

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 THE FALSE PROMISE OF ZERO-SHOT SUPER- RESOLUTION IN MACHINE-LEARNED OPERATORS

Anonymous authors

Paper under double-blind review

ABSTRACT

A core challenge in scientific machine learning, and scientific computing more generally, is modeling continuous phenomena which (in practice) are represented discretely. Machine-learned operators (MLO) have been introduced as a means to achieve this modeling goal, as this class of architecture can perform inference at arbitrary resolution. In this work, we evaluate whether this architectural innovation is sufficient to perform “zero-shot super-resolution,” namely to enable a model to serve inference on higher-resolution data than that on which it was originally trained. We comprehensively evaluate both zero-shot sub-resolution and super-resolution (i.e., multi-resolution) inference in MLOs. We decouple multi-resolution inference into two key behaviors: 1) extrapolation to varying frequency information; and 2) interpolating across varying resolutions. We empirically demonstrate that MLOs fail to do both of these tasks in a zero-shot manner. Consequently, we find MLOs are *not* able to perform accurate inference at resolutions different from those on which they were trained, and instead they are brittle and susceptible to aliasing. To address these failure modes, we propose a simple, computationally-efficient, and data-driven multi-resolution training protocol that overcomes aliasing and that provides robust multi-resolution generalization.

1 INTRODUCTION

Modeling physical systems governed by partial differential equations (PDEs) is critical to many computational science workflows:

$$S_2 = M(S_1), \quad (1)$$

where M is an approximation of the PDE’s solution operator, S_1 is the input state of the system, and S_2 is the predicted state. Central to this problem formulation is that *continuous* physical systems must be sampled and, therefore, modeled *discretely*. For a discrete model, M , to be useful in representing phenomena of different scales, scientists require the ability to use it at different resolutions accurately. For example, when modeling fluid flow, scientists often use adaptive mesh refinement (Berger & Oliger, 1984), a technique that increases simulation resolution in areas that require high accuracy (e.g., regions of turbulence), and coarsens it in less critical regions.

Traditionally, the approximation M is constructed by numerical methods which, by design, can be employed at arbitrary discretization (Forrester et al., 2008; Cozad et al., 2014; Asher et al., 2015; Sudret et al., 2017; Alizadeh et al., 2020; Kudela & Matousek, 2022). However, numerical methods are computationally expensive. Alternatively, machine-learned operators (MLOs), a class of data-driven machine learning (ML) models which parameterize the solution operator to families of PDEs, have been proposed (Raissi et al., 2019; Li et al., 2020a; Lu et al., 2021b; Kovachki et al., 2023; Raonic et al., 2023). Although querying MLOs at arbitrary discretization is computationally inexpensive, it is not obvious that this can be done *accurately*. The Fourier Neural Operator (FNO) (Li et al., 2020a), a specific MLO, claimed to address the discretization challenge in a *zero-shot* manner (Li et al., 2020a; Tran et al., 2021; George et al., 2024; Li et al., 2024b; Azizzadenesheli et al., 2024). The claim is that FNO can be trained at resolution m and then serve accurate inference at resolution $n > m$, without training on additional high resolution data e.g., zero-shot super-resolution. This claim of zero-shot super-resolution, if true, is especially attractive in settings where generating, and training on, high-resolution data is computationally expensive.

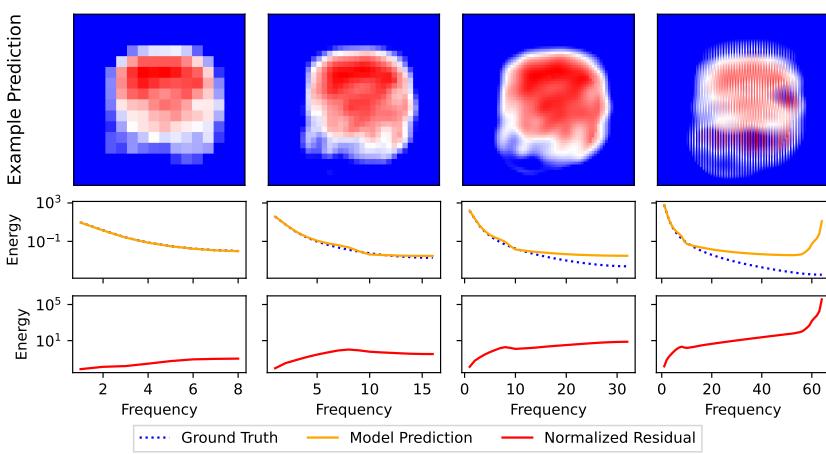


Figure 1: **Aliasing in zero-shot super-resolution.** Model trained on resolution 16 data, and evaluated at varying resolutions: 16, 32, 64, 128. **Top Row:** Sample prediction for Darcy flow; notice striation artifacts at resolution 128. **Middle Row:** Average test set 2D energy spectrum of label and model prediction. **Bottom Row:** Average residual spectrum normalized by label spectrum.

In this paper, we evaluate the claim of zero-shot super- (and sub-)resolution inference in MLOs. We document a substantial disparity in model performance across data of different discretizations, suggesting that MLOs are generally incapable of accurate inference at resolutions greater or less than their training resolution (i.e., zero-shot multi-resolution inference). Instead, we find that MLOs often misrepresent unseen frequencies and incorrectly infer their behavior in data whose discretization differs from its training discretization, i.e., they exhibit a form of aliasing (Fig. 1). In addition, we study two previously proposed solutions: (i) physics-informed optimization constraints (Li et al., 2024b) and (ii) band-limited learning (Raonic et al., 2023; Gao et al., 2025). We find that neither enables zero-shot multi-resolution, as they do *not* address the central issue: MLOs, like all machine-learned models, cannot typically generalize beyond their training data (Yang et al., 2023; Liu et al., 2020; Krueger et al., 2021). We establish that the discretization at which MLOs are trained impacts the discretization at which they accurately model the system.

To enable multi-resolution inference, we propose *multi-resolution training*, a simple, intuitive, and principled data-driven approach which trains models on data of multiple resolutions. We profile the impact of different multi-resolution training approaches, finding that optimal multi-resolution performance can often be achieved via training data sets that contain mostly low resolution (less expensive) data and very little high resolution (more expensive) data. This permits us to achieve low computational overhead, while also increasing the utility of a single MLO.

To summarize, the main contributions of our work are the following:

1. We assess the ability of trained MLOs to generalize beyond their training resolution. We demonstrate that MLOs struggle to perform accurate inference at resolutions higher or lower than which are they trained on, and instead they exhibit aliasing. Based on these results, we conclude that accurate *zero-shot* multi-resolution inference is unreliable (Sec. 3).
2. We evaluate two intuitive approaches—incorporating physics-informed constraints during training, as well as performing band-limited learning—and we find that neither approach enables reliable multi-resolution generalization. (Sec. 4).
3. We propose and test multi-resolution training, where we include training data of varying resolutions (in particular, a small amount of expensive higher-resolution data and a larger amount of cheaper lower-resolution data), and we show that multi-resolution inference improves substantially, without a significant increase in training cost (Sec. 5).

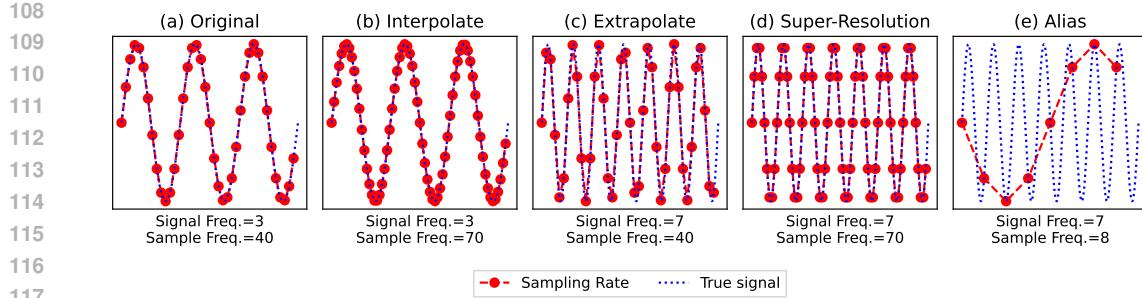


Figure 2: Accurate multi-resolution inference requires both interpolation & extrapolation.

Original: Signal is sampled at a rate greater than its Nyquist frequency. **Interpolation:** Adapting to new sampling rates of a given signal. **Extrapolation:** Adapting to new frequency information under constant sampling rate. **Super-Resolution:** Sampling a system at a higher rate which enables the capture of higher frequency information (interpolation & extrapolation). **Aliasing:** High-frequency information is misrepresented as a low-frequency information due to insufficient sampling.

2 BACKGROUND ON SIGNAL PROCESSING AND ALIASING

We start by discussing the practice of training ML models to represent continuous systems via discrete data. Next, we outline the implications of aliasing in ML as it relates to multi-resolution inference. Finally, we formally define “zero shot multi-resolution” inference in a discrete context.

Discrete Representations of Continuous Systems. The fundamental challenge in discretely representing continuous systems lies in the choice of sampling rate. The Whittaker–Nyquist–Shannon sampling theorem established that given a sampling rate r , the largest resolved frequency is $r/2$ (Unser, 2002; Shannon, 1949; Whittaker, 1928). Thus, ML models will be trained on discrete representations where only some of the frequencies are fully resolved. Resolving higher-order frequencies greater than $r/2$, consequently, becomes an out-of-distribution task. Aligning these discrete models’ predictions with the underlying continuous system is an open problem (Krishnapriyan et al., 2021; Queiruga et al., 2020; Ott et al., 2021; Ren et al., 2023; Takamoto et al., 2022; Subel et al., 2022; Chattopadhyay et al., 2024).

Aliasing. When sampling a continuous signal at rate r , aliasing occurs when frequency components greater than $r/2$ are projected onto lower frequency basis functions (Fig. 2) (Gruver et al., 2022). Thus, content with frequency $n > r/2$ is observed as a lower frequency contribution:

$$\text{Alias}(n) = \begin{cases} n \bmod r & \text{if } (n \bmod r) < r/2 \\ r - (n \bmod r) & \text{if } (n \bmod r) > r/2 \end{cases} \quad (2)$$

In an ML context, when inferring at different discretizations of a given signal, aliasing can manifest as the divergence between the energy spectrum of the model prediction and the expected output. Aliasing indicates a model’s failure to fit the underlying continuous system.

Zero-shot multi-resolution inference. We define *multi-resolution* inference as the ability to do inference at multiple resolutions (e.g., sub- and super-resolution). The *zero-shot* multi-resolution task employs an ML model, which is trained on data with some resolution and tested on data with a different resolution. Zero-shot multi-resolution inference raises two important questions with respect to the generalization abilities of trained models (see Fig. 2):

1. **Resolution Interpolation.** How do models behave when the frequency information in the data remains fixed, but its sampling rate changes from training to inference? Can the model *interpolate* the fully resolved signal to varying resolutions?
2. **Information Extrapolation.** How do models behave when the resolution remains fixed, but the number of fully resolved frequency components changes from training to inference? For super-resolution, this means can the model *extrapolate* beyond the frequencies in its training data and model higher frequency information?

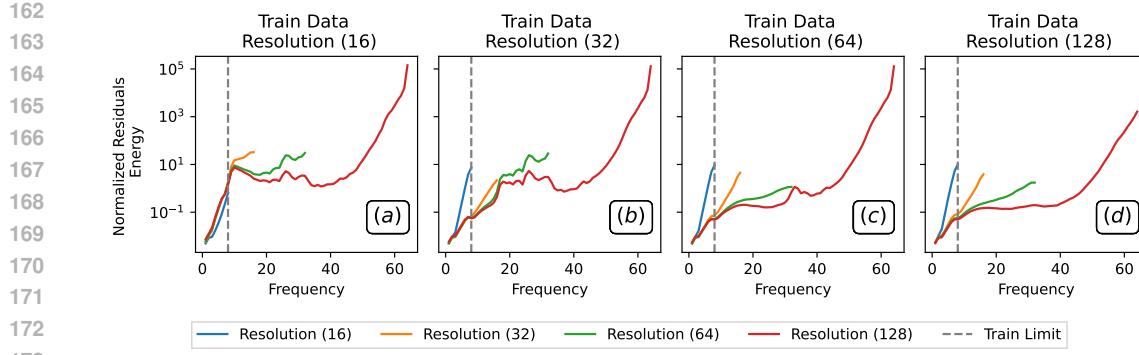


Figure 3: **Resolution Interpolation.** Four FNOs are trained on Darcy data at resolutions $\{16, 32, 64, 128\}$ from left to right with constant frequency information (low-pass limit of $8f$), and are tested on varying resolutions. We assess if each model can generalize to data with varying sampling rate. We visualize spectra of the normalized residuals across test data. Notice, residual spectra (error) increases substantially in the low frequencies. Lower residual energy at all frequencies is better.

3 ASSESSING MULTI-RESOLUTION GENERALIZATION

We examine the zero-shot multi-resolution abilities of FNO, an architecture for which this claim has been previously made (Li et al., 2020a). We study the multi-resolution inference task from an out-of-distribution data perspective by decoupling what it means for a model to perform inference at a resolution different from that used during training. Specifically, in Sec. 3.1, we assess whether models trained on a system sampled at rate r are capable of both interpolating accurately to **new** sampling rates (Fig. 2(b)) *and* extrapolating accurately to **additional/fewer** frequencies present in data (Fig. 2(c)). We systematically test an FNO’s ability to do both objectives and show neither are achieved. In light of these failure modes, in Sec. 3.2, we then assess the spatial zero-shot sub- and super-resolution capabilities of FNOs and show the claim does not hold (Fig. 2(d)).

We evaluate FNO on three standard scientific datasets: Darcy, Burgers, and Navier Stokes. For each dataset we optimize FNO hyperparameters via grid-search as described in Appendix A.

3.1 BREAKING DOWN MULTI-RESOLUTION CAPABILITIES

Resolution Interpolation. We assess if an FNO trained on data of a specific resolution can generalize to data of both lower and higher resolution under fixed frequency information. Specifically, we keep the set of populated frequencies constant in the train and test data, while varying the resolution of the test data. We do this by applying a low-pass filter to all data at the highest resolution, and then down-sampling as needed. The sampling rates of all data are large enough to resolve all remaining frequencies.

We begin with a simple experiment: we train an FNO on a Darcy flow dataset at resolution $N = 16$ and assess the trained model’s performance across test datasets at varying resolutions $\{16, 32, 64, 128\}$, all low-pass filtered with limit $8f$ where f is the frequency unit $2\pi/N$. In Fig. 3(a), we visualize the average spectral energy of the model predictions normalized by the spectral energy of the unfiltered ground truth for each test dataset. For the test data with resolutions that are different from the training data, we observe sharp increases in their residual’s energy spectrum in frequencies greater than $8f$. This is especially concerning since the model was never trained on nor shown inference data containing frequencies greater than $8f$. In other words, FNOs, trained in a zero-shot manner, fail to reliably *interpolate* to varying resolutions.

In Figs. 3(b-d), we repeat the same experiment at varying training resolutions (e.g. 32, 64, 128) with low-pass limit $8f$. We notice that at each training resolution, the model consistently and incorrectly assigns high energy in frequencies greater than $8f$ across all test resolutions. We perform corresponding experiments for the Burgers dataset with low-pass limit $64f$ and resolutions $\in \{128, 256, 512, 1024\}$ and Navier Stokes dataset with low-pass limit $32f$ and resolutions $\in \{64, 128, 255, 510\}$ and observe the same failure mode (Appendix D: Figs. 11-12). **We conclude that changing**

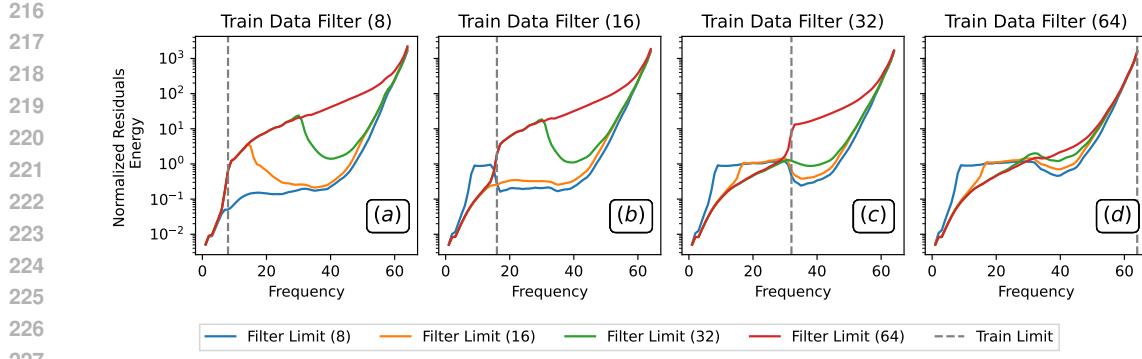


Figure 4: **Information Extrapolation.** Four FNOs trained on Darcy data of resolution 128 (constant sampling rate) and low-pass filtered with limits $\{8, 16, 32, 64\}f$ (varying frequency information) from left to right. Test if each model can generalize to data with varying frequency information. Visualizing spectrum of the normalized residuals across test data. Notice, residual spectra (error) increases substantially in the high frequencies. Lower residual energy at all frequencies is better.

resolution at test time is akin to out-of-distribution inference: the model is never trained on data with a broad range of sampling rates and consequently fails to generalize.

Information Extrapolation. We assess if an FNO trained on data containing a fixed set of frequencies can generalize to data containing both fewer and additional frequencies under fixed resolution. Specifically, we keep the resolution constant but vary the number of populated frequencies in the train and test datasets by applying varying low-pass filters to data at a fixed resolution.

We begin with a simple experiment: an FNO is trained on a Darcy flow dataset at resolution 128 which is low-pass filtered at limit $8f$. In Fig. 4(a), we assess the trained model’s performance across four versions of a test dataset, all of which have the same sampling resolution (128) but are filtered with low-pass limits $\{8, 16, 32, 64\}f$ (e.g., increasing amounts of frequency information). There is a sharp increase in the residual spectra in higher frequencies across all test sets; the residual error increases as the test and training filters diverge. In other words, FNOs, trained in a zero-shot manner, fail to *extrapolate* on data with previously unseen frequency information.

In Figs. 4(b-d), we repeat the same experiment at varying training data low-pass filter limits (e.g., 16, 32, 64f). Each model consistently and incorrectly assigns high energy in the high frequencies regardless of whether the test data contained any high-frequency information. This is a concerning failure mode indicating FNOs do not generalize both in the presence of frequencies greater than, and the absence of frequencies less than, what was present in their training data. We perform corresponding experiments for the Burgers dataset with resolution 1024 and low-pass limits $\in \{64, 128, 256, 512\}f$ and Navier Stokes dataset with resolution 510 and low-pass limits $\in \{8, 16, 32, 255\}f$ and observe the same failure mode (Appendix D: Figs. 13-14). **We conclude that varying the frequency information at test time is effectively out-of-distribution inference** as the model was not trained on data with such variation in frequency information and, therefore, fails to generalize.

3.2 ZERO-SHOT MULTI-RESOLUTION INFERENCE

We study whether FNOs are capable of spatial multi-resolution inference: simultaneously changing the sampling rate and frequency information. For each dataset in {Darcy, Burgers, Navier Stokes}, we train a model on data at resolutions $(16, 32, 64, 128), (1024, 512, 256, 128), (510, 255, 128, 64)$, respectively. In Fig. 1, we see that models trained at low resolutions do not generalize to high resolutions. Similarly, in Fig. 5, we again see a failure to generalize and instead observe high-frequency artifacts in model predictions in multi-resolution settings. Further, for time-varying PDEs, such as Navier Stokes, we observe that these high frequency aliasing artifacts compound across time steps (Fig. 15). In Fig. 16, we show that models trained at a given resolution do not achieve low loss across all test resolutions. In fact, losses vary by $1\times$, $2\times$, and $10\times$ across test resolutions for the Darcy, Burgers, and Navier Stokes datasets, respectively. **Therefore, we conclude that Fourier neural operators are not capable of consistent zero-shot super- or sub-resolution.**

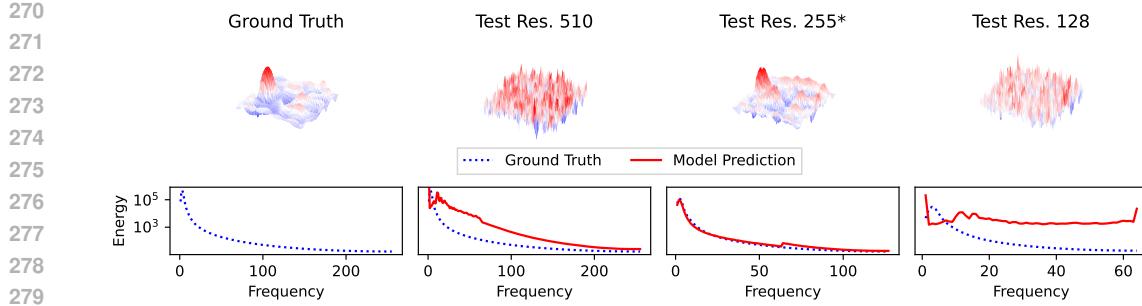


Figure 5: **FNOs do not generalize to higher or lower resolutions.** Model trained on Navier Stokes dataset at resolution 255 (indicated by *), evaluated resolutions 510, 255, 128. **Top Row:** Ground truth, prediction at resolutions 510 (super-resolution), resolution 255 (same as train resolution), and resolution 128 (sub-resolution). **Bottom Row:** Average energy spectra over test data.

4 EVALUATING POTENTIAL CORRECTIVE METHODS

Here we assess two previously proposed strategies for accurate zero-shot multi-resolution inference: physics-informed optimization objectives (Li et al., 2024b) and band-limited learning (Raonic et al., 2023; Gao et al., 2025). We find that neither enables accurate multi-resolution inference.

4.1 PHYSICAL OPTIMIZATION CONSTRAINTS

Physics-informed optimization constraints have been proposed as a means of achieving accurate inference in the *zero-shot* super-resolution setting (Li et al., 2024b). For each dataset in {Darcy, Burgers, Navier Stokes}, we train FNOs at all available resolutions. We optimize each set of model parameters θ with a dual optimization objective $\mathcal{L}(\theta) = (1 - w) * \ell_{\text{data}}(\theta) + w * \ell_{\text{phys}}(\theta)$, where ℓ_{data} is the original data-driven loss (mean squared error) with an additional physics-informed loss ℓ_{phys} , which explicitly enforces that the governing partial differential equation is satisfied. We use the physics losses of Li et al. (2024b) and detail the implementation in Appendix G.

We find that the data-driven loss always outperforms any training objective that includes a physics constraint (Fig. 19). We determine this by evaluating $w \in \{0, 0.1, 0.25, 0.5\}$ for {Darcy, Burgers, Navier Stokes} at resolutions 64, 512, 255, respectively. Among the objectives that contain a physics constraint, we observe a clear trend: the lower the physics constraint is weighted, the better test performance the model achieves. This indicates that physics constraints make it more challenging for the model to be trained optimally despite extensive hyperparameter tuning (perhaps due to practical reasons such as being difficult to optimize (Krishnapriyan et al., 2021; Subramanian et al., 2022; Gao & Wang, 2023; Wang et al., 2023)).

To further illustrate, we use the smallest $w = 0.1$ and compare the physics-informed optimization with solely data-driven optimization. In Fig. 6, the predicted spectra of data from models optimized with physics loss generally diverge more substantially across test resolutions than models optimized with only a data loss. Models optimized with physics constraints even fail to accurately fit their training distributions (Fig. 6(c)), and fail to generalize to both super- and sub-resolution data (Fig. 6(a,b,d)). See Appendix G for results on all datasets. We conclude that physics informed constraints do *not* reliably enable multi-resolution generalization.

4.2 BAND-LIMITED LEARNING

We study two approaches which propose learning band-limited representations of data: convolutional neural operators (CNO) (Raonic et al., 2023) and the Cross-Resolution Operator-Learning (CROP) pipeline (Gao et al., 2025). Both CNO and CROP have been proposed as alias-free solutions to enable multi-resolution inference (Bartolucci et al., 2023). CNOs are *fixed-resolution* models; to use them, one must first interpolate their input to the model’s training resolution, do a forward pass, and then interpolate back to the original dimension. The CROP pipeline is more gen-

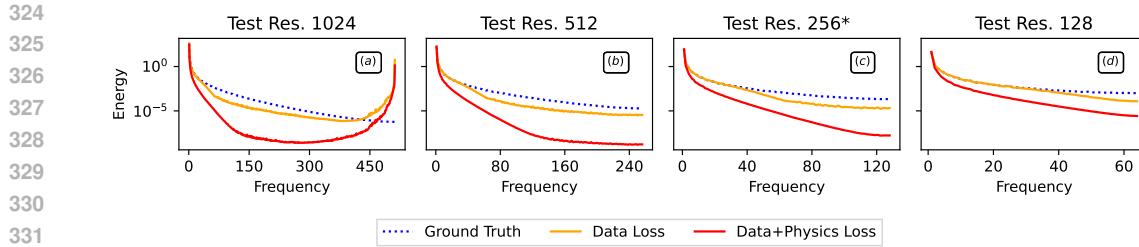


Figure 6: **(Physics-Informed) Optimization Evaluation.** FNO trained on Burgers data at resolution 256 (indicated by *) with and without physics optimization constraints. Average test spectra visualized. Spectra generated by the models trained with physics+data loss do not match the ground truth. Physics Loss term weighting $w = 10\%$. Full results in Appendix G, Figs. 21-23.

eral and extends to any class of model by interpolating inputs (both at training and inference) to and from a fixed-dimensional representation before and after a forward pass.

We train and evaluate a CNO and CROP+FNO model on all {Darcy, Burgers, Navier Stokes} datasets. Models are optimized using optimal hyper-parameters which were found via a grid search detailed in Appendix A. In Fig. 7, we visualize the predicted spectra of model's trained on resolution 16 data across test data of resolution {16, 32, 64, 128}. We observe that the CNO does accurately learn the band-limited representation of its training data: the spectra matches that of the ground truth until frequency $8f$ after which, by design, it drops sharply. This means that while CNO does not alias, it cannot predict frequencies higher than what was seen during training. Similarly, we observe that the CROP model, accurately fits lower frequencies, but struggles to fit high frequencies accurately across resolutions.

We note more broadly that the design of band-limiting a model's training data and predictive capacity is counter-intuitive to the goal of multi-resolution inference, in which, a broad range of frequencies must be modeled accurately. **Band-limiting a model's predictive capacity may enable accurate fixed-resolution representations, but ensures that high-frequency information is not predicted accurately (or at all).** We conclude that band-limited learning limits a model's utility for *multi-resolution* inference (full results in Appendix H).

5 MULTI-RESOLUTION TRAINING

We hypothesize that the reason models struggle to do *zero-shot* multi-resolution inference is because data representing a physical system at varying resolutions is sufficiently *out-of-distribution* to a model's fixed-resolution training data. To remedy this, we propose a data-driven solution: multi-resolution training (i.e., training on more than one resolution).

We compose multi-resolution datasets by randomly sampling different proportions of training data at varying resolutions $\{r_1, \dots, r_n\}$ where r_x is the proportion of x resolution training data and $n = 4$. We begin by evaluating dual-resolution training. Li et al. (2024a) have previously shown that dual-resolution active learning enables more accurate *high-resolution* inference. We extend this finding and evaluate if dual-resolution training can enable accurate *multi-resolution* inference. For each dataset in {Darcy, Burgers, Navier Stokes}, we combine data across resolutions in a pairwise manner creating $\frac{n(n-1)}{2}$ dual-resolution sets; the ratio of data samples between the two resolutions is varied over $p \in \{0.5, 0.1, 0.25, 0.5, 0.75, 0.9, 0.1\}$. In Fig. 8(a-f), we observe for pair-wise training, the test

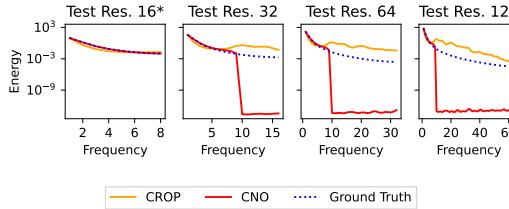


Figure 7: **Bandlimited Learning Evaluation.** Models trained on Darcy dataset at resolution 16 (indicated by *), evaluated resolutions 16, 32, 64, 128. Average test set 2D energy spectrum of model predictions and ground truth. Notice both model prediction spectrums diverge from ground truth after frequency 8.

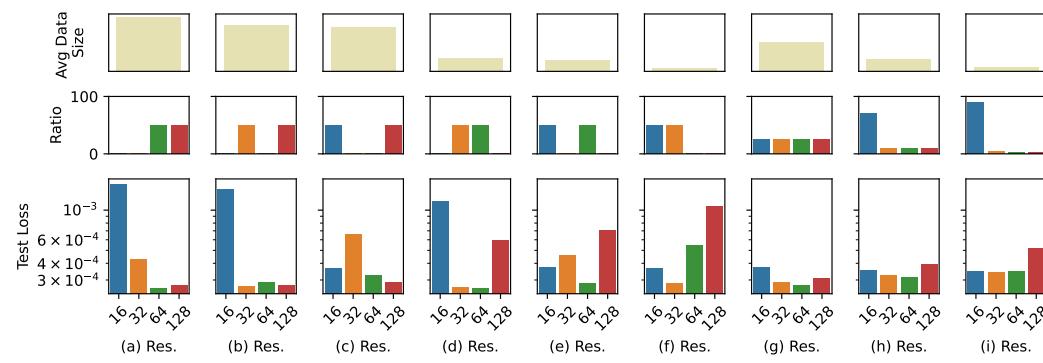


Figure 8: **Multi-resolution training.** FNO trained on multi-resolution Darcy data. **Top row:** Average number of pixels in a data sample in the training set. Lower number of pixels enables faster data generation and model training. **Middle row:** The ratios of data within each resolution bucket. **Bottom row:** The average test loss across different resolutions. Lower loss is better. Notice: mixed resolution datasets achieve both low average data size and low loss (ideal scenario).

performance for data that corresponds to the two training resolutions is generally better, but there are not consistent gains for the two non-training resolutions. This indicates that models perform best on the data resolutions on which they are trained.

To improve multi-resolution capabilities, we investigate the impact of including data from *all* resolutions. We first assess an equal number of samples across resolutions. In Fig. 8(g), the test performance across all resolutions improves which confirms that multi-resolution training benefits multi-resolution inference. Next, we ask: *Can we improve the computationally efficiency of multi-resolution training?* To do this, the training dataset must be composed of primarily low resolution data as it is both the cheapest to generate and train over (Fig. 33). We compose two additional multi-resolution datasets: $\{(0.7, 0.1, 0.1, 0.1), (0.9, 0.5, 0.3, 0.2)\}$. In Fig. 8(h, i), models remain competitive across test resolutions, even as we decrease the amount of high-resolution data. In Fig. 9, we observe the consistent trend for all datasets: models are able to achieve a balance between dataset size and multi-resolution test loss via multi-resolution training. Optimizing the ratios across all resolutions remains an exciting future direction. Full results in Appendix I.

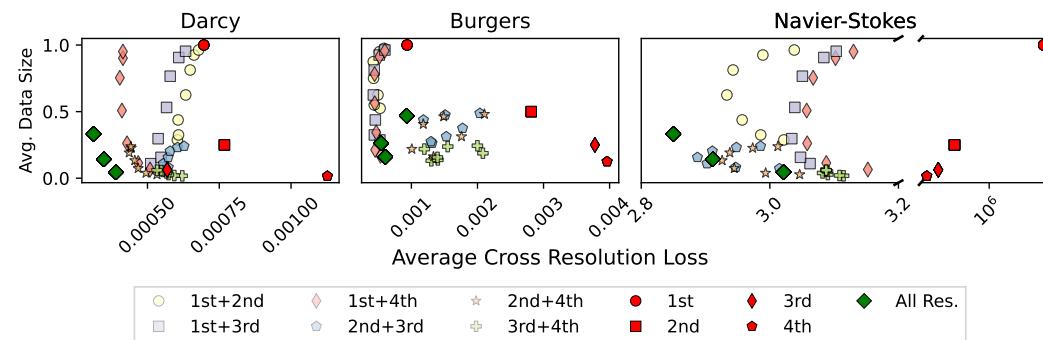


Figure 9: **Data Cost vs. Loss Tradeoff.** Lower data size and test loss are desirable (bottom left corner). We generally notice the “All Res.” datasets form the Pareto front of achieving optimal data size vs. test loss. “ m ” and “ $m+n$ ” indicate the one or two resolutions that were included in the training set. “All Res.” indicates that dataset contains points from all resolutions available. “Avg. Data Size” is the normalized average number of pixels in a data point.

432

6 RELATED WORK

433
 434
 435 **Modeling PDEs via Deep Learning.** Three prominent approach exist to discretely model PDEs
 436 with deep learning. **1) Mesh-Free** models learn the solution operator to a *specific* instance of a PDE
 437 (Yu et al., 2018; Raissi et al., 2019; Bar & Sochen, 2019; Smith et al., 2020; Pan & Duraisamy, 2020).
 438 Mesh-free models can be queried to return a measurement at any time and space coordinate. While,
 439 this approach means that a single model can resolve its solution at arbitrary discretization, it has
 440 two shortcomings: (i) Inference costs increases with number of points queried. (ii) Models cannot
 441 generalize beyond the specific PDE instance it was trained on. **2) Fixed-Mesh** models remedy both
 442 issues by learning a solution operator for a PDE *family* over a fixed-resolution mesh (Guo et al.,
 443 2016; Zhu & Zabaras, 2018; Adler & Öktem, 2017; Bhatnagar et al., 2019; Khoo et al., 2021). The
 444 inference costs of fixed-mesh models are lower than traditional numerical methods at corresponding
 445 mesh resolutions. However, fixed-mesh models fall short when one is interested in modeling scale
 446 phenomena that cannot be resolved via the fixed-mesh resolution (e.g., high-frequency information
 447 in turbulent systems). **3) Mesh-invariant** models, unlike fixed-mesh models, are capable of doing
 448 inference at arbitrary mesh resolutions (Li et al., 2020b;a; Lu et al., 2021b; Bhattacharya et al., 2021;
 449 Nelsen & Stuart, 2021; Patel et al., 2021; Rahman et al., 2022; Fanaskov & Oseledets, 2023). They
 450 have been proposed as a means to learn **mesh-invariant** solution operators to entire PDE families
 451 cheaply: train on low-resolution data, and use in a zero-shot fashion on arbitrary resolution data
 452 (e.g., zero-shot multi-resolution). In this work, we examine the zero-shot multi-resolution utility of
 453 mesh-invariant models.

454 **Aliasing (and corrective measures) in Deep Learning.** Sources of aliasing in deep learning include
 455 both artifacts of a pixel grid which models learn and amplify and the application of point-wise
 456 non-linearities to intermediate model representations (Karras et al., 2021; Gruver et al., 2022;
 457 Wilson, 2025). A straightforward approach, first introduced in generative adversarial networks, to
 458 stem nonlinearity-caused aliasing is to up sample a signal before applying a non-linearity followed
 459 by down sampling the signal (Karras et al., 2021). Bartolucci et al. (2023) and Raonic et al. (2023)
 460 extend the application of anti-alias activation function design to scientific machine learning; while
 461 this does prevent *aliasing*, it does not enable models trained at a specific resolution to resolve higher
 462 frequencies in higher resolution data. Gao et al. (2025) proposed a framework that “lifts” arbitrarily
 463 discretization data to a fixed-resolution band-limited space, to both train and do inference in. We
 464 investigate the sources of aliasing the context of *zero-shot* multi-resolution inference and show that
 465 proposed solutions fall short in remedying the core issue: out-of-distribution generalization.

466

7 CONCLUSION AND FUTURE WORK

467 For machine-learned operators to be as versatile as numerical methods-based approaches for mod-
 468 eling PDE’s they must perform accurate multi-resolution inference. To better understand an MLO’s
 469 abilities, we break down the task of multi-resolution inference and assess a trained model’s ability
 470 to both *extrapolate* to higher/lower frequency information in data and *interpolate* across varying
 471 data resolutions. We find that models trained on low resolution data and used for inference on
 472 high-resolution data can neither extrapolate nor interpolate, and therefore, more generally fail to do
 473 accurate multi-resolution inference. Changing the resolution of data at inference time is akin to out-
 474 of-distribution inference: models have not learned how to generalize in such settings. We document
 475 that models used in a *zero-shot* multi-resolution setting are prone to aliasing. We study the utility
 476 of two existing solutions—physics-informed constrains and learning band-limited learning—and find
 477 that neither enable accurate multi-resolution inference.

478 We introduce a simple and principled approach to enable accurate multi-resolution inference: multi-
 479 resolution training. We first show that models perform best at resolutions they have been trained
 480 on. We then extend this finding and demonstrate that one can computationally efficiently achieve
 481 the benefits of multi-resolution training via datasets composed with mostly low-resolution data and
 482 small amounts of high-resolution data. This enables accurate multi-resolution learning with the
 483 added benefit of low data-generation and model training cost. A promising future direction remains
 484 the automated selection of multi-resolution training data using strategies like active learning.

486 REPRODUCIBILITY STATEMENT
487

488 To enable reproduction of our results we will link to the code base upon acceptance. We addition-
489 ally include details about hyperparameter tuning in Appendix A, models (and their configura-
490 tion details) in Appendix B, datasets (and their configuration details) in Appendix C, data extrapola-
491 tion/interpolation implementation details in Appendix D.1. To ensure reliability of results we in-
492 clude an additional sensitivity analysis in Appendix L.

494 REFERENCES
495

496 Jonas Adler and Ozan Öktem. Solving ill-posed inverse problems using iterative deep neural net-
497 works. *Inverse Problems*, 33(12):124007, 2017.

498 Reza Alizadeh, Janet K Allen, and Farrokh Mistree. Managing computational complexity using
499 surrogate models: A critical review. *Research in Engineering Design*, 31(3):275–298, 2020.

500 Michael J Asher, Barry FW Croke, Anthony J Jakeman, and Luk JM Peeters. A review of surrogate
501 models and their application to groundwater modeling. *Water Resources Research*, 51(8):5957–
503 5973, 2015.

504 Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiavini, Jean Kossaifi, and
505 Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature
506 Reviews Physics*, 6(5):320–328, 2024.

508 Leah Bar and Nir Sochen. Unsupervised deep learning algorithm for PDE-based forward and inverse
509 problems. *arXiv preprint arXiv:1904.05417*, 2019.

510 Francesca Bartolucci, Emmanuel de Bezenac, Bogdan Raonic, Roberto Molinaro, Siddhartha
511 Mishra, and Rima Alaifari. Representation equivalent neural operators: A framework for alias-
512 free operator learning. *Advances in Neural Information Processing Systems*, 36:69661–69672,
513 2023.

515 Marsha J Berger and Joseph Oliger. Adaptive mesh refinement for hyperbolic partial differential
516 equations. *Journal of computational Physics*, 53(3):484–512, 1984.

517 Saakaar Bhatnagar, Yaser Afshar, Shaowu Pan, Karthik Duraisamy, and Shailendra Kaushik. Predic-
518 tion of aerodynamic flow fields using convolutional neural networks. *Computational Mechanics*,
519 64(2):525–545, 2019.

521 Kaushik Bhattacharya, Bamdad Hosseini, Nikola B Kovachki, and Andrew M Stuart. Model reduc-
522 tion and neural networks for parametric PDEs. *The SMAI journal of computational mathematics*,
523 7:121–157, 2021.

524 Ashesh Chattopadhyay, Y. Qiang Sun, and Pedram Hassanzadeh. Challenges of learning multi-
525 scale dynamics with AI weather models: Implications for stability and one solution, 2024. URL
526 <https://arxiv.org/abs/2304.07029>.

528 Alison Cozad, Nikolaos V Sahinidis, and David C Miller. Learning surrogate models for simulation-
529 based optimization. *AIChE Journal*, 60(6):2211–2227, 2014.

530 Vladimir Sergeevich Fanaskov and Ivan V Oseledets. Spectral neural operators. In *Doklady Math-
531 ematics*, volume 108, pp. S226–S232. Springer, 2023.

533 Alexander Forrester, Andras Sobester, and Andy Keane. *Engineering design via surrogate mod-
534 elling: A practical guide*. John Wiley & Sons, 2008.

535 Wenhao Gao and Chunmei Wang. Active learning based sampling for high-dimensional nonlinear
536 partial differential equations. *Journal of Computational Physics*, 475:111848, 2023.

538 Wenhao Gao, Ruichen Xu, Yuefan Deng, and Yi Liu. Discretization-invariance? On the discretiza-
539 tion mismatch errors in neural operators. In *The Thirteenth International Conference on Learning
Representations*, 2025.

540 Robert Joseph George, Jiawei Zhao, Jean Kossaifi, Zongyi Li, and Anima Anandkumar. Incremental
 541 spatial and spectral learning of neural operators for solving large-scale PDEs. *Transactions on*
 542 *Machine Learning Research*, 2024. ISSN 2835-8856. URL [https://openreview.net/](https://openreview.net/forum?id=xI6cPQObp0)
 543 [forum?id=xI6cPQObp0](https://openreview.net/forum?id=xI6cPQObp0).

544

545 Nate Gruver, Marc Finzi, Micah Goldblum, and Andrew Gordon Wilson. The Lie derivative for
 546 measuring learned equivariance. *arXiv preprint arXiv:2210.02984*, 2022.

547

548 Xiaoxiao Guo, Wei Li, and Francesco Iorio. Convolutional neural networks for steady flow ap-
 549 proximation. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge*
 550 *discovery and data mining*, pp. 481–490, 2016.

551

552 D Hendrycks. Gaussian error linear units (Gelus). *arXiv preprint arXiv:1606.08415*, 2016.

553

554 Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and
 555 Timo Aila. Alias-free generative adversarial networks. *Advances in neural information processing*
 556 *systems*, 34:852–863, 2021.

557

558 Yuehaw Khoo, Jianfeng Lu, and Lexing Ying. Solving parametric PDE problems with artificial
 559 neural networks. *European Journal of Applied Mathematics*, 32(3):421–435, 2021.

560

561 Jean Kossaifi, Nikola Kovachki, Zongyi Li, David Pitt, Miguel Liu-Schiaffini, Valentin Duruisseaux,
 562 Robert Joseph George, Boris Bonev, Kamyar Azizzadenesheli, Julius Berner, et al. A library for
 563 learning neural operators. *arXiv preprint arXiv:2412.10354*, 2024.

564

565 Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, An-
 566 drew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces
 567 with applications to PDEs. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

568

569 Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael W Mahoney. Char-
 570 acterizing possible failure modes in physics-informed neural networks. *Advances in neural infor-*
 571 *mation processing systems*, 34:26548–26560, 2021.

572

573 David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai
 574 Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapo-
 575 lation (rex). In *International conference on machine learning*, pp. 5815–5826. PMLR, 2021.

576

577 Jakub Kudela and Radomil Matousek. Recent advances and applications of surrogate models for
 578 finite element method computations: A review. *Soft Computing*, 26(24):13709–13733, 2022.

579

580 Shibo Li, Xin Yu, Wei Xing, Robert Kirby, Akil Narayan, and Shandian Zhe. Multi-resolution
 581 active learning of Fourier neural operators. In Sanjoy Dasgupta, Stephan Mandt, and Yingzhen Li
 582 (eds.), *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*,
 583 volume 238 of *Proceedings of Machine Learning Research*, pp. 2440–2448. PMLR, 02–04 May
 584 2024a. URL <https://proceedings.mlr.press/v238/li24k.html>.

585

586 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, An-
 587 drew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential
 588 equations. *arXiv preprint arXiv:2010.08895*, 2020a.

589

590 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, An-
 591 drew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differ-
 592 ential equations. *arXiv preprint arXiv:2003.03485*, 2020b.

593

594 Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar
 595 Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial
 596 differential equations. *ACM / IMS J. Data Sci.*, 1(3), May 2024b. doi: 10.1145/3648506. URL
 597 <https://doi.org/10.1145/3648506>.

598

599 Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detec-
 600 tion. *Advances in neural information processing systems*, 33:21464–21475, 2020.

594 Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning
 595 nonlinear operators via deeponet based on the universal approximation theorem of operators.
 596 *Nature machine intelligence*, 3(3):218–229, 2021a.
 597

598 Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning
 599 nonlinear operators via DeepONet based on the universal approximation theorem of operators.
 600 *Nature machine intelligence*, 3(3):218–229, 2021b.
 601

602 Nicholas H Nelsen and Andrew M Stuart. The random feature model for input-output maps between
 603 banach spaces. *SIAM Journal on Scientific Computing*, 43(5):A3212–A3243, 2021.
 604

605 Katharina Ott, Prateek Katiyar, Philipp Hennig, and Michael Tiemann. ResNet after all: Neural
 606 ODEs and their numerical solution. In *International Conference on Learning Representations*,
 607 2021.
 608

609 Shaowu Pan and Karthik Duraisamy. Physics-informed probabilistic learning of linear embeddings
 610 of nonlinear dynamics with guaranteed stability. *SIAM Journal on Applied Dynamical Systems*,
 611 19(1):480–509, 2020.
 612

613 Ravi G Patel, Nathaniel A Trask, Mitchell A Wood, and Eric C Cyr. A physics-informed operator re-
 614 gression framework for extracting data-driven continuum models. *Computer Methods in Applied
 615 Mechanics and Engineering*, 373:113500, 2021.
 616

617 Alejandro F Queiruga, N Benjamin Erichson, Dane Taylor, and Michael W Mahoney. Continuous-
 618 in-depth neural networks. *arXiv preprint arXiv:2008.02389*, 2020.
 619

620 Md Ashiqur Rahman, Zachary E Ross, and Kamyar Azizzadenesheli. U-NO: U-shaped neural
 621 operators. *arXiv preprint arXiv:2204.11127*, 2022.
 622

623 Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A
 624 deep learning framework for solving forward and inverse problems involving nonlinear partial
 625 differential equations. *Journal of Computational physics*, 378:686–707, 2019.
 626

627 Bogdan Raonic, Roberto Molinaro, Tim De Ryck, Tobias Rohner, Francesca Bartolucci, Rima Alai-
 628 fari, Siddhartha Mishra, and Emmanuel de Bézenac. Convolutional neural operators for robust
 629 and accurate learning of PDEs. *Advances in Neural Information Processing Systems*, 36:77187–
 77200, 2023.
 630

631 Pu Ren, N Benjamin Erichson, Shashank Subramanian, Omer San, Zarija Lukic, and Michael W
 632 Mahoney. Superbench: A super-resolution benchmark dataset for scientific machine learning.
 633 *arXiv preprint arXiv:2306.14070*, 2023.
 634

635 C.E. Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.
 636 doi: 10.1109/JRPROC.1949.232969.
 637

638 Jonathan D Smith, Kamyar Azizzadenesheli, and Zachary E Ross. Eikonet: Solving the eikonal
 639 equation with deep neural networks. *IEEE Transactions on Geoscience and Remote Sensing*, 59
 640 (12):10685–10696, 2020.
 641

642 A Subel, Y Guan, A Chattopadhyay, and P Hassanzadeh. Explaining the physics of transfer learning
 643 a data-driven subgrid-scale closure to a different turbulent flow. *arXiv preprint arXiv:2206.03198*,
 644 724, 2022.
 645

646 Shashank Subramanian, Robert M Kirby, Michael W Mahoney, and Amir Gholami. Adaptive self-
 647 supervision algorithms for physics-informed neural networks. *arXiv preprint arXiv:2207.04084*,
 648 2022.
 649

650 Bruno Sudret, Stefano Marelli, and Joe Wiart. Surrogate models for uncertainty quantification:
 651 An overview. In *2017 11th European conference on antennas and propagation (EUCAP)*, pp.
 652 793–797. IEEE, 2017.
 653

654 Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani,
 655 Dirk Pflüger, and Mathias Niepert. PDEBench: An extensive benchmark for scientific machine
 656 learning. *Advances in Neural Information Processing Systems*, 35:1596–1611, 2022.
 657

648 Alasdair Tran, Alexander Mathews, Lexing Xie, and Cheng Soon Ong. Factorized fourier neural
649 operators. *arXiv preprint arXiv:2111.13802*, 2021.
650

651 Michael Unser. Sampling-50 years after Shannon. *Proceedings of the IEEE*, 88(4):569–587, 2002.
652

653 Sifan Wang, Shyam Sankaran, Hanwen Wang, and Paris Perdikaris. An expert’s guide to training
654 physics-informed neural networks. *arXiv preprint arXiv:2308.08468*, 2023.
655

656 John M Whittaker. The “Fourier” theory of the cardinal function. *Proceedings of the Edinburgh
657 Mathematical Society*, 1(3):169–176, 1928.
658

659 Andrew Gordon Wilson. Deep learning is not so mysterious or different. *arXiv preprint
660 arXiv:2503.02113*, 2025.
661

662 Liu Yang, Siting Liu, Tingwei Meng, and Stanley J Osher. In-context operator learning with data
663 prompts for differential equation problems. *Proceedings of the National Academy of Sciences*,
664 120(39):e2310142120, 2023.
665

666 Bing Yu et al. The deep Ritz method: a deep learning-based numerical algorithm for solving varia-
667 tional problems. *Communications in Mathematics and Statistics*, 6(1):1–12, 2018.
668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 A HYPER-PARAMETER SEARCH
703

704 **FNO HP Tuning:** For each dataset $\in \{\text{Darcy, Burgers, Navier Stokes}\}$ we do a grid search for
705 optimal training hyper-parameters: learning rate $\in \{1e-2, 1e-3, 1e-4, 1e-5\}$ and weight decay $\in \{1e-5, 1e-6, 1e-7\}$ for both the data driven loss (i.e., mean squared error) and the respective data+physics
706 driven loss. Each model was trained for 150 epochs. For the models optimized with the data+physics
707 loss, we optimized the physics loss term's weighting coefficient $w \in \{0.1, 0.25, 0.5\}$. For Darcy,
708 Burgers, and Navier Stokes we do this hyper-parameter search for data at resolution 64, 512, 255
709 respectively, and then use the optimize parameter values for each dataset to train models at remaining
710 resolutions. See Tab. 1 for the optimal hyper-parameter for each dataset/loss combination.
711

712 **CROP/CNO HP Tuning:** For each dataset $\in \{\text{Darcy, Burgers, Navier Stokes}\}$ we do a grid search
713 for optimal training hyper-parameters: learning rate $\in \{1e-3, 1e-4, 1e-5\}$ for the data driven loss
714 (i.e., mean squared error). Each model was trained for 150 epochs. For Darcy, Burgers, and Navier
715 Stokes we do this hyper-parameter search for data at resolution 64, 512, 255 respectively, and then
716 use the optimize parameter values for each dataset to train models at remaining resolutions. See
717 Tab. 2 for the optimal hyper-parameter for each dataset/loss combination.
718

719 Table 1: **Optimal FNO hyper-parameters** from hyper-parameter search outlined in Appendix
720 A. *NS batch size had to be reduced to 1 for multi-resolution training experiments (see Sec. 5),
721 therefore we used a lower learning rate in that setting. w =Physic Loss Coefficient (see Sec. 4.1).

Data	Loss	w	Learning Rate	Weight Decay	Batch Size
Darcy	Data	-	1e-3	1e-5	128
Darcy	Data+Physics	0.1	1e-2	1e-5	128
Burgers	Data	-	1e-3	1e-5	64
Burgers	Data+Physics	0.1	1e-3	1e-5	64
Navier Stokes	Data	-	1e-2	1e-6	4
Navier Stokes	Data+Physics	0.1	1e-4	1e-5	4
Navier Stokes*	Data	-	1e-5	1e-6	1

731
732 Table 2: **Optimal CNO/CROP/DeepONet hyper-parameters** from hyper-parameter search out-
733 lined in Appendix A. *The original CROP implementation did not include a 1D version, so we omit
734 CROP for the 1D Burgers dataset.
735

Data	Loss	Model	Learning Rate	Weight Decay	Batch Size
Darcy	Data	CNO	0.0001	1e-5	128
Darcy	Data	CROP	0.001	1e-5	128
Darcy	Data	DeepONet	0.001	1e-5	128
Burgers*	Data	CNO	0.001	1e-5	64
Navier Stokes	Data	CNO	0.001	1e-6	1
Navier Stokes	Data	CROP	0.001	1e-6	1

756 B MODEL IMPLEMENTATIONS AND CONFIGURATIONS

758 The **Fourier Neural Operator** is described in detail in (Li et al., 2020a); we closely follow
 759 their implementation which can be found at <https://neuraloperator.github.io/dev/index.html>. We detail key model configuration parameter choices here:
 760

- 762 1. **Max Modes:** When training and evaluating models to do zero-shot super-resolution or sub-
 763 resolution and multi-resolution training/inference, we train with the max modes parameter
 764 set to half the maximum training resolution. Further, we do an ablation study of the max
 765 modes parameter in Appendix F to assess impact of FNO’s ability to do information ex-
 766 trapulation and resolution interpolation.
- 767 2. **Layer count:** 4 (standard implementation)
- 768 3. **Hidden channels:** 32 (standard implementation)
- 769 4. **Activation Function:** All main paper experiments utilize the Gelu nonlinearity
 770 (Hendrycks, 2016). Further, we assess the impact of utilizing anti-aliasing activation
 771 function on zero-shot super and sub-resolution, information extrapolation, and resolution
 772 interpolation in Appendix K (Karras et al., 2021).
- 773 5. **All other parameters:** All other FNO parameters were initialized to default values detailed
 774 in the standard `neuraloperator` Python package (Kossaifi et al., 2024).

775 The **CNO** is described in detail in (Raonic et al., 2023); we closely follow their
 776 implementation which can be found at <https://github.com/camlab-ethz/ConvolutionalNeuralOperator/tree/main>. We detail key model configuration
 777 parameter choices here:
 778

- 779 1. **Residual block count:** 6 (standard implementation)
- 780 2. **Latent size:** 64
- 781 3. **Input Resolution:** Varries based on training data resolution.
- 782 4. **All other parameters:** All other CNO parameters were initialized to default values de-
 783 tailed in the official CNO implementation: <https://github.com/camlab-ethz/ConvolutionalNeuralOperator/tree/main>.

784 The **CROP** pipeline is described in detail in (Gao et al., 2025); we closely follow their imple-
 785 mentation for CROP+FNO which can be found at <https://github.com/wenhangao21/ICLR25-CROP/tree/main>. We note that they did not include a 1D CROP implementation,
 786 so we exclude evaluation of CROP on the 1D Burgers dataset. We detail key model configuration
 787 parameter choices here:
 788

- 789 1. **Hidden channels:** 32 (standard implementation)
- 790 2. **Latent size:** 64
- 791 3. **Max Mode:** 32 (half latent size)
- 792 4. **All other parameters:** All other CROP parameters were initialized to default values de-
 793 tailed in the official CROP implementation: <https://github.com/wenhangao21/ICLR25-CROP/tree/main>.

794 The **DeepONet** model is described in detail in (Lu et al., 2021a); we closely follow the imple-
 795 mentation of (Gao et al., 2025) which can be found at <https://github.com/wenhangao21/ICLR25-CROP/tree/main>. We detail key model configuration parameter choices here:
 796

- 797 1. **Branch layer count:** 4 (standard implementation)
- 798 2. **Trunk layer count:** 4 (standard implementation)
- 799 3. **Branch layer width:** 128 (standard implementation)
- 800 4. **Trunk layer width:** 128 (standard implementation)
- 801 5. **All other parameters:** All other DeepONet parameters were initialized to default values
 802 detailed in the official implementation of (Gao et al., 2025): <https://github.com/wenhangao21/ICLR25-CROP/tree/main>.

C DATA

We study three standard scientific datasets: Darcy, Burgers, and turbulent incompressible Navier Stokes released in PDEBench ([Takamoto et al., 2022](#)); please refer to the original publication for how these datasets were created.

We summarize each dataset here, for full details refer to [Takamoto et al. \(2022\)](#):

Darcy: We study the steady-state solution of 2D Darcy Flow over the unit square with viscosity term $a(x)$ as an input of the system. We learn the mapping from $a(x)$ to the steady-state solution described by:

$$\begin{aligned} -\nabla(a(x)\nabla u(x)) &= f(x), \quad x \in (0, 1)^2 \\ u(x) &= 0, \quad x \in \partial(0, 1)^2 \end{aligned}$$

The force term is a constant value $f = 1$.

Burgers: We study Burgers' equation which is used to model the non-linear behavior and diffusion process in fluid dynamics:

$$\partial_t u(t, x) + \partial_x(u^2(t, x)/2) = v/\pi \partial_{xx} u(t, x), \quad x \in (0, 1), t \in (0, 2] \quad (3)$$

$$u(0, x) = u_0(x), \quad x \in (0, 1) \quad (4)$$

The diffusion coefficient is a constant value $f = 0.001$

(Turbulent) Inhomogeneous, Incompressible Navier Stokes: We study a popular variant of the Navier Stokes equation: the incompressible version. This equation is used to model dynamics far lower than the speed of propagation of waves in the medium:

$$\nabla \cdot v = 0, \rho(\partial_t v + v \cdot \nabla v) = -\nabla p + \eta \Delta v \quad (5)$$

Takamoto et al. (2022) employ an augmented form of (5) which includes a vector field forcing term u :

$$\rho(\partial_t v + v \cdot \nabla v) = -\nabla p + \eta \Delta v + u$$

The viscosity is a constant value $\nu=0.01$. We convert the incompressible Navier Stokes dataset to vorticity form to enable direct comparison with [Li et al. \(2020a\)](#).

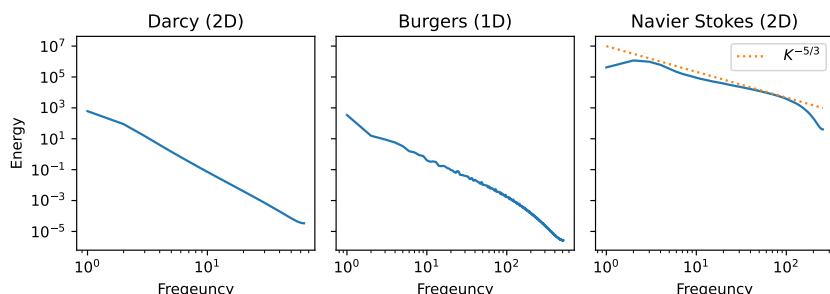
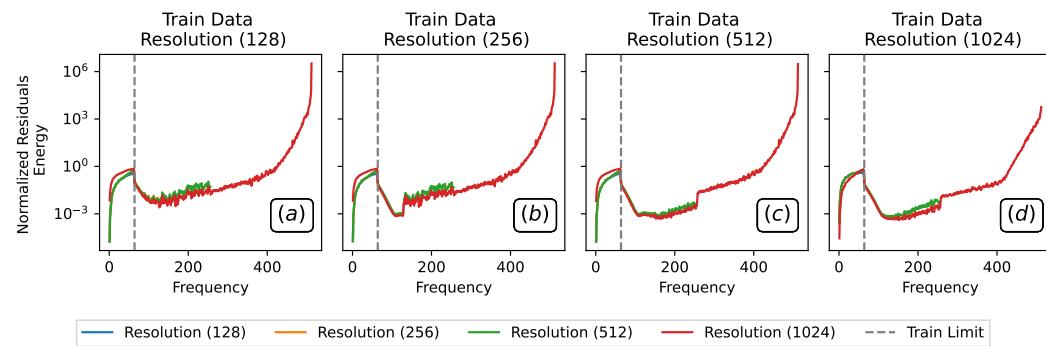
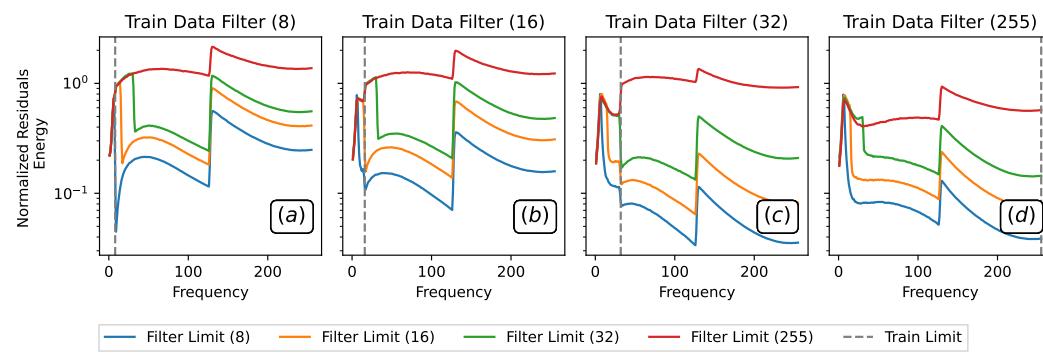
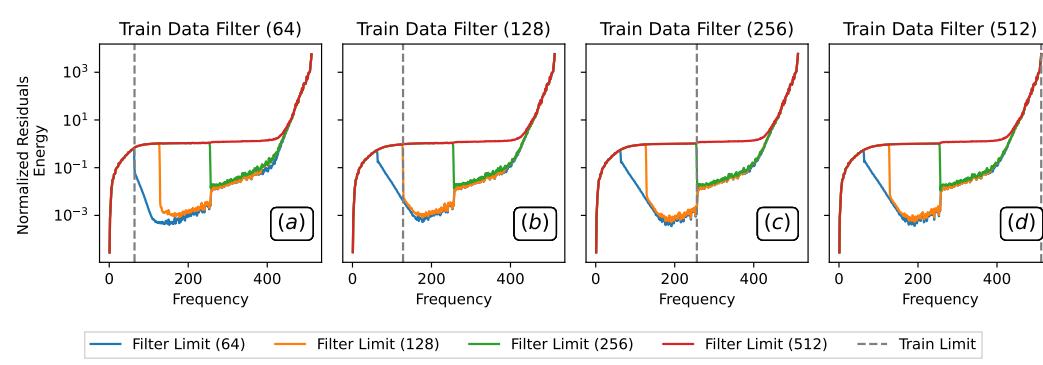
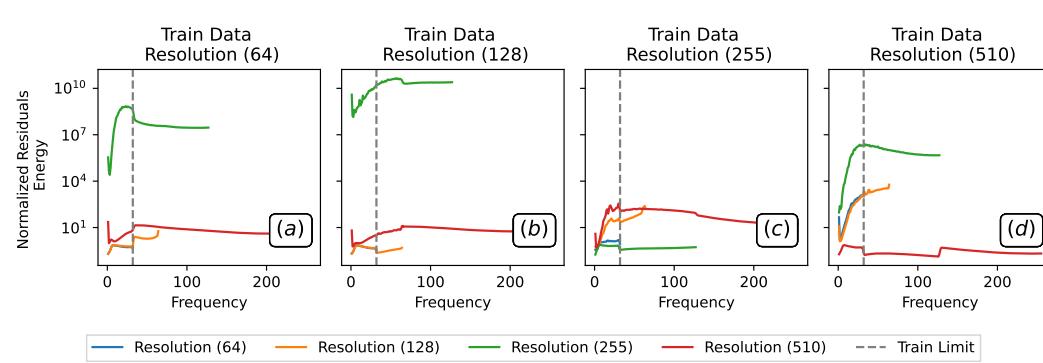
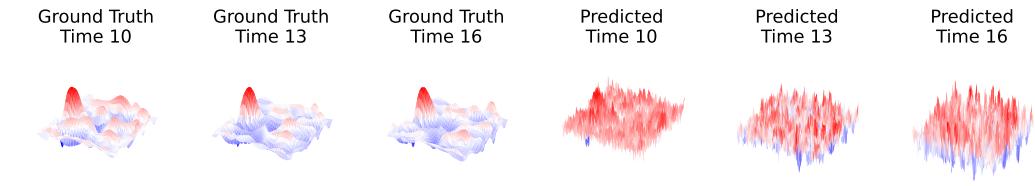


Figure 10: **Dataset Energy Spectrum.** Average energy spectrums over test datasets. Notice that Navier Stokes is in the turbulent regime. K = Kolmogorov coefficient.

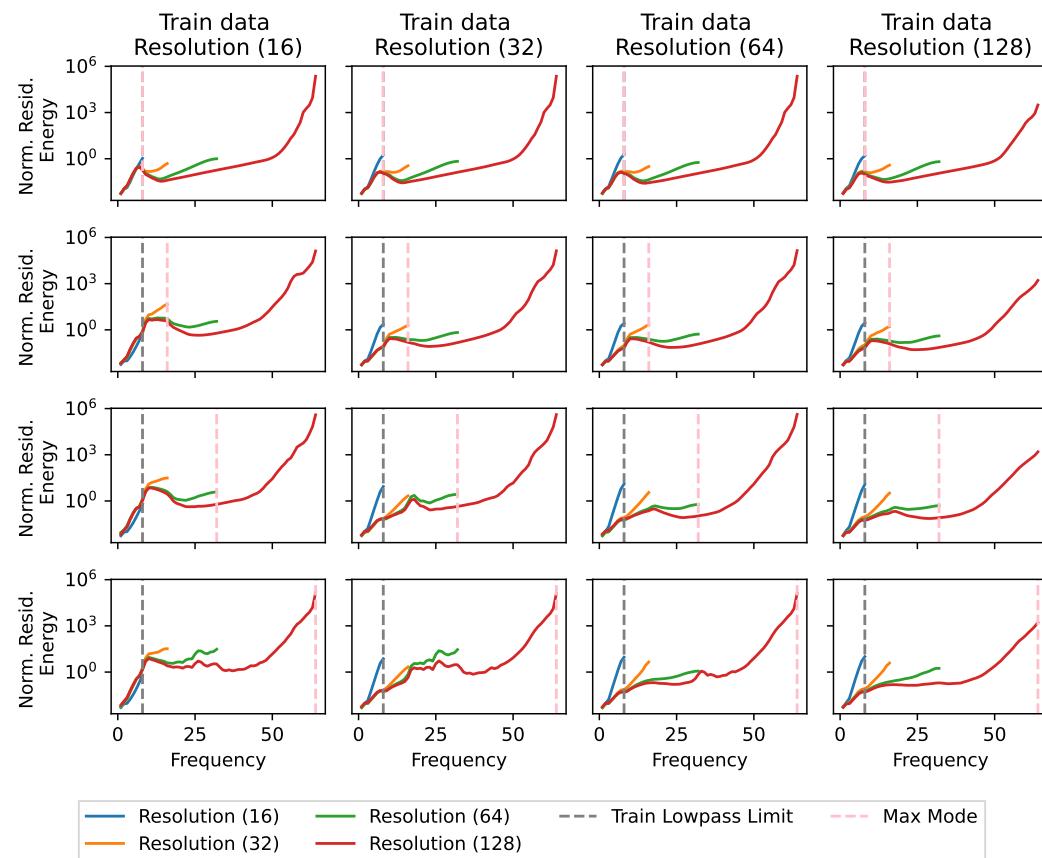
864 D INFORMATION EXTRAPOLATION AND RESOLUTION INTERPOLATION
865866 D.1 EXTRAPOLATION AND INTERPOLATION IMPLEMENTATION DETAILS
867868 **Down sampling Procedure:** For the resolution interpolation experiments, data was sampled at
869 different resolution by employing standard down sampling of the data (keeping every N th element)
870 from the maximal data resolution. This is the same down sampling procedure used by
871 PDEBench (Takamoto et al., 2022).872 **Low-pass Filtering Procedure:** For the information extrapolation experiments, data was filtered
873 with a standard stop-band filter which restricts any frequency components above the desired filter
874 limit N . Concretely, to low pass filter data, we transform the data to the frequency domain via
875 an FFT and shift it such that the lowest frequency component is the center of the spectrum. We
876 determine which frequency components are greater than our limit N by assessing which components
877 fall outside a radius of N from the center; we then set the energy of any frequency components
878 greater than N to zero. Finally, the (filtered) data is transformed back to the spatial domain via an
879 inverse FFT.880 D.2 FULL RESULTS
881882 Here we include the full experimental results of studying FNOs' abilities to do both information
883 extrapolation and resolution interpolations as described in Sec. 3.
884895
896
897 **Figure 11: Interpolation.** Four FNOs trained on Burgers data low-pass limit $64f$ (constant frequency
898 information) and down sampled to resolutions $\{128, 256, 512, 1024\}$ (varying sampling rate)
899 from left to right. We test if each model can generalize to data with varying sampling rate. We
900 visualize the normalized residual spectra across test data. Notice, residual spectra (error) increases
901 substantially in the low frequencies. Lower energy at all wave numbers is better.
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917



972 E EVALUATING SUB- AND SUPER-RESOLUTION
973974 Here we included the full experiment results of studying FNOs' abilities to do zero-shot multi-
975 resolution inference as described in Sec. 3.
976985 **Figure 15: Aliasing artifacts compound over time.** FNO trained on resolution 255 Navier-Stokes
986 data, evaluated at resolution 510. **Left:** Ground truth evolution of NS fluid flow. **Right:** Corresponding
987 FNO predictions at resolution 510. Notice, high frequency artifacts become more prevalent over
988 time.1001 **Figure 16: FNOs do not generalize to higher or lower resolutions.** Heatmaps of losses incurred
1002 by FNO trained and tested at varying resolutions (lower is better). When the test resolution varies
1003 from the training resolution, the models often incur a substantial increase in loss.1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

1026 F MAX MODES

1028 A key design decision in the FNO architecture is parameter m that indicates maximum frequencies
 1029 to keep along each dimension in the Fourier layer during the forward pass; this has implications
 1030 both during training (which frequencies are learned in the Fourier layers) and at inference (which
 1031 frequencies are predicted over in the Fourier layers). In Sec. 3, the FNO is always initialized such
 1032 that it can make use all frequencies in its input; this is especially critical in the multi-resolution
 1033 setting where data of varying discretization will have varying frequency information. Here we study
 1034 the impact of varying m . In the zero-shot multi-resolution inference setting, in Figs. 18 and 17, we
 1035 find that that across all variation in m , the models assign high energy in the high-frequencies (e.g.,
 1036 alias). More broadly, we comment that in the context of multi-resolution inference it does not make
 1037 sense to set m to a value less than the largest populated frequency in a model input, as it ensure that
 1038 the model cannot make use that frequency information greater than m in the Fourier layers. In the
 1039 event frequency information above m is not useful to prediction (e.g., noise), we advocate low-pass
 1040 filtering and down sampling the data to a more compressed representation of data prior to inference
 1041 to remove unwanted frequencies and ensure faster inference.



1070 Figure 17: **Resolution Interpolation.** Four FNOs trained on Darcy data low-pass limit = 8 (constant
 1071 frequency information) and down sampled to resolutions $\{16,32,64,128\}$ (varying sampling rate)
 1072 from left to right, and top to bottom with max modes $m \in \{8,16,32,64\}$. Test if each model can
 1073 generalize to data with varying sampling rate. Visualizing spectrum of the normalized residuals
 1074 across test data. Notice, residual spectra (error) increases substantially in the low frequencies. Lower
 1075 energy at all wave numbers is better. We notice that across all variation in m , the models assign high
 1076 energy in the high-frequencies.

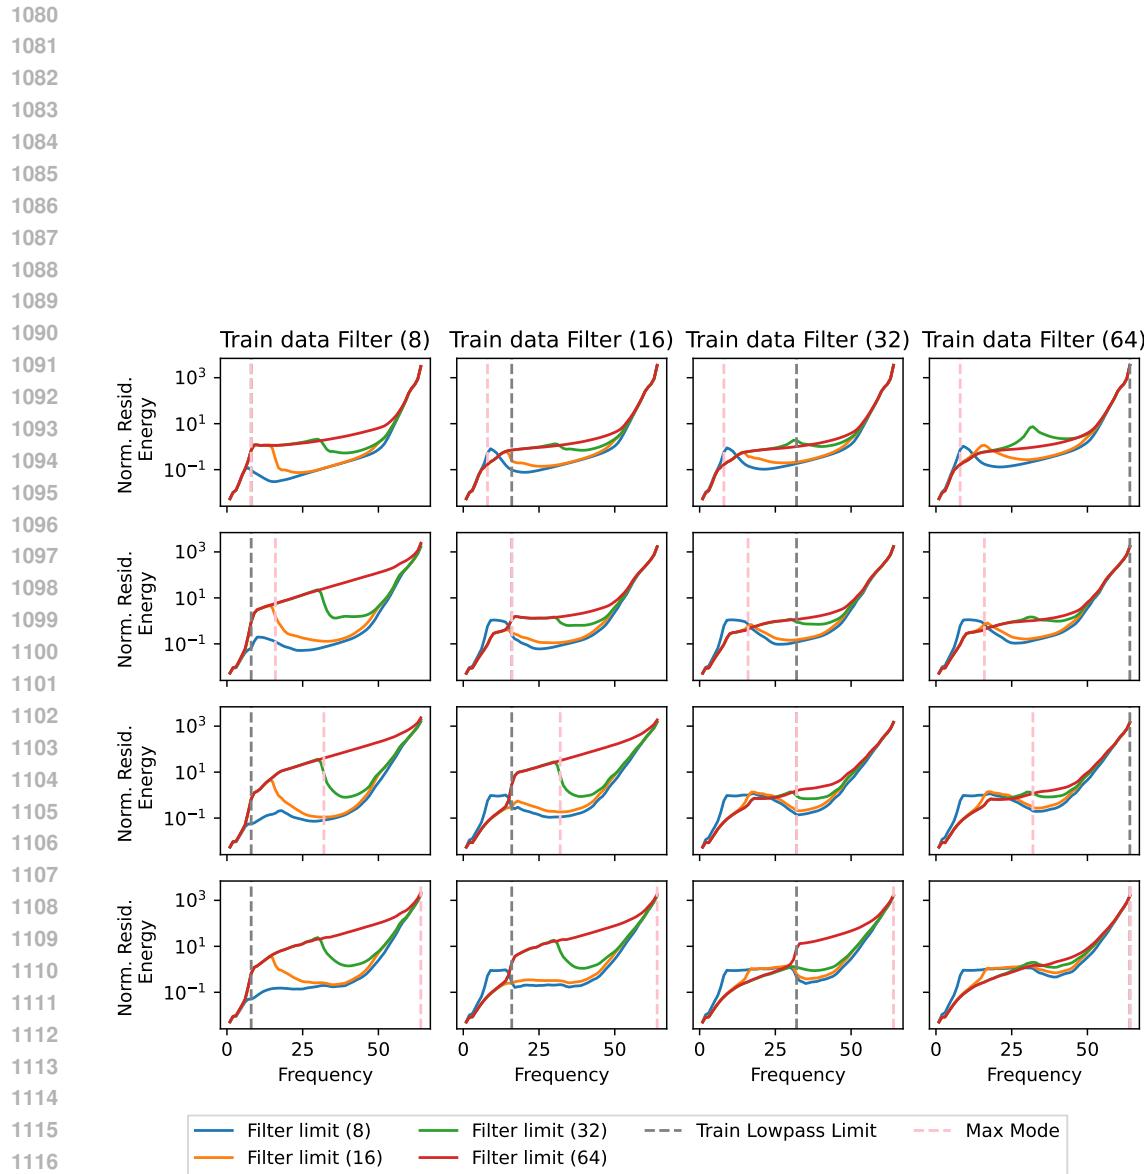
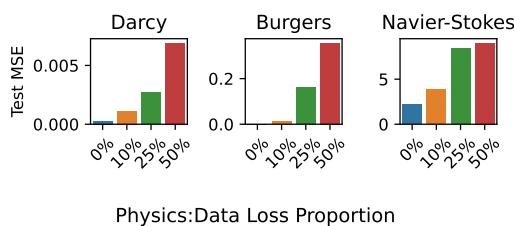
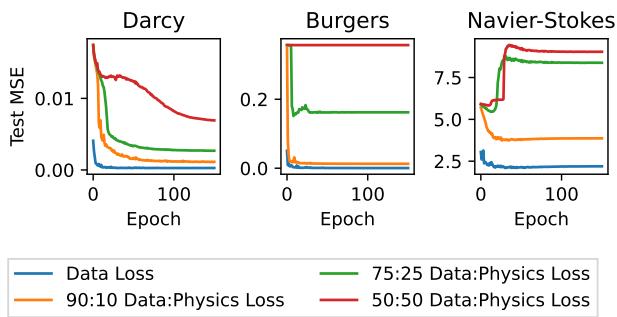
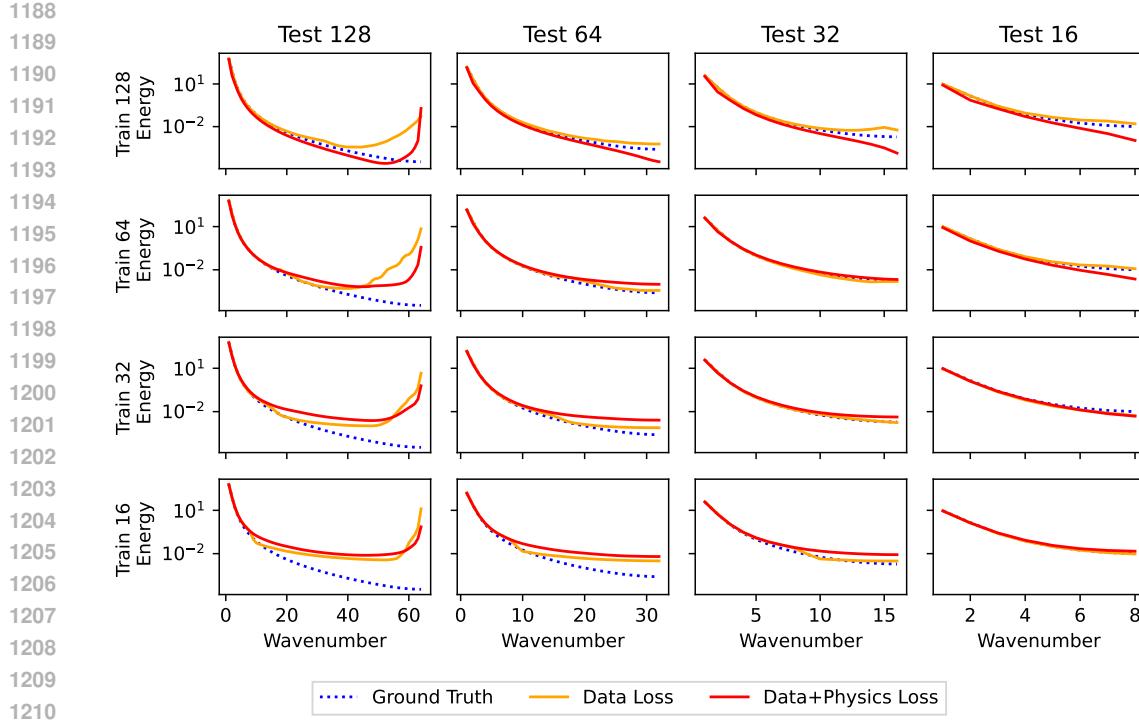


Figure 18: **Information Extrapolation.** Sixteen FNOs trained on Darcy data of resolution 128 (constant sampling rate) and low-pass filtered with limits $\{8, 16, 32, 64\}$ (varying frequency information) from left to right, and top to bottom with max modes $m \in \{8, 16, 32, 64\}$. Test if each model can generalize to data with varying frequency information. Visualizing spectrum of the normalized residuals across test data. Notice, residual spectra (error) increases substantially in the high frequencies. Lower residual energy at all wave numbers is better. We notice that across all variation in m , the models assign high energy in the high-frequencies.

1134 **G PHYSICS-INFORMED OPTIMIZATION**
11351136 We use the physics losses of [Li et al. \(2024b\)](#) which explicitly enforce that the governing partial
1137 differential equation is satisfied. The governing partial differential equations are detailed in Ap-
1138 pendix Appendix C.1139 Below we include the results of tuning the physics loss weighting coefficient w in Figs. 20 and 19.
1140 We then include the full comparisons of training each dataset (Darcy, Burgers, Navier Stokes) at
1141 each resolution with both $w \in \{0, 0.1\}$ in Figs. 21-22.
11421143
1144
1145
1146
1147
1148
1149
1150
1151
1152 **Figure 19: (Physics-Informed) Optimization.** Increasing the proportion of physics-informed loss
1153 in the optimization objective corresponds with increased test loss. Lower MSE is better. Darcy
1154 trained at resolution 64, Burgers trained at resolution 512, and Navier Stokes trained at resolution
1155 255.1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168 **Figure 20: (Physics-Informed) Optimization Objective.** The physics informed constraints never
1169 achieves better performance than pure data-driven constraints. Darcy trained at resolution 64, Burg-
1170 ers trained at resolution 512, and Navier Stokes trained at resolution 255.



1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Figure 21: **Darcy**. Energy spectra for models trained at a specific resolution (y-axis) and tested at multiple resolution (x-axis) with and without physics optimization constraint. The spectrums generated by the models trained with physics+data loss do not match ground truth. $w = 10\%$.

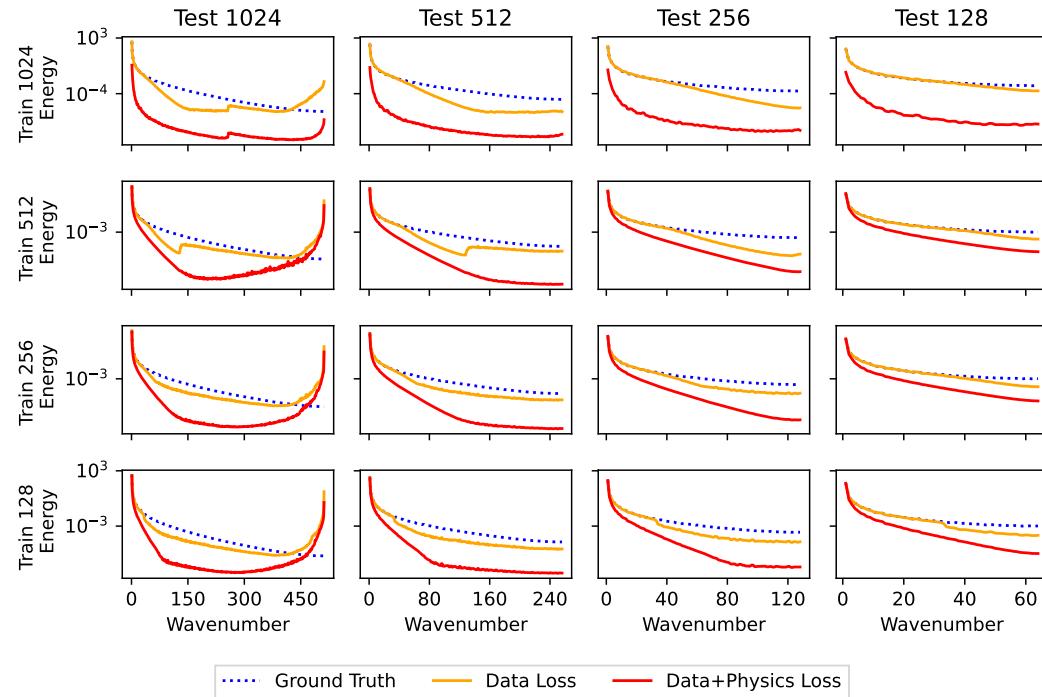


Figure 22: **Burgers**. Energy spectra for models trained at a specific resolution (y-axis) and tested at multiple resolution (x-axis) with and without physics optimization constraint. The spectrums generated by the models trained with physics+data loss do not match ground truth. $w = 10\%$.

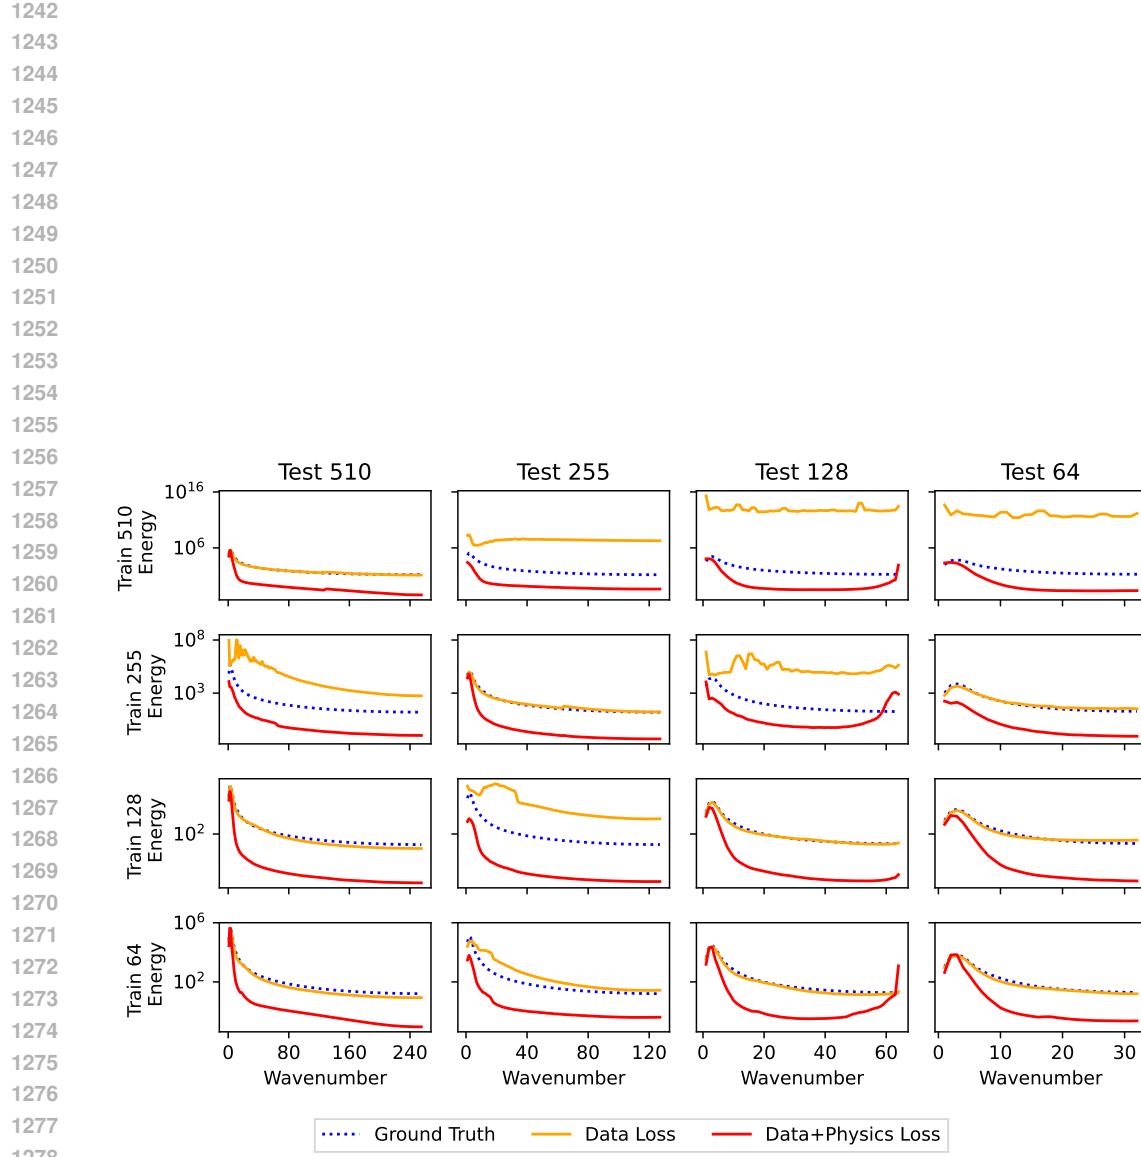


Figure 23: **Navier Stokes.** Energy spectra for models trained at a specific resolution (y-axis) and tested at multiple resolution (x-axis) with and without physics optimization constraint. The spectrums generated by the models trained with physics+data loss do not match ground truth. $w = 10\%$.

1296

H BAND-LIMITED LEARNING

1297

1298 Here we include the full experimental results of training models with mixed resolution datasets as
1299 described in Sec. 4.2. Note: [Gao et al. \(2025\)](#) did not release a 1D CROP pipeline, therefore, we
1300 were unable to test CROP with Burgers.

1301 First, we observe that band-limited learning, in which models are able to both infer and learn over
1302 band-limited representations of data, accurately learn the frequency region of the data included in
1303 the band-limit. However, they struggle to/do not learn anything outside this range (Figs. 7,28,26).
1304 The implication of this is that band-limited-approaches suffice for modeling data within a prespec-
1305 ified range as long as the band-limit range is wide enough, and the model will never need to be
1306 used to infer on data containing additional frequency information. However, we observe the band
1307 limited approach **under performs** multi-resolution training at fitting the full spectrum in the multi-
1308 resolution inference setting since the resolution of data, and consequently the resolved frequencies,
1309 are changing (Fig. 9). For both Darcy (Fig. 26) and Burgers (Fig. 27), we notice that multi-resolution
1310 training out-performs band-limited approaches. In Fig. 7, we observe that the predetermined band-
1311 limit leads to error in the high frequency range.

1312 A scenario in which CNO and CROP *appear* to perform well is on datasets in which the majority
1313 of the energy is concentrated in the predetermined band-limit (e.g., Navier Stokes, see Fig. 10). In
1314 this setting, we see that band-limited fit the lower frequencies in the spectrum very well (Fig. 28).
1315 However, we also see in Fig. 28, that multi-resolution training is the only method that consistently
1316 predicts both the correct amount of energy across the full spectrum. Band limited approaches *fail* to
1317 fit the high frequency range of the spectrum.

1318 We note that band-limited approaches do not accurately fit frequencies outside of their predeter-
1319 mined limit; thus making them effective for fixed-resolution inference but *ineffective* for multi-
1320 resolution inference. Alternatively, we demonstrated in Sec. 5 that multi-resolution inference can be
1321 scaled to new data resolutions (and therefore new parts of the spectral energy spectrum) via scaling
1322 up representative samples in the training dataset. We conclude that multi-resolution training is a
1323 more flexible and scalable approach to enabling accurate multi-resolution training at *all* parts of the
1324 energy spectrum.

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

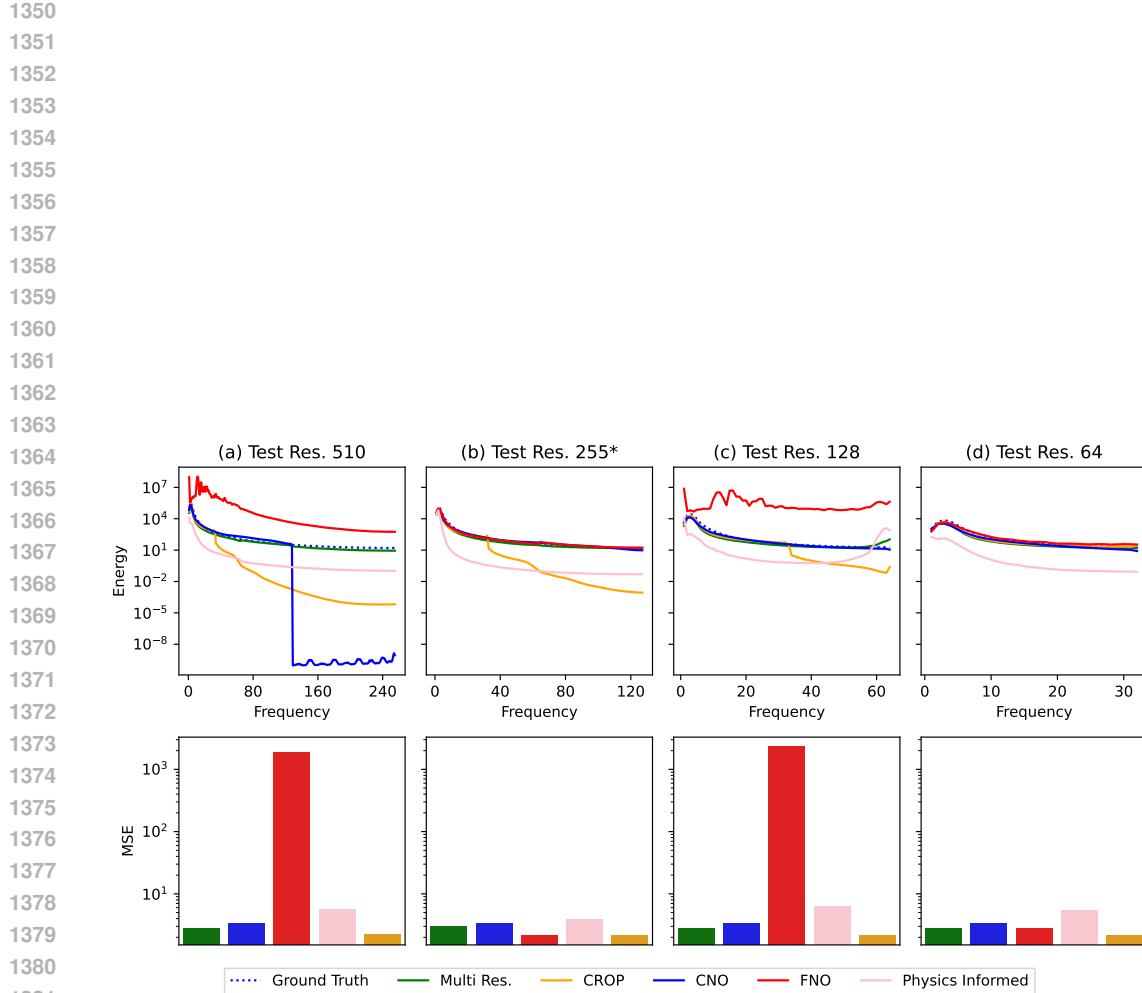


Figure 24: **Spectral Comparison Navier Stokes.** **Top Row:** Average predicted spectra for test data at varying resolutions across all methods. **Bottom row:** Average mean squared error loss over test data at varying resolutions across all methods. **Zero-shot methods:** CNO, FNO, Physics Informed and CROP are all zero-shot methods, meaning there are trained at a specific resolution (255, indicated by *), and evaluated at resolutions 510, 255, 128, 64. **Data-driven method:** Multi-resolution training; notice that multi-resolution training is the only method that consistently fits both the high and low parts of the spectra. **Band-limited methods:** CNO and CROP are both band-limited methods which are trained in a zero-shot manner at a fixed resolution; we observe that they only accurately fit the low frequencies.

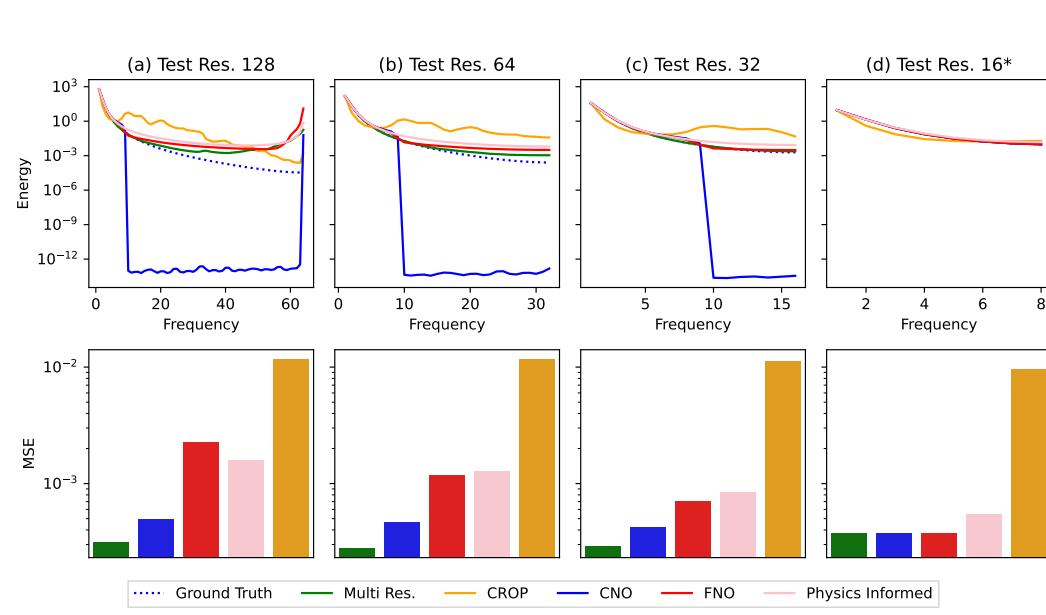


Figure 25: **Spectral Comparison Darcy.** **Top Row:** Average predicted spectra for test data at varying resolutions across all methods. **Bottom row:** Average mean squared error loss over test data at varying resolutions across all methods. **Zero-shot methods:** CNO, FNO, Physics Informed and CROP are all zero-shot methods, meaning they are trained at a specific resolution (16, indicated by *), and evaluated at resolutions 128, 64, 32, 16. **Data-driven method:** Multi-resolution training; notice that multi-resolution training is the only method that consistently fits both the high and low parts of the spectra. **Band-limited methods:** CNO and CROP are both band-limited methods which are trained in a zero-shot manner at a fixed resolution; we observe that they only accurately fit the low frequencies.

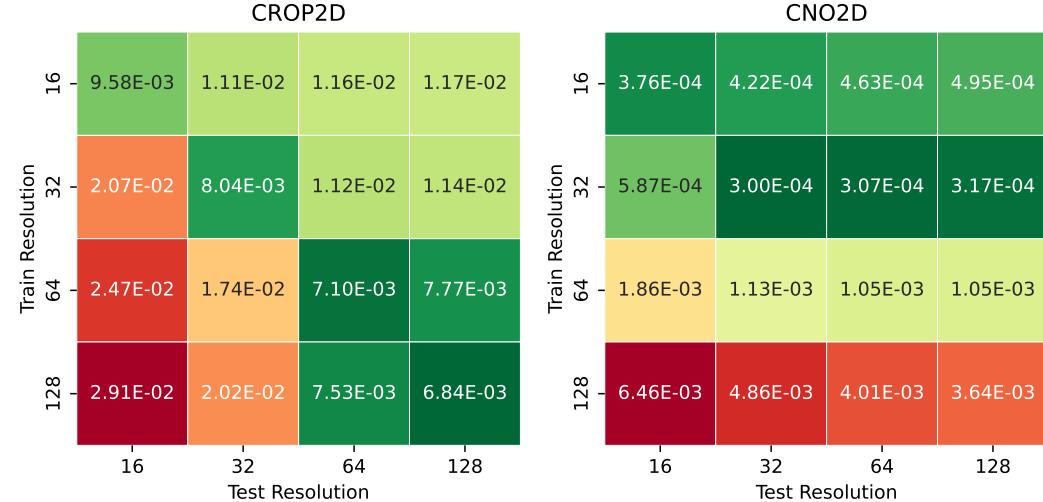


Figure 26: **CROP+FNO and CNO trained on Darcy.** On average both CROP+FNO and CNO incur higher losses across resolutions compared to both FNO (Fig. 16) and multi-resolution training (Fig. 9).

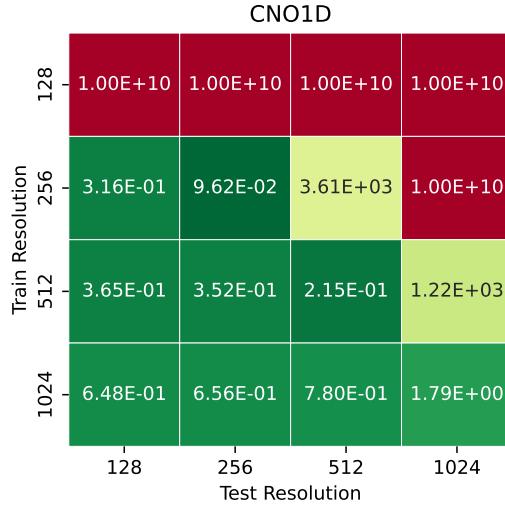


Figure 27: **CNO trained on Burgers.** On average, CNO incur higher losses across resolutions compared to both FNO (Fig. 16) and multi-resolution training (Fig. 9). We note that despite our hyperparameter search (Tab. 2) the CNO model trained on resolution 128 failed to converge.

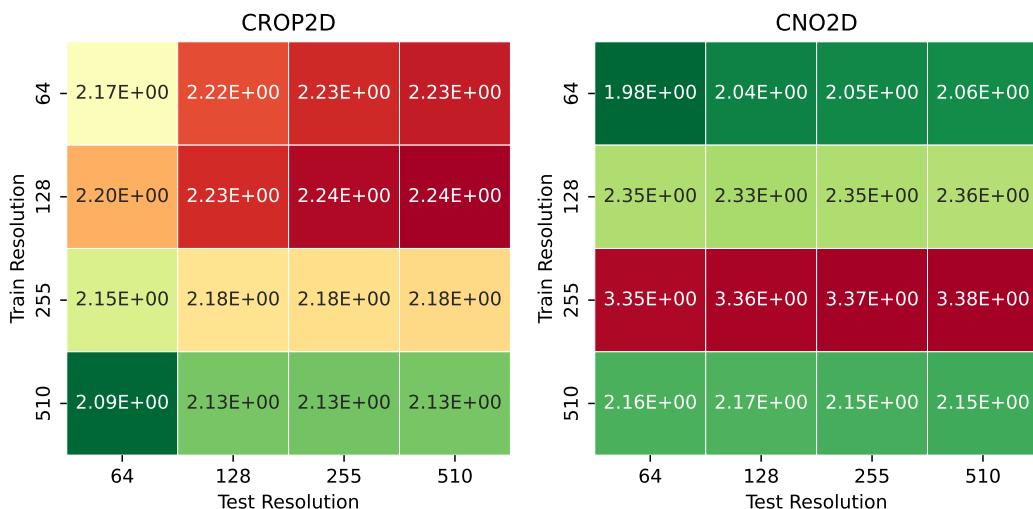
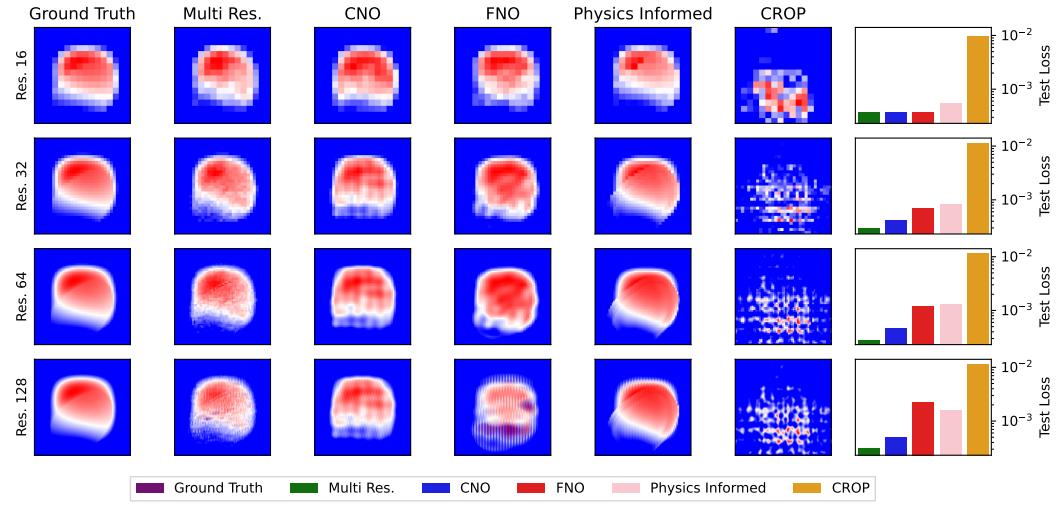


Figure 28: **CROP+FNO and CNO trained on Navier Stokes.** On average both CROP+FNO and CNO incur lower losses across resolutions compared to both FNO (Fig. 16) and multi-resolution training (Fig. 9).

1512 I MULTI-RESOLUTION TRAINING

1514 We provide an overview comparison across methods in Fig. 29. Here we include the full experimental
 1515 results of training models with mixed resolution datasets as described in Sec. 5 in Figs. 30-32.
 1516 We plot the association between increased dataset size and training time in Fig. 33.



1535 Figure 29: **Assessing multi-resolution inference.** Column 1: Expected prediction for Darcy flow
 1536 at varying resolutions. Columns 2-6: Sample prediction for Darcy flow at varying test resolutions.
 1537 Column 7: Average mean squared error test loss at each resolution (lower is better). **Zero-shot**
 1538 **methods:** CNO, FNO, Physics Informed and CROP are all zero-shot methods, meaning the model
 1539 was trained at a specific resolution (16) and evaluated at resolutions 16, 32, 64, 128. **Data-driven**
 1540 **method:** Multi-resolution training; notice that multi-resolution training consistently outperforms
 1541 zero-shot methods.

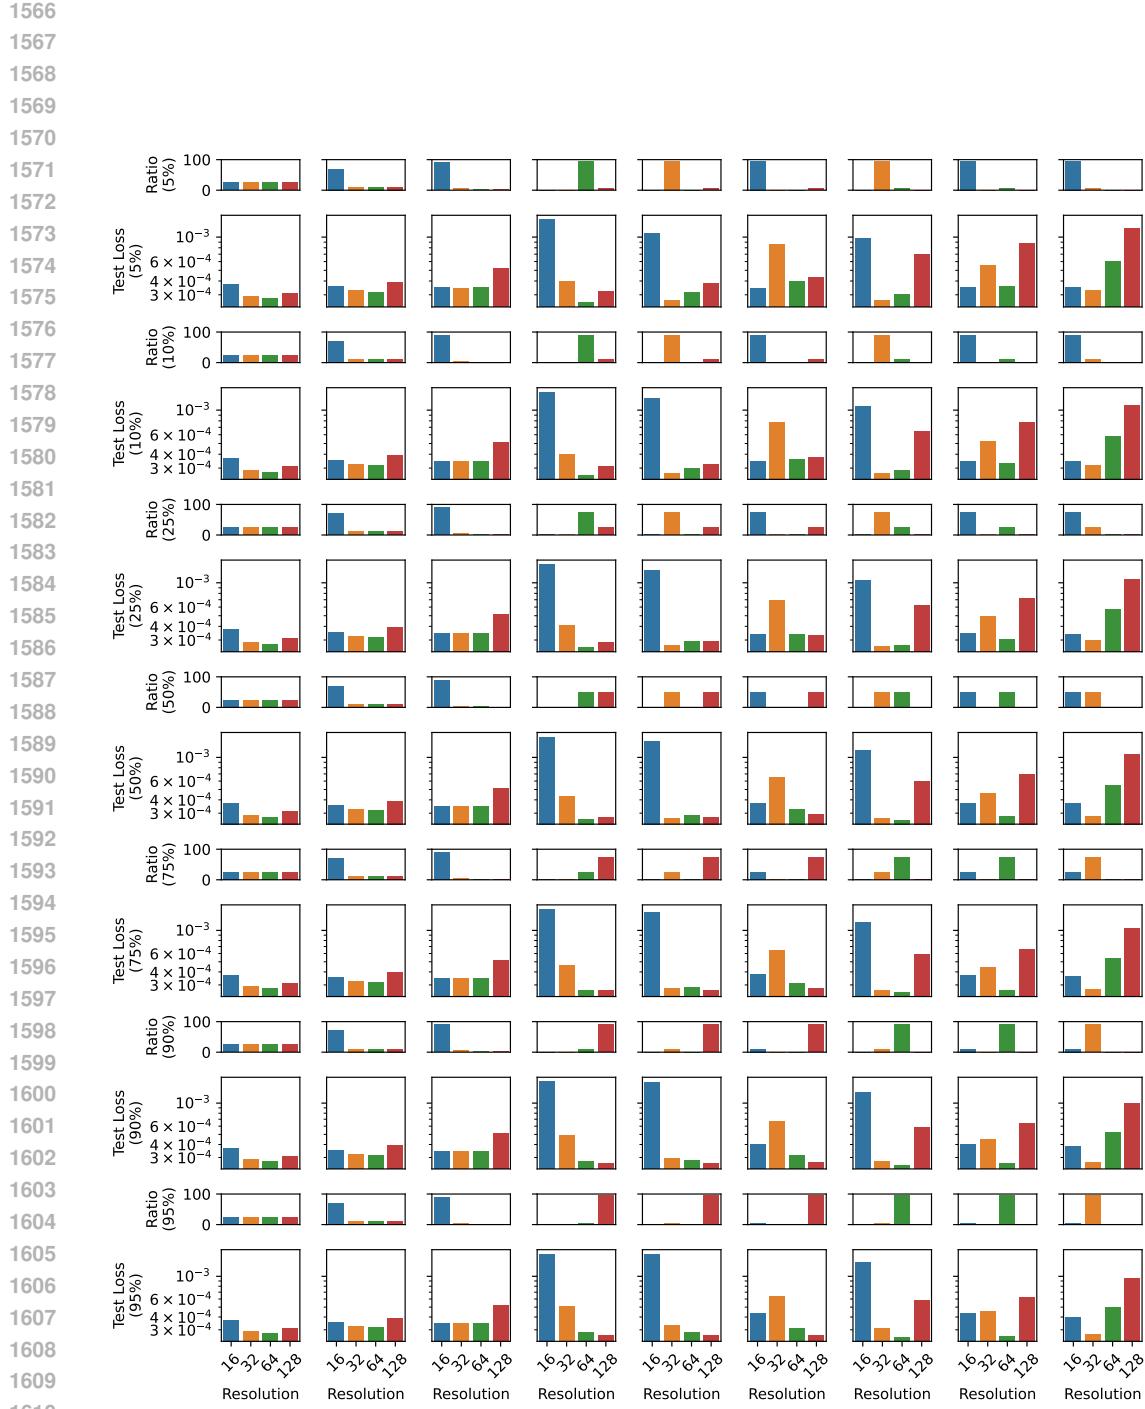


Figure 30: **Darcy Multi-Resolution Training.** FNO trained on multi-resolution Darcy data. Each row include two sub-rows; each row is delineated the the dual-resolution training ratio (indicated in y-axis label). The top sub-row illustrates the ratios of data within each resolution bucket. The bottom sub-row indicates the average test loss across different resolutions. Lower loss is better. Notice in the mixed resolution datasets achieve both low average data size and low loss (ideal scenario).

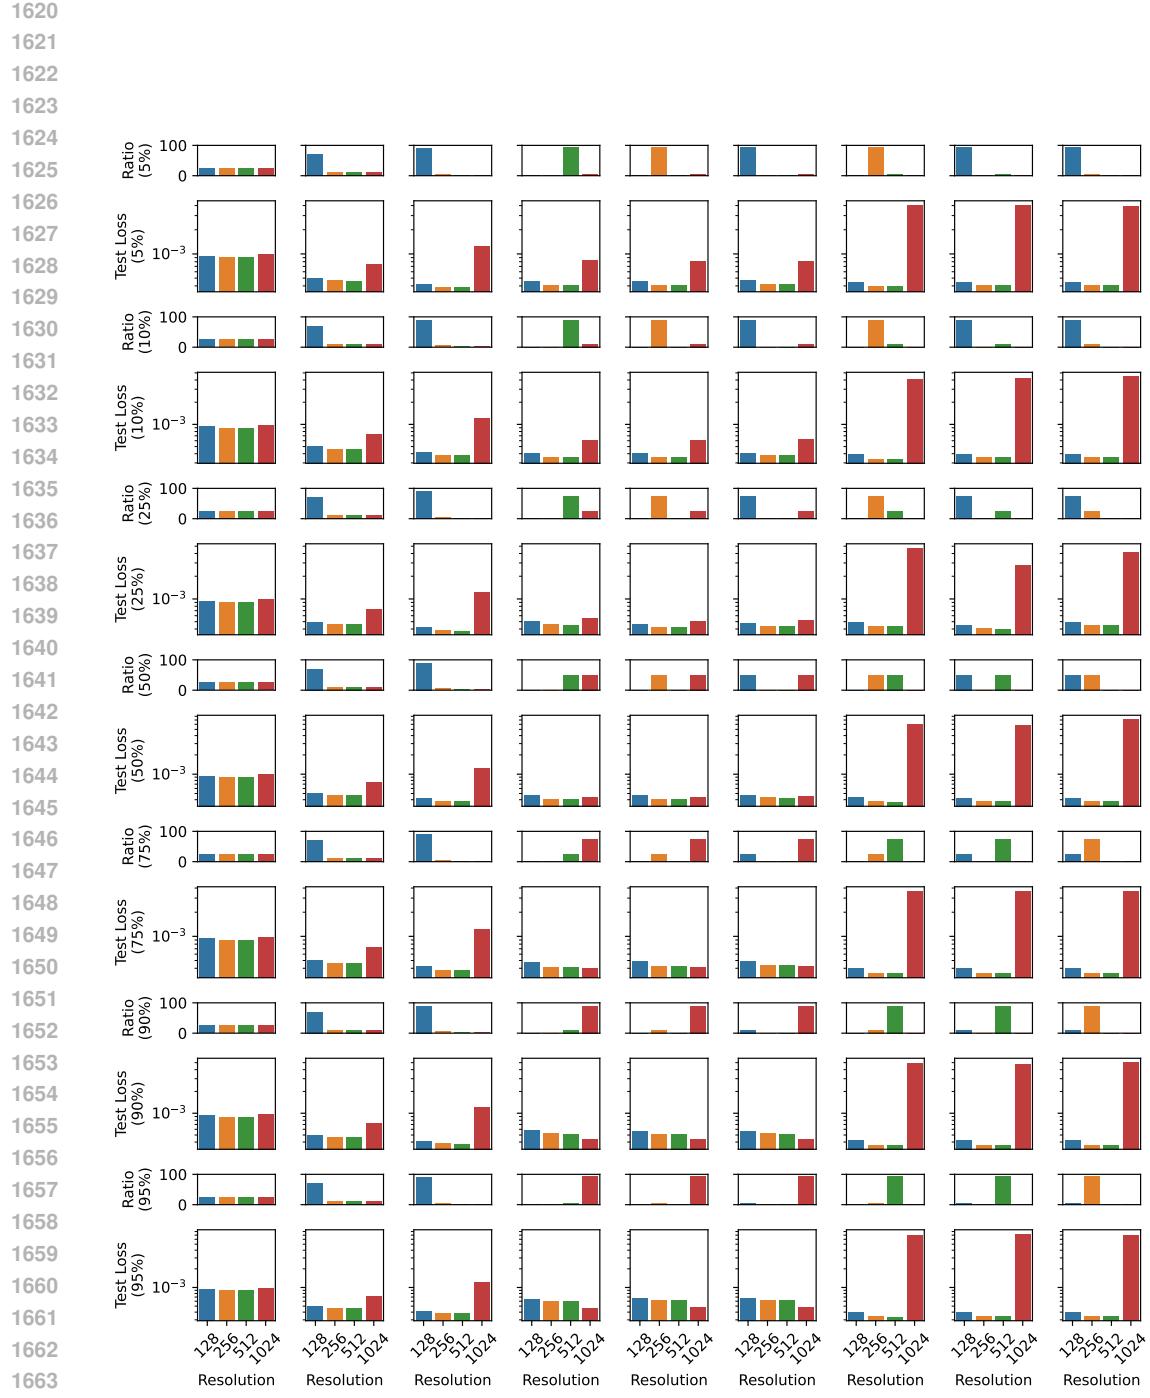


Figure 31: **Burgers Multi-Resolution Training.** FNO trained on multi-resolution Burgers data. Each row include two sub-rows; each row is delineated the the dual-resolution training ratio (indicated in y-axis label). The top sub-row illustrates the ratios of data within each resolution bucket. The bottom sub-row indicates the average test loss across different resolutions. Lower loss is better. Notice in the mixed resolution datasets achieve both low average data size and low loss (ideal scenario).

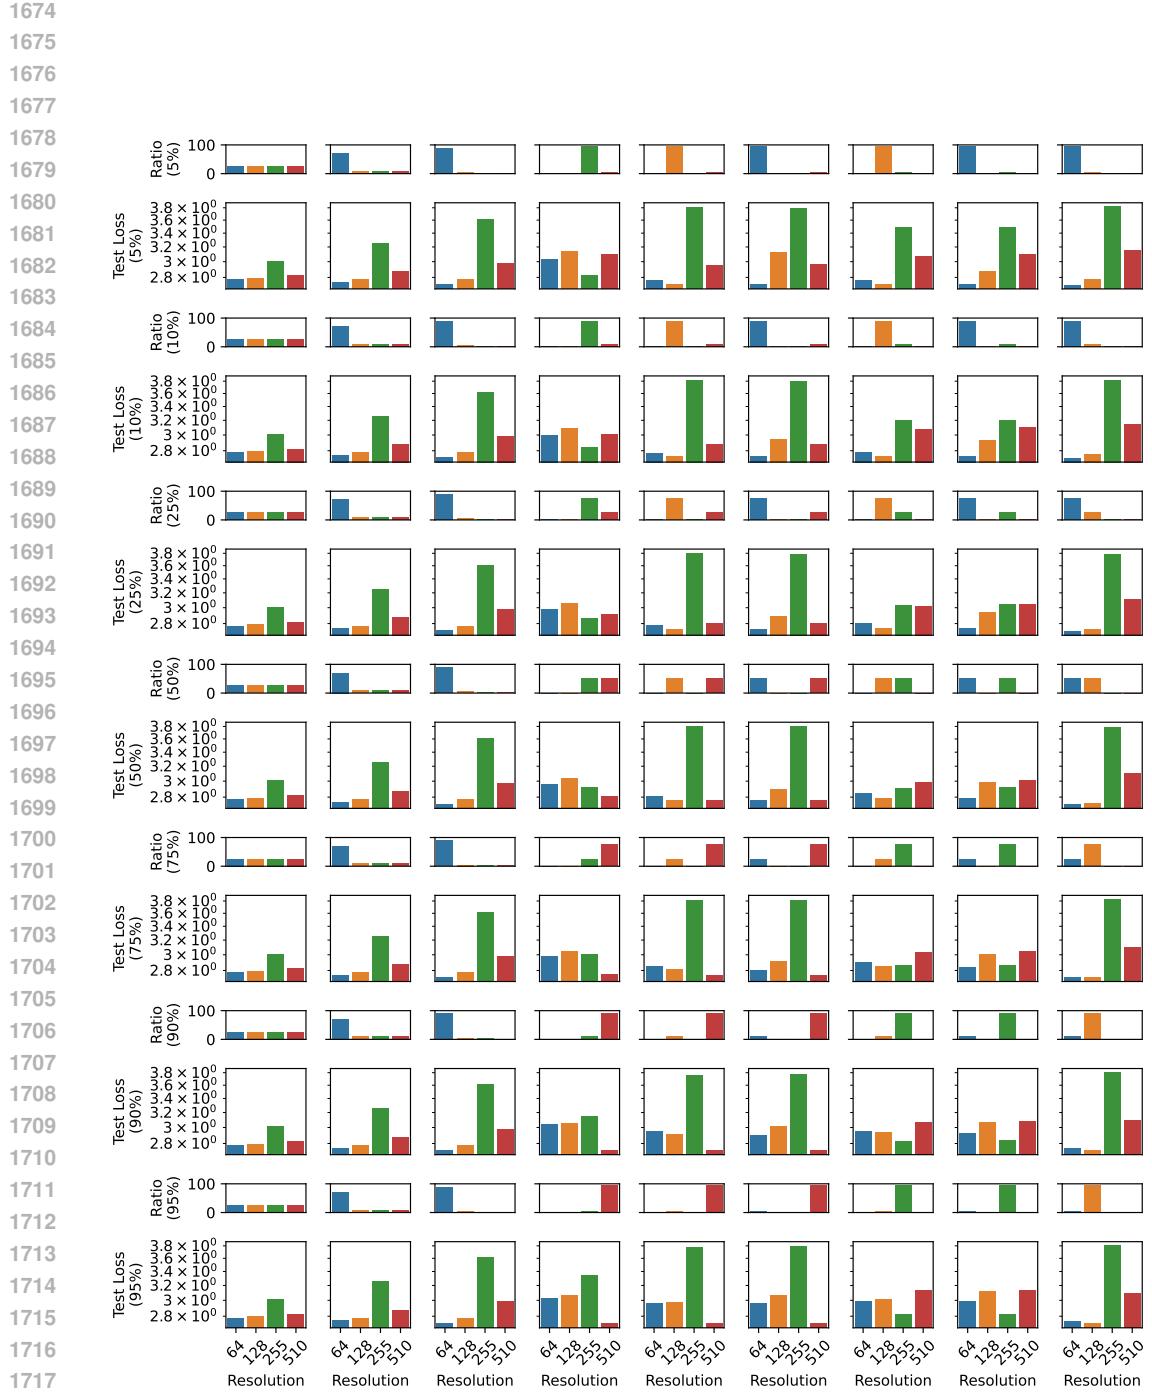


Figure 32: **Navier Stokes Multi-Resolution Training.** FNO trained on multi-resolution Navier Stokes data. Each row include two sub-rows; each row is delineated the the dual-resolution training ratio (indicated in y-axis label). The top sub-row illustrates the ratios of data within each resolution bucket. The bottom sub-row indicates the average test loss across different resolutions. Lower loss is better. Notice in the mixed resolution datasets achieve both low average data size and low loss (ideal scenario).

1728 I.1 MULTI-RESOLUTION TRAINING PERFORMANCE GAINS
1729

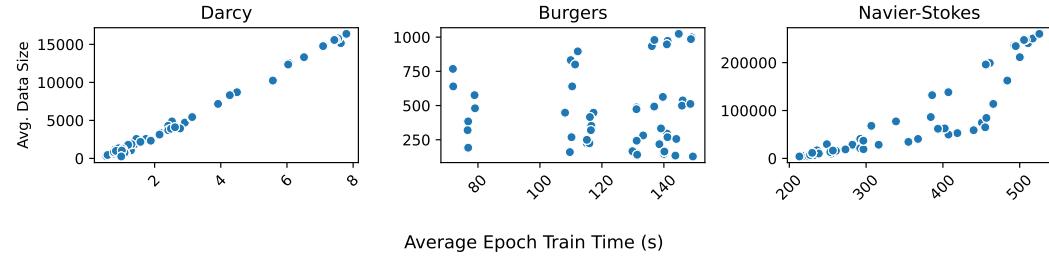
1730 Here we profile both the memory savings (see Tab. 3) and wall-clock training time (see Tab. 4)
 1731 saving due to the reduction in dataset size from multi-resolution training compared with the single-
 1732 resolution training counterpart. We empirically observe the smaller dataset size is associated with
 1733 reduction in training time in Fig. 33.

1734
 1735 **Table 3: Impact of Multi-Resolution Training on Memory.** Reporting the size of datasets in
 1736 Gigabytes (GB) for multi-resolution dataset with ratios (0.9, 0.5, 0.3, 0.2) vs. single-resolution
 1737 maximum resolution dataset along with the percent decrease in dataset size. The smaller the dataset,
 1738 the better.

1739 Dataset	1740 Single-Res. Size (GB)	1741 Multi-Res. Size (GB)	1742 % (↓)
1741 Darcy	0.54	0.02	96% ↓
1742 Burgers	0.34	0.005	98% ↓
1743 Navier Stokes	0.96	0.04	96% ↓

1744
 1745 **Table 4: Impact of Multi-Resolution Training on Wall-clock Training Time.** Reporting the
 1746 training time (over 150 epochs) for multi-resolution dataset with ratios (0.9, 0.5, 0.3, 0.2) vs. single-
 1747 resolution maximum resolution dataset along with the percent decrease in dataset size. The lower
 1748 the training time, the better.

1749 Dataset	1750 Single-Res. Training Time (hours)	1751 Multi-Res. Training Time (hours)	1752 % (↓)
1751 Darcy	0.33	0.05	86% ↓
1752 Burgers	6.03	5.8	3.2% ↓
1753 Navier Stokes	21.9	9.6	56% ↓



1765 **Figure 33: Data size corresponds with training time.** We notice a clear trend with Darcy and
 1766 Navier Stokes datasets: as data size increases, so does average training time per epoch. For Burgers,
 1767 this trend is less clear; however, we note that the Burgers dataset is several orders of magnitude
 1768 smaller and therefore can be used with a high batch size (see Tab. 1) which reduces the computational
 1769 gains achieved from a smaller sized dataset.

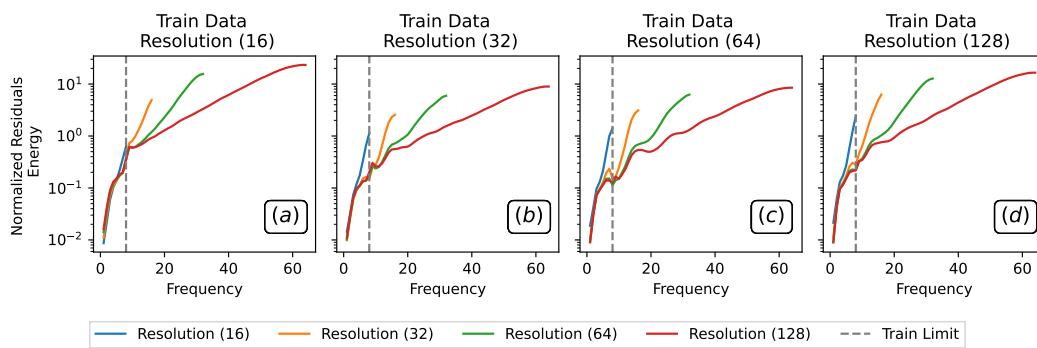
1770
 1771
 1772
 1773
 1774
 1775
 1776
 1777
 1778
 1779
 1780
 1781

1782 **J ASSESSING DEEPONET**
1783

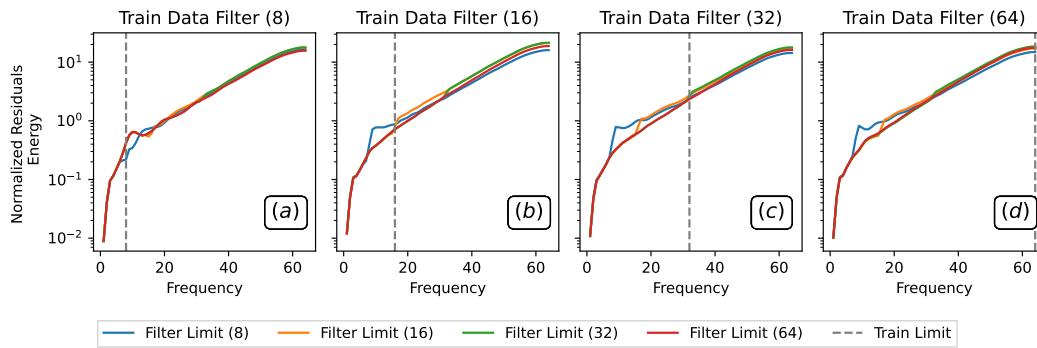
1784 We extend our investigation to include DeepONet (Lu et al., 2021a). DeepONet is a prominent MLO,
1785 however, it notably does not claim to perform accurate **zero-shot** super- and sub-resolution. Despite
1786 this, we still investigate the zero-shot super-resolution, sub-resolution, information extrapolation,
1787 and resolution interpolation abilities of DeepONet. We then assess if multi-resolution training can
1788 enable more accurate multi-resolution inference. Prior to the investigation, we note that DeepONet
1789 operates at a fixed resolution on the input and is only resolution independent on the output.

1790 We begin by performing a hyper-parameter search over learning rates $\in \{1e-2, 1e-3, 1e-4, 1e-5\}$; we set a static weight decay $wd = 1e-5$ as we do not notice significant performance changes
1791 w.r.t. weight decay for FNO/CNO/CROP (See Tables 1 and 2). We find the optimal learning rate is
1792 $1e-3$ which we report in Tab. 2.

1793 Next, we assess DeepONet’s ability to perform information extrapolation and resolution interpolation
1794 via experiments detailed in Sec. 3.1 and find that it cannot do either accurately (see Figures 34
1795 and 35). There is little variability in DeepONet’s information extrapolation performance across test
1796 datasets; there is a similar level of aliasing across all test datasets regardless of frequency content.
1797

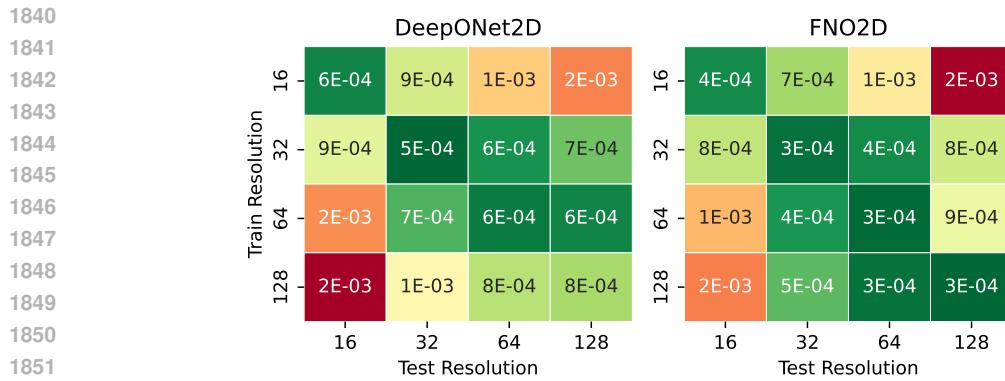


1811 **Figure 34: Resolution Interpolation w/ DeepONet.** Four DeepONets are trained on Darcy data at
1812 resolutions $\{16, 32, 64, 128\}$ from left to right with constant frequency information (low-pass limit
1813 of $8f$), and are tested on varying resolutions. We assess if each model can generalize to data with
1814 varying sampling rate. We visualize spectra of the normalized residuals across test data. Notice,
1815 residual spectra (error) increases substantially in the low frequencies. Lower residual energy at all
1816 frequencies is better.

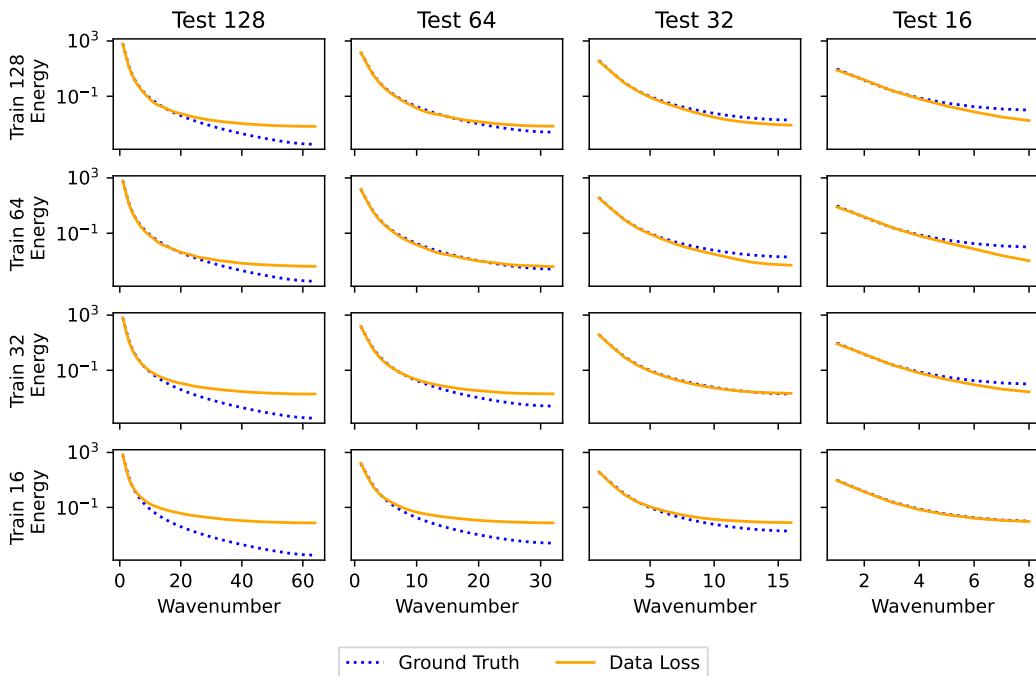


1831 **Figure 35: Information Extrapolation w/ DeepONet.** Four DeepONets trained on Darcy data of
1832 resolution 128 (constant sampling rate) and low-pass filtered with limits $\{8, 16, 32, 64\}f$ (varying
1833 frequency information) from left to right. Test if each model can generalize to data with varying
1834 frequency information. Visualizing spectrum of the normalized residuals across test data. Notice,
1835 residual spectra (error) increases substantially in the high frequencies. Lower residual energy at all
frequencies is better.

1836 We then assess if DeepONet can perform zero-shot super-resolution and sub-resolution. As we
 1837 found with other MLOs, we find that it cannot accurately perform zero-shot super-resolution and
 1838 sub-resolution (See Figures 36 and 37).
 1839



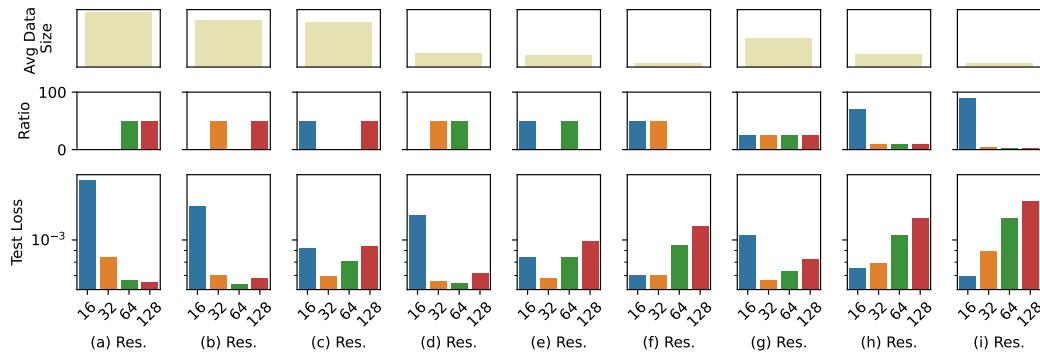
1853 **Figure 36: FNO and DeepONet trained on Darcy.** On average both models incur higher losses
 1854 across test resolutions that differ from the train resolution. This means both FNO and DeepONet
 1855 cannot accurately do zero-shot super-resolution or sub-resolution.
 1856



1880 **Figure 37: DeepONet for Darcy.** Energy spectra for models trained at a specific resolution (y-axis)
 1881 and tested at multiple resolution (x-axis). The spectrums generated by the models do not match
 1882 ground truth for test resolutions that diverge from the training resolution.
 1883

1884 Finally, we assess if multi-resolution training can improve DeepONet’s data-driven multi-resolution
 1885 performance. We employ the same experimental methodology outlined in Sec. 5. We find that
 1886 multi-resolution training **improves** DeepONet’s performance across test resolutions (see Fig. 38).
 1887 We observe that DeepONet is more sensitive than FNO w.r.t. the ratios of data across resolutions:
 1888 DeepONet benefits from having more data samples evenly spread out across resolutions. We hypothe-
 1889 size this is due to the fact that DeepONet is only resolution independent on the output, meaning
 that during training the model must learn mappings from exclusively low resolution inputs to

1890
 1891 varying resolution outputs (low through high). Because low resolution inputs contain less frequency
 1892 components than high-resolution inputs, it is harder learning task than learning between the same
 1893 resolution input output mappings (as FNO does).



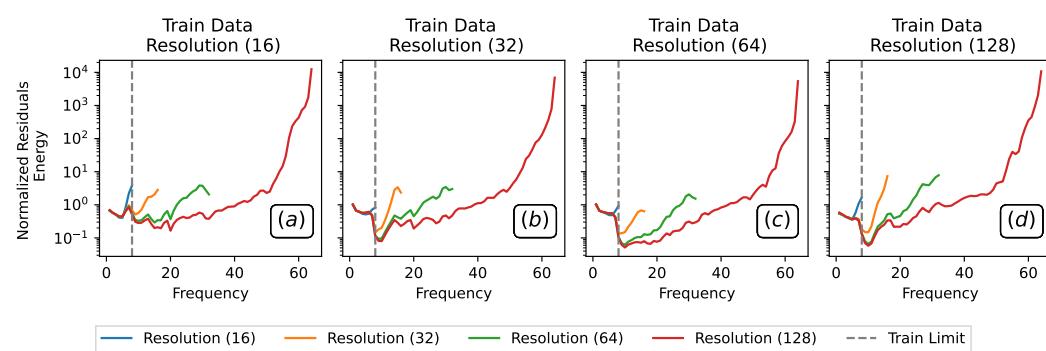
1906
 1907 **Figure 38: Multi-resolution training for DeepONet.** DeepONet trained on multi-resolution Darcy
 1908 data. **Top row:** Average number of pixels in a data sample in the training set. Lower number of
 1909 pixels enables faster data generation and model training. **Middle row:** The ratios of data within
 1910 each resolution bucket. **Bottom row:** The average test loss across different resolutions. Lower loss
 1911 is better. Notice: mixed resolution datasets achieve both low average data size and low loss (ideal
 1912 scenario).

1913
 1914
 1915
 1916
 1917
 1918
 1919
 1920
 1921
 1922
 1923
 1924
 1925
 1926
 1927
 1928
 1929
 1930
 1931
 1932
 1933
 1934
 1935
 1936
 1937
 1938
 1939
 1940
 1941
 1942
 1943

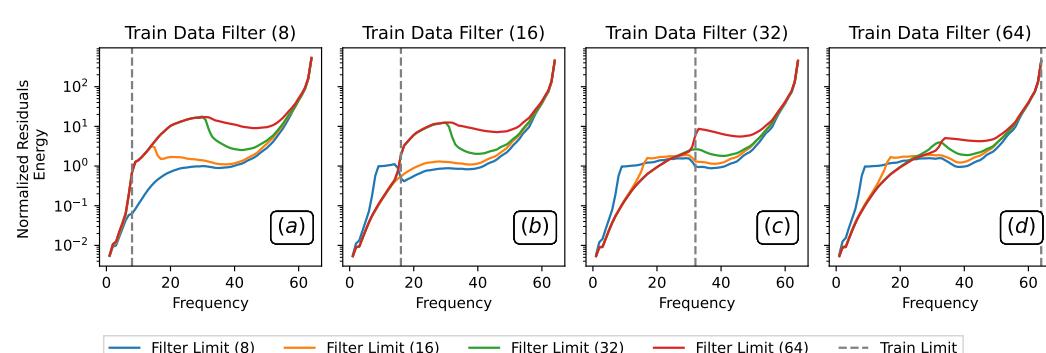
1944 K ASSESSING ANTI-ALIAS NONLINEARITIES

1946 We extend our investigation of FNO to assess the impact nonlinearities play on model aliasing in
 1947 the multi-resolution inference setting. As we discuss in Sec. 6, [Karras et al. \(2021\)](#) observed that
 1948 the application of point-wise nonlinearities to intermediate model representations introduces high-
 1949 frequency components that cannot be resolved via the given resolution. To mitigate this source of
 1950 aliasing, they propose *anti-aliasing* nonlinearities which simply interpolate data to a greater reso-
 1951 lution by a scalar factor, apply the desired non-linearity to the higher resolution representation,
 1952 and then interpolate the signal back down to the original resolution (thus removing many of the
 1953 introduced high-frequency components). The CNO architecture utilizes these anti-alias nonlinear-
 1954 ities. We now investigate the impact of anti-alias nonlinearities to disentangle architecture- vs.
 1955 data-pipeline-driven aliasing for FNO.

1956 We begin by assessing the FNO with anti-aliasing nonlinearities’ ability to perform information
 1957 extrapolation and resolution interpolation via experiments detailed in Sec. 3.1 and find that it cannot
 1958 do either accurately (see Figures 39 and 40). This confirms that the out-of-distribution nature of data
 1959 plays a significant role in aliasing for FNO.



1973 **Figure 39: Resolution Interpolation w/ FNO+Anti-Aliasing Activation.** Four FNOs w/ anti-
 1974 aliasing activation functions are trained on Darcy data at resolutions $\{16, 32, 64, 128\}$ from left to
 1975 right with constant frequency information (low-pass limit of $8f$), and are tested on varying reso-
 1976 lutions. We assess if each model can generalize to data with varying sampling rate. We visualize
 1977 spectra of the normalized residuals across test data. Notice, residual spectra (error) increases sub-
 1978 stantially in the low frequencies. Lower residual energy at all frequencies is better.



1993 **Figure 40: Information Extrapolation w/ FNO+Anti-Aliasing Activation.** Four FNOs w/ anti-
 1994 aliasing activation functions are trained on Darcy data of resolution 128 (constant sampling rate)
 1995 and low-pass filtered with limits $\{8, 16, 32, 64\}f$ (varying frequency information) from left to right.
 1996 Test if each model can generalize to data with varying frequency information. Visualizing spectrum
 1997 of the normalized residuals across test data. Notice, residual spectra (error) increases substan-
 1998 tially in the high frequencies. Lower residual energy at all frequencies is better.

We then assess if FNO with anti-aliasing nonlinearities can perform zero-shot super-resolution and sub-resolution. As we found with other MLOs, we find that it cannot accurately perform zero-shot super-resolution and sub-resolution (See Figures 41 and 42). Therefore, we conclude that while anti-aliasing nonlinearities stem architecture-driven aliasing artifacts in the model forward pass, they do not address the data-driven aliasing artifacts introduced during inference on data of resolutions that differ from a model’s training resolution.

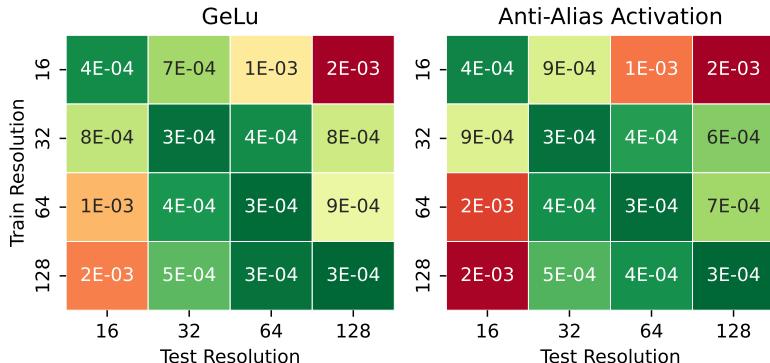


Figure 41: **FNO w/ Anti-Aliasing vs. GeLu activation functions trained on Darcy.** Both models incur higher losses across test resolutions that differ from the train resolution. This means that activation functions alone are not the source of FNO aliasing in the zero-shot multi-resolution setting; this confirms that data at resolutions different than the model’s training resolution is sufficiently out-of-distribution. Zero-shot multi-resolution inference is unreliable regardless of activation function choice.

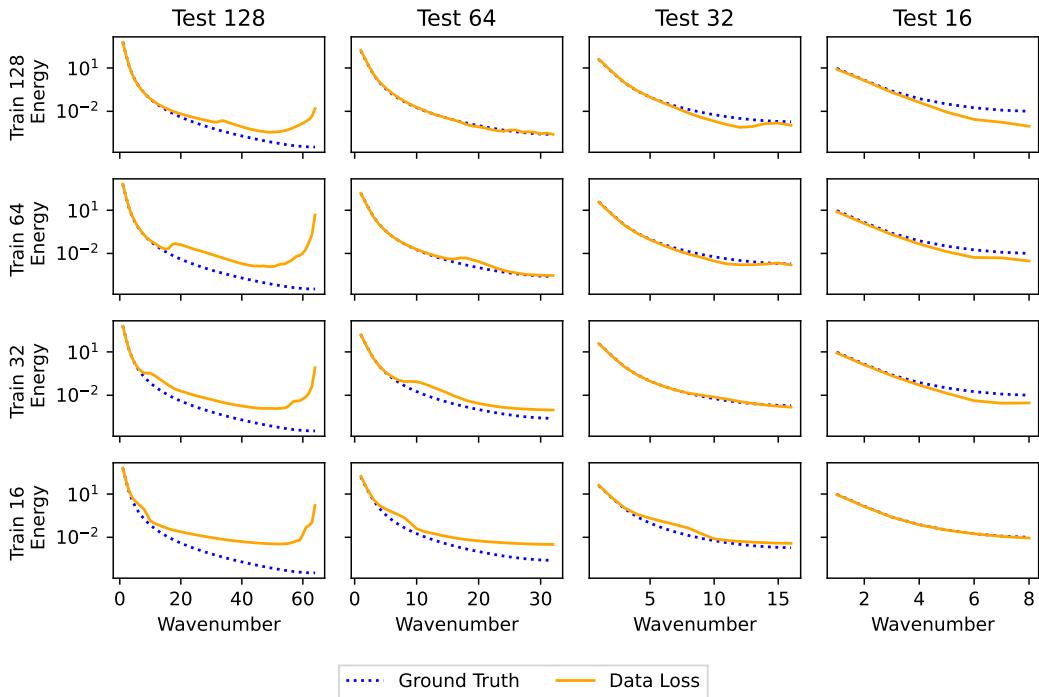
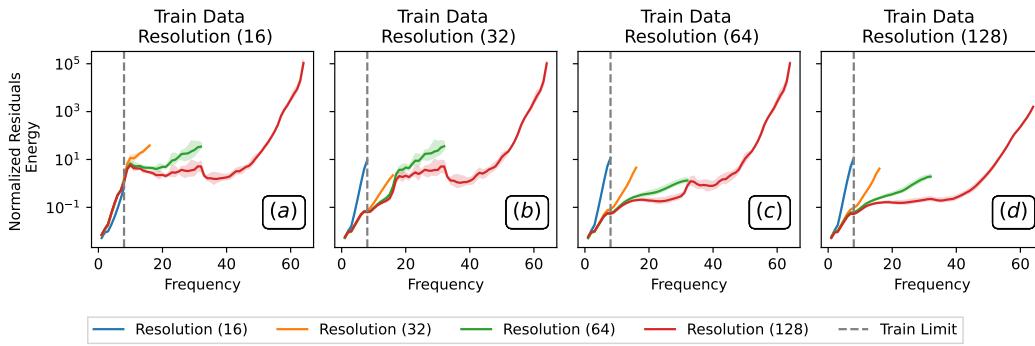


Figure 42: **FNO w/ Anti-Aliasing activation function for Darcy.** Energy spectra for models trained at a specific resolution (y-axis) and tested at multiple resolution (x-axis). The spectrums generated by the models do not match ground truth for test resolutions that diverge from the training resolution.

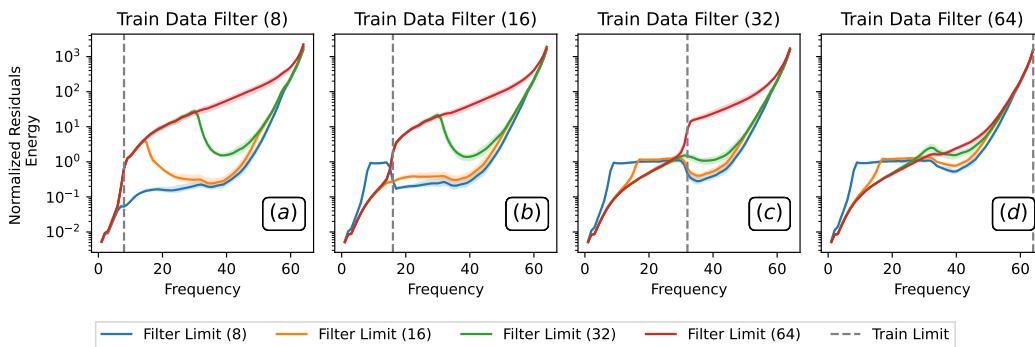
2052 L SENSITIVITY ANALYSIS

2053
 2054 We include an extended sensitivity analysis of our resolution interpolation, information extraction
 2055 and zero-shot super-/sub-resolution experiments for the Darcy dataset. We report the mean and stan-
 2056 dard deviation over these experiments over three random seeds and show consistent performance
 2057 across all seeds. In the main paper, we ensured the robustness of our results by replicating observed
 2058 trends across a wide array of datasets (Darcy, Burgers, Navier-Stokes) as a proxy for generaliz-
 2059 ability of results. Doing multi-seeded experiments was infeasible for the larger datasets (Burgers
 2060 and Navier-Stokes) due to computational cost. By including both multi-seeded experiments and
 2061 experiments over multiple datasets, we ensure that results are generalizable.

2062 We begin by assessing the FNO’s ability to perform information extrapolation and resolution in-
 2063 terpolation via experiments detailed in Sec. 3.1 over three random seeds and find that it cannot do
 2064 either accurately (see Figures 43 and 44). Results are consistent over multiple seeds.



2065
 2066
 2067
 2068
 2069
 2070
 2071
 2072
 2073
 2074
 2075
 2076
 2077 **Figure 43: Resolution Interpolation over 3 Random Seeds.** Four FNOs are trained on Darcy data
 2078 at resolutions $\{16, 32, 64, 128\}$ from left to right with constant frequency information (low-pass
 2079 limit of $8f$), and are tested on varying resolutions. We assess if each model can generalize to data
 2080 with varying sampling rate. We visualize spectra of the normalized residuals across test data. Notice,
 2081 residual spectra (error) increases substantially in the low frequencies. Lower residual energy at all
 2082 frequencies is better. Results averaged over three random seeds; shaded region depicts standard
 2083 deviation.



2084
 2085
 2086
 2087
 2088
 2089
 2090
 2091
 2092
 2093
 2094
 2095
 2096
 2097 **Figure 44: Information Extrapolation over 3 Random Seeds.** Four FNOs are trained on Darcy
 2098 data of resolution 128 (constant sampling rate) and low-pass filtered with limits $\{8, 16, 32, 64\}f$
 2099 (varying frequency information) from left to right. Test if each model can generalize to data with
 2100 varying frequency information. Visualizing spectrum of the normalized residuals across test data.
 2101 Notice, residual spectra (error) increases substantially in the high frequencies. Lower residual energy
 2102 at all frequencies is better. Results averaged over three random seeds; shaded region depicts standard
 2103 deviation.

2104 We then assess if FNO can perform zero-shot super-resolution and sub-resolution over three random
 2105 seeds (see Fig. 45). We observe very similar performance across all seeds: FNO consistently fails

2106 to perform accurate zero-shot super-/sub-resolution. We report mean and standard deviation of this
 2107 experiment across all three seeds in Tab. 5.
 2108

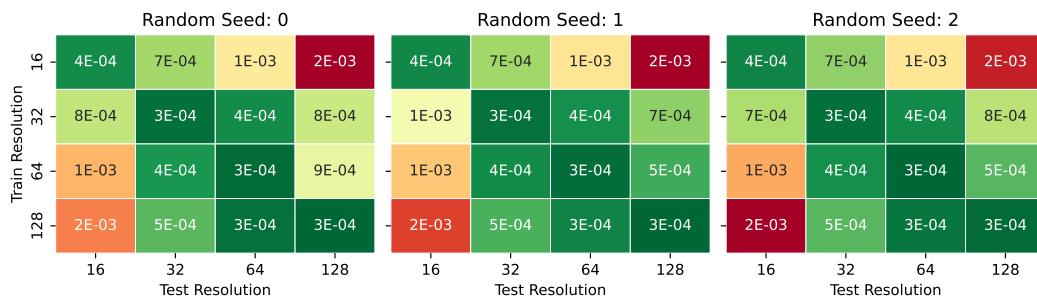


Figure 45: **FNO on Darcy over three random seeds.** On average FNO incurs higher losses across test resolutions that differ from the train resolution across all three random seeds.

Table 5: **Assessing FNO zero-shot super-/sub-resolution over multiple seeds.** Reporting mean and standard deviation of MSE over 3 random seeds.

Train Res. \ Test Res.	16	32	64	128
16	$3.6e-4 \pm 9e-6$	$6.8e-4 \pm 2e-5$	$1.2e-3 \pm 3e-5$	$2.0e-3 \pm 3e-4$
32	$8.2e-4 \pm 1e-4$	$2.9e-4 \pm 2e-6$	$3.8e-4 \pm 9e-6$	$7.4e-4 \pm 8e-5$
64	$1.4e-3 \pm 8e-5$	$3.9e-4 \pm 2e-5$	$2.7e-4 \pm 1e-5$	$6.4e-4 \pm 2e-4$
128	$1.8e-3 \pm 8e-5$	$5.2e-4 \pm 1e-5$	$3.2e-4 \pm 3e-5$	$2.7e-4 \pm 8e-6$

2126
 2127
 2128
 2129
 2130
 2131
 2132
 2133
 2134
 2135
 2136
 2137
 2138
 2139
 2140
 2141
 2142
 2143
 2144
 2145
 2146
 2147
 2148
 2149
 2150
 2151
 2152
 2153
 2154
 2155
 2156
 2157
 2158
 2159