
Physics-Informed CNNs for Super-Resolution of Sparse Observations on Dynamical Systems

Daniel Kelshaw

Department of Aeronautics
Imperial College London
djk21@imperial.ac.uk

Georgios Rigas

Department of Aeronautics
Imperial College London
g.rigas@imperial.ac.uk

Luca Magri

Department of Aeronautics
Imperial College London,
Alan Turing Institute
l.magri@imperial.ac.uk

Abstract

In the absence of high-resolution samples, super-resolution of sparse observations on dynamical systems is a challenging problem with wide-reaching applications in experimental settings. We showcase the application of physics-informed convolutional neural networks for super-resolution of sparse observations on grids. Results are shown for the chaotic-turbulent Kolmogorov flow, demonstrating the potential of this method for resolving finer scales of turbulence when compared with classic interpolation methods, and thus effectively reconstructing missing physics.




Figure 1: Diagrammatic overview of the super-resolution task.

1 Introduction

The application of machine learning for dynamical systems is gaining traction, providing means to extract physical information from data, and reducing the dependence on running computationally expensive simulations [1]. In many cases, access to only sparse or partial observations of a dynamical system is a limiting factor, obscuring the underlying dynamics and providing a challenge for system identification [2]. Super-resolution methods offer the means for high-fidelity state reconstruction from limited observations, a problem of fundamental importance in the physical sciences.

Convolutional neural networks (CNNs) are prominent in the domain of image reconstruction due to their inherent ability to exploit spatial correlations [3–5]. For the classic data-driven approach, there is a dependency on access to samples of high-resolution data. In the absence of ground-truth labels, as with observations on a dynamical system, a common approach is to impose prior knowledge of the physics, regularising predictions with respect to known governing equations [6]. The introduction of physics-informed neural networks (PINNs) has provided new tools for physically-motivated problems, exploiting the automatic-differentiation paradigm provided by neural networks to constrain gradients of the physical system [7–9]. Applications of PINNs for super-resolution show promising results for simple systems, using sparse observations for accurate reconstruction of high-resolution fields [10].

Applications of physics-informed CNNs for super-resolution are less prevalent. While Liu et al. [11] demonstrate a CNN for super-resolution of observations on a dynamical system, there is a strict dependence on high-resolution examples to train the model. Gao et al. [12] explore physics-informed CNNs for the super-resolution of a steady flow field, producing a one-to-one mapping for stationary solutions of the Navier-Stokes equations. Physics-based regularisation of the steady solution neglects any temporal component, introducing further complexities when considering dynamical systems.

Our work extends on the current use of physics-informed CNNs for super-resolution, considering the application to dynamical systems where the temporal nature cannot be neglected. We showcase the super-resolution of sparse, spatial observations of the chaotic-turbulent Kolmogorov flow, highlighting the ability to recover finer scales of turbulence through reconstruction of the missing physics.

2 Defining the Dynamical System

Our work considers super-resolution of sparse observations on dynamical systems of the form

$$\partial_t \mathbf{u} - \mathcal{N}(\mathbf{u}; \lambda) = 0 \quad \text{with} \quad \mathbf{u} = u(\mathbf{x}, t), \quad \mathbf{x} \in \Omega \subset \mathbb{R}^n, \quad t \in [0, T] \subset \mathbb{R}^+, \quad (1)$$

where $u : \Omega \times [0, T] \rightarrow \mathbb{R}^n$, \mathcal{N} is a sufficiently smooth differential operator, and λ are the physical parameters of the system. We define the residual of the system as the left-hand side of equation (1)

$$\mathcal{R}(\mathbf{u}; \lambda) \triangleq \partial_t \mathbf{u} - \mathcal{N}(\mathbf{u}; \lambda), \quad (2)$$

such that $\mathcal{R}(\mathbf{u}; \lambda) = 0$ when $u(\mathbf{x}, t)$ is a solution to the partial differential equation (PDE).

3 The Super-Resolution Task

Given sparse observations on a low-resolution grid, we aim to reconstruct the underlying solution to the PDE on a high-resolution grid. Mathematically, we denote this process by the mapping

$$f_\theta : u(\Omega_L, t) \rightarrow u(\Omega_H, t), \quad (3)$$

where the domain Ω is discretised on uniform, structured grids $\Omega_L \subset \mathbb{R}^{N^n}$, $\Omega_H \subset \mathbb{R}^{M^n}$ such that $\Omega_L \cap \Omega_H = \Omega_L$, and $M = \kappa N$ where $\kappa \in \mathbb{N}^+$ is the up-sampling factor. An overview is provided in Figure 1. We further discretise the time-domain, providing $\mathcal{T} = \{t_i \in [0, T]\}_{i=0}^{N_t}$ for N_t samples. Approximating the mapping f_θ as a CNN, we optimise weights θ of the network to minimise the loss

$$\mathcal{L}_\theta = \alpha \mathcal{L}_O + \mathcal{L}_R, \quad (4)$$

where α represents a fixed, empirical weighting parameter. We further define each loss term as

$$\begin{aligned} \mathcal{L}_O &= \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \|f_\theta(u(\Omega_L, t))|_{\Omega_L} - u(\Omega_L, t)\|_{\Omega_L}^2, \\ \mathcal{L}_R &= \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \|\mathcal{R}(f_\theta(u(\Omega_L, t)); \lambda)\|_{\Omega_H}^2, \end{aligned} \quad (5)$$

where $f_\theta(\cdot)|_{\Omega_L}$ denotes the corestriction of Ω_H on Ω_L , and $\|\cdot\|_\Omega$ represents the ℓ^2 -norm over the given domain. The observation-based loss \mathcal{L}_O seeks to minimise the distance between known observations and their corresponding predictions - not accounting for additional high-resolution information unavailable to the system. We regularise network predictions with the residual-based term \mathcal{L}_R , seeking to ensure that realisations of high-resolution fields satisfy the governing PDEs. As a consequence of imposing prior knowledge of the governing equations through the residual-based loss term, we are effectively able to condition the underlying physics on the observed data, allowing us to recover the underlying solution on the grid Ω_H .¹

¹All code is available on GitHub: <https://github.com/magrilab/PISR>

4 Results

We consider super-resolution for the Kolmogorov flow, a solution of the incompressible Navier-Stokes equations. The system is evaluated on the domain $\Omega \in [0, 2\pi) \subset \mathbb{R}^2$ with periodic boundary conditions applied on $\partial\Omega$, and a stationary, spatially-varying sinusoidal forcing term $g(\mathbf{x})$. The Kolmogorov flow is prominent in studies of turbulence, providing a suitable case study for the super-resolution task; this allows us to evaluate the quality of predictions with respect to the turbulent energy cascade [13]. The standard continuity and momentum equations are

$$\begin{aligned}\nabla \cdot \mathbf{u} &= 0, \\ \partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} &= -\nabla p + \nu \Delta \mathbf{u} + g(\mathbf{x}),\end{aligned}\tag{6}$$

where p, ν denote the scalar pressure field and kinematic viscosity respectively. We take $\nu = 1/34$ to ensure chaotic-turbulent dynamics, and prescribe the standard forcing $g(\mathbf{x}) = [\sin(4\mathbf{x}_2), 0]^\top$.

4.1 Differentiable Pseudospectral Discretisation

We utilise a differentiable pseudospectral spatial discretisation for the Kolmogorov flow, enabling backpropagation for the residual-based loss, $\mathcal{L}_{\mathcal{R}}$. By eliminating the pressure term, our discretisation handles the continuity constraint implicitly, allowing us to neglect the term in the loss [14]. Data is generated by time-integration of the dynamical system with the forward-Euler scheme, taking a time-step $\Delta t = 0.005$ to ensure numerical stability according to the Courant–Friedrichs–Lewy condition. We evaluate the residual-based loss $\mathcal{L}_{\mathcal{R}}$ in the Fourier domain, re-defining the loss

$$\mathcal{L}_{\mathcal{R}} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \|\partial_t \hat{f}_\theta(u(\Omega_{\mathbf{L}}, t)) - \hat{\mathcal{N}}(\hat{f}_\theta(u(\Omega_{\mathbf{L}}, t)))\|_{\hat{\Omega}_{\mathbf{k}}}^2,\tag{7}$$

where $\hat{\Omega}_{\mathbf{k}} \in \mathbb{Z}^{K^n}$ is the discretised Fourier wavespace grid, $\hat{f}_\theta = \mathcal{F} \circ f_\theta$ where \mathcal{F} is the Fourier operator, and $\hat{\mathcal{N}}$ denotes the Fourier differential operator; represented by the differentiable discretisation. We consider evaluation of the loss over discrete time-windows, each consisting of $\tau \geq 2$ consecutive time-steps. This windowing procedure allows for localised computation of the residual in the time domain, alleviating the requirement for the observations to be totally contiguous in time.

4.2 Numerical Experiments

We solve the flow in the Fourier domain on a grid $\hat{\Omega}_{\mathbf{k}} \in \mathbb{Z}^{61 \times 61}$. Data is generated on the high-resolution grid $\Omega_{\mathbf{H}} \in \mathbb{R}^{150 \times 150}$ prior to extracting the low-resolution grid $\Omega_{\mathbf{L}} \in \mathbb{R}^{10 \times 10}$ of observations. A total of 2048 time-windows are used for training, with a further 256 for validation; taking $\tau = 2$ in each instance. We employ the VDSR architecture [15], a variation on the common VGG-net [16]. By prepending the convolutional layers with a bi-cubic upsampling layer, we reformulate the problem in-terms of learning the residual between the interpolation, and the high-resolution field; exploiting benefits associated with residual learning problems [17]. Periodic boundary conditions are embedded in the model through the use of periodic padding in the convolutional layers. For training we use the Adam optimizer with a learning rate of 3×10^{-4} , weighting the loss with $\alpha = 1 \times 10^5$.

Results are compared with bi-linear (*BL*) and bi-cubic (*BC*) interpolation, measuring performance by the relative ℓ^2 -error between network predictions $f_\theta(\Omega_{\mathbf{L}}, t)$, and the true high-resolution fields, $u(\Omega_{\mathbf{H}}, t)$. Figure 2 shows a snapshot of the results, with network predictions yielding qualitatively superior results when compared with the interpolated alternatives. Quantitatively we observe a lower error for the network predictions, achieving a relative ℓ^2 -error of 0.0872, compared with 0.2091 for bi-linear interpolation and 0.1717 for bi-cubic interpolation; averaging results over available data.²

The energy spectrum is characteristic of the direct energy cascade observed in turbulent flows, a multi-scale phenomenon which sees energy content decaying with increasing wavenumber. In Figure 2, we observe that the energy content of the low-resolution field diverges from that of the high-resolution field; a consequence of spectral aliasing. Applying the $2/3$ de-aliasing rule to the Nyquist

²All experiments were run on a single NVIDIA Quadro RTX 8000.

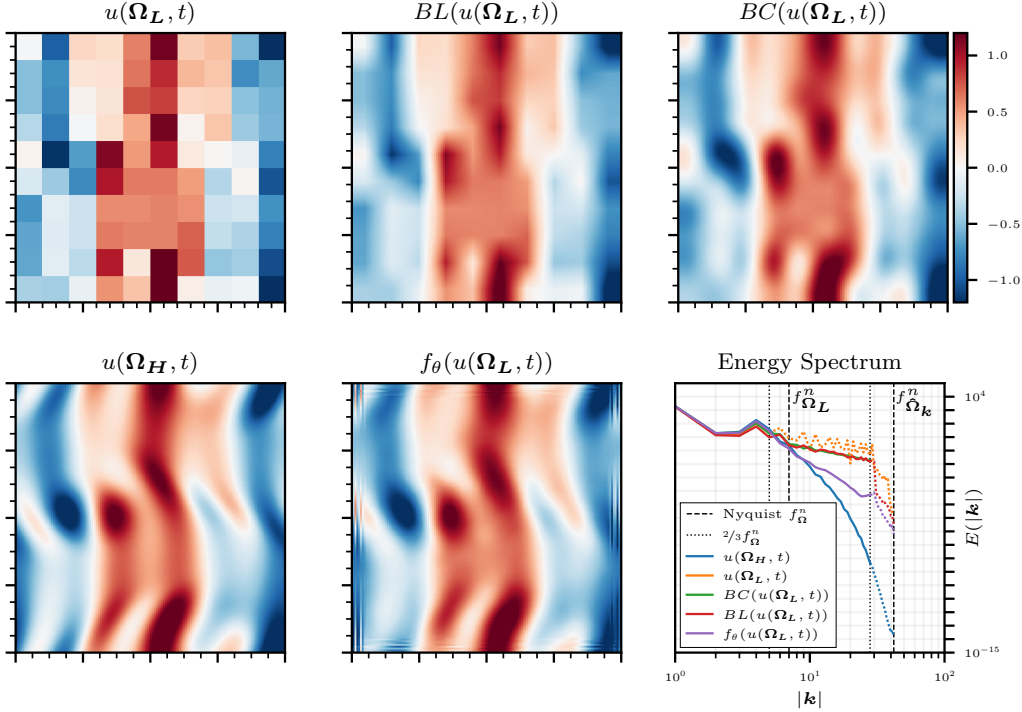


Figure 2: Super-resolution of $u(\Omega_L, t) \in \mathbb{R}^{10 \times 10}$ to $u(\Omega_H, t) \in \mathbb{R}^{150 \times 150}$, demonstrating: bi-linear interpolation, $BL(u(\Omega_L, t))$; bi-cubic interpolation, $BC(u(\Omega_L, t))$; and predictions, $f_\theta(u(\Omega_L, t))$.

frequencies, we see that aliasing is introduced at $|k| = 5, 28$ for the low- and high-resolution fields respectively. This correlates with the divergence in energy content observed in the spectrum. We find that fields super-resolved by the network are capable of capturing finer scales of turbulence compared to low-resolution and interpolation approaches, prior to diverging at $|k| = 10^1$. This result is indicative of the impact of the residual-based loss, signifying the ability of the network to act beyond simple interpolation. The network is capable of de-aliasing, or inferring missing physics.

Despite the use of periodic padding in the convolutional layers, we observe artefacts on the boundaries of the predicted field. These artefacts are largely responsible for the divergence of the predicted energy spectrum from the true spectrum and provide a consideration for future work.

5 Conclusion

In this work, we have demonstrated the application of physics-informed CNNs for the super-resolution of the chaotic-turbulent Kolmogorov flow. Our results show improved performance when compared with standard interpolation methods. Evaluation of the turbulent energy spectrum highlights the ability of the network to resolve turbulent structures on a finer scale than available in the sparse observations. We note that this reconstruction of turbulence showcases the ability to resolve missing physics in the sparse observations, a consequence of embedding prior knowledge of the physics in the loss term. With a view on the spectral representation, we find that the network effectively mitigates the aliasing introduced by taking sparse observations. This work opens opportunities for the accurate reconstruction of dynamical systems from sparse observations, as is often the case in experimental settings for the physical sciences.

Acknowledgments and Disclosure of Funding

D. Kelshaw. and L. Magri. acknowledge support from the UK EPSRC. L. Magri gratefully acknowledges financial support from the ERC Starting Grant PhyCo 949388.

References

- [1] S. Brunton, B. Noack, and P. Koumoutsakos, “Machine learning for fluid mechanics,” *Annual Review of Fluid Mechanics*, vol. 52, pp. 477–508, 1 2020.
- [2] S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Discovering governing equations from data by sparse identification of nonlinear dynamical systems,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 15, pp. 3932–3937, 2016.
- [3] C. Dong, C. C. Loy, K. He, and X. Tang, “Learning a deep convolutional network for image super-resolution,” in *European Conference on Computer Vision*, 2014, pp. 184–199.
- [4] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, “Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network,” *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 9 2016.
- [5] W. Yang, X. Zhang, Y. Tian, W. Wang, J.-H. Xue, and Q. Liao, “Deep learning for single image super-resolution: A brief review,” *IEEE Transactions on Multimedia*, vol. 21, pp. 3106–3121, 12 2019.
- [6] I. E. Lagaris, A. Likas, and D. I. Fotiadis, “Artificial neural networks for solving ordinary and partial differential equations,” *IEEE Transactions on Neural Networks*, vol. 9, pp. 987–1000, 1998.
- [7] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics*, vol. 378, pp. 686–707, 2 2019.
- [8] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, and S. Ho, “Lagrangian neural networks,” in *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2019.
- [9] S. Cai, Z. Mao, Z. Wang, M. Yin, and G. E. Karniadakis, “Physics-informed neural networks (pinns) for fluid mechanics: a review,” *Acta Mechanica Sinica*, vol. 37, pp. 1727–1738, 12 2021.
- [10] H. Eivazi and R. Vinuesa, “Physics-informed deep-learning applications to experimental fluid mechanics,” 2022. [Online]. Available: <https://arxiv.org/abs/2203.15402>
- [11] B. Liu, J. Tang, H. Huang, and X.-Y. Lu, “Deep learning methods for super-resolution reconstruction of turbulent flows,” *Physics of Fluids*, vol. 32, p. 25105, 2020.
- [12] H. Gao, L. Sun, and J.-X. Wang, “Super-resolution and denoising of fluid flow using physics-informed convolutional neural networks without high-resolution labels,” *Physics of Fluids*, vol. 33, p. 073603, 7 2021.
- [13] E. D. Fylladitakis, “Kolmogorov flow: Seven decades of history,” *Journal of Applied Mathematics and Physics*, vol. 6, pp. 2227–2263, 2018.
- [14] C. Canuto, M. Y. Hussaini, A. Quarteroni, and T. A. Zang, *Spectral Methods in Fluid Dynamics*. Springer Berlin Heidelberg, 1988.
- [15] J. Kim, J. K. Lee, and K. M. Lee, “Accurate image super-resolution using very deep convolutional networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [16] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *International Conference on Learning Representations*, 2015.
- [17] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, “Visualizing the loss landscape of neural nets,” in *Neural Information Processing Systems*, 2018, pp. 6391–6401.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?
[Yes] We state the ability of the network to resolve finer scales of turbulence than classical interpolation methods as a result of the super-resolution; this is demonstrated in Figure 2 and discussed in the results section. We show quantitative improvements in terms of relative ℓ^2 -error, as well as justifying results with respect to the turbulent energy cascade.
 - (b) Did you describe the limitations of your work?
[Yes] We discuss the presence of artefacts on the boundaries of the predicted field in the final paragraph of the results, §4. This is largely responsible for the divergence in the energy spectrum, as seen in Figure 2.
 - (c) Did you discuss any potential negative societal impacts of your work?
[No] We do not for-see any potential negative societal impacts.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[N/A]**
 - (b) Did you include complete proofs of all theoretical results? **[N/A]**
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)?
[Yes] We include a link to the code in the footprint of §3.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)?
[Yes] We provide all necessary training details in §4, including: architecture, hyperparameters, amount of data, physical parameters for simulation results. Further details are available on the link to the repository.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)?
[No] We do not believe the inclusion of error-bars is warranted for this study. Results in Figure 2 show snapshots of the super-resolution, and the energy spectrum plot is not well-suited to error bars. Repeats were indeed run, but not reported.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)?
[Yes] We provide a footnote in §4.2 which states the computational resource used.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[N/A]**
 - (b) Did you mention the license of the assets? **[N/A]**
 - (c) Did you include any new assets either in the supplemental material or as a URL?
[Yes] We provide a differentiable pseudospectral spatial discretisation for the Kolmogorov flow, a link to the repository is available in §3 - currently redacted.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[N/A]**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**