

SPUS: A LIGHTWEIGHT AND PARAMETER-EFFICIENT FOUNDATION MODEL FOR PDEs

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010

ABSTRACT

011 We introduce Small PDE U-Net Solver (SPUS), a compact and efficient foun-
012 dation model (FM) designed as a unified neural operator for solving a wide
013 range of partial differential equations (PDEs). Unlike existing state-of-the-art
014 PDE FMs—primarily based on large complex transformer architectures with high
015 computational and parameter overhead—SPUS leverages a lightweight residual
016 U-Net-based architecture that has been largely underexplored as a foundation
017 model architecture in this domain. To enable effective learning in this minimalist
018 framework, we utilize a simple yet powerful auto-regressive pretraining strategy
019 which closely replicates the behavior of numerical solvers to learn the underly-
020 ing physics. SPUS is pretrained on a diverse set of fluid dynamics PDEs and
021 evaluated across 7 challenging unseen downstream PDEs spanning various phys-
022 ical systems. Experimental results demonstrate that SPUS using residual U-Net
023 based architecture achieves state-of-the-art generalization on these downstream
024 tasks while requiring significantly fewer parameters and minimal fine-tuning data,
025 highlighting its potential as a highly parameter-efficient FM for solving diverse
026 PDE systems.

1 INTRODUCTION

027 Partial differential equations (PDEs) are fundamental mathematical tools for modeling a wide range
028 of complex spatio-temporal phenomena in science and engineering, including fluid dynamics, elec-
029 tronics, materials science, and climate systems (Neumann et al., 2012; Schaa et al., 2016;
030 Müller & Scheichl, 2014). Traditional numerical solvers—such as finite difference and finite ele-
031 ment methods—are widely used for PDE simulation but often come with high computational costs,
032 especially when repeated simulations are required for varying coefficients or boundary conditions
033 (Herde et al., 2024). To address these limitations, deep learning-based approaches like the Fourier
034 neural operator (Li et al., 2021), convolutional neural operator (Raonic et al., 2023), and DeepONet
035 (Lu et al., 2021) have been proposed. While these models have shown promising performance, they
036 are typically designed for specific PDE families and require retraining when applied to new classes
037 of governing equations, leading to significant computational overhead.

038 Some simulation data is more computationally expensive to produce from numerical solvers than
039 others; and so multiphysics PDE FMs take advantage of pretraining on large benchmark PDE data
040 to finetune on limited PDE data from more expensive simulations. PDE FMs—including MPP (Mc-
041 McCabe et al., 2024), POSEIDON (Herde et al., 2024), PROSE FD (Liu et al., 2024), DISCO (Morel
042 et al., 2025) and DPOT (Hao et al., 2024)—have emerged as a promising paradigm. These models
043 aim to learn unified representations by incorporating multiple physical systems into a single frame-
044 work, demonstrating the ability to generalize to unseen PDE families using limited data. However,
045 current state-of-the-art FM approaches predominantly utilize transformer-based architectures with
046 high parameter counts in the hundreds of millions, resulting in increased computational and data
047 demands (McCabe et al., 2024; Herde et al., 2024; Hao et al., 2024). To overcome these limita-
048 tions, in this work, we propose a efficient yet effective Small PDE U-Net Solver (SPUS), with an
049 order of magnitude fewer parameters, for PDE foundation modeling. To the best of our knowledge,
050 this is the first work to explore a residual U-Net architecture as an FM pretrained on a large and
051 diverse PDE dataset, beyond single-family PDE prediction. U-Net has been shown to significantly
052 outperform neural operators such as FNO in solving PDEs (Gupta & Brandstetter, 2023). How-
053 ever, recent transformer-based FM approaches primarily compare against U-Net on a single family

of PDEs (McCabe et al., 2024; Hao et al., 2024; Shen et al., 2024), overlooking its potential as a foundation model architecture—particularly given the availability of large-scale PDE datasets from diverse systems.

PDE FMs are formulated in various ways with different assumptions on the form of input data and output predictions. Other FMs rely on multiple previous timesteps as input and predict output trajectories, or include temporal information in the input (McCabe et al., 2024; Herde et al., 2024; Hao et al., 2024). SPUS is trained autoregressively as an operator predicting a single timestep; which closely replicates the behavior of numerical solvers; potentially learning the underlying physics (Lippe et al., 2023). We demonstrate that the autoregressive training produces a foundation model that generalizes to time-independent PDEs via finetuning.

This work addresses the following key questions regarding FMs for PDEs:

- (a) Rather than designing a new and complex architecture, can we utilize a simple, existing one—such as residual U-Net—as an FM for PDEs?
- (b) Can a lightweight, low-parameter FM achieve state-of-the-art generalization on unseen PDEs?
- (c) Can pretraining on a set of simpler PDEs but exhibiting diverse physical behaviors (e.g., shocks, shear, vorticity) enable effective transfer to downstream tasks governed by complex PDEs and dominant dynamics, such as vortex evolution from piecewise-constant or shear-layer initial conditions?
- (d) Can an FM be pretrained to emulate the behavior of numerical solvers by autoregressively predicting the next time step from the current one, thereby potentially learning the underlying physics?

Finally, we demonstrate that SPUS, built on a simple residual U-Net and pretrained to emulate the behavior of numerical solvers, achieves state-of-the-art generalization with a lightweight design, transfers knowledge effectively from simpler PDEs to more complex ones, and thereby establishes a path toward efficient, generalizable PDE foundation models.

2 PRELIMINARIES

PDEs model a wide range of physical phenomena and include equations such as Navier-Stokes, compressible Euler, the Wave equation and others. The general form of a time-dependent PDE is:

$$\begin{aligned} \delta_t u(p, t) + L(u, \nabla_p u, \nabla_p^2 u, \dots) &= 0, & \forall p \in D \subset \mathbb{R}^d, t \in (0, T), \\ \mathcal{B}(u) &= 0, & \forall (p, t) \in \delta D \times (0, T), \\ u(p, 0) &= u_0(p), & p \in D. \end{aligned} \tag{1}$$

for given boundary conditions \mathcal{B} and initial conditions u_0 . Note that we use atypical PDE notation with p for the position variable (reserving x for the input as is typical in machine learning). Many PDE datasets are discretized in space and time. We denote the discretized spatial state at each timestep as $u_t = \{(p^j, u_t^j) : p^j \in \mathcal{P}\}$, $t \in [0, 1, \dots, n]$ where \mathcal{P} is the discretized spatial mesh and n is the number of discretized timesteps. Initial conditions are given by $u_{t=0}$ and each $u_t \in \mathbb{R}^d$ where d is the dimensionality of system variables.

3 RELEVANT WORK

The closest PDE FMs to ours fall into three distinct formulations.

- (a) PDE FMs which take $\{u_{t=[0,m]}\}$ of a PDE trajectory as input and autoregressively predict $\{u_{t=[m+1,n]}\}$; where $m = 15$ for MPP (McCabe et al., 2024), and $m = 9$ for DPOT (Hao et al., 2024). MPP projects normalized field variables from diverse physical systems into a unified latent space and utilizes an axial attention vision transformer-based architecture to perform autoregressive prediction over multiple systems. On the other hand, DPOT injects small-scale noise to $\{u_{t=[0,m]}\}$ and utilize a Fourier attention based transformer architecture to perform autoregressive prediction over multiple systems.

108 (b) PROSE FD (Liu et al., 2024) which takes $\{u_{t=[0,m]}\}$ of a PDE trajectory as input and si-
 109 multaneously predict $\{u_{t=[m+1,n]}\}$ as a trajectory where $m = 9$ and $n = 19$. PROSE FD
 110 introduces a multimodal transformer framework which takes $\{u_{t=[0,m]}\}$ of a PDE trajec-
 111 tory and mathematical description of the physical behavior as input and performs simulta-
 112 neous prediction for multi-physics systems.

113 (c) POSEIDON (Herde et al., 2024) which takes $(u_{t=0}, \Delta t)$ as input and predicts
 114 $\{u_{t=\Delta t}\} \forall \Delta t \in [1, T]$ where $T = 14$. POSEIDON proposes a multiscale operator trans-
 115 former architecture enhanced with time-conditioned normalization to perform prediction
 116 on multiple physical systems. Similar to SPUS, POSEIDON uses only a single time step
 117 (rather than a trajectory) as input; however, instead of performing autoregressive rollout,
 118 it predicts arbitrary future time steps directly. For a dataset with n time steps, POSEI-
 119 DON trains on $O(n^2)$ input-output pairs, whereas our approach is more sample-efficient,
 120 requiring only $O(n)$ sequential pairs.

121 **4 METHODS**

122 SPUS is a lightweight, low-parameter residual U-Net architecture designed for modeling PDE dy-
 123 namics. To enable effective learning within this compact model, we utilize an auto-regressive pre-
 124 training scheme. This method facilitates the efficient modeling of temporal dynamics of PDEs with
 125 reasonable accuracy and low computational overhead.

126 **Problem statement** Given an initial state $u_{t=0}$ of a trajectory governed by a specific PDE, where
 127 $u_t \in \mathbb{R}^d$ represents the system state at time step t with d variables, our objective is to predict the
 128 future states $u_{t=1}, u_{t=2}, \dots, u_{t=n}$.

129 **Auto-regressive pretraining and finetuning** We formulate the problem as a *first-order Markov*
 130 process (Pillai, 2002), in which the evolution of the system depends only on the immediately pre-
 131 ceding state. That is, the prediction of u_{t+1} is conditioned solely on u_t , satisfying the Markov
 132 property:

$$P(u_{t+1} | u_t, u_{t-1}, u_{t-2}, \dots, u_0) = P(u_{t+1} | u_t). \quad (2)$$

133 This formulation allows the system dynamics to be modeled using an autoregressive framework
 134 consistent with the Markov assumption.

135 The proposed auto-regressive training methodology for the U-Net-based FM is illustrated in Figure 1. During pretraining, the proposed FM takes a randomly sampled ground-truth state u_t from a
 136 PDE trajectory in the pretraining dataset and predicts the next state u'_{t+1} . More specifically, during
 137 pretraining, only ground-truth states are used as inputs; predicted states are not used to generate
 138 future predictions. During finetuning, the pretrained model is adapted to a specific downstream PDE
 139 using the same input-output structure as in pretraining: the model receives u_t and predicts u'_{t+1} .
 140 At inference time, however, we provide the model with the initial state $u_{t=0}$ and auto-regressively
 141 generate predictions $u'_{t=1}, u'_{t=2}, \dots, u'_{t=n}$, where each prediction u'_{t+1} is based on the previously
 142 predicted state u'_t as shown in Figure 1.

143 **4.1 MODEL ARCHITECTURE**

144 Figure 2 shows the residual U-Net architecture (Lan & Zhang, 2020; Ronneberger et al., 2015) with
 145 36 million parameters we have utilized for designing the FM for PDEs. The U-Net model takes any
 146 current state u_t of shape $d \times 128 \times 128$, where d is the number of system variables and applies a 3×3
 147 convolutional layer to project it into a 32-dimensional feature space. The residual encoder path
 148 comprises four hierarchical levels, each of which processes features through two residual blocks. Each
 149 residual block includes two 3×3 convolutional layers, with batch normalization (Bjorck et al., 2018)
 150 and GELU activation (Hendrycks & Gimpel, 2016) applied after each convolution, and incorporates
 151 a skip connection to preserve feature integrity and support gradient flow. Strided convolution is
 152 applied for downsampling at each level of encoder except the last. The encoder processes features
 153 through 32-channel blocks at Level 0, increases to 64 channels at Levels 1 and 2, and reaches 128
 154 channels at Level 3. The residual bottleneck consists of two residual blocks that operate on 128-
 155 channel feature maps, effectively capturing high-level representations. The residual decoder mirrors
 156

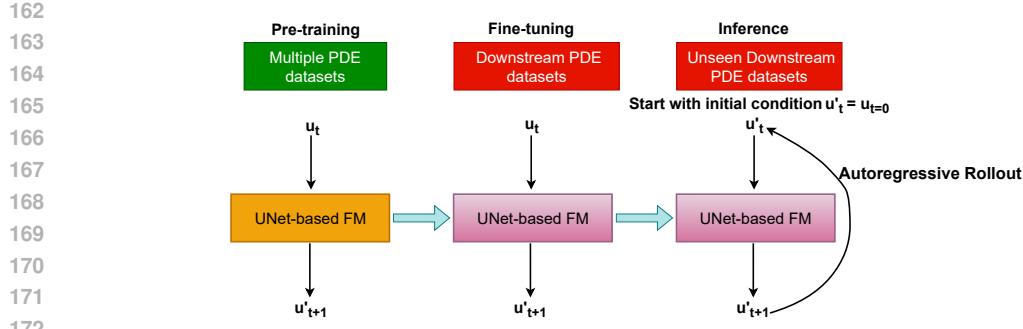


Figure 1: Proposed auto-regressive training methodology for the U-Net-based FM. During both pre-training and finetuning, the FM randomly samples a ground truth state u_t , where $u_t \in \mathbb{R}^d$ represents the system variables at time step t , and learns to predict the next state u'_{t+1} . During inference, the full trajectory is predicted autoregressively from the initial condition $u_{t=0}$. The FM takes $u'_t = u_{t=0}$ as input and recursively predicts subsequent states based on its own previous outputs for $t = 1, \dots, n$, where n is the length of the trajectory under consideration.

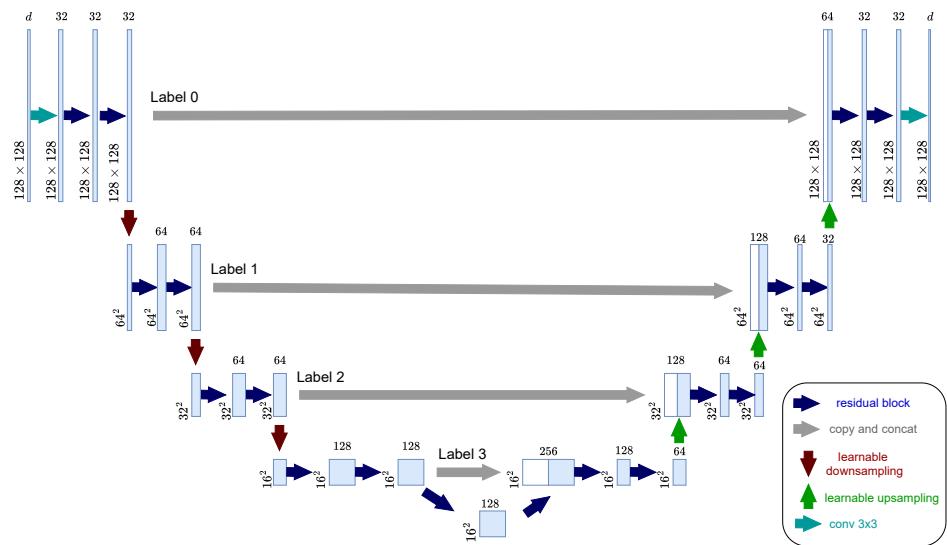


Figure 2: Illustration of the residual U-Net based FM architecture for PDEs with 36M parameters. The network takes an input of shape $d \times 128 \times 128$, representing the current time step of a PDE trajectory, and predicts the next time step of the same shape. It employs an encoder-decoder structure with residual blocks, skip connections, and progressive downsampling and upsampling to preserve spatial and contextual information.

the encoder with three upsampling stages implemented with transposed convolution layers. At each stage, the upsampled features are concatenated with the corresponding encoder feature maps via skip connections, facilitating the recovery of spatial details. After concatenation, the features are passed through two residual blocks, each followed by batch normalization and GELU activation. The decoder progressively reduces the feature dimensionality from 128 to 64, and subsequently from 64 to 32 across its stages. Finally, a 3×3 convolution maps the decoder output back to the number of system variable d .

During pretraining, the model was trained for 200 epochs using the Adam optimizer (Kingma & Ba, 2014) with a linear learning rate schedule starting from 10^{-4} , and a batch size of 10. The learning rate decreased linearly over the course of training. The model achieving the best performance on the evaluation set of the pretraining dataset was saved for downstream use.

216 4.2 THEORETICAL FOUNDATIONS OF RESIDUAL U-NET EFFICIENCY AND
217 GENERALIZATION IN PDE LEARNING
218219 Design choices are driven by theoretical motivation: a) convolutional layers for spatial bias, b)
220 residual blocks for numerical stability, c) U-Net architecture for encoder-bottleneck-decoder com-
221 pact representation, and d) multi-PDE pre-training for generalization.222 **Convolutional Layers** The step-to-step evolution of nonlinear PDEs on uniform grids exhibits a
223 locally translation-equivariant structure: each grid cell interacts primarily with its spatial neighbors
224 through couplings that decay smoothly with distance (LeVeque, 2007). This locality aligns naturally
225 with the inductive bias of the residual U-Net architecture, whose convolutional encoder captures
226 spatially local dependencies and whose decoder reconstructs fine-scale features through upsampling
227 and skip connections (Ronneberger et al., 2015; Ruthotto & Haber, 2020). By progressively reducing
228 and restoring spatial resolution, the U-Net builds a **multiscale hierarchy** that aggregates coarse and
229 fine spatial information, enabling efficient representation of both local interactions and long-range
230 dependencies. This hierarchical design allows the network to approximate global spatial couplings
231 without the quadratic computational overhead of transformer-based attention mechanisms.
232233 **Residual Blocks** Residual connections enhance both stability and learning efficiency by predicting
234 incremental field updates rather than full mappings, analogous to time-stepping schemes in numeri-
235 cal PDE solvers. These design principles collectively make the residual U-Net a parameter-efficient
236 and numerically stable architecture for learning complex, nonlinear PDE operators across diverse
237 physical regimes.238 The following section formalizes a theoretical error decomposition, illustrating how the residual
239 encoder-bottleneck-decoder hierarchy reduces approximation, projection, and statistical errors, and
240 how multi-PDE pretraining further enhances cross-family generalization.241 4.2.1 THEORY
242243 **Encode-Bottleneck-Decoder Architecture** Let $\Psi : X \rightarrow Y$ denote the one-step solution operator
244 of a time-dependent PDE. Given a field state $u_t \in \mathbb{R}^d$ at time t , representing the spatial distributions
245 of system variables (e.g., density, velocity, or pressure), the operator Ψ produces the evolved field
246 $u_{t+1} = \Psi(u_t) \in Y$ at the next timestep. For notational convenience, we denote $x = u_t \in X$ as the
247 input field and $y = \Psi(x) = u_{t+1} \in Y$ as the corresponding evolved field. In the residual U-Net
248 framework, this mapping is approximated by a multiscale neural surrogate $F_\theta : X \rightarrow Y$, parame-
249 terized by weights θ , where each stage of the network (encoder, bottleneck, and decoder) performs
250 local residual updates that emulate incremental field evolution consistent with the underlying PDE
251 dynamics across spatial scales. Let μ denote the probability measure associated with the sampling
252 of input fields x at time t within the function space X . Each sample $x \sim \mu$ therefore represents a
253 physical state of the PDE system at the previous timestep, and the network is trained to predict the
254 corresponding evolved field $y = \Psi(x)$ at time $t + 1$, yielding the expected mean-squared prediction
255 error

256
$$\mathcal{E}(\theta) = \mathbb{E}_{x \sim \mu} [\|\Psi(x) - F_\theta(x)\|_2^2], \quad (3)$$

257 which quantifies the average discrepancy between the true PDE evolution and the residual U-Net
258 prediction across all field states encountered in the data. Following the generalization frame-
259 work of Bhattacharya et al. (2021), this error can be decomposed into three fundamental contribu-
260 tions—neural approximation, model reduction, and finite-sample statistics:

261
$$\mathcal{E}(\theta) \lesssim \underbrace{\varepsilon_{\text{approx}}^2}_{\text{neural approximation}} + \underbrace{R^\mu(V_x) + R^{\Psi^\# \mu}(V_y)}_{\text{model reduction / projection}} + \underbrace{\mathcal{O}\left(\frac{1}{\sqrt{N}}\right)}_{\text{finite-sample term}}. \quad (4)$$

262
263

264 Here, neural approximation denotes the best achievable error of the residual U-Net within the learned
265 subspaces; $V_x \subset X$ and $V_y \subset Y$ denote the encoder and decoder subspaces, respectively. The pro-
266 jection terms quantify the residual energy of the data distributions outside these learned subspaces.
267 Specifically, $R^\mu(V_x) = \mathbb{E}_{x \sim \mu} [\|x - \Pi_{V_x}(x)\|_2^2]$ measures the expected reconstruction error of input
268 fields $x \in X$, where Π_{V_x} denotes the projection or learned feature mapping induced by the encoder
269 onto the subspace V_x . Likewise, $R^{\Psi^\# \mu}(V_y) = \mathbb{E}_{y \sim \Psi^\# \mu} [\|y - \Pi_{V_y}(y)\|_2^2]$ measures the correspond-
ing residual energy of output fields $y = \Psi(x) \in Y$, where $\Psi^\# \mu$ is the output distribution induced

270 by the PDE operator Ψ , that is, the distribution of evolved fields obtained by mapping input samples
 271 $x \sim \mu$ through Ψ . Here, Π_{V_y} denotes the decoder-side projection onto the subspace V_y , and N
 272 represents the number of training samples. In a residual U-Net, the encoder, bottleneck, and decoder
 273 all contain residual blocks of the form

$$274 \quad h^{(l+1)} = h^{(l)} + R_{\theta_l}^{(h)}(h^{(l)}), \quad (5)$$

276 where $h^{(l)}$ denotes the feature tensor at depth l *within the current stage*, and each residual operator
 277 $R_{\theta_l}^{(h)}$ is a local convolutional transformation acting on the corresponding feature space. The
 278 **residual encoder** E progressively constructs a hierarchical latent subspace V_x that captures coarse-
 279 to-fine spatial and dynamical modes. By composing small residual updates, the encoder acts as a
 280 data-driven reduced-basis generator, minimizing the input projection error $R^\mu(V_x)$ in equation 4.
 281 At the network’s core, the **residual bottleneck** B captures cross-scale and nonlocal interactions in
 282 the latent feature space, providing an efficient, low-dimensional approximation of global coupling
 283 that further reduces the intrinsic neural approximation error $\varepsilon_{\text{approx}}$. The **residual decoder** D then
 284 reconstructs the output subspace V_y through a sequence of incremental updates, which refine hier-
 285 archical features and approximate stable step-to-step updates of the underlying PDE evolution. The
 286 complete network composition is thus

$$287 \quad F_\theta = D \circ B \circ E = (\text{Id} + R_\theta^{(D)}) \circ (\text{Id} + R_\theta^{(B)}) \circ (\text{Id} + R_\theta^{(E)}), \quad (6)$$

288 where Id denotes the identity mapping on the corresponding feature space of each component (en-
 289 coder, bottleneck, or decoder). This composition can be interpreted as a **multiscale residual in-**
 290 **tegrator** in feature space, in which each component performs a small, well-conditioned update
 291 aligned with the underlying PDE dynamics. Combining Equations 5–6 with the error decomposi-
 292 tion in Equation 4, we find that the residual encoder–bottleneck–decoder hierarchy reduces all three
 293 error sources simultaneously: (a) the multiscale residual encoder–decoder minimizes the projection
 294 errors $R^\mu(V_x)$ and $R^{\Psi\#\mu}(V_y)$; (b) the bottleneck and residual formulation lower the intrinsic ap-
 295 proximation error $\varepsilon_{\text{approx}}$ by capturing local-to-global nonlinear interactions; and (c) convolutional
 296 weight sharing limits parameter growth, thereby reducing the statistical term $\mathcal{O}(N^{-1/2})$. In contrast
 297 to transformer-based operator learners that rely on global self-attention—incurred quadratic com-
 298 putational and memory complexity with respect to input size and requiring large parameter counts
 299 to infer spatial locality from data—the residual U-Net embeds these spatial priors directly through
 300 its convolutional and multiscale hierarchical design. This inductive bias enables the residual U-
 301 Net to achieve comparable or superior accuracy with substantially fewer parameters and improved
 302 numerical stability in PDE operator learning.

303 **Multi-PDE Pretraining** When the residual U-Net is pretrained across multiple PDE families
 304 $\{\Psi_m\}_{m=1}^M$ defined on the same grid class, the model learns a **shared multiscale encoder–**
 305 **bottleneck–decoder representation** that captures structural invariants common to these operators,
 306 including spatial locality, smooth spectral decay, and hierarchical coupling across scales. Assuming
 307 that these structural invariants exist in the downstream finetuning PDEs, then model error will be
 308 reduced compared to training from scratch (Bhattacharya et al., 2021; Ben-David & Borbely, 2008).
 309 Let μ_m denote the probability measure associated with sampling the variable fields at time t from
 310 trajectories of the m -th PDE family within the function space X , and let $\Psi_{m\#}\mu_m$ denote the cor-
 311 responding measure of their evolved fields at time $t + 1$. Define $\bar{\mu} = \frac{1}{M} \sum_{m=1}^M \mu_m$ as the mixture
 312 of input field distributions and $\bar{\Psi}_{\#\mu} = \frac{1}{M} \sum_{m=1}^M \Psi_{m\#}\mu_m$ as the corresponding mixture of output
 313 field measures. Under the same decomposition as in equation 4, the expected joint training error
 314 satisfies

$$315 \quad \frac{1}{M} \sum_{m=1}^M \mathbb{E}_{x \sim \mu_m} [\|\Psi_m(x) - F_\theta(x)\|_2^2] \lesssim \varepsilon_{\text{approx}}^2 + R^{\bar{\mu}}(V_x) + R^{\bar{\Psi}_{\#\mu}}(V_y) + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right). \quad (7)$$

318 Here, the projection terms depend on the shared mixture distributions rather than on any single
 319 PDE family. By exposing the network to diverse yet structurally related dynamics, multi-PDE pre-
 320 training enables the residual encoder, bottleneck, and decoder to learn common multiscale feature
 321 subspaces that align with recurring spatial and spectral patterns across PDE families. This shared
 322 representation captures a greater portion of the underlying functional variability with the same latent
 323 dimensionality, thereby reducing the average projection error. As a result, **during downstream fine-**
 324 **tuning**, the pretrained residual U-Net requires only minor residual adaptations in its local blocks,

324 leading to improved numerical stability, and strong cross-family generalization without increasing
 325 parameter count.
 326

327 **4.3 PRETRAINING AND FINETUNING DATASET**
 328

329 SPUS is pretrained on a diverse set of PDE types from the PDEGYM dataset (Herde et al., 2024),
 330 which includes four operators derived from the compressible Euler (CE) equations:
 331

- 332 • **CE-RP**, containing trajectories initialized with four-quadrant Riemann problems;
- 333 • **CE-CRP**, initialized with multiple curved Riemann problems;
- 334 • **CE-KH**, representing shear-driven Kelvin–Helmholtz instabilities; and
- 335 • **CE-Gauss**, featuring initial conditions with Gaussian vorticity profiles.
 336

337 Each dataset consists of 10,000 trajectories. Each trajectory has 21 time steps and each time step
 338 consists of five physical fields: density ρ , horizontal velocity u , vertical velocity v , pressure p , and
 339 energy E with spatial grid of resolution 128×128 .
 340

341 We fine-tune SPUS on seven previously unseen downstream PDEs from the PDEGYM dataset, using
 342 128 trajectories for each PDE task. These downstream PDEs include three operators governed by
 343 the CE equations, three operators governed by the incompressible Navier-Stokes (NS) equations,
 344 and one based on the wave equation:
 345

- 346 • **CE-RPUI**: consisting of trajectories initialized with four-quadrant Riemann problems fea-
 347 turing uncertain interfaces;
- 348 • **CE-RM**: representing the Richtmyer-Meshkov instability problem;
- 349 • **NS-PwC**: initialized from piecewise-constant vorticity fields;
- 350 • **NS-SL**: initialized with double shear layer conditions;
- 351 • **FNS-KF**: also initialized from piecewise-constant vorticity fields;
- 352 • **Wave-Gauss**: containing trajectories initialized as a sum of Gaussians that are propagated
 353 by the spatially varying wave speed.
- 354 • **SE-AF**: contains contains the steady-state density over airfoils
 355

356 Each CE-RPUI and CE-RM trajectory contains 21 time steps. Each time step has five physical
 357 fields: density ρ , horizontal velocity u , vertical velocity v , pressure p , and energy E . On the other
 358 hand, each trajectory in the three NS datasets also has 21 time steps but only two physical fields:
 359 horizontal velocity u , and vertical velocity v . For the Wave-Gauss dataset, trajectories have 15 time
 360 steps with one physical field, spatially varying wave speed w . For the SE-AF dataset, the samples
 361 are time-independent and solution operator maps a shape coefficient into the steady state solution.
 362 All fine-tuning datasets share a common spatial resolution of 128×128 grid points.
 363

364 **4.4 FINETUNING STRATEGIES AND BASELINE MODELS**
 365

366 In downstream tasks, the number of variables per time step may differ from those used during
 367 pretraining. To adapt the pretrained model to downstream tasks with different input and output
 368 dimensions than pretraining, we introduce lightweight input and output adapters. Specifically, we
 369 use 1×1 convolutional layers as adapters:
 370

- 371 • The **InputAdapter** maps the task-specific input (e.g., 2 fields for NS-SL) to the 5-field
 372 format expected by the pretrained SPUS model.
- 373 • The **OutputAdapter** maps the model’s 5-field output back to the task-specific output di-
 374 mensionality (e.g., 2 fields for NS-SL).
 375

376 These adapters are simple, efficient, and allow the pretrained model to be flexibly applied to a
 377 variety of downstream tasks without modifying its internal architecture. For each downstream task,
 378 we fine-tuned either the pretrained model or the pretrained model with adapters (if the number of
 379 the fields differed from five) using 128 trajectories. The model was fine-tuned for 200 epochs using
 380

378 Table 1: Comparison of model performance (average MSE over all predicted timesteps from the
 379 initial conditions of the trajectories) on six unseen downstream PDE datasets fine-tuned with 128
 380 trajectories. Lower is better. The U-Net* is trained from scratch using 128 trajectories for each
 381 downstream PDE dataset.

Dataset	SPUS (Ours, 36M)	DPOT (122M)	POSEIDON (158M)	U-Net* (36M)
CE-RPUI	0.0054	0.0570	0.0085	0.0337
CE-RM	0.0159	0.0222	0.4181	0.0218
NS-PwC	0.0048	0.0294	0.0004	0.0048
FNS-KF	0.0015	0.0301	0.0017	0.0047
NS-SL	0.0163	0.1461	0.0163	0.0165
SE-AF	0.0006	-	0.0031	0.0040
Wave-Gauss	0.0069	0.0107	0.0068	0.0071

391
 392
 393 the Adam optimizer with a linear learning rate schedule starting from 10^{-4} , and a batch size of 10.
 394 The learning rate decreased linearly over the course of training.
 395

396 To ensure a fair comparison, we fine-tune two baseline FMs: DPOT “M” (122M parameters) (Hao
 397 et al., 2024) and POSEIDON “B” (158M parameters) (Herde et al., 2024). DPOT was pretrained on
 398 12 PDE datasets governed by the Navier-Stokes, diffusion-reaction, and shallow-water equations,
 399 whereas POSEIDON was pretrained on 6 PDE datasets governed by the compressible Euler and
 400 Navier-Stokes equations. For both baselines, we adopt the exact hyperparameter settings recom-
 401 mended in their original papers and accompanying code repositories (Hao et al., 2024; Herde et al.,
 402 2024). All models, including SPUS, are fine-tuned separately on each downstream PDE task using
 403 the same set of 128 trajectories for 200 epochs with MSE loss. Performance is evaluated on testing
 404 dataset corresponding to each PDE task.

405 DPOT recommends a context window of 10 timesteps. Accordingly, to predict trajectories from
 406 their initial conditions, we follow the same fine-tuning methodology described in Herde et al. (2024),
 407 padding input sequences with timestep 0 when predicting steps earlier than the 10th. For instance,
 408 to predict the state at timestep 4, the input sequence is padded as follows:

$$[ts_0, ts_0, ts_0, ts_0, ts_0, ts_0, ts_0, ts_1, ts_2, ts_3].$$

409 POSEIDON, on the other hand, is designed to take a single timestep as input, along with the corre-
 410 sponding Δt , and directly predict any future frame within the trajectory. This allows POSEIDON to
 411 predict any timestep (using only the initial timestep as context) without requiring an autoregressive
 412 rollout. In practice, such “direct” predictions result in higher average accuracy compared to predic-
 413 tions generated via autoregressive rollout. Therefore, we report POSEIDON’s performance based
 414 on its direct prediction accuracy.

415 To further assess the parameter efficiency and architectural simplicity of SPUS, we trained a U-
 416 Net model with 36M parameters—sharing the same architecture as SPUS—from scratch using the
 417 same set of 128 trajectories per downstream dataset on which SPUS, POSEIDON, and DPOT were
 418 fine-tuned, and compared its performance with these foundation models.

419 The comparison of model performance—measured as average mean squared error (MSE) across all
 420 predicted timesteps from the initial condition of the trajectories—on seven unseen downstream PDE
 421 datasets fine-tuned with 128 trajectories is presented in Table 1.

422 5 EXPERIMENTS

423
 424 **425 Is SPUS an effective lightweight PDE emulator? Does SPUS with only 36 million parameters
 426 generalize as accurately as larger models?** To address these questions, we design and evaluate
 427 the following three experiments.

428
 429 **(A). Does SPUS generalize to unseen systems governed by the compressible Euler (CE) equa-
 430 tions, consistent with its pretraining?** We investigate whether SPUS can generalize to previously

unseen physical systems that are governed by CE equations, consistent with its pretraining. To evaluate this, we fine-tune the pretrained SPUS model on the CE-RPUI dataset. While this dataset adheres to the same underlying physical laws, its distribution of initial conditions differs from those seen during pretraining, presenting a clear out-of-distribution (OOD) generalization challenge (Herde et al., 2024). As shown in Table 1, SPUS achieves strong performance in autoregressively predicting full trajectories from initial conditions, despite having only 36 million parameters. Notably, it outperforms both the substantially larger POSEIDON model (158 million parameters) and the DPOT model (122 million parameters) in terms of average mean squared error (MSE) across 240 test trajectories. A randomly selected trajectory prediction from the CE-RPUI test set is shown in Figure A.2, where the SPUS predictions closely match the ground truth at each time step. These results demonstrate the effectiveness and computational efficiency of the lightweight SPUS model relative to significantly larger architectures.

We also fine-tune the pretrained SPUS model on the CE-RM dataset, which exhibits significantly more complex dynamics compared to CE-RPUI. SPUS demonstrates strong generalization capability in predicting entire trajectories from initial conditions, as illustrated in Figure A.3. Furthermore, as shown in Table 1, SPUS achieves a lower average MSE across 130 test trajectories compared to both the POSEIDON and DPOT models, despite their substantially larger parameter counts.

(B). Does SPUS generalize to systems governed by equations different from those used in pre-training? We investigate the ability of SPUS to generalize to previously unseen physical systems governed by equations different from those used during pretraining. Specifically, we fine-tune the pretrained SPUS model on three datasets governed by incompressible NS equations that were not part of the pretraining data: NS-PwC, NS-SL, and FNS-KF. As shown in Table 1, despite not being exposed to incompressible NS dynamics during pretraining, surprisingly, SPUS achieves superior time-step prediction performance compared to DPOT across all three datasets—even though DPOT was pretrained on operators of both compressible and incompressible NS equations. For the POSEIDON model, whose pretraining data includes two operators governed by NS equations, SPUS outperforms it on FNS-KF, matches its performance on NS-SL, and is outperformed on NS-PwC, as summarized in Table 1. These results demonstrate the strong generalization capability of SPUS to new physical regimes outside its pretraining distribution and highlight its effective transferability to downstream tasks governed by equations different from those seen during pretraining. Randomly selected trajectory predictions from the test datasets of NS-PwC, NS-SL, and FNS-KF are shown in Figure A.4 and Figure A.5 (in Appendix). As observed, SPUS demonstrates strong generalization performance on each of the NS dataset; however, the predicted variables gradually deviate from the ground truth over time. We also fine-tuned SPUS on the Wave-Gauss dataset, which is governed by the wave equation. As shown in Table 1, SPUS outperforms DPOT and is narrowly outperformed by POSEIDON for 240 Wave-Gauss testing trajectories.

(C). Does SPUS generalize to time-independent PDEs although it is pretrained on time-dependent PDEs? We fine-tune the pretrained SPUS model on the time-independent SE-AF trajectories to evaluate its generalization capability beyond time-dependent PDEs. Although SPUS is originally pretrained on time-dependent PDEs, it can be readily adapted to time-independent problems because it does not explicitly take time as an input variable. Instead, the model simply learns the mapping between input and output fields required for prediction in static PDEs. As shown in Table 1, SPUS achieves better performance than both POSEIDON and the unpretrained U-Net on the SE-AF testing dataset. A randomly selected prediction from the SE-AF test set is illustrated in Figure A.7.

Does SPUS show scalability with dataset size? Table 4 reports the average MSE of predicting the time steps of entire trajectories from the initial condition of the trajectories across six downstream PDEs using SPUS, finetuned with 32, 128, and 256 trajectories. As shown, increasing the fine-tuning set size reduces the MSE across the six downstream PDEs, demonstrating that SPUS scales favorably with additional data.

Summaries of the experiments Based on the above experiments, we observe that SPUS, built on a residual U-Net architecture with only 36 million parameters, achieves state-of-the-art generalization on downstream tasks, outperforming significantly larger models such as POSEIDON (158 million parameters) and DPOT (122 million parameters). These results highlight that a simple, well-

486 established architecture—specifically, a residual U-Net—can be effectively leveraged as a founda-
 487 tion model (FM) for PDEs. Despite its architectural simplicity and relatively small parameter count,
 488 SPUS is capable of capturing complex dynamics and performs competitively with more sophisti-
 489 cated, larger models. We also observe that SPUS, when pretrained on a diverse set of simpler PDEs
 490 (such as CE), demonstrates strong performance on complex downstream PDEs (such as NS). This
 491 indicates the effective transferability of SPUS across distinct physical regimes, despite differences
 492 in the underlying governing equations. Furthermore, this suggests that even when the pretraining
 493 data are derived from PDEs governed by simple CE equations, a sufficiently diverse pretraining
 494 dataset—spanning variations in initial and boundary conditions, domain geometries, and external
 495 forcing—can enable the FM to generalize effectively. Moreover, SPUS is pretrained to emulate the
 496 behavior of numerical solvers by autoregressively predicting the next time step from the current one.
 497 The results on downstream tasks for SPUS suggest that this pretraining strategy helps the FM learn
 498 the underlying physics of PDEs, enabling more accurate and physically consistent predictions.
 499

500 Additional results on the performance evaluation of SPUS on downstream tasks, as well as visual
 501 comparisons of SPUS’s performance on trajectory prediction from initial conditions with larger
 502 models (POSEIDON and DPOT), are presented in Appendices A.1–A.6.

503 6 LIMITATIONS

504 SPUS is pretrained on a limited set of PDE families, using only four CE datasets. Despite this
 505 narrow pretraining scope, SPUS shows promising transferability across equation types, generalizing
 506 from compressible Euler to incompressible Navier–Stokes and wave equations. Currently, SPUS
 507 cannot simultaneously predict multiple future timesteps from an initial condition. In future work, we
 508 plan to extend SPUS to support both direct and autoregressive temporal prediction. We also expect
 509 that pretraining on a broader range of governing equations would further enhance its generalization
 510 capability. Moreover, the current architecture is restricted to regular geometries, and extending
 511 SPUS to irregular domains remains an important direction for future research.
 512

514 7 CONCLUSIONS

516 We propose SPUS, a compact and lightweight FM for PDEs, capable of handling a broad range of
 517 physical systems. The model is based on a simple residual U-Net architecture and is trained us-
 518 ing a straightforward autoregressive pretraining strategy. Despite its relatively small size—only 36
 519 million parameters—SPUS demonstrates strong generalization capabilities across six diverse down-
 520 stream PDE tasks. SPUS consistently outperforms the significantly larger DPOT model across all
 521 downstream datasets. When compared to the POSEIDON model, which also has substantially more
 522 parameters, SPUS achieves superior performance on three datasets, matches performance on one,
 523 is narrowly outperformed on another (MSE: 0.0069 vs. 0.0068), and is outperformed on one task.
 524 These results establish SPUS as a highly parameter-efficient foundation model, capable of solving a
 525 wide range of complex PDE systems with competitive accuracy. Furthermore, we demonstrate that
 526 pretraining SPUS on simpler PDEs (such as CE) with autoregressive training to emulate a numerical
 527 solver enables effective transfer to more complex PDEs (such as NS), reducing the amount of data
 528 required for finetuning even when the downstream task involves more complex dynamics than those
 529 seen during pretraining.

531 REPRODUCIBILITY STATEMENT

533 All datasets used in this work are publicly available. The code will be released at the time of
 534 publication.
 535

537 REFERENCES

538 Shai Ben-David and Reba Schuller Borbely. A notion of task relatedness yielding provable multiple-
 539 task learning guarantees. *Mach. Learn.*, 73(3):273–287, 2008.

540 Kaushik Bhattacharya, Bamdad Hosseini, Nikola B Kovachki, and Andrew M Stuart. Model reduc-
 541 tion and neural networks for parametric pdes. *The SMAI journal of computational mathematics*,
 542 7:121–157, 2021.

543

544 Nils Bjorck, Carla P Gomes, Bart Selman, and Kilian Q Weinberger. Understanding batch normal-
 545 ization. *Advances in Neural Information Processing Systems*, 31, 2018.

546 Jayesh K Gupta and Johannes Brandstetter. Towards multi-spatiotemporal-scale generalized PDE
 547 modeling. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.

548

549 Zhongkai Hao, Chang Su, Songming Liu, Julius Berner, Chengyang Ying, Hang Su, Anima Anand-
 550 kumar, Jian Song, and Jun Zhu. DPOT: Auto-regressive denoising operator transformer for large-
 551 scale PDE pre-training. In *Proceedings of the 41st International Conference on Machine Learn-
 552 ing*, 2024.

553 Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (GELUs). *arXiv preprint
 554 arXiv:1606.08415*, 2016.

555

556 Maximilian Herde, Bogdan Raonic, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Emmanuel
 557 de Bézenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for PDEs. *Advances
 558 in Neural Information Processing Systems*, 37:72525–72624, 2024.

559

560 Z Huang, H Wang, Z Deng, J Ye, Y Su, H Sun, J Junjun He, Y Gu, L Gu, S Zhang, et al. Stu-
 561 net: Scalable and transferable medical image segmentation models empowered by large-scale
 562 supervised pre-training. *arXiv preprint arXiv:2304.06716*, 2023.

563 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint
 564 arXiv:1412.6980*, 2014.

565

566 Yancheng Lan and Xuming Zhang. Real-time ultrasound image despeckling using mixed-attention
 567 mechanism based residual UNet. *IEEE Access*, 8:195327–195340, 2020.

568

569 Randall J LeVeque. *Finite difference methods for ordinary and partial differential equations: steady-
 570 state and time-dependent problems*. SIAM, 2007.

571

572 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, An-
 573 drew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential
 574 equations. In *International Conference on Learning Representations*, 2021.

575

576 Phillip Lippe, Bas Veeling, Paris Perdikaris, Richard Turner, and Johannes Brandstetter. PDE-
 577 refiner: Achieving accurate long rollouts with neural PDE solvers. *Advances in Neural Infor-
 578 mation Processing Systems*, 36:67398–67433, 2023.

579

580 Yuxuan Liu, Jingmin Sun, Xinjie He, Griffin Pinney, Zecheng Zhang, and Hayden Schaeffer.
 581 PROSE-FD: A multimodal PDE foundation model for learning multiple operators for forecasting
 582 fluid dynamics. *arXiv preprint arXiv:2409.09811*, 2024.

583

584 Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning
 585 nonlinear operators via DeepONet based on the universal approximation theorem of operators.
 586 *Nature Machine Intelligence*, 3(3):218–229, 2021.

587

588 Michael McCabe, Bruno Régaldo-Saint Blancard, Liam Parker, Ruben Ohana, Miles Cranmer, Al-
 589 berto Bietti, Michael Eickenberg, Siavash Golkar, Geraud Krawezik, Francois Lanusse, et al.
 590 Multiple physics pretraining for spatiotemporal surrogate models. *Advances in Neural Infor-
 591 mation Processing Systems*, 37:119301–119335, 2024.

592

593 Rudy Morel, Jiequn Han, and Edouard Oyallon. Disco: learning to discover an evolution operator
 594 for multi-physics-agnostic prediction. *arXiv preprint arXiv:2504.19496*, 2025.

595

596 Eike H Müller and Robert Scheichl. Massively parallel solvers for elliptic partial differential equa-
 597 tions in numerical weather and climate prediction. *Quarterly Journal of the Royal Meteorological
 598 Society*, 140(685):2608–2624, 2014.

594 Philipp Neumann, Hans-Joachim Bungartz, Miriam Mehl, Tobias Neckel, and Tobias Weinzierl. A
 595 coupled approach for fluid dynamic problems using the PDE framework Peano. *Communications*
 596 *in Computational Physics*, 12(1):65–84, 2012.

597

598 S Unnikrishna Pillai. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, 2002.

599

600 Bogdan Raonic, Roberto Molinaro, Tim De Ryck, Tobias Rohner, Francesca Bartolucci, Rima Alai-
 601 fari, Siddhartha Mishra, and Emmanuel de Bézenac. Convolutional neural operators for robust
 602 and accurate learning of PDEs. *Advances in Neural Information Processing Systems*, 36:77187–
 603 77200, 2023.

604

605 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical
 606 image segmentation. In *International Conference on Medical Image Computing and
 Computer-Assisted Intervention*, pp. 234–241. Springer, 2015.

607

608 Lars Ruthotto and Eldad Haber. Deep neural networks motivated by partial differential equations.
 609 *Journal of Mathematical Imaging and Vision*, 62(3):352–364, 2020.

610

611 R Schaa, L Gross, and J Du Plessis. PDE-based geophysical modelling using finite elements: Ex-
 612 amples from 3d resistivity and 2d magnetotellurics. *Journal of Geophysics and Engineering*, 13
 (2):S59–S73, 2016.

613

614 Junhong Shen, Tanya Marwah, and Ameet Talwalkar. UPS: Efficiently building foundation models
 615 for PDE solving via cross-modal adaptation. *Transactions on Machine Learning Research*, 2024.

616

617 Anthony Zhou and Amir Barati Farimani. Masked autoencoders are pde learners. *arXiv preprint*
 arXiv:2403.17728, 2024.

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

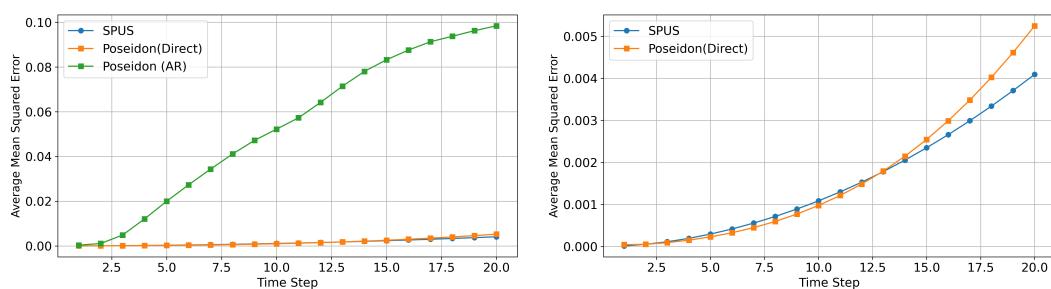
647

648 **A APPENDIX**
649650 **A.1 ERROR GROWTH OVER TIME**
651

652 Figure A.1 presents the average mean squared error (MSE) of trajectory predictions over time for the
 653 FNS-KF test datasets. As shown, the SPUS model—with only 36 million parameters—exhibits an
 654 approximately linear increase in prediction error over time, a behavior consistently observed across
 655 our downstream datasets. POSEIDON, a larger FM with 158 million parameters, demonstrates a
 656 similar error growth pattern as shown in Figure A.1b. As can be seen, SPUS achieves comparable
 657 performance to POSEIDON (Direct) at early time steps and surpasses it at later time steps. These
 658 results highlight the potential of SPUS to deliver accurate long-term predictions despite having sig-
 659 nificantly fewer parameters.

660 Table 2: Comparison of model performance (average MSE across all predicted 20 timesteps from the
 661 initial conditions of the trajectories) on the CE-RM, NS-SL, and FNS-KF PDE datasets, fine-tuned
 662 with 128 trajectories using three different pretrained model sizes. SPUS demonstrates scalability
 663 with increasing model size.

Dataset	SPUS-9M	SPUS-36M	SPUS-76M
CE-RM	0.0203	0.0159	0.0129
NS-SL	0.0247	0.0163	0.0027
FNS-KF	0.0045	0.0015	0.0008



670 (a) Prediction error over time for FNS-KF with Posei- 671 (b) Prediction error over time for FNS-KF with Posei-
 672 don (Direct), Poseidon (AR), and SPUS 673 don (Direct), and SPUS
 674
 675
 676
 677
 678

679 Figure A.1: Average MSE of trajectory predictions over time for the FNS-KF test datasets. SPUS
 680 (36M) shows approximately linear error growth similar to POSEIDON (Direct). It surpasses PO-
 681 SEIDON (direct) at later steps on FNS-KF, highlighting its efficiency and long-term accuracy as
 682 shown in Figure (b).
 683

687 Table 3: Performance analysis of SPUS with different numbers of residual blocks at each stage.
 688 The table reports the average MSE across all 20 predicted timesteps from the initial timestep on the
 689 CE-RM, NS-SL, and FNS-KF test datasets. Using five residual blocks per stage in SPUS reduces
 690 the prediction MSE across all downstream tasks.

Model	Number of Residual Blocks	CE-RM	NS-SL	FNS-KF
SPUS-36M	2	0.0159	0.0163	0.0015
SPUS-76M	5	0.0129	0.0027	0.0008

695 **A.2 SPUS SCALES WITH MODEL SIZE**
696

697 We pretrain SPUS with three model sizes to study scalability:

698 **• SPUS-9M:** This variant contains approximately 9 million parameters. Its architecture is
 699 identical to SPUS-36M except that the number of channels in each encoder and decoder
 700 stage is reduced by half.
 701

702 Table 4: Evaluation (average MSE on all predicted timesteps from initial condition) of SPUS on
 703 downstream datasets under different numbers of finetuned trajectories
 704

Downstream Dataset	Number of Trajectories		
	32	128	256
CE-RPUI	0.0057	0.0054	0.0041
CE-RM	0.0246	0.0159	0.0130
NS-PwC	0.0076	0.0048	0.0025
NS-SL	0.0286	0.0163	0.0004
FNS-KF	0.0098	0.0015	0.0012
Wave-Gauss	0.0097	0.0069	0.0068

713

- 714 • **SPUS-36M:** The baseline configuration illustrated in Figure 2. It includes two residual
 715 blocks at each level of the encoder, decoder, and bottleneck.
- 716 • **SPUS-76M:** This larger variant follows the same architecture as SPUS-36M but increases
 717 the number of residual blocks in each encoder, decoder, and bottleneck stage from two to
 718 five following Huang et al. (2023) .

719

720 All three pretrained models are fine-tuned on three downstream datasets: CE-RM, NS-SL, and FNS-
 721 KF. Table 2 summarizes the average test MSE across all predicted timesteps for each dataset and
 722 model configuration. As shown, increasing model size consistently reduces the test error across for
 723 all three downstream datasets, demonstrating that SPUS scales effectively with model size.

724

725 **A.3 PERFORMANCE EVALUATION OF SPUS WITH DIFFERENT NUMBERS OF RESIDUAL
 726 BLOCKS PER STAGE**

727

728 As shown in Table 3, increasing the number of residual blocks in each stage of SPUS improves
 729 model accuracy. Specifically, the larger variant, SPUS-76M with five residual blocks per encoder,
 730 decoder, and bottleneck stage, achieves lower prediction errors on all three downstream PDE datasets
 731 compared to the configuration with two residual blocks (SPUS-36M), demonstrating the benefit of
 732 deeper residual refinement within each stage.

733

734 **A.4 PERFORMANCE EVALUATION OF SPUS WITH AUTOREGRESSIVE TRAINING UNDER A
 735 FIRST-ORDER MARKOV ASSUMPTION**

736

737 We formulate the autoregressive training of SPUS as a first-order Markov process, where the evo-
 738 lution of the system depends only on its immediately preceding state. This formulation contrasts
 739 with DPOT, which conditions on the previous ten timesteps to predict the next one. Relying on
 740 multiple past states, as in DPOT, to start inference, it requires the 10 past states to be generated from
 741 expensive numerical simulator. On the other hand, SPUS only need the initial condition which has
 742 no computational cost to start the inference.

743

744 Relying on multiple past states, as in DPOT, can accumulate redundant temporal information and
 745 potentially introduce compounding errors during long rollouts. As shown in Table 1, our first-
 746 order Markov formulation enables SPUS to achieve lower prediction errors and improved long-term
 747 stability across all downstream datasets compared to DPOT.

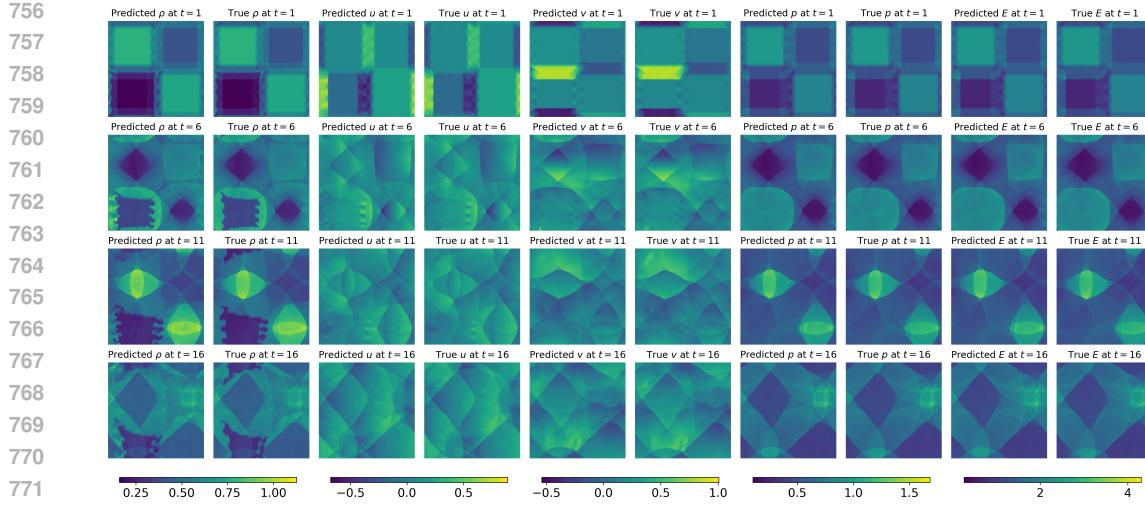
748

749 **A.5 PERFORMANCE EVALUATION OF SPUS ON THE CE-RPUI, CE-RM, FNS-KF AND
 750 WAVE-GAUSS DATASETS**

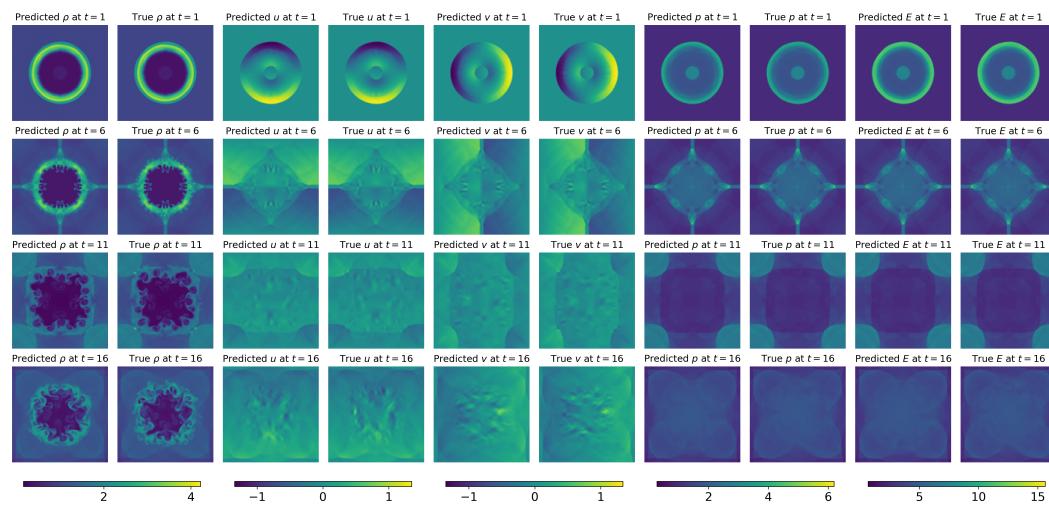
751

752 Figure A.2 presents randomly selected trajectory prediction from the CE-RPUI test dataset. Notably,
 753 the predicted variables closely match the ground truth at each time step, although the deviation
 754 between prediction and ground truth increases more noticeably over time for CE-RM compared to
 755 CE-RPUI due to its more complex dynamics.

756 Figure A.3 presents randomly selected trajectory prediction from the CE-RM test dataset. Notably,
 757 the predicted variables closely match the ground truth at each time step.

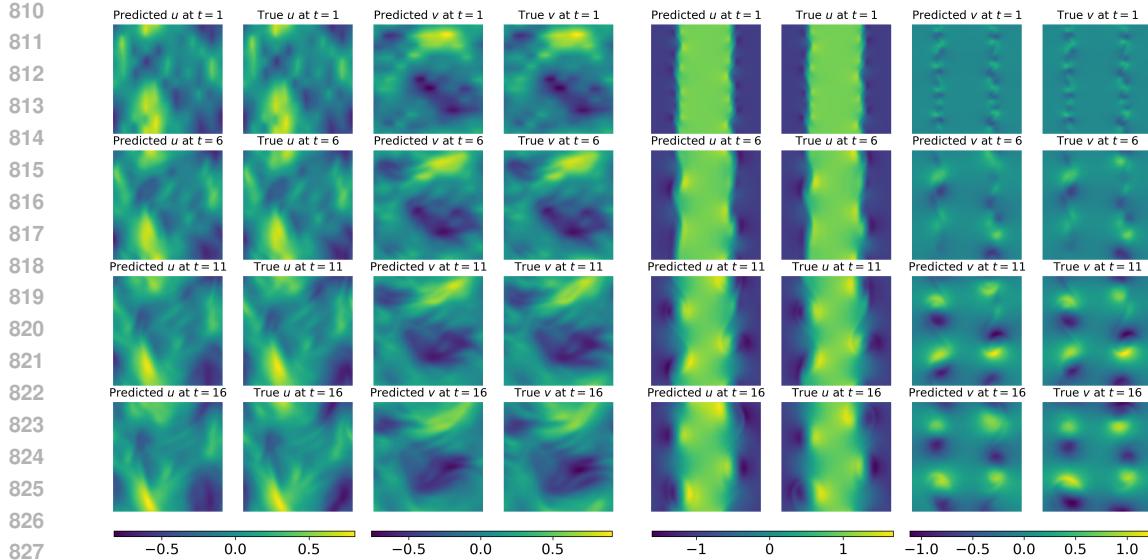


773 Figure A.2: Autoregressive trajectory prediction by SPUS from the initial condition of a randomly
774 selected trajectory in the CE-RPUI testing dataset (240 test trajectories). The figure shows example
775 results at time steps $t = 1, 6, 11, 16$ for five system variables: density ρ , horizontal velocity u ,
776 vertical velocity v , pressure p , and energy E . SPUS takes the initial condition $u'_t = u_{t=0}$ as input
777 and recursively predicts subsequent states based on its own previous outputs for $t = 1, \dots, 20$, as
778 described in Figure 1 (inference step). As shown, the predicted variables closely match the ground
779 truth at each time step.



797 Figure A.3: Autoregressive trajectory prediction by SPUS from the initial condition of a randomly
798 selected trajectory in the CE-RM testing dataset (130 test trajectories). The figure shows example
799 results at time steps $t = 1, 6, 11, 16$ for five system variables: density ρ , horizontal velocity u ,
800 vertical velocity v , pressure p , and energy E . SPUS takes the initial condition $u'_t = u_{t=0}$ as input
801 and recursively predicts subsequent states based on its own previous outputs for $t = 1, \dots, 20$, as
802 described in Figure 1 (inference step). As shown, the predicted variables closely match the ground
803 truth at each time step, although the deviation between prediction and ground truth increases more
804 noticeably over time for CE-RM compared to CE-RPUI due to its more complex dynamics.

805
806
807 Figure A.4 presents randomly selected trajectory prediction from the NS-PwC and NS-SL test
808 dataset. Notably, although SPUS was not pretrained on incompressible Navier–Stokes dynamics,
809 its predictions closely follow the ground truth variables at each time step, demonstrating robust
generalization to the NS-PwC and NS-SL system.



(a) Autoregressive trajectory prediction from initial state for NS-PwC

(b) Autoregressive trajectory prediction from initial state for NS-SL

Figure A.4: Autoregressive trajectory prediction by SPUS from the initial condition of a randomly selected trajectory in the NS-PwC and NS-SL testing datasets (each with 240 test trajectories). The figure shows example results at time steps $t = 1, 6, 11, 16$ for two system variables: horizontal velocity u , vertical velocity v . As shown, despite not being exposed to incompressible NS dynamics during pretraining, the predicted variables by SPUS closely match the ground truth at each time step. The deviation between prediction and ground truth (GT) increases more noticeably over time for NS-PwC compared to NS-SL.

Figure A.5 presents randomly selected trajectory prediction from the FNS-KF test dataset. Notably, although SPUS was not pretrained on incompressible Navier–Stokes dynamics, its predictions closely follow the ground truth variables at each time step, demonstrating robust generalization to the FNS-KF system.

Figure A.6 illustrates a representative trajectory prediction from the Wave-Gauss test set, demonstrating SPUS’s ability to generalize to the Wave-Gauss system. However, in comparison to the Navier–Stokes downstream tasks, the deviation between the predicted variables and the ground truth increases more noticeably over time for Wave-Gauss.

A.6 VISUALIZATION OF SPUS PERFORMANCE ON TRAJECTORY PREDICTION COMPARED WITH POSEIDON AND DPOT

Figure A.8 shows a random trajectory predictions for FNS-KF made by SPUS (36M), POSEIDON (158M), and DPOT (122M). Each model is finetuned with 128 trajectories.

Figure A.9 shows a random trajectory predictions for NS-SL made by SPUS (36M), POSEIDON (158M), and DPOT (122M). Each model is finetuned with 128 trajectories.

A.7 CAN SPUS GENERALIZE TO DIFFERENT INPUT RESOLUTIONS?

We fine-tuned SPUS-36M using 128 PDE trajectories governed by the Burgers equation (Zhou & Farimani, 2024). Each trajectory consists of 100 time steps, with the velocity field represented on a spatial grid of size 64×64 . For comparison, we also trained an unpretrained U-Net model with 36M parameters—sharing the same architecture as SPUS—from scratch, using the same 128 trajectories.

We evaluated both models on 64 test trajectories, predicting all 99 future time steps from the initial condition. The average mean squared error (MSE) across all 99 predicted time steps for the 64 trajectories was 0.0035 for SPUS and 0.0079 for the unpretrained U-Net. These results demonstrate

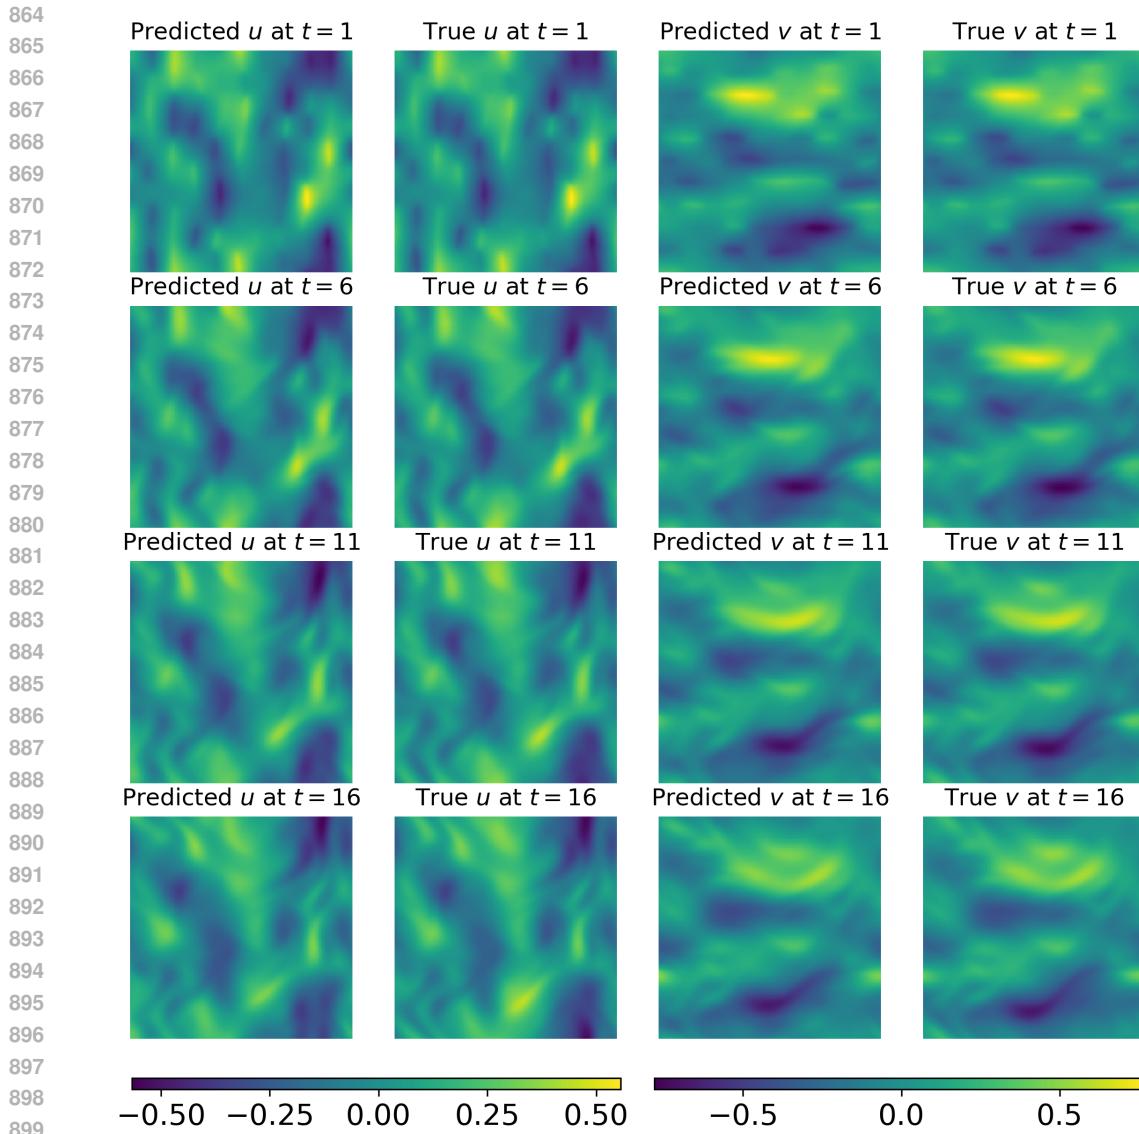
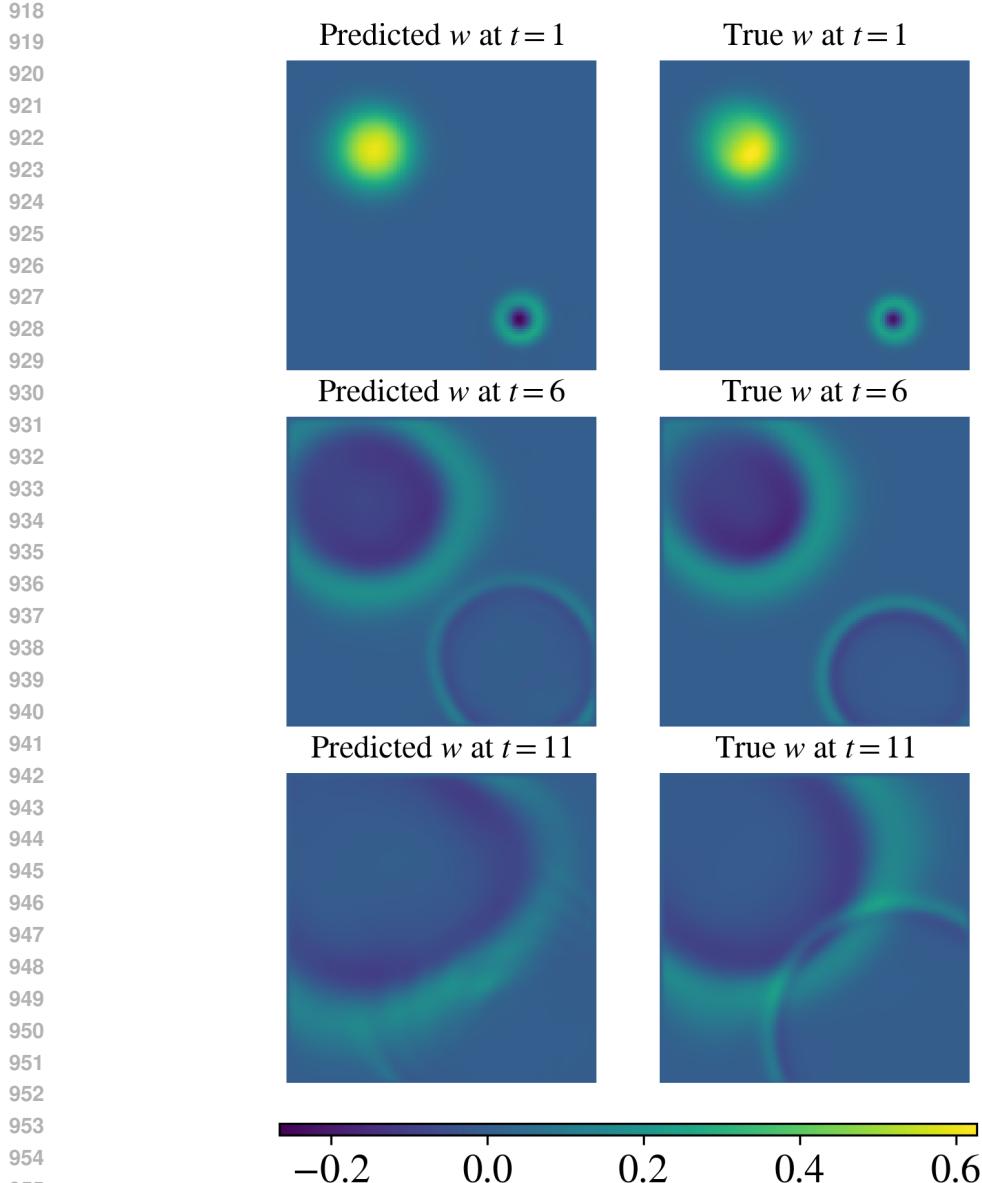


Figure A.5: Autoregressive trajectory prediction by SPUS from the initial condition of a randomly selected trajectory in the FNS-KF testing dataset (240 trajectories). The figure shows example results at time steps $t = 1, 6, 11, 16$ for two system variables: horizontal velocity u , vertical velocity v . SPUS takes the initial condition $u'_t = u_{t=0}$ as input and recursively predicts subsequent states based on its own previous outputs for $t = 1, \dots, 20$, as described in Figure 1 (inference step). As shown, the predicted variables closely match the ground truth at each time step.

the strong generalization capability of SPUS to different spatial resolutions and its robustness for longer rollouts. Figure A.10 shows a randomly selected trajectory prediction for PDEs governed by the Burgers equation, generated by SPUS (36M) and an unpretrained U-Net (36M). The figure shows example results at time steps $t = 21, 31, 41, 51, 99$ for the velocity field on a 64×64 spatial grid. As shown, the unpretrained U-Net begins to deviate from the ground truth (GT) after $t = 41$, whereas SPUS remains close to the GT across all time steps, demonstrating its potential for longer rollouts. Figure A.11 shows the average MSE over 99 predicted time steps for 64 test trajectories of Burgers equation. As shown, SPUS (36M) maintains a nearly constant MSE after timestep 40, remaining stable through the final prediction. In contrast, the unpretrained U-Net (36M) exhibits a steadily increasing MSE.



957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
Figure A.6: Autoregressive trajectory prediction by SPUS from the initial condition of a randomly selected trajectory in the Wave-Gauss testing dataset (240 trajectories). The figure shows example results at time steps $t = 1, 6, 11$ for one system variable: wave speed u . SPUS takes the initial condition $u'_t = u_{t=0}$ as input and recursively predicts subsequent states based on its own previous outputs for $t = 1, \dots, 14$, as described in Figure 1 (inference step). As shown, the predicted variables deviates very quickly from ground truth at each time step for Wave-Gauss compared to other downstream tasks.

A.8 CAN SPUS SCALE TO PDES WITH HIGHER SPATIAL RESOLUTIONS THAN THOSE USED IN PRETRAINING?

SPUS is pretrained on PDEs with a spatial resolution of 128×128 . To evaluate its scalability to higher-resolution PDEs, we finetuned the pretrained SPUS model using 128 trajectories governed by the Burgers equation (Zhou & Farimani, 2024). Each trajectory contains 20 time steps, with the velocity field represented on a 256×256 spatial grid. For comparison, we also trained a 36M-parameter U-Net—matching the architecture of SPUS—from scratch using the same 128 trajectories.

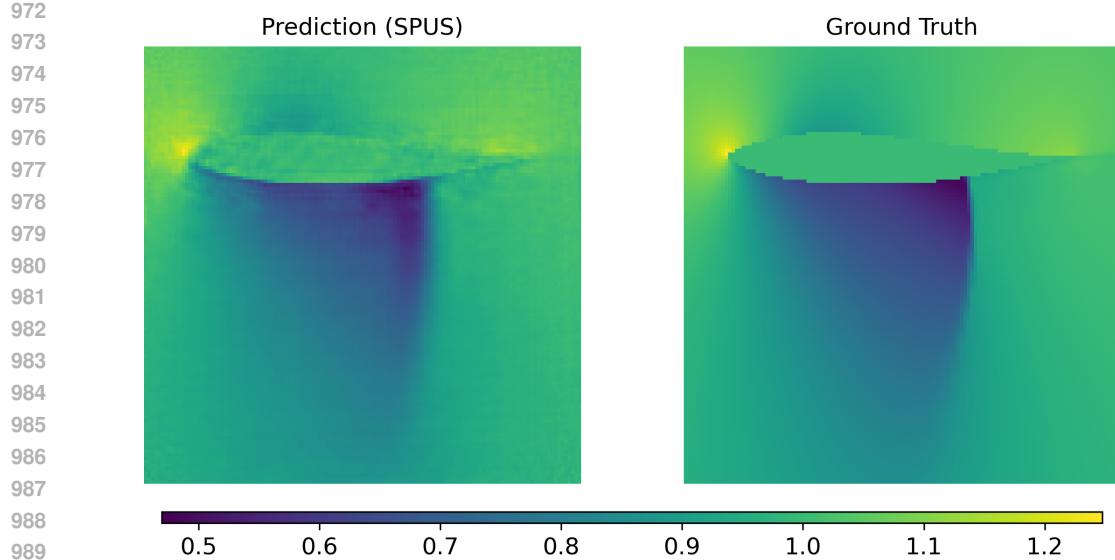
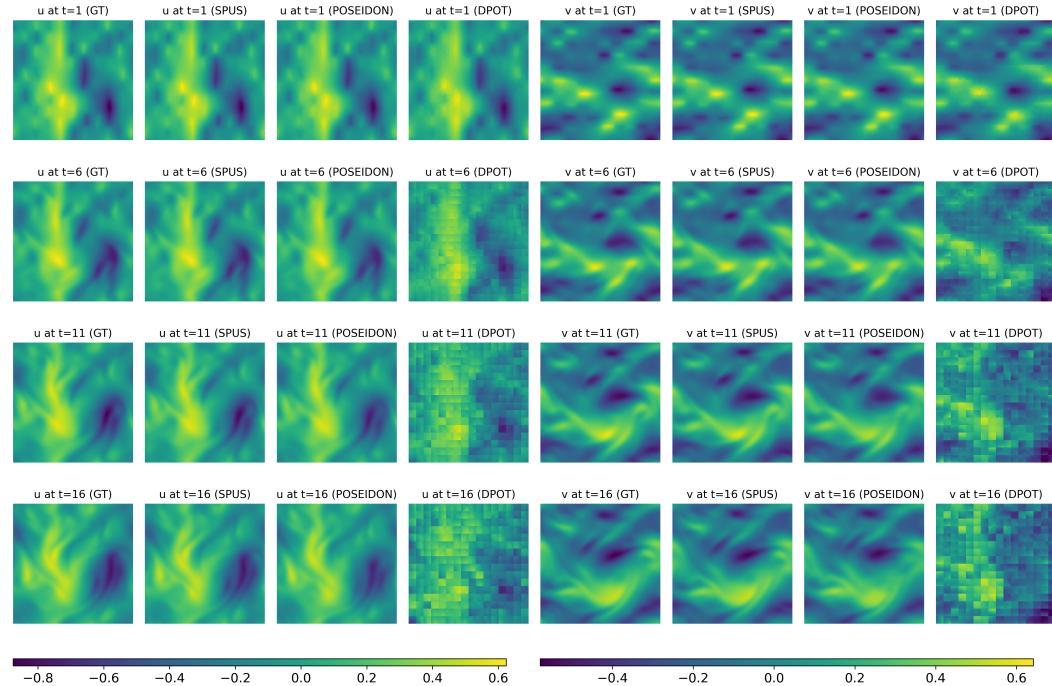


Figure A.7: A randomly selected predicted samples by SPUS for SE-AF

Figure A.8: A random trajectory predictions for FNS-KF made by SPUS (36M), POSEIDON (158M), and DPOT (122M). The figure shows example results at time steps $t = 1, 6, 11, 16$ for two system variables: horizontal velocity u , vertical velocity v .

1019
1020
1021
1022
1023
1024
1025

We evaluated both models on 64 test trajectories, predicting all 19 future time steps from the initial condition. The average mean squared error (MSE) across the 19 predicted time steps was 0.0017 for SPUS and 0.0023 for the U-Net trained from scratch. These results demonstrate the strong generalization capability of SPUS to spatial resolutions higher than those used during pretraining.

Figure A.12 shows a randomly selected trajectory prediction for the Burgers equation generated by SPUS (36M) and the unpretrained U-Net (36M). Example velocity-field snapshots at time steps

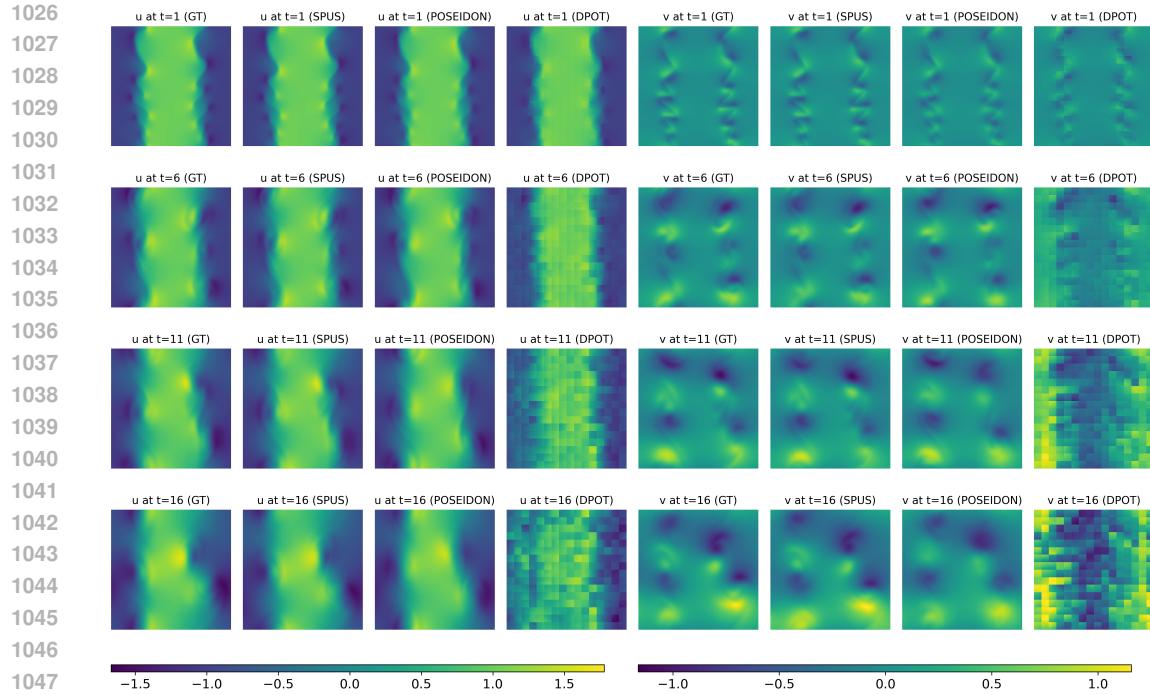


Figure A.9: A random trajectory predictions for NS-SL made by SPUS (36M), POSEIDON (158M), and DPOT (122M). The figure shows example results at time steps $t = 1, 6, 11, 16$ for two system variables: horizontal velocity u , vertical velocity v .

$t = 2, 6, 11, 16, 19$ on the 256×256 grid are presented. As illustrated, SPUS remains close to the ground truth across all time steps, whereas the U-Net trained from scratch develops noticeable periodic striping artifacts at later time steps that are absent in the ground-truth (GT) solution. These findings show that SPUS successfully scales to PDEs with higher spatial resolutions than those seen during pretraining.

1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

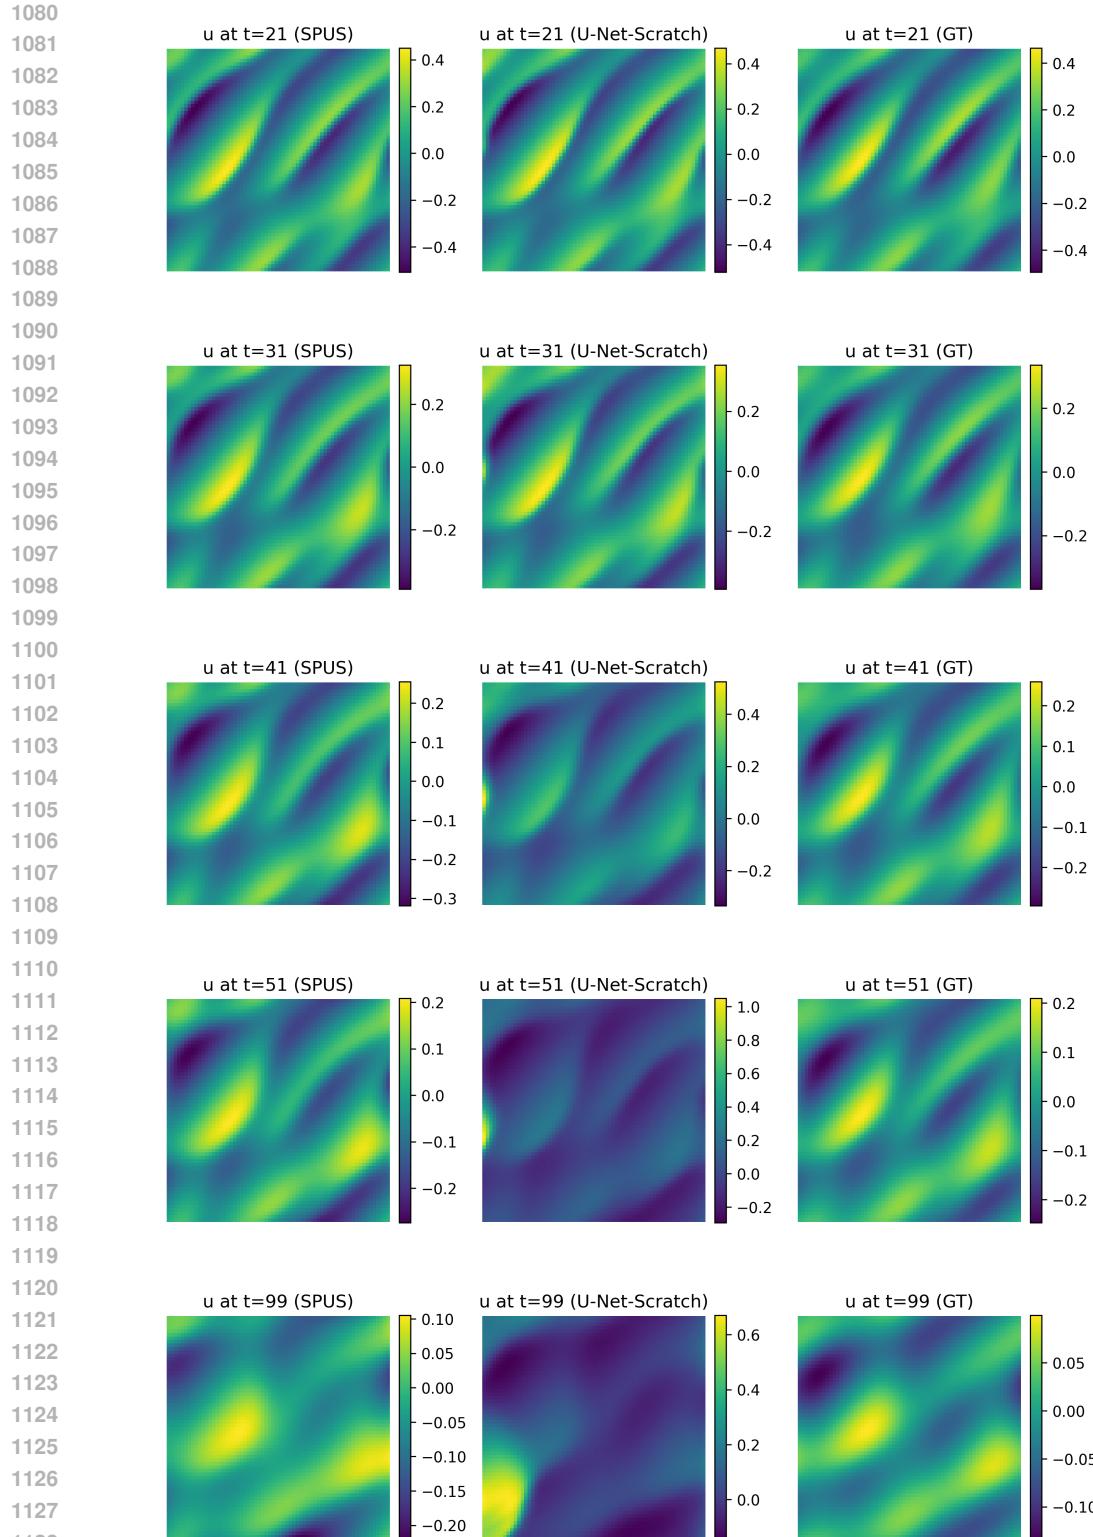


Figure A.10: A randomly selected trajectory prediction for PDEs governed by the Burgers equation, generated by SPUS (36M) and an unpretrained U-Net (36M). The figure shows example results at time steps $t = 21, 31, 41, 51, 99$ for the velocity field on a 64×64 spatial grid. As shown, the unpretrained U-Net begins to deviate from the ground truth (GT) after $t = 41$, whereas SPUS remains close to the GT across all time steps, demonstrating its potential for longer rollouts.

```

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

```

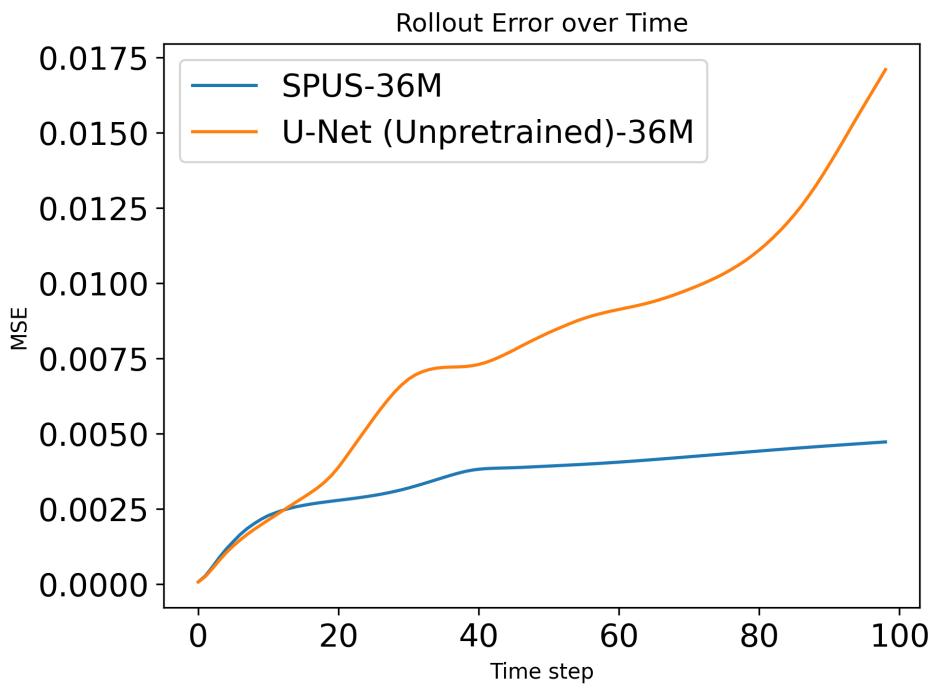


Figure A.11: Average MSE over 99 predicted time steps for 64 test trajectories of Burgers equation. As shown, SPUS (36M) maintains a nearly constant MSE after timestep 40, remaining stable through the final prediction. In contrast, the unpretrained U-Net (36M) exhibits a steadily increasing MSE.

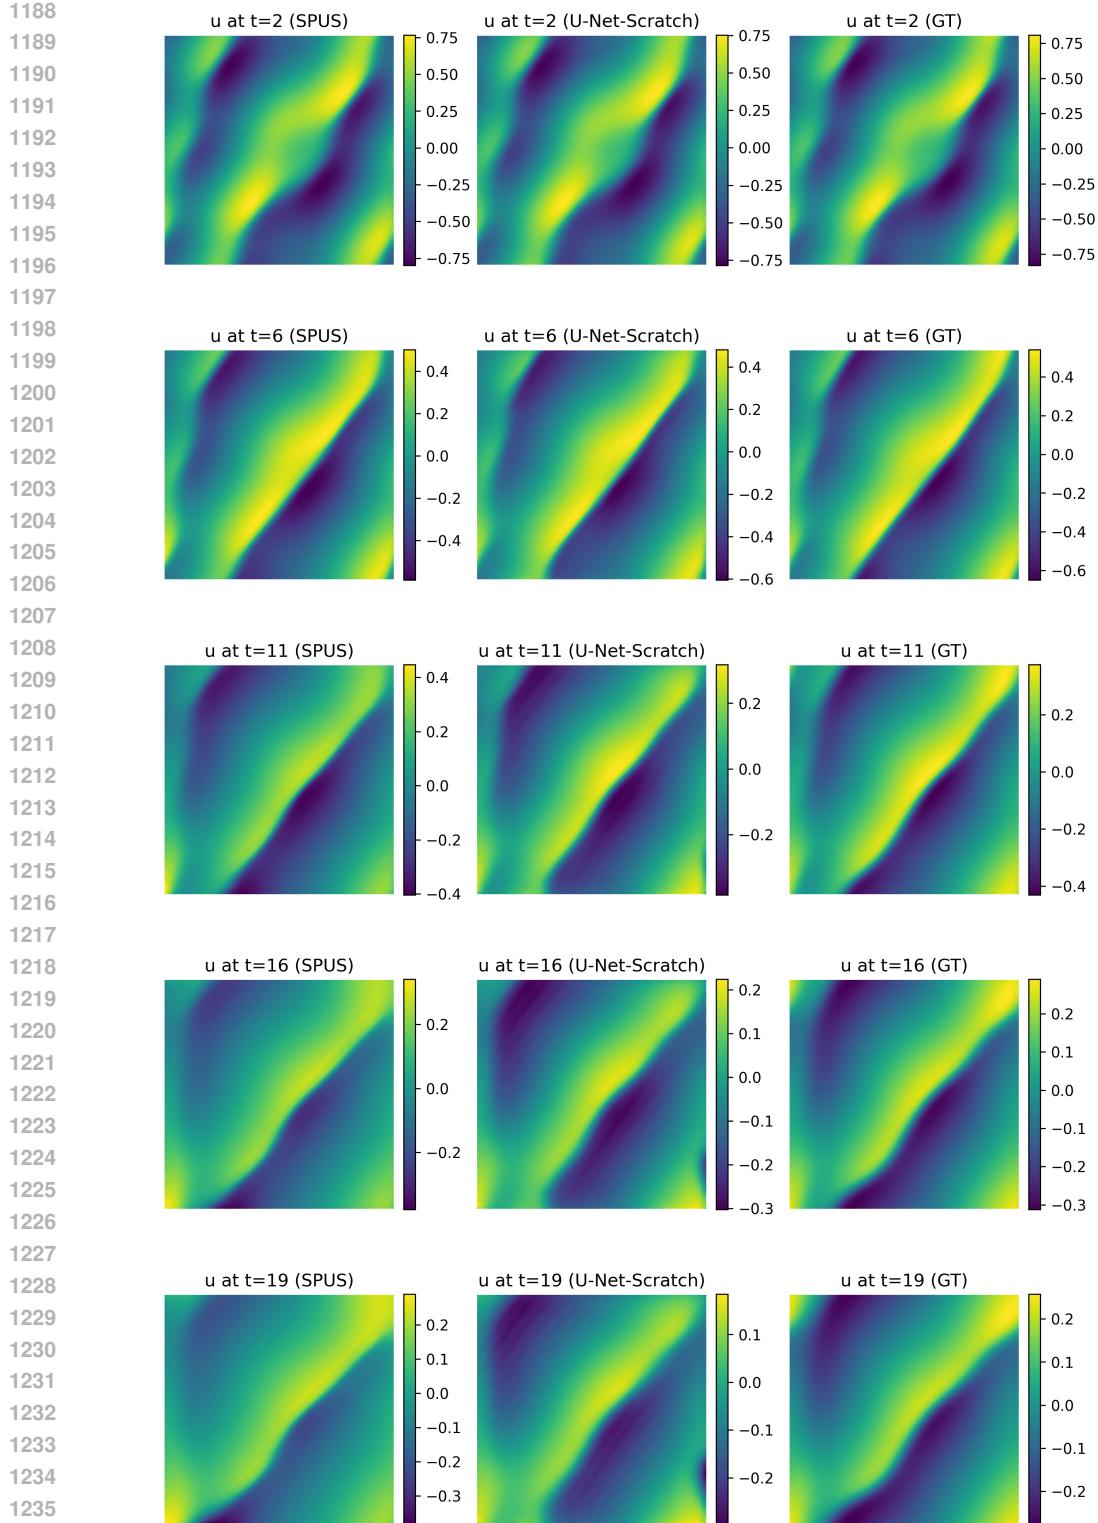


Figure A.12: Scalability of SPUS to PDE resolutions higher than those used in pretraining (128×128). Shown is a randomly selected trajectory prediction for PDEs governed by the Burgers equation, generated by SPUS (36M) and an unpretrained U-Net (36M). The figure shows example results at time steps $t = 2, 6, 11, 16, 19$ for the velocity field on a 256×256 spatial grid. As illustrated, SPUS remains close to the ground truth across all time steps, whereas the U-Net trained from scratch develops noticeable periodic striping artifacts at later time steps that are absent in the ground-truth (GT) solution.