s
P3D- <i>B</i>	46.1M	2330.6	2.1G	30.9it/s

2160 C.2 VORTICITY
2161

2162 The vorticity ω , is a measure of the local rotation in a fluid flow. It is defined as the curl of the velocity
2163 vector field $\mathbf{u} = (u_x, u_y, u_z)$, where u_x, u_y and u_z are the velocity fields in x, y and z direction. The
2164 vorticity is defined as

$$2165 \omega = \nabla \times \mathbf{u} = \left(\frac{\partial u_z}{\partial y} - \frac{\partial u_y}{\partial z} \right) e_1 + \left(\frac{\partial u_x}{\partial z} - \frac{\partial u_z}{\partial x} \right) e_2 + \left(\frac{\partial u_y}{\partial x} - \frac{\partial u_x}{\partial y} \right) e_3, \quad (9)$$

2166 where e_1, e_2 and e_3 represent the unit vectors in x, y and z direction.
2167

2168 The vorticity is not part of the data. Only the velocity channels in X/Y/Z are available.
2169

2170 **Approximation by finite differences** To compute the components of the vorticity numerically
2171 from the velocity, we use finite difference approximations for the partial derivatives in Equation (9).
2172

2173 For a grid point (i, j, k) located in the interior of the discretized simulation domain, we approximate
2174 the partial derivatives as
2175

$$\begin{aligned} 2176 \frac{\partial u_x}{\partial x} &\approx \frac{u_x^{i+1,j,k} - u_x^{i-1,j,k}}{2\Delta x} \\ 2177 \frac{\partial u_x}{\partial y} &\approx \frac{u_x^{i,j+1,k} - u_x^{i,j-1,k}}{2\Delta y} \\ 2178 \frac{\partial u_x}{\partial z} &\approx \frac{u_x^{i,j,k+1} - u_x^{i,j,k-1}}{2\Delta z} \end{aligned}$$

2183 and analogously for u_y and u_z .
2184

2185 C.3 ENSTROPHY GRAPH
2186

2187 We consider the magnitude $|\omega|$ of the vorticity ω , which is a scalar field defined as
2188

$$2189 |\omega| = \sqrt{\sum_{i=1}^n |\omega_i|^2}. \quad (10)$$

2190 The enstrophy of the magnitude of the vorticity field $|\omega|$ at wavenumber $k \in \mathbb{R}_+$ is defined as
2191

$$2193 \text{Enstrophy}(k) = \sum_{k \leq |m| \leq k+1} |\hat{\omega}(m)|^2, \quad (11)$$

2195 where $\hat{\omega}$ with $m \in \mathbb{Z}^2$ are the Fourier coefficients of $|\omega|$, see Chen et al. (2024, C.2) for reference.
2196 The enstrophy spectrum is the graph of the function $k \mapsto \text{Enstrophy}(k)$. We average the enstrophy
2197 spectrum over different vorticity fields.
2198

2199 We compute the L2 enstrophy error between the reference enstrophy spectrum and the enstrophy
2200 spectrum of generated vorticities as
2201

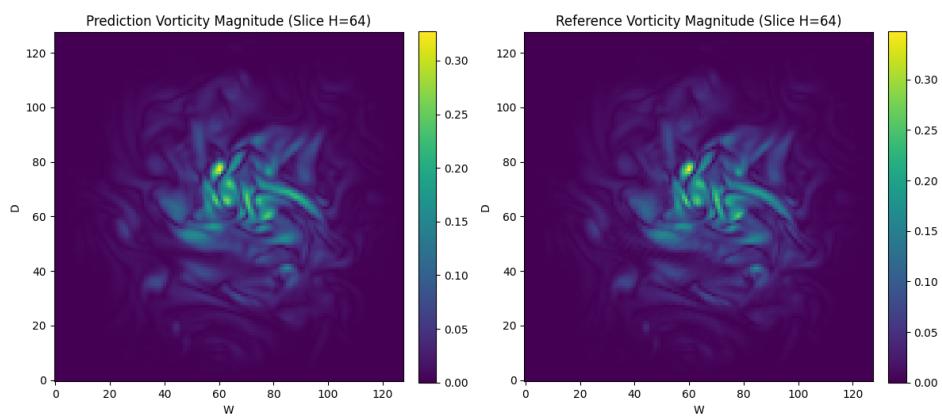
$$2201 L_{\text{Enstrophy}}^2 := \frac{1}{K+1} \sum_{0 \leq k \leq K} \left[\text{Enstrophy}_{\text{reference}}^{\text{avg}}(k) - \text{Enstrophy}_{\text{generated}}^{\text{avg}}(k) \right]^2 \quad (12)$$

2203 with $K = 16$.
2204

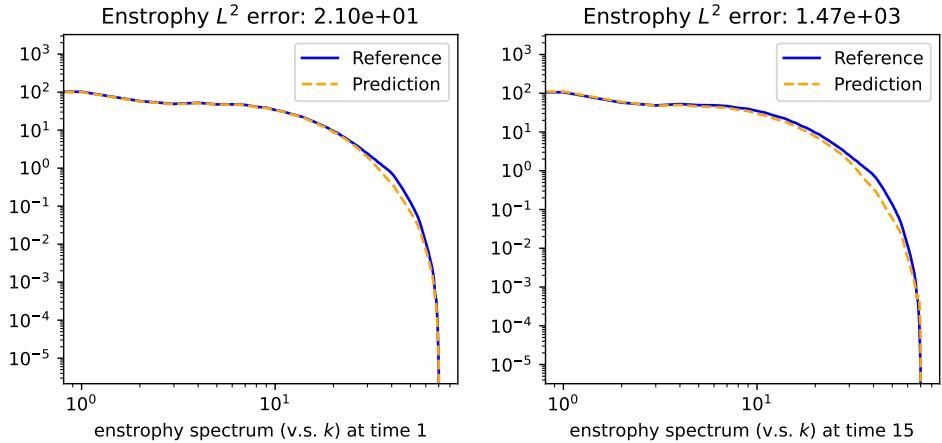
2205 **Hanning filter** The data on the cropped domain of size 128^3 is not periodic, thus there are artifacts
2206 at the boundary of the crop. This leads to problems when calculating the Fourier coefficients $\hat{\omega}$.
2207 We therefore smoothen $|\omega|$ towards the boundary by multiplying the data with the Hann window as
2208 defined in `torch.hann_window`.
2209

2210 In Figure C2 we show the magnitude of the vorticity $|\omega|$ generated by P3D-*B* and the reference after
2211 applying the Hanning filter ($t = 1$ autoregressive steps).
2212

2213 See Figure C3 for the enstrophy spectrum $\text{Enstrophy}_{\text{generated}}^{\text{avg}}$ of the predicted vorticities by P3D-*B*
2214 and the enstrophy spectrum $\text{Enstrophy}_{\text{reference}}^{\text{avg}}$ of the reference for autoregressive unrolling steps
2215 $t = 1$ and $t = 15$ on the test set.



2230 Figure C2: Generated vorticity magnitude $|\omega|$ by P3D-*B* (left) and reference (right) after applying
 2231 the Hanning filter.



2248 Figure C3: Enstrophy spectrum of P3D-*B* and the reference simulation at resolution 128^3 for $t = 1$
 2249 and $t = 15$ autoregressive prediction steps.

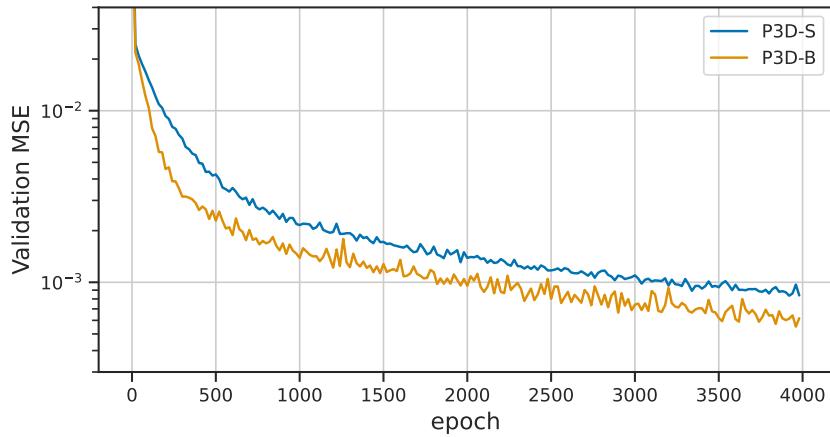
C.4 TRAINING AND EVALUATION

2250
 2251
 2252
 2253
 2254
 2255
 2256
 2257
 2258
 2259
 2260
 2261
 2262
 2263
 2264
 2265
 2266
 2267

The model weights for the evaluation in Table 4 are the EMA weights at epoch 2000. We train both the *S* and *B* configurations of P3D on crop size 128^3 . Training P3D-*S* for 4000 epochs took 11h 48m and training P3D-*B* took 20h 25m on four A100 GPUs. See Figure C4 for the validation loss curve.

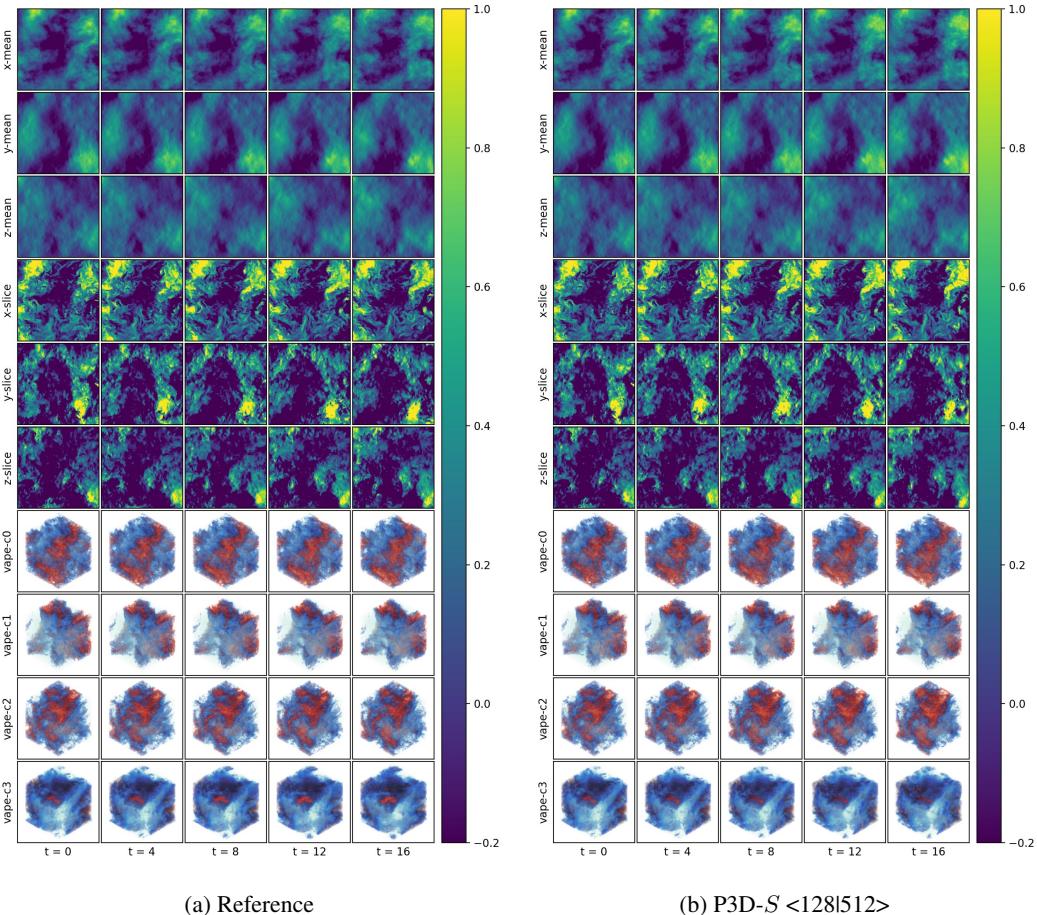
In Figure C5 we show a comparison between the reference and P3D-*S* $<128|512>$ with an autoregressive rollout until $t = 16$ on the test set at resolution 512^3 .

2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283



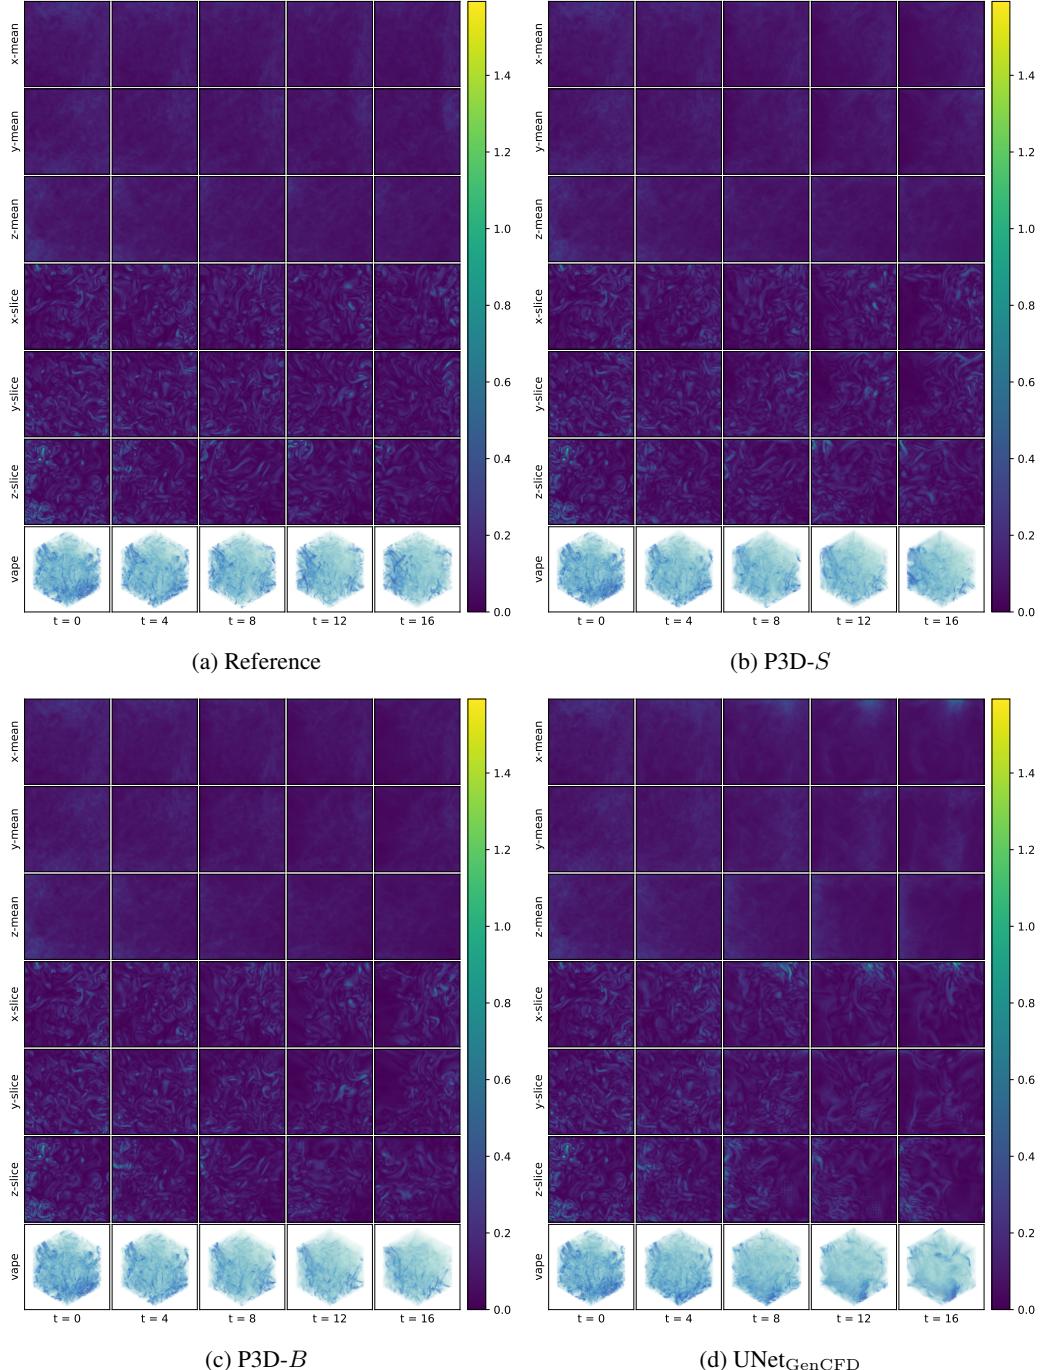
2284 Figure C4: Isotropic Turbulence. Validation MSE for P3D-*S* and P3D-*B* during training.
2285
2286
2287
2288

2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321



2319 Figure C5: Isotropic Turbulence. Autoregressive prediction for $t = 16$ steps with P3D-*S* <128|512>
2320 on the test set at resolution 512^3 . The first six rows are 2D mean projections and slices of the velocity
2321 in x -direction. The 3D renderings show velocity X/Y/Z and pressure.
2321

2322
 2323
 2324
 2325
 2326
 2327
 2328
 2329
 2330
 2331
 2332
 2333
 2334
 2335
 2336
 2337
 2338
 2339
 2340
 2341
 2342
 2343
 2344
 2345
 2346



2369
 2370
 2371 Figure C6: Vorticity magnitude $|\omega|$ calculated from the predicted velocity X/Y/Z for different models
 2372 (1/2).

2373
 2374
 2375

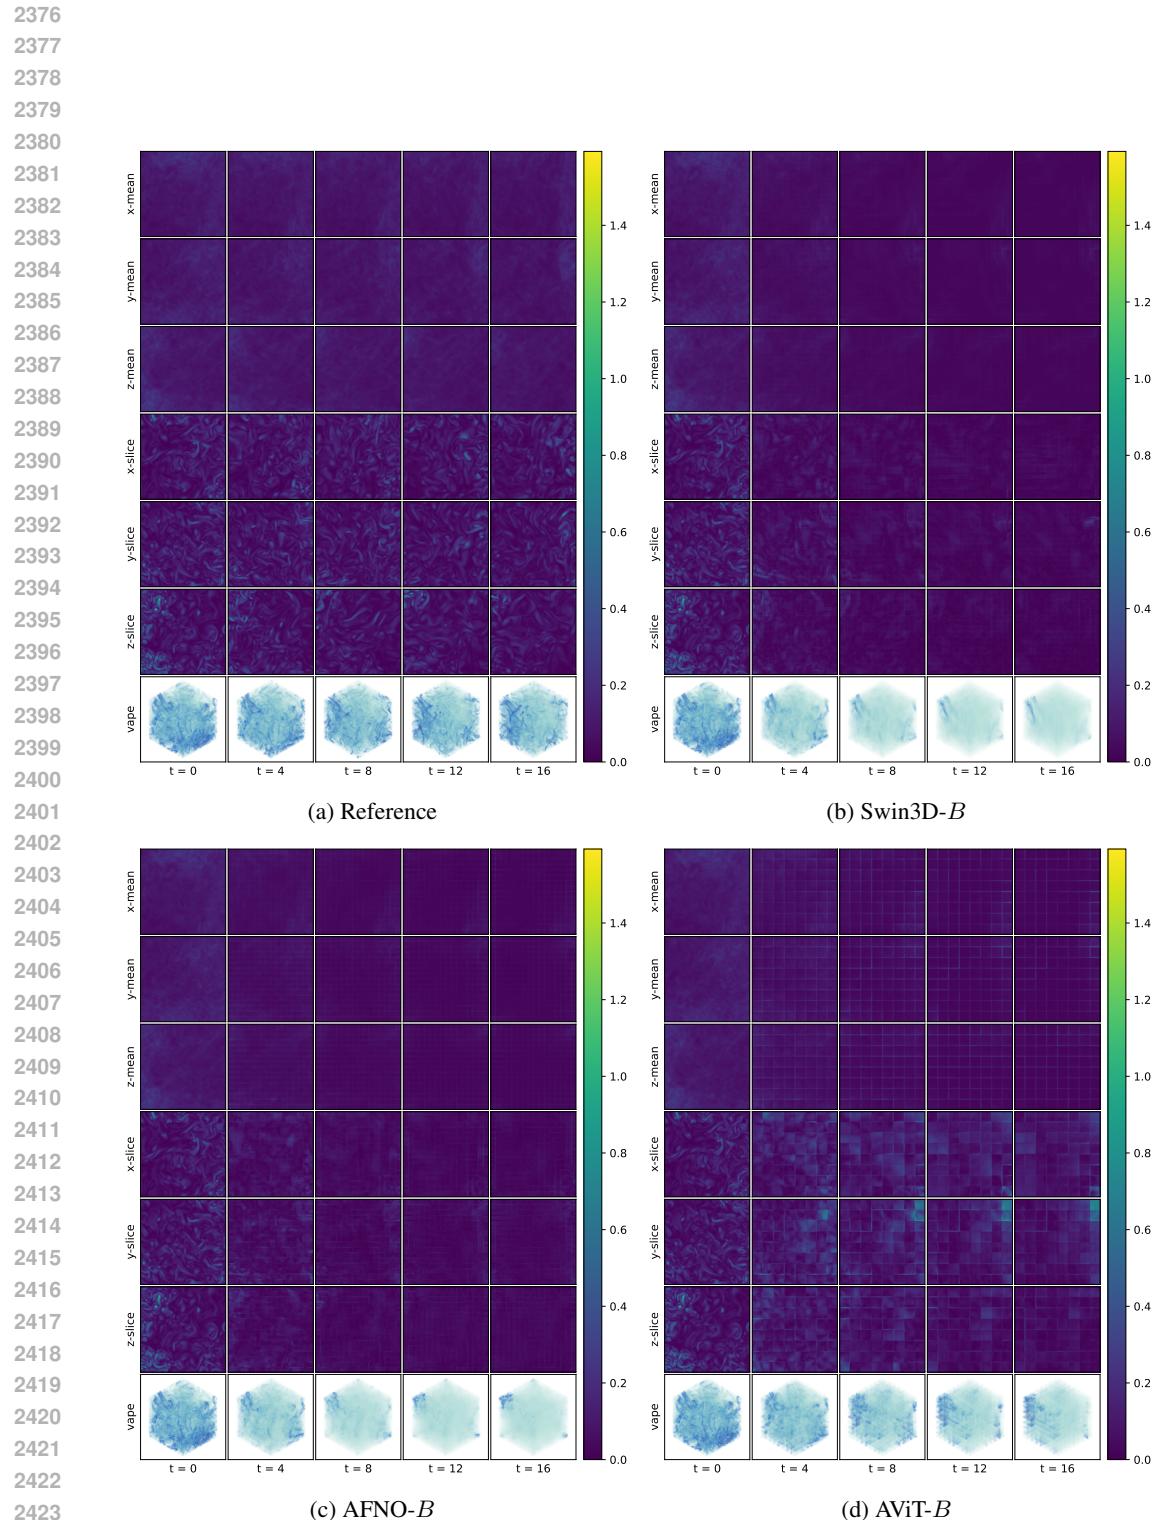


Figure C7: Vorticity magnitude $|\omega|$ calculated from the predicted velocity X/Y/Z for different models (2/2).

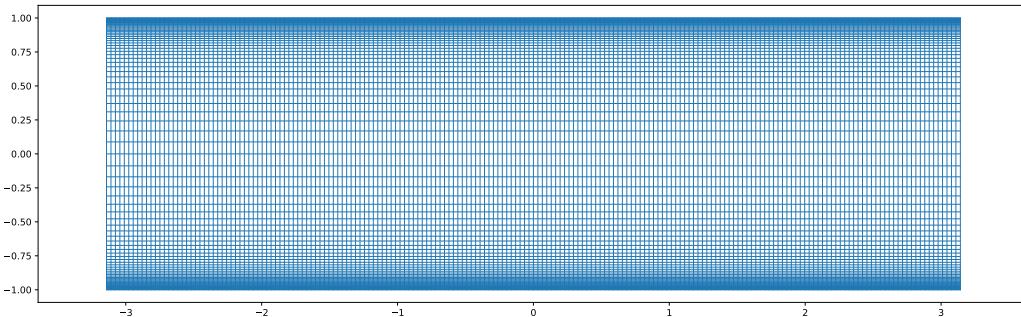
2430 **D EXPERIMENT 3: TURBULENT CHANNEL FLOW**
2431

2432 The dataset for the turbulent channel flow (TCF) represents a periodic channel with no-slip
 2433 boundaries at $\pm y$ that is driven by a dynamic forcing to re-inject energy lost due to wall friction, and
 2434 prevent the flow from slowing down. This results in a continuous production of vortex structures at
 2435 the walls, which have a very characteristic and well-studied, spatially-varying distribution (Hoyas &
 2436 Jiménez, 2008). Due to the complexity of the flow, these flows require very long transient phases to
 2437 develop the characteristic structures. We target this scenario by learning with a generative model, in
 2438 the context of which the TCF problem represents a probabilistic learning problem to infer turbulent
 2439 stats from the equilibrium phase, bypassing the costly transient warm-up phase.

2440
2441 **D.1 DATASET**
2442

2443 We generate a dataset comprising 20 simulations with Reynolds numbers within the interval $[400, 800]$
 2444 spaced equidistantly. After the initial-warmup phase, we simulate $ETT = 20$ eddy turnover cycles,
 2445 which we save in 200 snapshots with $\Delta t = 0.1$. The computational grid comprises $96 \times 96 \times 192$
 2446 spatially adaptive cells with a finer discretization near the wall. The data contains channels for the
 2447 velocity in X/Y/Z direction as well as pressure. The computational grid is shown in Figure D1.
 2448 In Figures D2 and D3 we show visualizations of the turbulent channel flow for Reynold numbers
 2449 $Re = 400$ and $Re = 640$ respectively.

- 2450 • Dimensionality: $s = 20, t = 200, f = 4, x = 96, y = 96, z = 192$
- 2451 • Initial conditions: noise
- 2452 • Boundary conditions: periodic (x), wall (y,z)
- 2453 • Time step of stored data: 0.1
- 2454 • Number of warmup steps (discarded, in time step of data storage): 200
- 2455 • Spatial domain: $[-1, 1] \times [-1, 1] \times [-\pi, \pi]$
- 2456 • Fields: velocity X/Y/Z, pressure
- 2457 • Varied parameters: Reynolds number $\in [400, 800]$
- 2458 • Validation set: random 15% split of Reynolds number



2473 Figure D1: Computational grid of the turbulent channel flow simulation. The spatial discretization is
 2474 refined in the near-wall region to resolve the boundary layer.

2475
2476
2477
2478
2479
2480
2481
2482
2483

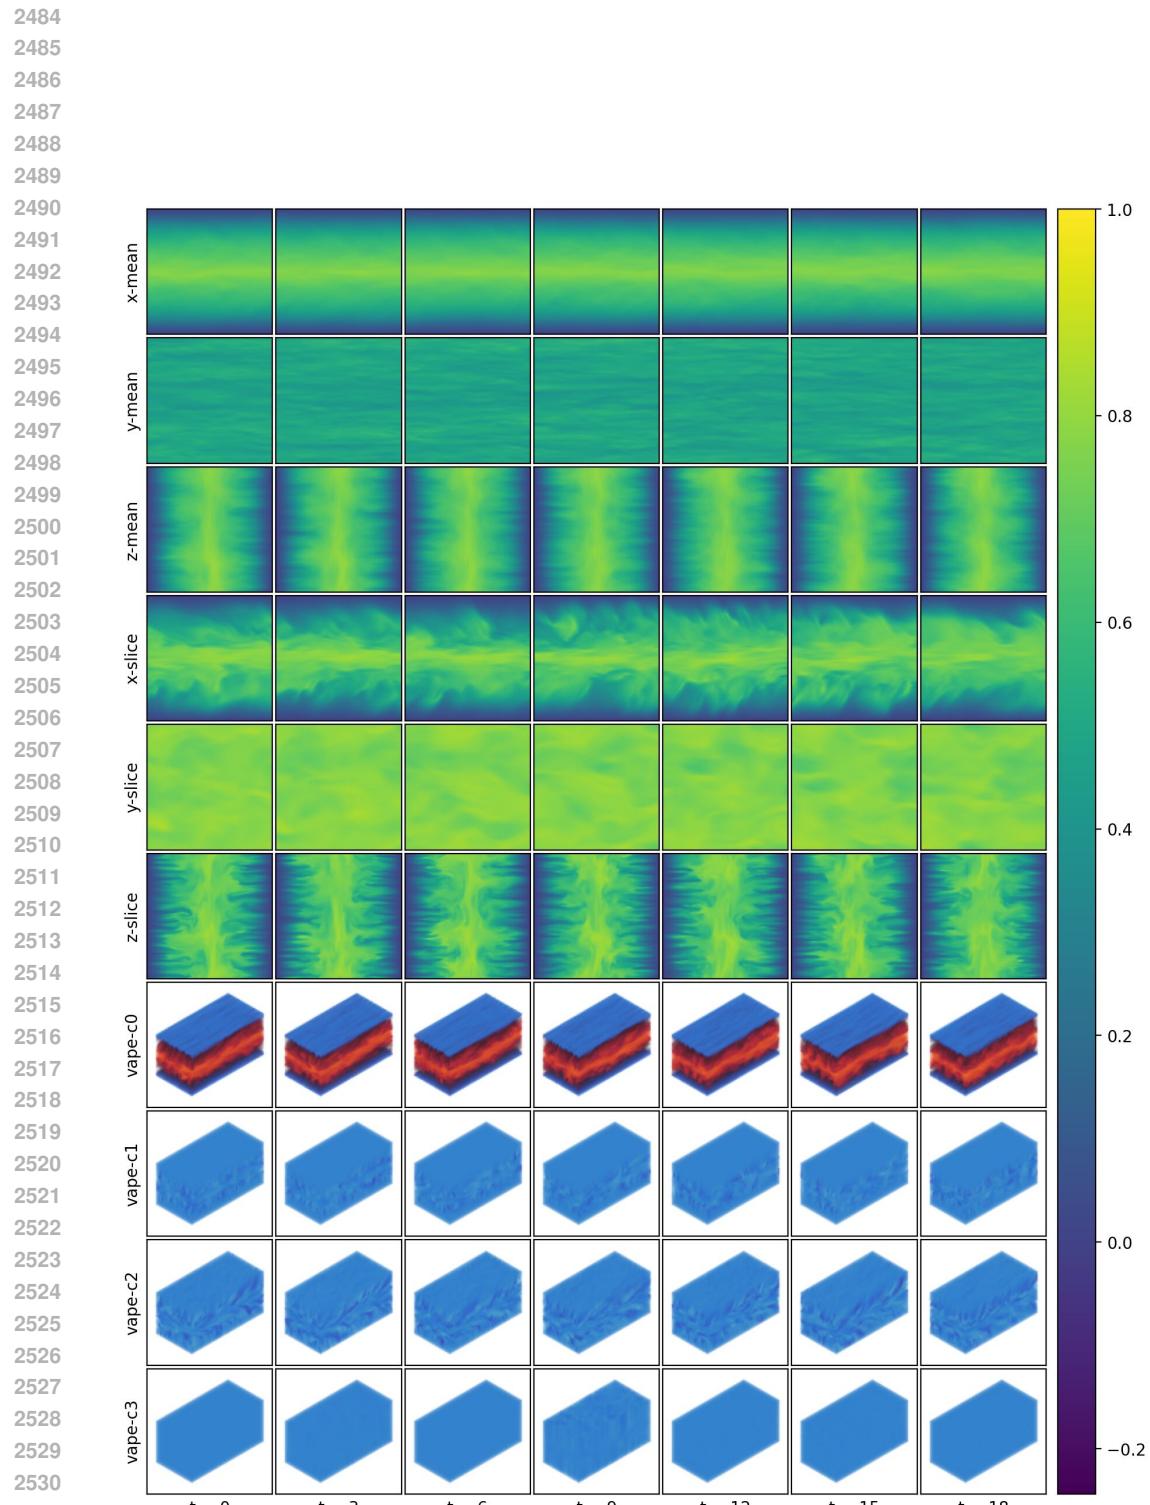


Figure D2: Turbulent channel flow with Reynolds number 400.

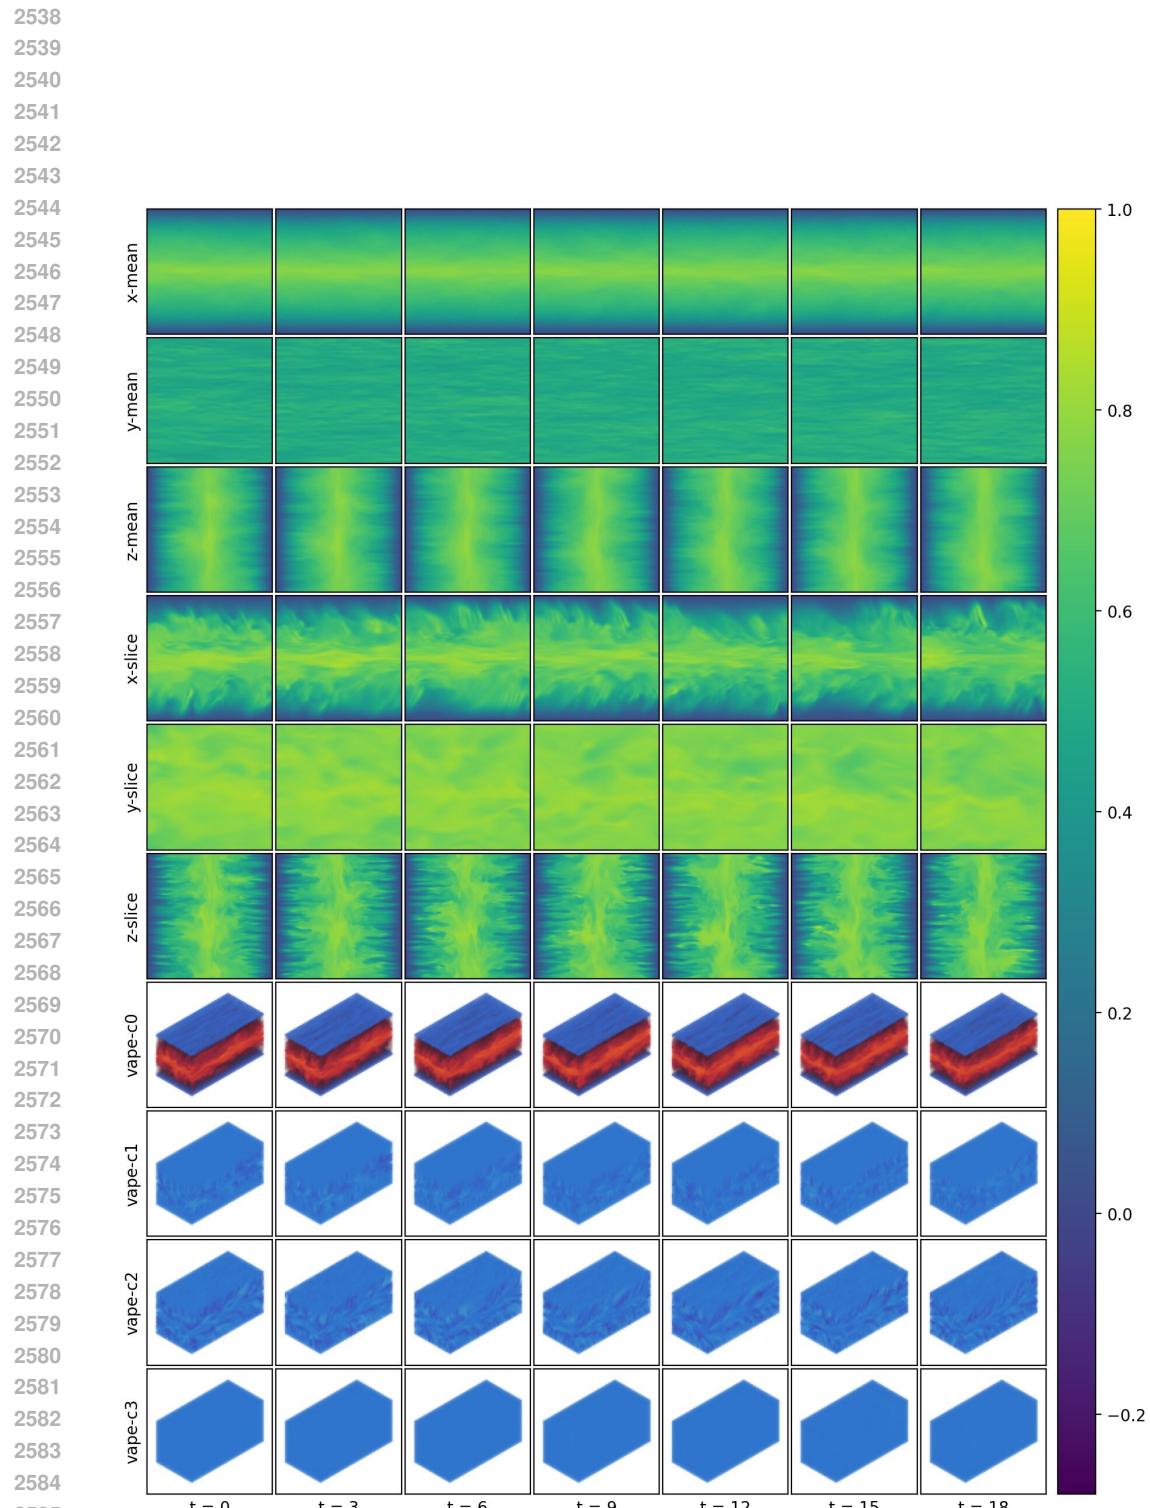
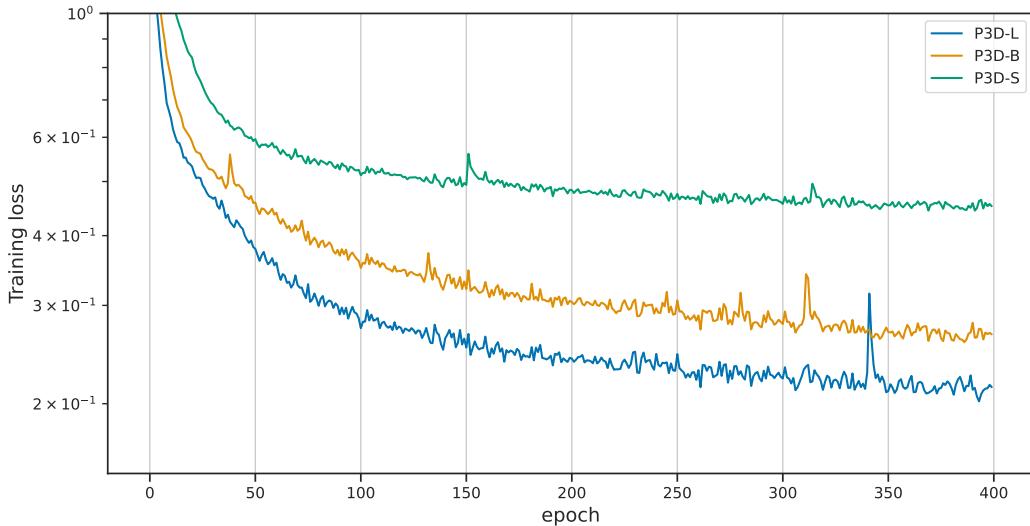


Figure D3: Turbulent channel flow with Reynolds number 640.

2592 **Training S, B and L configs** We train P3D with the configurations S , B and L for 400 epochs on
 2593 the full domain of size $96^2 \times 192$. The training loss is shown in Figure D4. It is important to choose
 2594 large architectures in our generative modeling setup based on flow matching. The network size,
 2595 specifically the embedding dimension is critical for this task with the L config reaching significantly
 2596 lower loss values compared to the S config trained with the same number of epochs. All models were
 2597 trained on 4 A100 GPUs with 80GB VRAM. Training took 11h 4m, 14h 43m and 27h 55m for the
 2598 S , B and L configs respectively. **We train P3D-L both with conditioning on the Reynolds number
 2599 enabled and with conditioning disabled. We use the former for the visualization in Figure 9 and the
 2600 latter in the comparison against baselines.**



2619 Figure D4: Turbulent channel flow. Training curve for different configurations of P3D.
 2620

2621 D.2 STATISTICAL EVALUATION

2623 Let $\mathbf{x} = (x_1, x_2, x_3)$ denote the spatial coordinates and $u(\mathbf{x}, t, \text{Re})$ denote the velocity of the flow
 2624 direction \mathbf{X} . The reference simulations reach an equilibrium phase after the initial transient phase of
 2625 the warmup. Therefore, for the reference simulations, the moments $u_m(\mathbf{x}, t, \text{Re})$ should be the same
 2626 for all t . Additionally, the setup combining periodic boundary conditions and no-slip boundaries
 2627 for the wall implies that the statistics *only* depend on the distance to the wall on the flow axis x_1 ,
 2628 i.e., $u_m((x_1, x_2, x_3), t, \text{Re})$ is the same for all x_2 and x_3 inside the domain. Thus it is reasonable
 2629 to consider $u_m(x_1, \text{Re})$ and calculate the moments by sampling over x_2, x_3, t . Since the baseline
 2630 methods cannot be conditioned on the Reynolds number Re , we also compute the velocity profile
 2631 over Re . We compute the L2 distance between mean μ ($m = 0$) and variance σ^2 ($m = 1$) of the
 2632 velocity profile graph $x_1 \mapsto u_m(x_1)$ of the time-resolved DNS reference and generated samples.
 2633 The resulting L2 distance for the mean μ and variance σ^2 is shown in Table 5, which provides a
 2634 meaningful statistical metric to evaluate accuracy. We additionally train two baselines, AFNO and
 2635 UNetGenCFD with identical training setups as P3D-L on the full domain.

2636 **Moments of the flow field** We included comparisons of the first three moments (mean, variance,
 2637 skewness) u_m of the flow direction (velocity in x -direction) averaging over x_1 as well. We report
 2638 the standard deviation when estimating the moments of the reference based on randomly drawing 20
 2639 simulations states of the equilibrium phase per Reynolds number as done for the velocity profiles to
 2640 properly assess how close the generated samples should match the reference. For the finetuned P3D-L
 2641 <X48|X48>, we picked the best model out of the five finetuned models. See Table D1. Overall,
 2642 P3D-L trained on the full domain and P3D-L <X48|X48> achieve the best results.
 2643

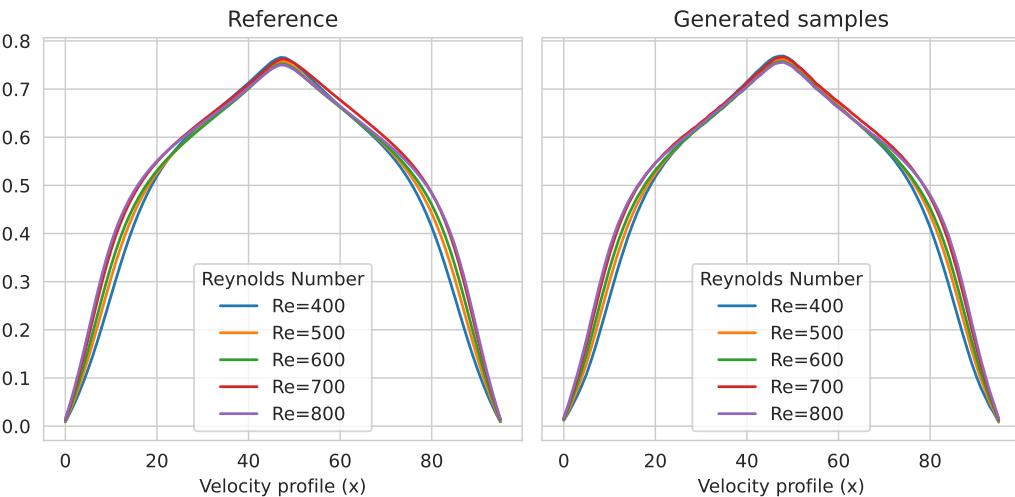
2644 **Velocity profile comparison** See Figure D5 for a comparison between the mean channel flow
 2645 of the reference simulation for different Reynolds number and the mean flow from the generative
 2646 model P3D-L trained on the full domain. We show additional comparisons between the different

2646 Table D1: First three moments of the velocity field in the flow direction for the reference, P3D-*L* and
 2647 the two baseline methods.

Moment	Mean	Variance	Skewness
Reference	0.5034±0.0007	0.0511±0.0001	-0.776±0.007
UNet _{GenCFD} full domain	0.5002	0.0532	-0.723
AFNO full domain	0.5040	0.0930	-0.361
P3D- <i>L</i> full domain	0.5009	0.0513	-0.789
P3D- <i>L</i> <X48 X48>	0.5044	0.0510	-0.802

2657 training and inference strategies in Figure D6. Scaling the P3D-*L* network trained on the small
 2658 crops of size 48^3 to the full domain does not work well and results in incorrect velocity profiles.
 2659 By finetuning with the context network, region crops can coordinate and obtain information about
 2660 their relative position to each other as well as to the wall. As a results, the flow statistics improve
 2661 significantly, more closely matching the reference and samples from P3D-*L* trained on the full domain.

2662 We show samples from P3D-*L* trained on the full domain in Figure D7, when applying
 2663 P3D-*L* pretrained on the full domain without any finetuning in Figure D8 and with finetuning via the
 2664 context network and learned region-dependent conditioning in Figure D9 respectively.



2683 Figure D5: Comparison of the mean channel flow of the reference and of generated samples from
 2684 P3D-*L* trained on the full domain.

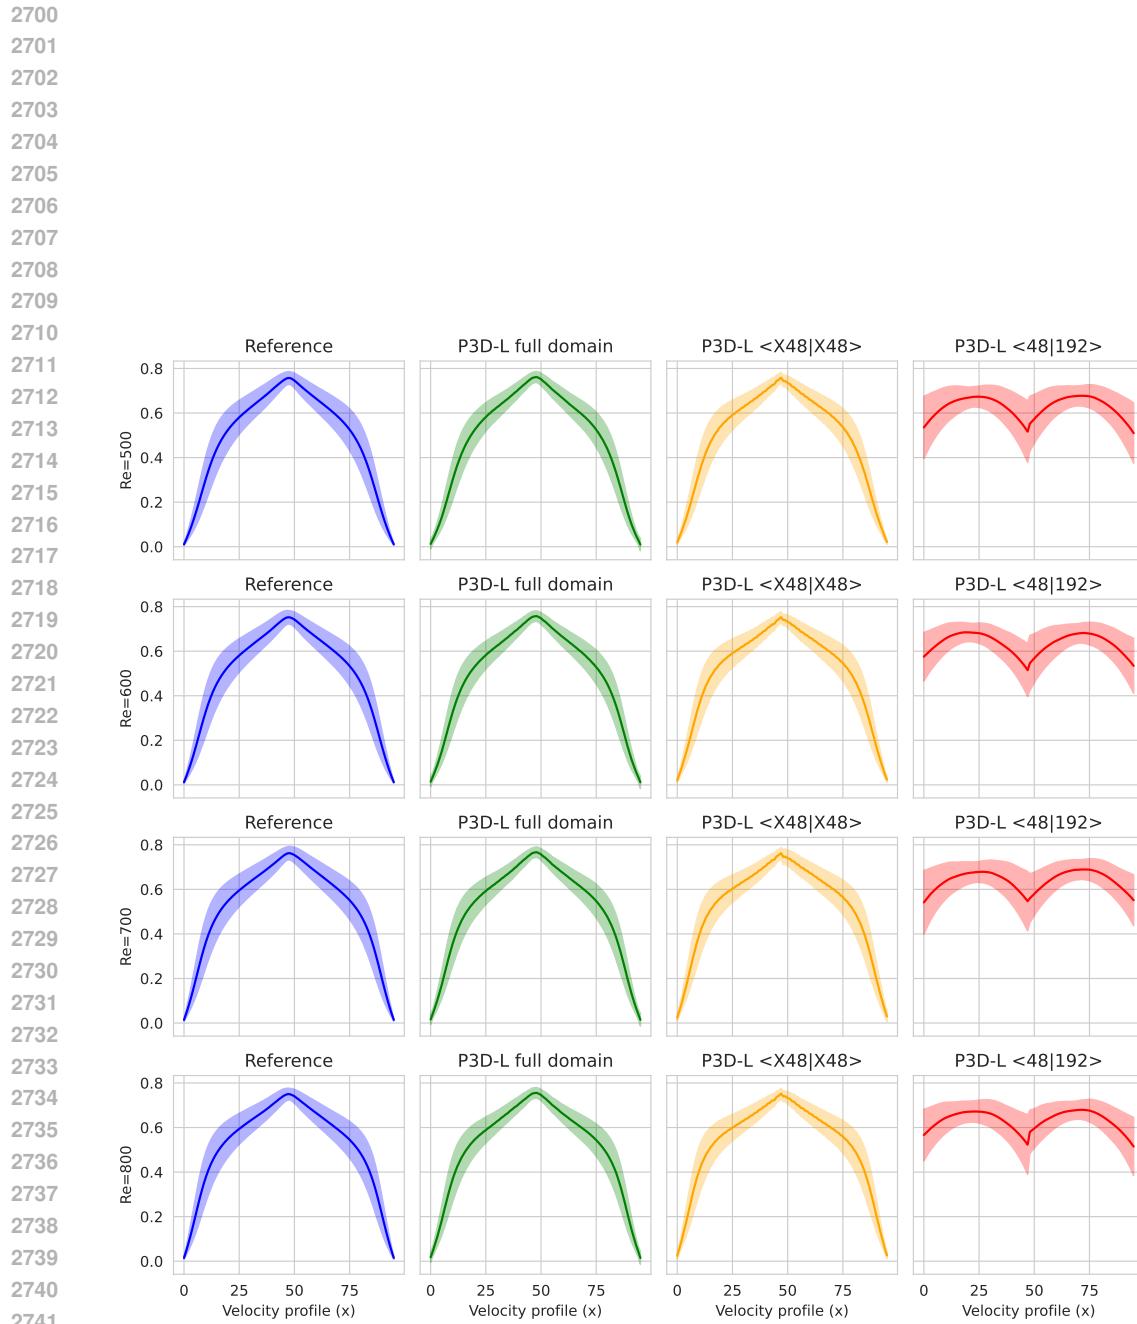
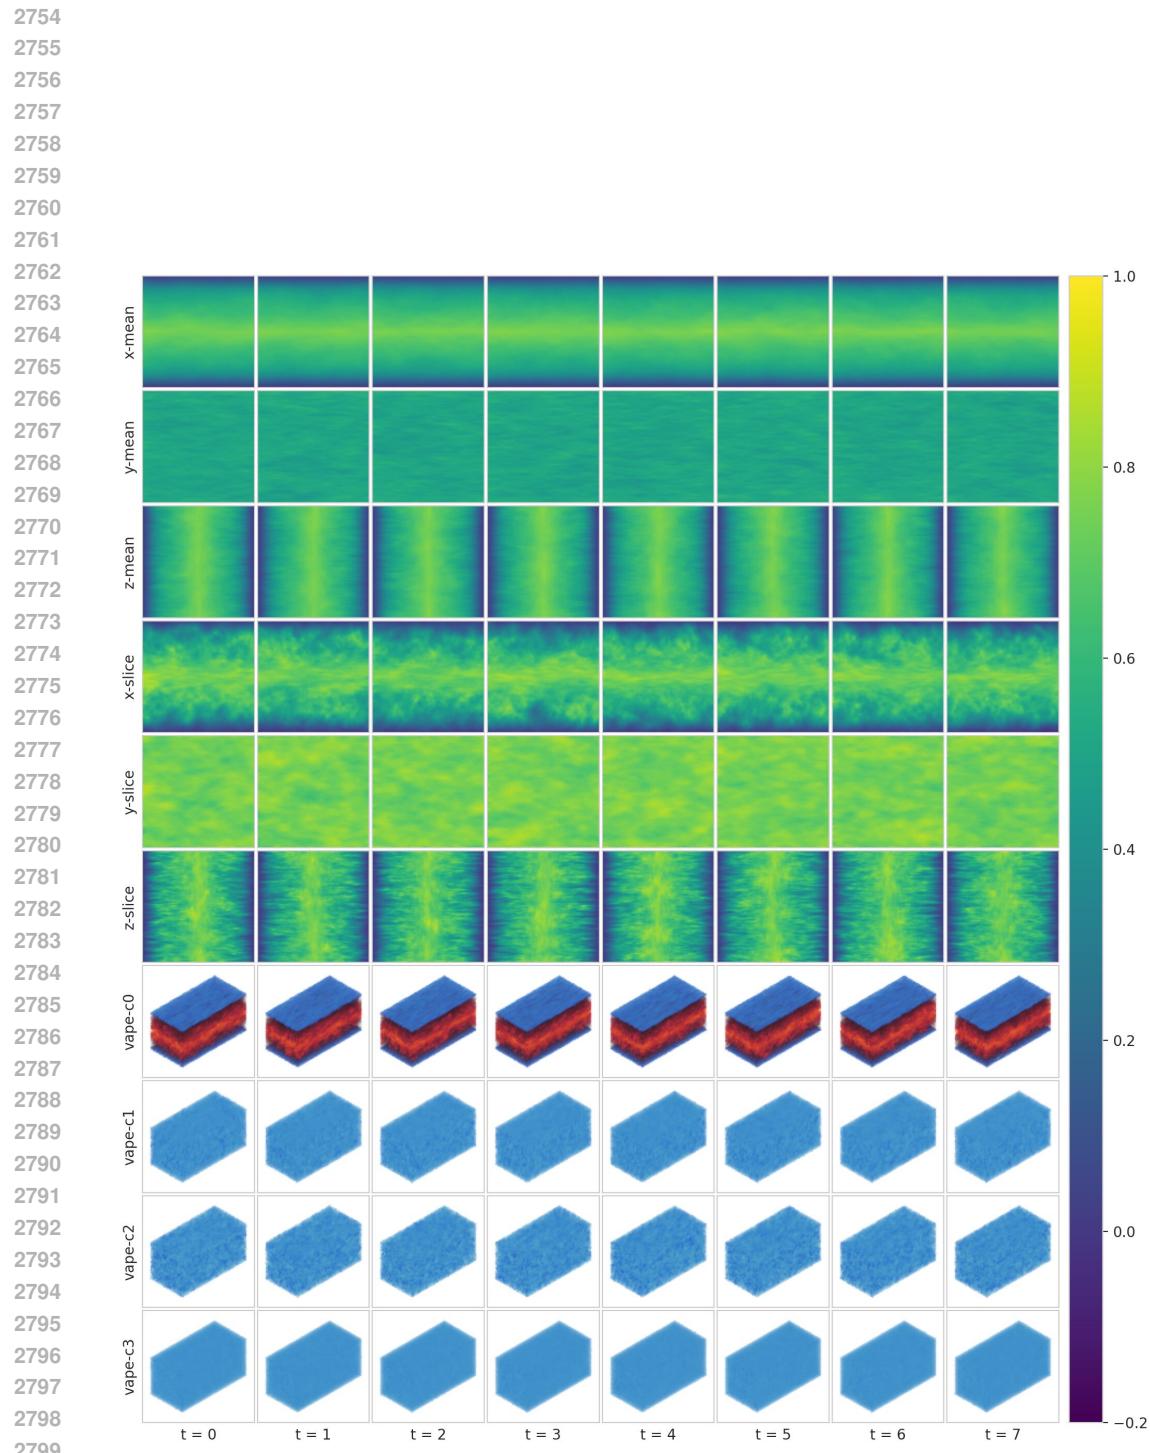
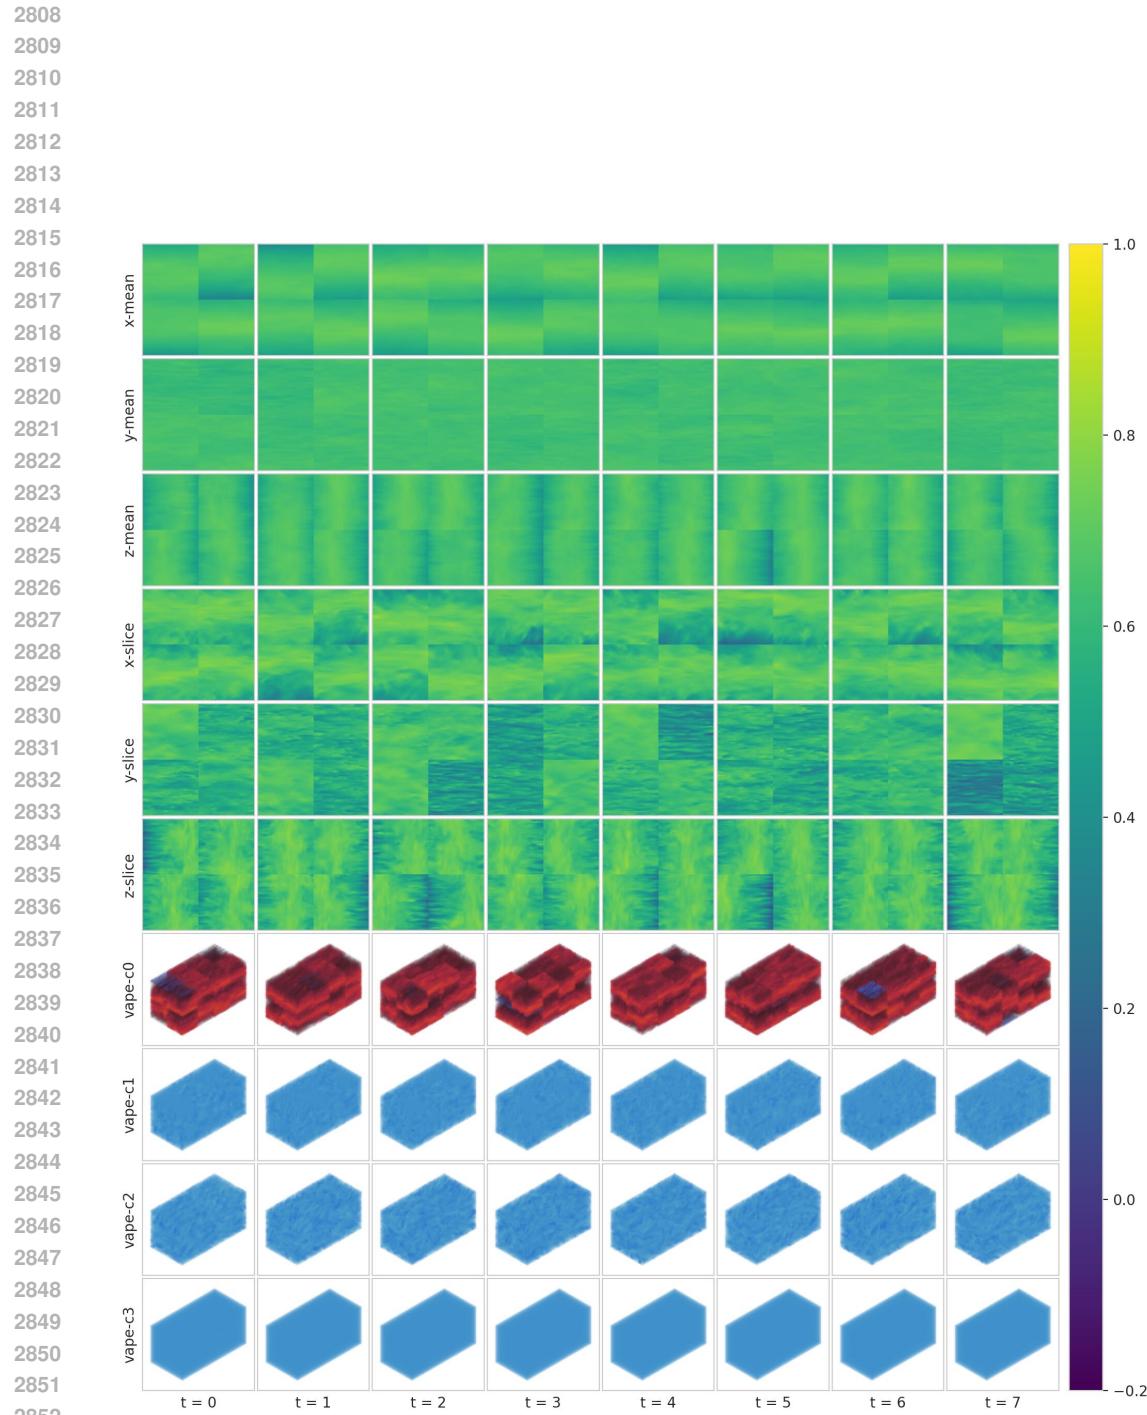


Figure D6: Comparison of the first two moments of the velocity profile between the reference DNS and generated samples from P3D-L with different training and inference strategies.

Figure D7: Samples from P3D-L trained on the full domain at $Re = 800$.



2853 Figure D8: Samples from P3D-L $<48|192>$ pretrained on crops of size 48^3 . Inference on the full
2854 domain at $Re = 800$ produces incorrect samples, as information on the relative positions between
2855 region crops is not available.

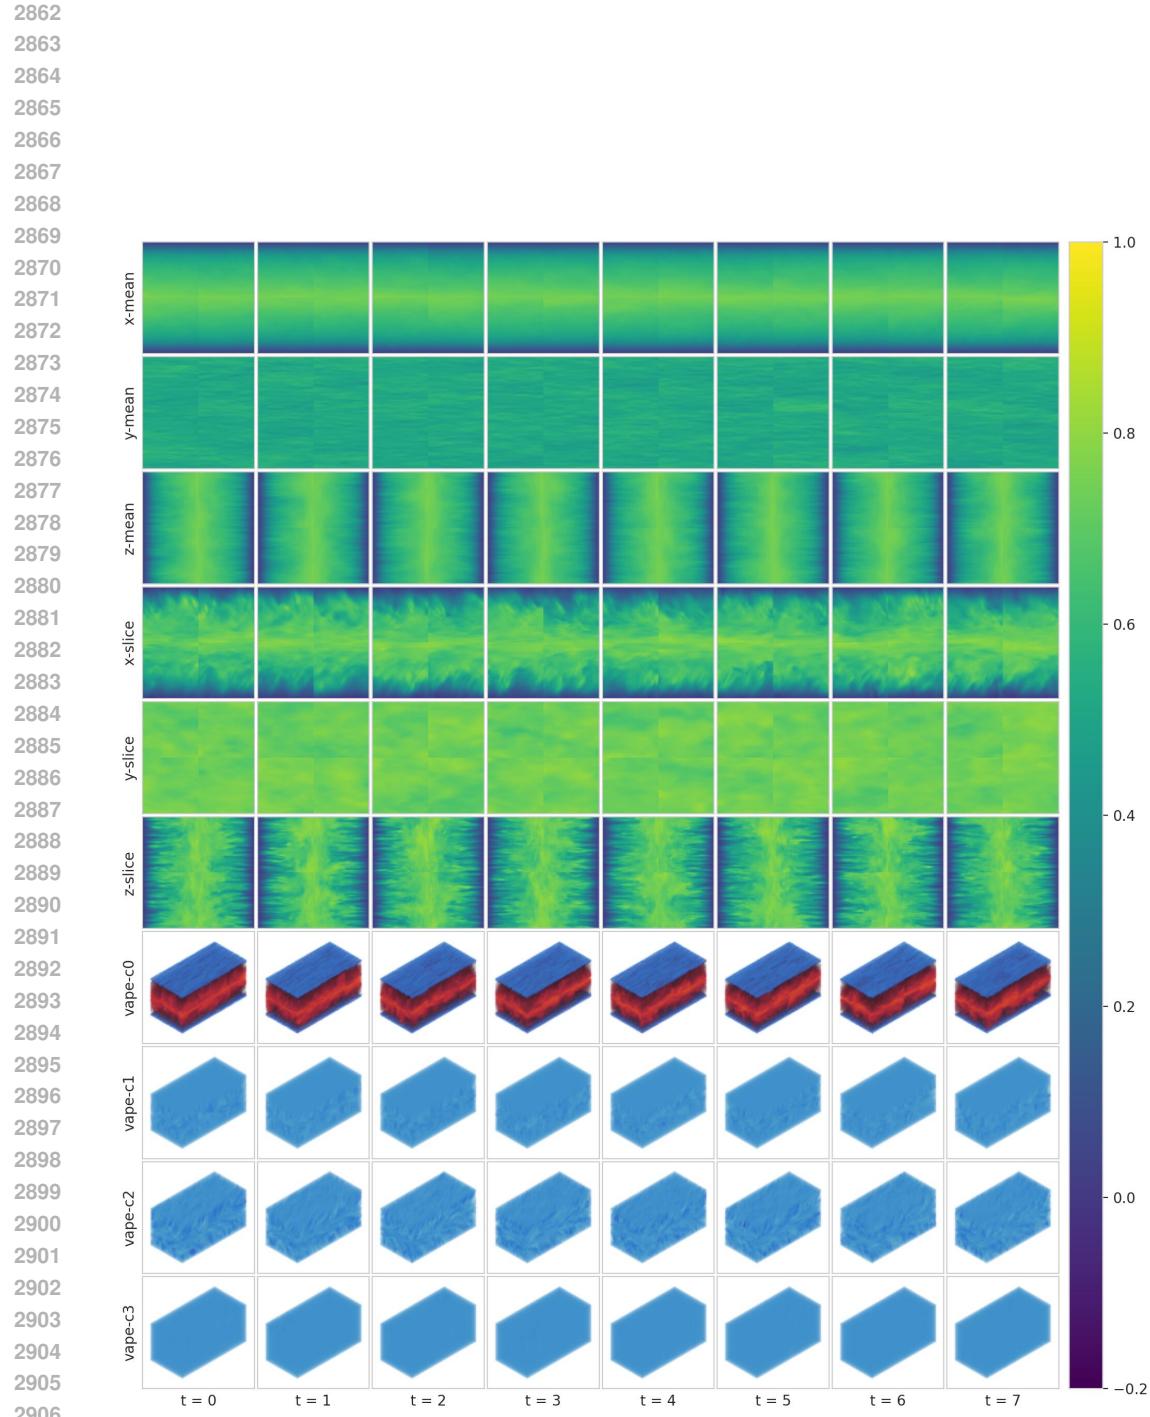


Figure D9: Samples from P3D-L $\langle X48|X48 \rangle$ pretrained on crops of size 48^3 and finetuned with the context network. Inference on the full domain at $Re = 800$ produces samples that exhibit the correct flow statistics.