

Physics-Informed Deep Learning and Compressive Collocation for High-Dimensional Diffusion-Reaction Equations: Practical Existence Theory and Numerics

Simone Brugiapaglia

SIMONE.BRUGIAPAGLIA@CONCORDIA.CA

*Department of Mathematics and Statistics
Concordia University
1400 De Maisonneuve Blvd. W.
Montréal, QC H3G 1M8, Canada*

Nick Dexter

NICK.DEXTER@FSU.EDU

*Department of Scientific Computing
Florida State University
400 Dirac Science Library
Tallahassee, FL 32306-4120, USA*

Samir Karam

SAMIR.KARAM@MAIL.CONCORDIA.CA

*Department of Mathematics and Statistics
Concordia University
1400 De Maisonneuve Blvd. W.
Montréal, QC H3G 1M8, Canada*

WeiQi Wang

WEIQIWANG@UVIC.CA

*Department of Physics & Astronomy
University of Victoria
3800 Finnerty Road
Victoria, BC V8P 5C2, Canada*

Editor: Animashree Anandkumar

Abstract

On the forefront of scientific computing, Deep Learning (DL), i.e., machine learning with Deep Neural Networks (DNNs), has emerged a powerful new tool for solving Partial Differential Equations (PDEs). It has been observed that DNNs are particularly well suited to weakening the effect of the *curse of dimensionality*, a term coined by Richard E. Bellman in the late '50s to describe challenges such as the exponential dependence of the sample complexity, i.e., the number of samples required to solve an approximation problem, on the dimension of the ambient space. However, although DNNs have been used to solve PDEs since the '90s, the literature underpinning their mathematical efficiency in terms of numerical analysis (i.e., stability, accuracy, and sample complexity), is only recently beginning to emerge. In this paper, we leverage recent advancements in function approximation using sparsity-based techniques and random sampling to develop and analyze an efficient high-dimensional PDE solver based on DL. We show, both theoretically and numerically, that it can compete with a novel stable and accurate compressive spectral collocation method for the solution of high-dimensional, steady-state diffusion-reaction equations with periodic boundary conditions. In particular, we demonstrate a new *practical existence theorem*, which establishes the existence of a class of trainable DNNs with suitable bounds on the

network architecture and a sufficient condition on the sample complexity, with logarithmic or, at worst, linear scaling in dimension, such that the resulting networks stably and accurately approximate a diffusion-reaction PDE with high probability.

Keywords: physics-informed neural networks, compressive Fourier collocation, numerical methods for high-dimensional PDEs, practical existence theorem

1. Introduction

PDEs over high-dimensional domains are a powerful mathematical modelling tool adopted in a variety of applications including molecular dynamics, computational finance, optimal control, and statistical mechanics. Important high-dimensional Partial Differential Equation (PDE) models in these areas are the many-electron Schrödinger equation, the Hamilton–Jacobi–Bellman equation, the Fokker–Planck equation and the Black-Scholes model. Analytic solutions to these equations are in general not available and, hence, it is necessary to design efficient numerical PDE solvers to approximate their solutions. A crucial challenge that immediately arises in this context is the *curse of dimensionality*, see Bellman (1957, 1961). This refers to the tendency of numerical methods for solving high-dimensional problems to exhibit a computational cost or require an amount of data that scales exponentially with the problem’s dimension.

Recent work has shown that *compressive sensing* and *Deep Learning (DL)* are promising techniques to develop efficient high-dimensional PDE solvers and lessen the curse. This success is part of a larger research trend in the area of *scientific machine learning* (Baker et al., 2019), where state-of-the-art techniques from machine learning are applied to solve challenging scientific computing problems, including the numerical solution of PDEs. One of the most popular recent examples in this area are *Physics-Informed Neural Networks (PINNs)*, see Raissi et al. (2019) and earlier studies on the topic by Lagaris et al. (1998), which recently gained an impressive amount of attention in the scientific computing community. In particular, DL based methods have shown great promise for high-dimensional PDEs, see Han et al. (2018) and the recent review paper by E et al. (2021), and PDEs on domains with complex geometries, see, e.g., Chen et al. (2022).

Concurrently, recent advancements involving compressive sensing in scientific computing include the adoption of sparsity-based techniques for function approximation from random samples, see Rauhut and Ward (2012), whose initial success was due to their application in the field of *Uncertainty Quantification (UQ)* of parametric PDEs by Doostan and Owhadi (2011), see also Adcock et al. (2022b) for a comprehensive review of the topic. Here we focus on *Compressive Fourier Collocation (CFC)*, a method proposed in Wang and Brugiapaglia (2024) as an improvement of the compressive spectral collocation method from Brugiapaglia (2020) and able to lessen the curse of dimensionality in the number of collocation points. A detailed literature review on compressive sensing and DL methods for PDEs can be found in §1.2.

Motivated by these recent advances, in this paper we study and compare numerical solvers for high-dimensional PDEs based on compressive sensing and DL from both the theoretical and the numerical viewpoint, focusing on steady-state diffusion-reaction equations with periodic boundary conditions. The choice of the time-independent setting introduces a gap with real-world applications. Yet, it allows us to carry out a rigorous analysis while

keeping the level of technical difficulty moderate and focus on the impact of the physical domain’s dimension, which is the main subject of this work. Our methodological approach is inspired by the recent paper by Adcock and Dexter (2021), where a similar practical and theoretical study was made in the context of high-dimensional function approximation from pointwise samples.

1.1 Main Contributions

Our main contributions, of both theoretical and computational nature, are summarized below.

i) New convergence theorem for high-dimensional periodic PINNs. The main theoretical contribution of the paper is a new convergence theorem for periodic PINNs applied to (possibly high-dimensional) steady-state diffusion-reaction problems with periodic boundary conditions. Our theoretical guarantee is a convergence result in the form of *practical existence theorem*, a theoretical approach recently developed in the context of scalar-valued function approximation via DL in Adcock and Dexter (2021) and Hilbert-valued function approximation relevant to parametric PDEs and UQ in Adcock et al. (2022a). In contrast with better-known universal approximation theorems (see, e.g., Elbrächter et al. (2021) and references therein) that prove the existence of networks with suitable approximation properties, a *practical* existence theorem also addresses the issues of network training and sample complexity. Our result, stated in Theorem 1 and proved in Section 5, establishes the existence of a class of trainable periodic PINNs with explicit architecture bounds and shows that networks in this class can achieve near-optimal approximation rates for PDE solutions that are sparse with respect to the Fourier basis through training using a number of samples that is only mildly affected by the curse of dimensionality (i.e., that scales logarithmically or, at worst, linearly with the domain’s dimension d). Despite the presence of the adjective “practical” in the terminology used to refer to Theorem 1, there are still gaps between this result and common implementation practice. For example, Theorem 1 is a result about networks obtained as minimizers of a certain regularized training loss. However, it does not involve any specific optimization method that may be used to approximate such minimizers (e.g., stochastic gradient descent or Adam). Further limitations of Theorem 1 are discussed extensively in §6.

ii) Numerical study of CFC and periodic PINNs in high dimensions. Our second contribution is of computational nature and it is the implementation and numerical study of CFC and periodic PINNs for high-dimensional, steady-state diffusion-reaction problems on the torus \mathbb{T}^d , up to dimension $d = 30$. We also compare these two methods by studying their accuracy as a function of the training set size (i.e., the number of collocation points). The code to reproduce our experiments can be found in the GitHub repository https://github.com/WeiqiWangMath/PINN_high_dimensional_PDE.

iii) Improved implementation of CFC. We propose a new variant of CFC, based on *adaptive lower Orthogonal Matching Pursuit (OMP)* recovery (see Algorithm 1). This variant improves the accuracy of the CFC method from Wang and Brugiapaglia (2024) under suit-

able structural assumptions on the PDE solution and allows the method to scale in higher dimensions. This is due to the way it constructs the approximation support set by adaptively exploring the set of possible candidates to be added in an iterative procedure, thereby avoiding dealing with a large *a priori* truncation set.

iv) CFC convergence theorem for diffusion-reaction problems. We show the convergence of CFC for steady-state, periodic diffusion-reaction problems in Theorem 8. Although this is mainly an auxiliary step needed to prove our main theoretical result (Theorem 1) it is an extension of the analysis in Wang and Brugiapaglia (2024) (relying on the same techniques employed therein) of independent interest.

1.2 Contextualization of Our Contributions

Numerical methods for PDEs based on compressive sensing and sparse recovery have been considered in tandem with different discretization techniques. This includes Galerkin (Jokar et al., 2010), Petrov-Galerkin (Brugiapaglia et al., 2021b, 2015), Fourier-Galerkin (Gross and Iwen, 2025), isogeometric analysis (Brugiapaglia et al., 2020), spectral collocation approaches (Brugiapaglia, 2020; Wang and Brugiapaglia, 2024) and methods based on the sparse Fourier transform (Daubechies et al., 2007). In this paper, we are interested in methods that can be applied to high-dimensional domains. In particular, we focus on *Compressive Fourier Collocation (CFC)* (see Wang and Brugiapaglia, 2024), which aims to compute a sparse approximation to the PDE solution with respect to the Fourier basis from random collocation points. This method will be reviewed in detail in §3. The Wavelet-Fourier CORSING method (Brugiapaglia et al., 2021b) can in principle be implemented in domains in arbitrary dimension, although making practical numerical implementations scale in high dimensions is nontrivial. The sparse spectral method proposed in Gross and Iwen (2025), based on Fourier-Galerkin discretization and sublinear time algorithms, can scale to extremely high-dimensional problems.

The design of numerical PDE solvers based on neural networks dates back to the 1990s (Lagaris et al., 1998; Lee and Kang, 1990). More recently, this field became extremely popular thanks to the introduction of *Physics-Informed Neural Networks (PINNs)* (see Raissi et al., 2019; Karniadakis et al., 2021). Solvers based on DL have shown great promise specifically in the case of high-dimensional PDEs. In this direction, approaches proposed in the literature include the Deep Galerkin Method (Sirignano and Spiliopoulos, 2018), methods based on the reformulation of high-dimensional PDEs as Stochastic Differential Equations (SDEs) (Han and Jentzen, 2017; Han et al., 2018), PINNs for high-dimensional problems (Hu et al., 2024; Zeng et al., 2022), and deep genetic algorithms (Putri et al., 2024). In this paper, we will consider PINNs combined with a periodic layer (Dong and Ni, 2021) to solve high-dimensional PDEs on the d -dimensional torus. This approach will be presented in detail in §2.1.

Other approaches for the numerical solution of PDEs able to scale to moderately high dimensions include sparse grid methods (see, e.g., Shen and Yu (2010, 2012)), methods based on tensor-based approximation and low rank structures (Bachmayr and Dahmen, 2015; Bachmayr et al., 2016; Dahmen et al., 2016), and sparse grid spectral methods (Kupka, 1997).

Our main theoretical result (Theorem 1) is a convergence theorem for PINNs over high-dimensional periodic domains. Currently, the convergence analysis of PINNs is an active research area and several studies have appeared in the literature. However, the theory is arguably far from being fully developed. The analysis in Shin et al. (2020) shows the convergence of PINNs for linear second-order elliptic problems, but it is based on Hölder-type regularization that is in general not implementable. The study in Shin et al. (2023) provides *asymptotic* convergence results (i.e., with training set size $m \rightarrow \infty$) relying on assumptions in terms of Bernstein-type inequalities or Rademacher’s complexity that could be challenging to verify in practice. The analysis in Doumèche et al. (2025) shows convergence results for PINNs of asymptotic type and nonasymptotic convergence rates for the expected squared L^2 -error of the form $\log^2(m)/m^{1/2}$ for Sobolev-type regularized loss functions, where m is the number of training data points used to collocate the PDE in the PINN’s loss function. The work (De Ryck and Mishra, 2022) proves generalization bounds for PINNs for Kolmogorov equations showing that PINNs can lessen the curse of dimensionality, under the assumption that the network’s weights are bounded. Compared to these results, our practical existence theorem (Theorem 1) has the advantage of being of nonasymptotic type and the corresponding error bound leads to fast convergence rates with respect to the size of the training set when the PDE solution is sparse or compressible with respect to the Fourier basis.

Finally, the framework of practical existence theorems developed in Adcock and Dexter (2021); Adcock et al. (2022a) for scalar- and Hilbert-valued high-dimensional function approximation was recently extended to Banach-valued function approximation in Adcock et al. (2025) and reduced-order modelling of parametric PDEs based on convolutional autoencoders in Franco and Brugiapaglia (2025). This framework relies on recently proposed convergence results for compressive sensing-based approximation in high dimensions by Adcock et al. (2022b). For a review of practical existence theory, see Adcock et al. (2024). All the practical existence theorems developed so far rely on the emulation of orthogonal polynomials with neural networks (see Daws and Webster, 2019; De Ryck et al., 2021; Opschoor et al., 2022) and are only applied to approximating Hilbert or Banach-valued functions relevant to parametric PDEs or scalar-valued functions. In this paper, we extend the scope of practical existence theorems by emulating Fourier basis functions and considering the case of (non-parametric) PDE solvers.

1.3 Outline of the Paper

We briefly outline the organization of the paper. We start by illustrating the model problem (a high-dimensional periodic diffusion-reaction equation), periodic PINNs and our main theoretical result (Theorem 1) in §2. Then, we present CFC and adaptive lower OMP in §3. §4 contains an extensive numerical study of CFC and periodic PINNs. The proof of our convergence result is presented in §5. Finally, we draw some conclusions and describe possible directions of future work in §6. Additional experiments investigating the numerical impact of certain components of Theorem 1 are presented in Appendix A.

2. Problem Setting

Notation. We start by recalling some standard mathematical notation that will be employed throughout the paper. We denote the d -dimensional torus by \mathbb{T}^d , with $d \in \mathbb{N}$, where $\mathbb{T} := [0, 1] / \sim$ and \sim is the equivalence relation on $[0, 1]$ defined by $x \sim y$ if and only if $x - y \in \mathbb{Z}$. $L^2(\mathbb{T}^d)$ denotes the space of square-integrable functions, equipped with inner product $\langle v, w \rangle := \int_{\mathbb{T}^d} v(\mathbf{x})\overline{w(\mathbf{x})} d\mathbf{x}$ and norm $\|v\|_{L^2} = \langle v, v \rangle^{1/2}$. Moreover, $L^\infty(\mathbb{T}^d)$ denotes the space of functions such that $\|v\|_{L^\infty} := \text{ess sup}_{\mathbf{x} \in \mathbb{T}^d} |v(\mathbf{x})| < \infty$. We also consider the Sobolev spaces $H^k(\mathbb{T}^d)$, $k = 1, 2$ equipped with norms $\|v\|_{H^1} := (\|v\|_{L^2}^2 + \|\nabla v\|_{L^2}^2)^{1/2}$ and $\|v\|_{H^2} := (\|v\|_{H^1}^2 + \|\nabla^2 v\|_{L^2}^2)^{1/2}$, respectively. Here ∇ and ∇^2 denote the gradient and the Hessian operators, respectively. Moreover, $\|\nabla v\|_{L^2}^2 = \int_{\mathbb{T}^d} \|\nabla v(\mathbf{x})\|_2^2 d\mathbf{x}$ and $\|\nabla^2 v\|_{L^2}^2 = \int_{\mathbb{T}^d} \|\nabla^2 v(\mathbf{x})\|_F^2 d\mathbf{x}$, where $\|\cdot\|_2$ is the discrete 2-norm and $\|\cdot\|_F$ is the Frobenius norm. We will also consider the Sobolev spaces $W^{k,\infty}(\mathbb{T}^d)$, $k = 1, 2$, equipped with norms $\|v\|_{W^{1,\infty}} := \max\{\|v\|_{L^\infty}, \|\nabla v\|_{L^\infty}\}$ and $\|v\|_{W^{2,\infty}} := \max\{\|v\|_{W^{1,\infty}}, \|\nabla^2 v\|_{L^\infty}\}$, respectively. Here $\|\nabla v\|_{L^\infty} = \text{ess sup}_{\mathbf{x} \in \mathbb{T}^d} \|\nabla v(\mathbf{x})\|_\infty$ and $\|\nabla^2 v\|_{L^\infty} = \text{ess sup}_{\mathbf{x} \in \mathbb{T}^d} \max_{i,j} |(\nabla^2 v(\mathbf{x}))_{ij}|$. $C^k(\mathbb{T}^d)$ denotes the space of k -times continuously differentiable functions over \mathbb{T}^d . We denote the set of first n positive integers by $[n] := \{1, \dots, n\}$. The cardinality of a set X is denoted as $|X|$. For a vector $\mathbf{z} \in \mathbb{C}^N$, we define its support as $\text{supp}(\mathbf{z}) = \{j \in [N] : z_j \neq 0\}$.

Model problem. Our model problem is a periodic diffusion-reaction equation over the d -dimensional torus \mathbb{T}^d . In this paper, we are interested in the scenario where $d \gg 1$. Moreover, we consider the following steady-state, periodic diffusion-reaction equation as our model problem:

$$-\nabla \cdot (a(\mathbf{x})\nabla u(\mathbf{x})) + \rho u(\mathbf{x}) = f(\mathbf{x}), \quad \forall \mathbf{x} \in \mathbb{T}^d, \quad (1)$$

where $u : \mathbb{T}^d \rightarrow \mathbb{R}$ is the PDE solution, and the diffusion coefficient $a : \mathbb{T}^d \rightarrow \mathbb{R}$, the reaction term $\rho \in \mathbb{R}$ and the forcing term $f : \mathbb{T}^d \rightarrow \mathbb{R}$ are assumed to satisfy

$$a \in C^1(\mathbb{T}^d), \quad \min_{\mathbf{x} \in \mathbb{T}^d} a(\mathbf{x}) \geq a_{\min} > 0, \quad \rho > 0, \quad \text{and} \quad f \in L^2(\mathbb{T}^d). \quad (2)$$

These conditions are sufficient for the problem (1) to be well-posed (see, e.g., (Brugiapaglia et al., 2021b, Proposition 2.1)) and guarantee that its weak solutions belong to the Sobolev space $H^2(\mathbb{T}^d)$ (see, e.g., (Evans, 2010, §6.3)). Throughout the paper, we also assume f to be regular enough for pointwise evaluations to be well-defined. Moreover, we define the PDE operator

$$\mathcal{L}[u](\mathbf{x}) := -\nabla \cdot (a(\mathbf{x})\nabla u(\mathbf{x})) + \rho u(\mathbf{x}), \quad \forall \mathbf{x} \in \mathbb{T}^d, \quad \forall u \in H^2(\mathbb{T}^d). \quad (3)$$

Despite its simplicity, (1) is an interesting model problem since it shares the same second-order diffusion term with more complex PDEs such as the Black-Scholes, Schrödinger and Fokker-Planck models.

2.1 Physics-Informed Neural Networks (PINNs) with Periodic Layer

We start by illustrating the PINN setting adopted to solve the periodic high-dimensional diffusion-reaction problem (1). First, we describe the Deep Neural Network (DNN) architecture employed in the method, then illustrate the training strategy. The framework

presented here will encompass both our main theoretical result (presented in §2.2) and the numerical experiments in §4.

To enforce periodic boundary conditions, we consider an approach proposed in Dong and Ni (2021). This is achieved by adding a C^∞ *periodic layer* to the DNN as the first layer. The periodic layer contains, in turn, two layers of width dl , for some $l \in \mathbb{N}$, denoted as $\mathbf{q}^{(1)} = (q_{ij}^{(1)})_{i \in [d], j \in [l]}$ and $\mathbf{v}^{(2)} = (v_{ij}^{(2)})_{i \in [d], j \in [l]}$, respectively. The neuron $q_{ij}^{(1)}$ operates on the i th component of the input vector $\mathbf{x} \in \mathbb{R}^d$ and applies a cosine transformation to enforce periodicity while adding a phase shift parameter ϕ_{ij} :

$$q_{ij}^{(1)}(\mathbf{x}) = \cos(2\pi x_i + \phi_{ij}), \quad \forall \mathbf{x} \in \mathbb{R}^d, \quad \forall i \in [d], \forall j \in [l], \quad (4)$$

where $\phi = (\phi_{ij})_{i \in [d], j \in [l]} \in \mathbb{R}^{dl}$ are trainable parameters and l controls the number of neurons per dimension. Then, the neurons in $\mathbf{v}^{(2)}$ collect the outputs of the corresponding nodes in $\mathbf{q}^{(1)}$, and apply an affine transformation and (possibly nonlinear) activation function $\sigma^{(2)} : \mathbb{R}^{dl} \rightarrow \mathbb{R}^{dl}$, i.e.,

$$\mathbf{v}^{(2)}(\mathbf{x}) = \sigma^{(2)} \left(\text{diag}(\mathbf{w}^{(2)})\mathbf{x} + \mathbf{b}^{(2)} \right), \quad \forall \mathbf{x} \in \mathbb{R}^{dl}, \quad (5)$$

where $\mathbf{w}^{(2)}$ and $\mathbf{b}^{(2)}$ are dl -dimensional vectors of trainable parameters and $\text{diag}(\mathbf{w}^{(2)})$ is a $dl \times dl$ matrix with $\mathbf{w}^{(2)}$ on the main diagonal and zeros elsewhere. After the periodic layer, the DNN has $h \in \mathbb{N}$ traditional hidden layers having width $w \in \mathbb{N}$. Each of these takes the output of the previous layer as input and applies a trainable affine transformation and a componentwise activation function σ to it. These hidden layers are denoted by

$$\mathbf{v}^{(k)}(\mathbf{x}) = \sigma^{(k)}(W^{(k)}\mathbf{x} + \mathbf{b}^{(k)}), \quad \forall \mathbf{x} \in \begin{cases} \mathbb{R}^{dl} & \text{if } k = 3 \\ \mathbb{R}^w & \text{otherwise} \end{cases}, \quad \forall k = 3, \dots, h+2, \quad (6)$$

where the weight matrix $W^{(k)}$ is $w \times dl$ if $k = 3$ and $w \times w$ otherwise, $\mathbf{b}^{(k)} \in \mathbb{R}^w$ and $\sigma^{(k)} : \mathbb{R}^w \rightarrow \mathbb{R}^w$. §4.2.3 discusses the selection of the hyper-parameters l (number of neurons per dimension in the periodic layer), h and w defining the DNN architecture. The last hidden layer of the network $\mathbf{v}^{(h+2)}$ activates linearly into one output neuron so that the output of the network is scalar-valued. That is,

$$v^{(h+3)}(\mathbf{x}) = W^{(h+3)}\mathbf{x}, \quad \forall \mathbf{x} \in \mathbb{R}^w,$$

where $W^{(h+3)}$ is $1 \times w$. In this paper, we consider componentwise activations $\sigma^{(k)}$ of the form $\sigma^{(k)}(\mathbf{x}) = (\sigma_j^{(k)}(x_j))$, where $\sigma_j^{(k)}$ could be either a linear activation (i.e., $\sigma_j^{(k)}(x) = x$) or a nonlinear activation of the following two types: the hyperbolic tangent (i.e., $\sigma_j^{(k)}(x) = \tanh(x)$) or the *Rectified Power Unit (RePU)*, defined by

$$\text{RePU}_\ell(x) := \max\{0, x^\ell\}, \quad \ell \in \mathbb{N}. \quad (7)$$

Note that when $\ell = 1$ the function RePU_ℓ is the Rectified Linear Unit (ReLU). In our numerical experiments (see §4) we will consider \tanh activations. However, RePU and linear activations will be used to derive our theoretical result (see §2.2).

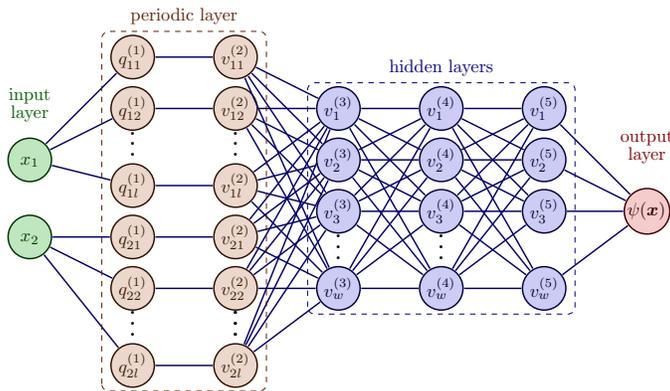


Figure 1: Architecture of the neural network with the periodic layer ($d = 2$). The number in superscript represents the layer number. In this case, $h = 3$ and $\text{depth}(\psi) = 7$.

In summary, we consider DNNs $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$ of the form

$$\psi = v^{(h+3)} \circ v^{(h+2)} \dots \circ v^{(2)} \circ q^{(1)}.$$

Fig.1 depicts the architecture of a DNN with periodic layer in dimension $d = 2$. We also define the depth and width of the network as

$$\text{width}(\psi) = \max\{dl, w\} \quad \text{and} \quad \text{depth}(\psi) = h + 4, \quad (8)$$

where we included the input and the output layer in the depth count.

The DNN $\psi(\mathbf{x})$ is then trained to approximate the solution $u(\mathbf{x})$ of the high-dimensional PDE. Given collocation points drawn independently and uniformly at random from \mathbb{T}^d , i.e.,

$$\mathbf{x}_1, \dots, \mathbf{x}_m \stackrel{\text{i.i.d.}}{\sim} \text{Uniform}(\mathbb{T}^d), \quad (9)$$

the parameters ϕ , $\mathbf{w}^{(2)}$, $W^{(k)}$, $\mathbf{b}^{(k)}$ of the DNN are learned by minimizing a regularized *Root Mean Squared Error (RMSE)* loss, i.e.,

$$\min_{\psi} \sqrt{\frac{1}{m} \sum_{i=1}^m |\mathcal{L}[\psi](\mathbf{x}_i) - f(\mathbf{x}_i)|^2 + \lambda \mathcal{R}(\psi)}, \quad (10)$$

where \mathcal{L} is the PDE operator defined in (3), $\lambda \geq 0$ is a tuning parameter, and \mathcal{R} is a regularization term that usually involves the ℓ^2 - or ℓ^1 -norm of the networks' weights (corresponding to the *weight decay* or *sparse regularization* strategies, respectively). Minimizing the loss function (10) corresponds to finding the network ψ that minimizes the PDE residual at the collocation points in a (regularized) least-squares sense. In the general PINN setting for solving stationary PDEs, the loss function usually consists of the sum of two components: the first one aims at minimizing the PDE residual (like in (10)) and the second component enforces boundary conditions. However, in our setting the periodic layer forces ψ to be periodic, hence automatically enforcing boundary conditions. For this reason, the training

loss does not contain a boundary condition term. The loss is then minimized by stochastic gradient descent methods. More technical details on the training procedure can be found in §4.2. We conclude by noting that other losses different from the RMSE can be considered for the PDE residual minimization. These include losses based on, e.g., L^p -norms (Wang et al., 2022) or Sobolev norms (Son et al., 2021). Here we limit our attention to the regularized RMSE loss for the sake of simplicity and also because our main theoretical result holds for this loss. The presence of the regularization term \mathcal{R} will be crucial for our convergence theorem, illustrated in the next subsection. In our numerical results we will train by simply minimizing the unregularized RMSE (or, equivalently, unregularized MSE) loss.

2.2 A Practical Existence Theorem for Periodic PINNs

Before presenting our main result, namely a convergence theorem for periodic PINNs based on the framework of practical existence theory (Theorem 1), we need to introduce some definitions and further technical ingredients. In a nutshell, our main result shows that trained periodic PINNs are able to achieve the same accuracy as a sparse Fourier approximation of the PDE solution using a training set of collocation points whose size scales logarithmically or, at worst, linearly with the dimension d . This is a natural choice given the presence of periodic boundary conditions in (1). In addition, we will require some technical conditions on the PDE coefficients a and ρ .

Target accuracy: sparse Fourier approximation. The elements of the L^2 -orthonormal Fourier basis are defined as

$$F_{\boldsymbol{\nu}}(\boldsymbol{x}) = \exp(2\pi i \boldsymbol{\nu} \cdot \boldsymbol{x}), \quad \forall \boldsymbol{\nu} \in \mathbb{Z}^d, \forall \boldsymbol{x} \in \mathbb{T}^d. \quad (11)$$

In particular, we will focus on *sparse* Fourier approximations supported on *hyperbolic crosses* (see, e.g., Dũng et al. (2018); Temlyakov (2018) and references therein). The *hyperbolic cross* of order n is a multi-index set of \mathbb{Z}^d defined as

$$\Lambda_{d,n}^{\text{HC}} = \left\{ \boldsymbol{\nu} \in \mathbb{Z}^d : \prod_{k=1}^d (|\nu_k| + 1) \leq n \right\}. \quad (12)$$

The hyperbolic cross is a convenient choice in high-dimensional approximation since its cardinality grows moderately with respect to d and n when compared to other standard multi-index set choices such as the *tensor product* and the *total degree* sets, see, e.g., (Adcock et al., 2022b, §2.3). Specifically, the following upper bound on the cardinality of the hyperbolic cross holds (see Kühn et al. (2015); Chernov and Dũng (2016), and also Eq. (2.10) in (Wang and Brugiapaglia, 2024)):

$$|\Lambda_{d,n}^{\text{HC}}| \leq \min\{4n^5 16^d, e^2 n^{2+\log_2(d)}\}. \quad (13)$$

In order to leverage the CFC convergence theory needed for our practical existence theorem, we consider a rescaled version of the system $\{F_{\boldsymbol{\nu}}\}_{\boldsymbol{\nu} \in \mathbb{Z}^d}$, defined by

$$\Psi_{\boldsymbol{\nu}} = \frac{1}{4\pi^2 \|\boldsymbol{\nu}\|_2^2 + \rho/a_0} F_{\boldsymbol{\nu}}, \quad \forall \boldsymbol{\nu} \in \mathbb{Z}^d, \quad (14)$$

where $a_{\mathbf{0}} = \langle a, F_{\mathbf{0}} \rangle = \int_{\mathbb{T}^d} a(\mathbf{x}) \, d\mathbf{x}$. As explained in detail in §5.1, this rescaling ensures that $\{\mathcal{L}[\Psi_{\nu}]\}_{\nu \in \mathbb{Z}^d}$ is a *bounded Riesz system* (Brugiapaglia et al., 2021a) of $L^2(\mathbb{T}^d)$ under sufficient conditions on a and ρ , see (16) and (17) below, which is a crucial property needed to show the convergence of the CFC method. To gain some intuition about this rescaling, consider the simple case of constant diffusion $a \equiv 1$. Applying the PDE operator to the rescaled Fourier system yields $\mathcal{L}[\Psi_{\nu}] = (-\Delta + \rho)\Psi_{\nu} = F_{\nu}$ for all $\nu \in \mathbb{Z}^d$. In this simple scenario, $\{\mathcal{L}[\Psi_{\nu}]\}_{\nu \in \mathbb{Z}^d}$ is a bounded orthonormal system and, as such, is ideally suited for compressive sensing, (see Foucart and Rauhut, 2013, §12).

In this setting, we consider a finite-dimensional truncation of the solution u to a finite multi-index set $\Lambda \subseteq \mathbb{Z}^d$ (expressed with respect to the rescaled Fourier basis), i.e.,

$$u_{\Lambda} = \sum_{\nu \in \Lambda} c_{\nu} \Psi_{\nu}, \quad \text{with } c_{\nu} = (4\pi^2 \|\nu\|_2^2 + \rho/a_{\mathbf{0}}) \cdot \langle u, F_{\nu} \rangle, \quad (15)$$

and let $\mathbf{c}_{\Lambda} = (c_{\nu})_{\nu \in \Lambda} \in \mathbb{C}^N$, where $N = |\Lambda|$. An s -sparse approximation to u is then obtained by only keeping s terms in the expansion (15). In general, we recall that a vector is said to be s -sparse if it has at most s nonzero entries. The best possible accuracy of such an approximation is measured by the *best s -term approximation error*, see (Cohen et al., 2009) and references therein, defined as

$$\sigma_s(\mathbf{c}_{\Lambda})_p = \min_{\mathbf{z} \in \mathbb{C}^N} \{\|\mathbf{c}_{\Lambda} - \mathbf{z}\|_p : \mathbf{z} \text{ is } s\text{-sparse}\}.$$

Sufficient conditions on the PDE coefficients. We also introduce a technical condition on the PDE coefficients a and ρ necessary for the convergence result. We consider diffusion coefficients $a : \mathbb{T}^d \rightarrow \mathbb{R}$ having a sparse Fourier expansion, namely,

$$a = a_{\mathbf{0}} + \sum_{\nu \in T} a_{\nu} F_{\nu}, \quad \text{with } a_{\nu} = \langle a, F_{\nu} \rangle, \quad \forall \nu \in T \cup \{\mathbf{0}\}, \quad (16)$$

for some $T \subseteq \mathbb{Z}^d \setminus \{\mathbf{0}\}$. In our main result, we will assume that a and ρ satisfy the following assumption:

$$\sqrt{|T|} \cdot \|a - a_{\mathbf{0}}\|_{H^1} < \sqrt{\left(a_{\mathbf{0}} + \frac{\rho}{4\pi^2}\right)^2 + a_{\mathbf{0}}^2} - \left(a_{\mathbf{0}} + \frac{\rho}{4\pi^2}\right). \quad (17)$$

Note that in the case of constant diffusion (i.e., $a \equiv a_{\mathbf{0}}$) the condition above is always satisfied under (2). In general, condition (17) controls the oscillatory behaviour of a . We emphasize that the sparsity of a and condition (17) are sufficient for our convergence theorem to hold, but we do not claim (nor believe) they are necessary. We also observe that our analysis does not take into account the potential sample complexity needed to estimate the term $a_{\mathbf{0}} = \int_{\mathbb{T}^d} a(\mathbf{x}) \, d\mathbf{x}$. In order to check the validity of condition (17), one needs direct access to this quantity. If it is not already available, the extra samples needed to compute it via, e.g., Monte Carlo quadrature or other high-dimensional integration methods would have to be taken into account.

We are now in a position to state our main result, which provides the existence of a class of trained neural networks with architecture of the form described in §2.1 able to approximate the solution to (1) with accuracy comparable to that of an s -sparse Fourier approximation with high probability.

Theorem 1 (Practical existence theorem for high-dimensional periodic PINNs)

Given a dimension $d \in \mathbb{N}$, target sparsity $s \in \mathbb{N}$, hyperbolic cross order $n \in \mathbb{N}$, RePU power $\ell \in \mathbb{N}$, with $\ell \geq 2$, and probability of failure $\varepsilon \in (0, 1)$, there exist:

- (i) a class of neural networks \mathcal{N} of the form described in §2.1 with d -dimensional input and 1-dimensional output layers, RePU $_\ell$ or linear activations, complex-valued weights and biases, and such that, for all $\psi \in \mathcal{N}$,

$$\text{width}(\psi) \leq c_\ell^{(1)} \cdot \min \left\{ 4n^5 16^d, e^2 n^{2+\log_2 d} \right\} \cdot d \cdot \min \{2^d, n\}, \quad (18)$$

$$\text{depth}(\psi) \leq c^{(2)} \cdot (\log_2(n) + \min \{\log_2 d, n\}), \quad (19)$$

- (ii) a regularization function $\mathcal{R} : \mathcal{N} \rightarrow [0, \infty)$,

- (iii) a choice of tuning parameter λ depending only on a , ρ and s ,

such that the following holds with probability $1 - \varepsilon$. For all $a : \mathbb{T}^d \rightarrow \mathbb{R}$ and $\rho \in \mathbb{R}$ satisfying (2), (16) and (17), let

$$m \geq c_{a,\rho}^{(3)} \cdot s \cdot \log^2 \left(c_{a,\rho}^{(4)} \cdot s \right) \cdot (\min \{\log(n) + d, \log(2n) \log(2d)\} + \log(\varepsilon^{-1})), \quad (20)$$

$\Lambda = \Lambda_{d,n}^{\text{HC}}$ as in (12), and consider collocation points $\mathbf{x}_1, \dots, \mathbf{x}_m$ randomly and independently sampled from the uniform measure on \mathbb{T}^d . Then, every minimizer $\hat{\psi}$ of the training program

$$\min_{\psi \in \mathcal{N}} \sqrt{\frac{1}{m} \sum_{i=1}^m |\mathcal{L}[\psi](\mathbf{x}_i) - f(\mathbf{x}_i)|^2} + \lambda \mathcal{R}(\psi) \quad (21)$$

satisfies

$$\|u - \hat{\psi}\|_{L^2} + \|(\Delta - \rho)(u - \hat{\psi})\|_{L^2} \leq C_{a,\rho}^{(1)} \cdot \frac{\sigma_s(\mathbf{c}_\Lambda)_1}{\sqrt{s}} + C_{a,d,\rho}^{(2)} \cdot \left(\frac{\|u - u_\Lambda\|_{W^{2,\infty}}}{\sqrt{s}} + \|u - u_\Lambda\|_{H^2} \right). \quad (22)$$

Moreover, if $\rho < 1$, we also have

$$\|u - \hat{\psi}\|_{H^2} \leq C_{a,\rho}^{(3)} \cdot \frac{\sigma_s(\mathbf{c}_\Lambda)_1}{\sqrt{s}} + C_{a,d,\rho}^{(4)} \cdot \left(\frac{\|u - u_\Lambda\|_{W^{2,\infty}}}{\sqrt{s}} + \|u - u_\Lambda\|_{H^2} \right). \quad (23)$$

Here each constant depends only on the subscripted parameters. Moreover, the dependence of each constant on d is at most linear (when present).

Proof (sketch) A complete proof of Theorem 1 is given in §5.2. It leverages the convergence theory of CFC (presented and proved subsequently in §5.1) thanks to an *ad hoc* construction of the network class \mathcal{N} . The idea is to construct networks $\psi \in \mathcal{N}$ so as to explicitly replicate linear combinations of Fourier functions $\{F_\nu\}_{\nu \in \Lambda}$ supported on the hyperbolic cross Λ , recall (15). In this construction, only the last layer, corresponding to the coefficients of the linear combination is trained. The rest of the network is explicitly constructed to replicate (a suitably rescaled version of) the Fourier basis functions. In this

setting, minimizing the regularized loss (21) where $\mathcal{R}(\psi)$ is the 1-norm of the weights in the last layer of the network ψ is equivalent to solving a sparse regularization problem, specifically, a square-root LASSO problem (Adcock et al., 2019; Belloni et al., 2011), which allows to rigorously connect periodic PINNs’ training with the CFC convergence theory (see Theorem 8). \blacksquare

We conclude this section by highlighting some important features of Theorem 1.

- (i) Theorem 1 is called a *practical* existence theorem since, as opposed to standard neural network existence results such as universal approximation theorems, see, e.g., (Elbrächter et al., 2021) and references therein, it not only guarantees the existence of neural networks (in this case, periodic PINNs) with favorable approximation properties, but also establishes that such networks can be computed by training a regularized RMSE loss and provides a sufficient condition on the minimum number of samples (i.e., collocation points) for the training process to be successful.
- (ii) One of the key benefits of Theorem 1 is that the minimum number of training samples m needed to successfully train a periodic PINN from the class \mathcal{N} to solve a d -dimensional reaction-diffusion problem scales logarithmically (when $d \gg n$) or, at worst, linearly (when $n \gg d$) in d . This indicates that periodic PINNs are provably able to alleviate the curse of dimensionality.
- (iii) As established by the error bound (22), the periodic PINN approximation accuracy guaranteed by Theorem 1 is controlled by the best s -term approximation error $\sigma_s(\mathbf{c}_\Lambda)_1$ and by the truncation error $u - u_\Lambda$ (measured with respect to the $W^{2,\infty}$ - and the H^2 -norm). This accuracy is inherited by the CFC convergence result that Theorem 1’s proof relies on (see Theorem 8). In this paper we do not assume that u belongs to a specific function class, but we note that the best s -term approximation and the hyperbolic cross truncation error could be bounded for functions satisfying suitable mixed regularity conditions, see (Dũng et al., 2018; Temlyakov, 2018) and the discussion in (Wang et al., 2022, §2.3). In general, $\sigma_s(\mathbf{c}_\Lambda)_1$ can be estimated via Stechkin’s inequality, see, e.g., (Adcock et al., 2022b, Lemma 3.5).
- (iv) Theorem 1 also provides explicit bounds on the architecture of the periodic PINNs from class \mathcal{N} . The networks’ depth scales logarithmically in the dimension d and the networks’ width scales polynomially in d (note that $n^{\log_2 d} = d^{\log_2 n}$), for sufficiently large d .

The main limitations of Theorem 1 and related open problems are discussed in §6. The numerical impact of some of the theoretical ingredients of Theorem 1 is investigated in Appendix A.

3. Compressive Fourier Collocation (CFC)

In this section, we illustrate the CFC method and its efficient numerical implementation via adaptive lower OMP. For more details, we refer to Wang and Brugiapaglia (2024). Similarly to §2.2, we consider a finite multi-index set $\Lambda \subset \mathbb{Z}^d$, the rescaled Fourier basis $\{\Psi_\nu\}_{\nu \in \mathbb{Z}^d}$

defined in (14) and a finite-dimensional expansion u_Λ of the form (15). Then, similarly to the PINN approach, we collocate the diffusion-reaction equation (1) by means of Monte Carlo sampling. Hence, we randomly generate m i.i.d. uniform points $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathbb{T}^d$, as in (9). Letting $N = |\Lambda|$, we assume to have an ordering for the multi-indices in $\Lambda = \{\nu_1, \dots, \nu_N\}$ (e.g., the lexicographic ordering). The PDE collocation process leads to the linear system

$$Az = \mathbf{b}, \quad (24)$$

where $A \in \mathbb{C}^{m \times N}$ and $\mathbf{b} \in \mathbb{C}^m$ are defined by

$$A_{ij} = \frac{1}{\sqrt{m}}[-\nabla \cdot (a\nabla\Psi_{\nu_j}) + \rho\Psi_{\nu_j}](\mathbf{x}_i) \quad \text{and} \quad b_i = \frac{1}{\sqrt{m}}f(\mathbf{x}_i), \quad \forall i \in [m], j \in [N]. \quad (25)$$

We refer to A as the CFC matrix. This collocation method is *compressive* since we choose $m \ll N$. This makes the linear system (24) underdetermined.

Computing the CFC solution. Following the approach in Wang and Brugiapaglia (2024), a CFC approximation \hat{u} to u can be computed by approximately solving the underdetermined linear system (24) via sparse recovery techniques. Therefore, we need to determine (i) a suitable truncation multi-index set Λ and (ii) a sparse recovery method to approximately solve the linear system (24). In this paper, the truncation set Λ is chosen as a *hyperbolic cross* of order n , i.e., $\Lambda = \Lambda_{d,n}^{\text{HC}}$, recall (12). The hyperbolic cross offers a twofold advantage. First, as discussed in §2.2, the cardinality of $\Lambda_{d,n}^{\text{HC}}$ grows moderately in n and d when compared to other standard multi-index families (e.g., tensor product and total degree set). Second, a hyperbolic cross can be characterized as the union of *lower* (or, equivalently, *downward closed* or *monotone*) sets of a given cardinality (*cp.* Definition 2).

After fixing Λ , we compute an approximate solution $\hat{\mathbf{c}} = (\hat{c}_\nu)_{\nu \in \Lambda} \in \mathbb{C}^N$ to the underdetermined linear system (24). To do so, we employ sparse recovery techniques, such as *Orthogonal Matching Pursuit (OMP)* or ℓ^1 minimization (Foucart and Rauhut, 2013). Finally, we define the CFC approximation as

$$\hat{u} = \sum_{\nu \in \Lambda} \hat{c}_\nu \Psi_\nu. \quad (26)$$

In Wang and Brugiapaglia (2024) it was shown that \hat{u} is an accurate and stable approximation to u for high-dimensional diffusion equations, under sufficient conditions on the diffusion term and for a number of collocation points that scales only logarithmically with the ambient dimension d . In Theorem 8, we will extend the CFC convergence analysis from Wang and Brugiapaglia (2024) to diffusion-reaction problems. From the computational viewpoint, if one wants to compute a very sparse CFC approximation to the solution u , then OMP typically offers a faster reconstruction than solving an ℓ^1 minimization problem via a convex optimization solver.

3.1 Adaptive Lower Orthogonal Matching Pursuit (OMP)

For very high-dimensional domains (say, $d > 20$), the cardinality of the hyperbolic cross, despite being moderate with respect to other multi-index choices, becomes considerably large.

This can make storing A and computing \hat{c} genuinely challenging. To deal with higher dimensions, we need a more efficient recovery procedure that does not rely on choosing a large *a priori* truncation set Λ , but constructs it iteratively. This can be achieved by considering a reconstruction strategy called *adaptive lower Orthogonal Matching Pursuit (OMP)*, that we now illustrate. The adaptive lower OMP algorithm presented here is based on analogous techniques employed in adaptive high-dimensional approximation, in particular in sparse grids (Gerstner and Griebel, 2003) and least squares methods (Migliorati, 2014, 2019).

First, we define the notion of lower set. Lower sets are an important class of multi-index sets in approximation theory and we refer to, e.g., Cohen and Migliorati (2018) or (Adcock et al., 2022b, §1.5) and references therein for further reading. Note that for polynomial approximations, lower sets are typically defined in \mathbb{N}_0^d . However, since we employ the complex Fourier basis here we consider lower sets in \mathbb{Z}^d . Before defining lower sets, we introduce a convenient notation to compare multi-indices of \mathbb{Z}^d . For $\boldsymbol{\mu}, \boldsymbol{\nu} \in \mathbb{Z}^d$, we use $\boldsymbol{\mu} \preceq \boldsymbol{\nu}$ to indicate that $|\mu_i| \leq |\nu_i|$ for every $i \in [d]$. Moreover, if $\boldsymbol{\mu} \preceq \boldsymbol{\nu}$ and there exists an $i \in [d]$ such that $|\mu_i| < |\nu_i|$, then we say that $\boldsymbol{\mu} \prec \boldsymbol{\nu}$.

Definition 2 (Lower set of \mathbb{Z}^d) *A multi-index set $\Lambda \subseteq \mathbb{Z}^d$ is said to be lower if the following holds for all $\boldsymbol{\nu} \in \Lambda$ and $\boldsymbol{\mu} \in \mathbb{Z}^d$: if $\boldsymbol{\mu} \preceq \boldsymbol{\nu}$, then $\boldsymbol{\mu} \in \Lambda$.*

An example of a lower set is given in Fig. 2 (blue dots). An equivalent condition for $\Lambda \subseteq \mathbb{Z}^d$ to be a lower set is the following: if $\boldsymbol{\nu} \in \Lambda$, then the box $\prod_{k=1}^d [-|\nu_k|, |\nu_k|] \subseteq \Lambda$. Note that lower sets of \mathbb{Z}^d are symmetric with respect to all coordinate hyperplanes $\{\boldsymbol{\nu} \in \mathbb{Z}^d : \nu_k = 0\}$, with $k \in [d]$. This symmetry can be justified as follows. For real-valued solutions $u : \mathbb{T}^d \rightarrow \mathbb{R}$ one could employ a real Fourier expansion with respect to basis functions of the form $\prod_{k=1}^d \xi_k(2\pi\nu_k x_k)$, where $\xi_k \in \{\sin, \cos\}$ and $\boldsymbol{\nu} \in \mathbb{N}_0^d$. It is a simple exercise to verify that the complex Fourier expansion of each of these real-valued trigonometric functions is supported on a set symmetric with respect to every coordinate hyperplane.

In order to create a mechanism able to iteratively enlarge a multi-index set while preserving the lower set structure, we introduce the notion of reduced margin.

Definition 3 (Reduced margin) *The reduced margin of a (nonempty) lower set $\Lambda \subseteq \mathbb{Z}^d$ is the multi-index set $\mathcal{R}(\Lambda) := \{\boldsymbol{\nu} \in \mathbb{Z}^d : \boldsymbol{\nu} \notin \Lambda \text{ and } \forall \boldsymbol{\mu} \prec \boldsymbol{\nu}, \boldsymbol{\mu} \in \Lambda\}$. Moreover, we let $\mathcal{R}(\emptyset) := \{\mathbf{0}\}$.*

We provide an illustration of the reduced margin in Fig. 2 (red dots). A key property of the reduced margin is that if $\Lambda \subseteq \mathbb{Z}^d$ is a lower set and $\boldsymbol{\nu} \in \mathcal{R}(\Lambda)$ then $\Lambda \cup \{\boldsymbol{\mu} \in \mathbb{Z}^d : |\mu_i| = |\nu_i|, \forall i \in [d]\}$ is also a lower set.

The adaptive lower OMP algorithm is a structured variant of OMP, see, e.g., (Foucart and Rauhut, 2013, §3.2) and references therein, with the only difference that the greedy search is restricted to iteratively enlarged lower sets. The algorithm generates a nested sequence of lower sets $\emptyset = \Lambda^{(0)} \subseteq \Lambda^{(1)} \subseteq \dots \subseteq \Lambda^{(K)}$ and then computes a sequence of least-squares solutions $\mathbf{z}^{(n)}$ to the linear system (24) such that $\text{supp}(\mathbf{z}^{(n)}) \subseteq \Lambda^{(n)}$. At each iteration, adaptive lower OMP extends the existing multi-index set Λ_n by picking an element $\boldsymbol{\nu}^{(n)} \in \mathcal{R}(\Lambda^{(n)})$ corresponding to the largest absolute residual $|(A^*(\mathbf{b} - A\mathbf{z}^{(n)}))_{\boldsymbol{\nu}}|$. Then, it also adds all reflections of $\boldsymbol{\nu}^{(n)}$ with respect to the coordinate hyperplanes $\{\boldsymbol{\nu} \in \mathbb{Z}^d : \nu_k = 0\}$ to preserve the lower structure. The greedy selection criterion of adaptive lower OMP (Line 5 in Algorithm 1) is illustrated in Fig. 2. The method is summarized in Algorithm 1.

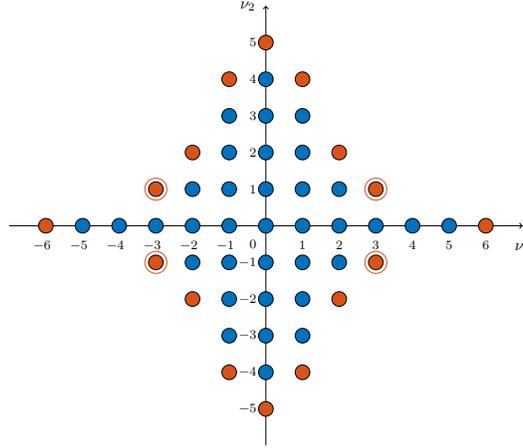


Figure 2: Illustration of the greedy selection criterion of adaptive lower OMP (Line 5 in Algorithm 1). Considering a given lower set $\Lambda^{(n)} \subset \mathbb{Z}^2$ (blue), its reduced margin $\mathcal{R}(\Lambda^{(n)})$ is drawn in red. If the quantity $|A^*(\mathbf{b} - A\mathbf{z}^{(n)})_{\nu}|$ is minimized at $\nu = (3, 1)$, then $\Lambda^{(n+1)}$ is constructed by adding the red circled dots to $\Lambda^{(n)}$.

Algorithm 1 Adaptive lower Orthogonal Matching Pursuit (OMP)

Require: $A \in \mathbb{C}^{m \times N}$, $\mathbf{b} \in \mathbb{C}^m$, and number of iterations $K \in \mathbb{N}$

Ensure: A K -sparse vector $\hat{\mathbf{c}} \in \mathbb{C}^N$, approximately solving (24)

- 1: $\Lambda^{(0)} = \emptyset$, $\mathbf{z}^{(0)} = \mathbf{0}$,
 - 2: **for** $n = 0, \dots, K - 1$ **do**
 - 3: Compute the columns of A corresponding to $\mathcal{R}(\Lambda^{(n)})$ (see Definition 3)
 - 4: $d_{\nu} \leftarrow 1/\sqrt{\sum_{i=1}^m |A_{i\nu}|^2}$, $A_{i\nu} \leftarrow A_{i\nu}d_{\nu}$, for all $\nu \in \mathcal{R}(\Lambda^{(n)})$ (ℓ^2 -normalize the new columns of A)
 - 5: $\nu^{(n)} \leftarrow \arg \max_{\nu \in \mathcal{R}(\Lambda^{(n)})} |(A^*(\mathbf{b} - A\mathbf{z}^{(n)}))_{\nu}|$ (greedy multi-index selection)
 - 6: $\Lambda^{(n+1)} \leftarrow \Lambda^{(n)} \cup \{\boldsymbol{\mu} \in \mathbb{Z}^d : |\mu_i| = |\nu_i^{(n)}|, \forall i \in [d]\}$
 - 7: $\mathbf{z}^{(n+1)} \leftarrow \arg \min_{\mathbf{z} \in \mathbb{C}^N} \left\{ \|\mathbf{b} - A\mathbf{z}\|_2^2, \text{supp}(\mathbf{z}) \subseteq \Lambda^{(n+1)} \right\}$
 - 8: **end for**
 - 9: $\hat{\mathbf{c}} \leftarrow D\mathbf{z}^{(K)}$, with $D = \text{diag}(\mathbf{d})$
-

4. A Numerical Study of CFC and Periodic PINNs

In this section, we present numerical experiments on the high-dimensional periodic diffusion-reaction equation (1). After illustrating the numerical setup in §4.1, we conduct tests on PINNs with periodic layer in §4.2 and the (adaptive) CFC method in §4.3. Then, we compare the two approaches in §4.4.

4.1 Numerical Setup

We start by introducing the setup of our numerical experiments.

Measurement of errors. In all numerical experiments, we use the relative L^2 -error to measure the approximation error. It is defined as

$$\text{relative } L^2\text{-error} = \frac{\|u - \hat{u}\|_{L^2}}{\|u\|_{L^2}},$$

where \hat{u} is the computed approximation to an exact solution u . Computing the L^2 -norms is a challenge in high dimension, therefore the norms $\|\cdot\|_{L^2}$ are approximated using Monte Carlo integration, i.e.,

$$\|u\|_{L^2} \approx \sqrt{\frac{1}{M} \sum_{i=1}^M |u(\mathbf{x}_i)|^2},$$

where the $\mathbf{x}_1, \dots, \mathbf{x}_M$ are $M \gg 1$ random independent points uniformly distributed in \mathbb{T}^d . Throughout this section, we choose $M = 10000$.

Diffusion coefficient. As shown in Wang and Brugiapaglia (2024), the choice of the type of diffusion coefficient (e.g., constant, sparse, or non-sparse) does not impact the numerical results for the compressive Fourier collocation method solving high-dimensional PDEs. Thus, in all experiments, we use the following diffusion coefficient:

$$a(\mathbf{x}) = 1 + 0.25 \sin(2\pi x_1) \sin(2\pi x_2).$$

Exact solutions. In our experiments, we will consider the problems having the following three high-dimensional functions as their solution:

$$u_1(\mathbf{x}) = \sin(4\pi x_1) \sin(2\pi x_2), \quad (\text{Example 1}) \quad (27)$$

$$u_2(\mathbf{x}) = \exp(\sin(2\pi x_1) + \sin(2\pi x_2)), \quad (\text{Example 2}) \quad (28)$$

$$u_3(\mathbf{x}) = \exp\left(\sum_{k=1}^d \frac{1}{k^2} \sin(2\pi x_k)\right). \quad (\text{Example 3}) \quad (29)$$

In the following, we will use Examples 1, 2, and 3 to refer to diffusion-reaction equations with exact solutions given by (27), (28), and (29), respectively. Note that for any given coefficients a and ρ , we will enforce the PDE to have a prescribed exact solution by suitable choice of the forcing term f . Here, u_1 is defined in terms of only one real-valued trigonometric function. This corresponds to a 4-sparse solution with respect to the complex Fourier system. Hence, in this case we have $\sigma_s(\mathbf{c}_\Lambda) = 0$ and $u - u_\Lambda = 0$ as soon as the truncation set Λ contains the four multi-indices $(\pm 2, \pm 1, 0, 0, \dots)$. Note, though, that these four multi-indices do not form a lower set. The solution u_2 is a smooth periodic function active only in the variables x_1 and x_2 . Hence, it exhibits a highly anisotropic behaviour. Finally, u_3 is another smooth solution defined using all variables from x_1, \dots, x_d . The behaviour of this solution is also anisotropic and the solution is uniformly bounded by

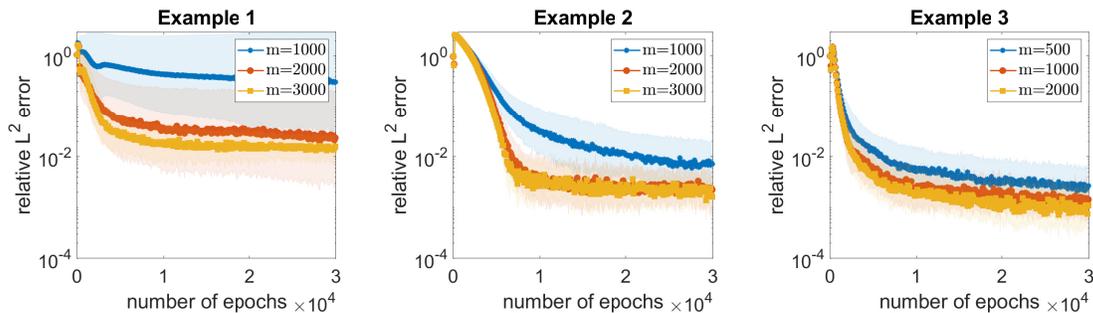


Figure 3: (Impact of the number of samples) Relative L^2 -error versus the number of epochs for approximating the exact solutions defined in (27)-(29) with $d = 6$, where m is the number of sample points.

$$\exp\left(\sum_{k=1}^{\infty} 1/k^2\right) = \exp(\pi^2/6).$$

Randomization of experiments and visualization. Due to the random nature of collocation points in the CFC and the periodic PINN method, we consider 25 random runs for each test. However, for the experiments in §4.2.3, where we examine the effect of the architecture on performance, we use 10 random runs for each width-depth combination. In our plots, the main curves represent the sample geometric mean of the error, and the size of the lightly shaded areas corresponds to its corrected geometric standard deviation. Note that calculating the geometric mean and standard deviation of the errors is equivalent to compute the classical mean and standard deviation of the log transformed errors. For more details on the visualization strategy, we refer to (Adcock et al., 2022b, §A.1.3).

4.2 Physics-Informed Neural Networks

We run tests using PINNs to solve the PDE defined in (1) with the various exact solutions from (27)-(29). For the experiments in §4.2.1 and §4.2.2, we use DNNs with $l = 11$ nodes in the periodic layer, depth $h = 3$ hidden layers, and width to depth ratio $r = 10$ so that the networks have 30 nodes per hidden layer. Moreover, we consider a tanh activation in every layer. In §4.2.3, we examine the impact of the architecture on performance by comparing results obtained by varying the number of nodes on the periodic layer and on the hidden layers of the fully connected part of the network, see Fig. 1. In this section, we train periodic PINNs without regularization, i.e., we let $\lambda = 0$ in (10). Moreover, we simply minimize the MSE loss (which is equivalent to minimizing the RMSE in the non-regularized setting). Additional experiments involving a regularized loss function can be found in Appendix A.

4.2.1 PERFORMANCE AND THE IMPACT OF THE NUMBER OF SAMPLES

We test the PINNs with the exact solutions defined in (27)-(29) and refer to them as Example 1, 2, 3, respectively. Fig. 3 presents the performance of the PINNs over 30000 epochs for each exact solution, with parameters $d = 6$, $\rho = 0.5$. The PINNs have 11 nodes on the periodic layer and three non-periodic layers with 30 nodes each. We choose the

commonly used tanh activation function and the Adam optimizer (Kingma and Ba, 2017) with early stopping. We use this neural network to conduct all numerical tests except those in §4.2.3. As shown in Fig. 3, the PINN model approximates all three exact solutions with L^2 -error on the order of 10^{-2} to 10^{-3} within the 30000 epoch training budget. Notably, the error decreases rapidly in the first 10000 epochs and can be seen to saturate for the remaining epochs. Also, in Examples 1 and 3, we see some small improvement in the error as the size of the training set is increased, while for Example 2, we note that moving from $m = 2000$ to $m = 3000$ points does not noticeably improve the error.

These differences in results obtained for our examples illustrate that the number of samples needed to reach a minimum error varies depending on the complexity of the exact solution. This difference in behaviour is not directly explained by Theorem 1. In fact, Example 1 is the the easiest one to approximate in the Fourier basis as it is only a linear combination of 4 Fourier basis functions, but it corresponds to the worst performance for periodic PINNs. The difficulty in approximating Example 1 could be heuristically explained by the presence of a higher frequency 4π (as opposed to 2π) associated with the variable x_1 . Yet, a rigorous explanation of this phenomenon remains an open problem.

4.2.2 IMPACT OF THE DIMENSION

We now consider the performance of PINNs in higher-dimensional settings. We present the results for exact solutions with $d = 6, 10, 20$ in Fig. 4. With 3000 sample points, the PINNs approximate Example 1's exact solution with relative L^2 -error on the order of 10^{-2} when $d = 6$. However, we observe that the error increases with increasing dimension with the error for $d = 10$ above that for $d = 6$, while for $d = 20$ the PINNs do not converge at all within the budget of epochs. Increasing the number of samples to 10000 (result in the top right of Fig. 4), the PINNs can be seen to be converging to the exact solution, saturating at an identical error around 10^{-2} in dimensions $d = 6, 10$, and 20 after 30000 epochs. This highlights the dependence on the problem dimension in the number of samples required to achieve a given accuracy. This can also be observed for Examples 2 and 3, where we see that with fixed $m = 3000$ samples the error increases as the dimension increases.

As discussed in the introduction, we seek methods to solve high-dimensional PDEs which overcome or substantially mitigate the curse of dimensionality. The experiments in Fig. 5 aim to illuminate the relationship between the dimension of the problem and the number of samples required to achieve a given error. There we plot the relative L^2 -error vs. the number of sample points in $d = 6, 10$, and 20 dimensions for each of our examples. We observe in all three cases that the number of samples required to achieve a given error increases with the dimension. However, focusing on the point at which the PINNs begin to saturate in their relative L^2 -error, this scaling does not appear to be exponential in the dimension which would certainly be the case if the method was not capable of mitigating the curse. Rather, the scaling in the required number of samples appears to be approximately linear. Comparing the results for all three examples when $d = 6, 10$ and 20, we observe that the number of samples required for the PINNs to reach saturation at relative L^2 -error around 10^{-2} roughly doubles moving from 6 to 10 dimensions and again moving from 10 to 20 dimensions. This empirical observation is in accordance with Theorem 1, which establishes that m scales at worst linearly in d (see (20)). It is however worth stressing that PINNs are

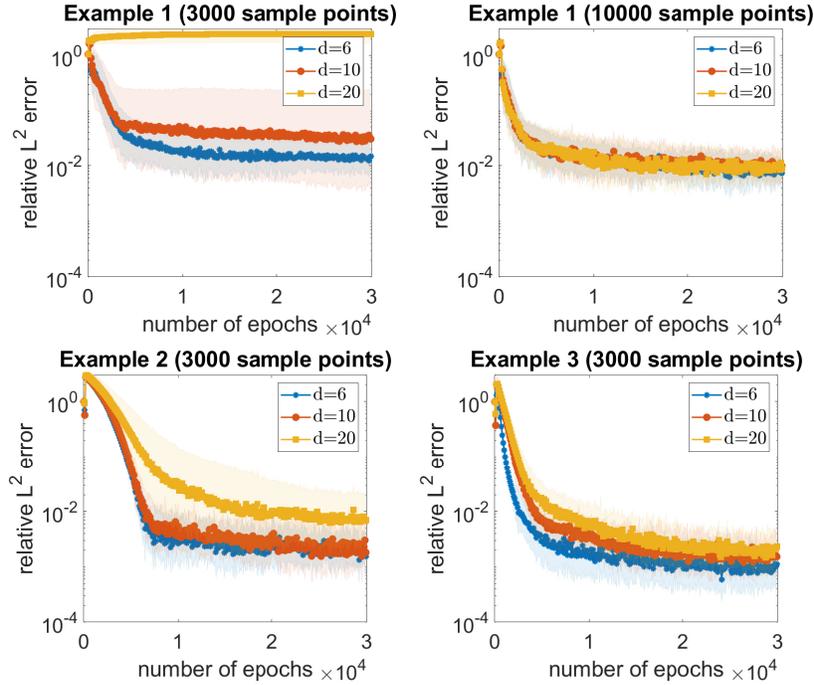


Figure 4: (Impact of the dimension) Relative L^2 -error versus the number of epochs for approximating the exact solutions defined in (27)-(29) with (top left, bottom left, bottom right) $m = 3000$ and (top right) $m = 10000$ samples, in $d = 6, 10, 20$ dimensions.

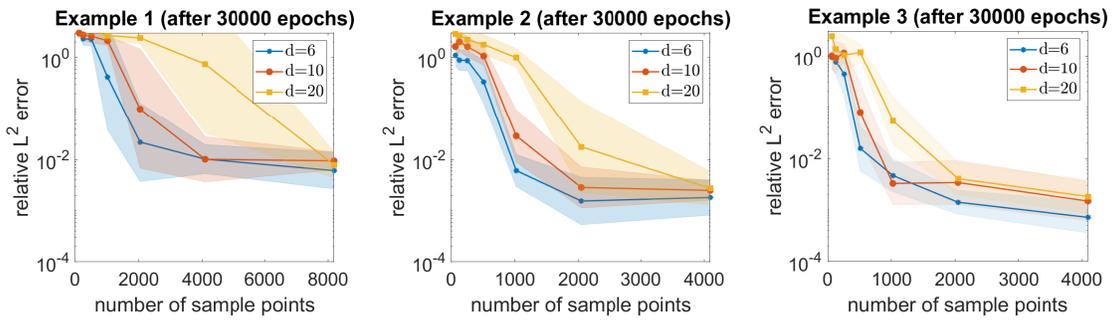


Figure 5: (Dimensionality and number of samples) Relative L^2 -error after 30000 epochs versus number of sample points m for the exact solutions defined in Eq. (27)-(29), where $d = 6, 10, 20$ are the dimension of the problem. Left: Example 1. Middle: Example 2. Right: Example 3.

not able to reach very high accuracy (an L^2 -error of around 10^{-2} might not be considered fully satisfactory by numerical PDE experts). This fact could depend on the optimizer or

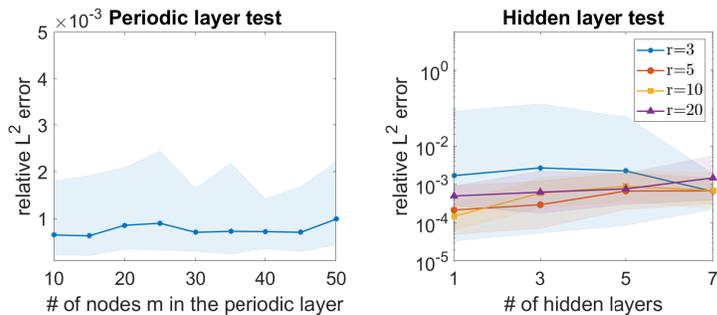


Figure 6: (Impact of the architecture) Left: Relative L^2 -error versus number of nodes in the periodic layer. Right: Relative L^2 -error versus the number of hidden layers. Different curves represent the different width-depth ratios of the hidden layers. For both figures, we test the neural networks on Example 3, in $d = 6$ dimensions using 5000 sample points, and the relative L^2 -error are measured after 30000 epochs.

the choice of architecture (one could consider more sophisticated options, such as networks where hidden layers do not have constant width, or networks with residual layers).

4.2.3 IMPACT OF THE ARCHITECTURE

We also examine the impact of the neural network’s architecture on performance. Specifically, we consider varying the number of nodes in the periodic layer and the width-depth ratio $r = w/h$ of the hidden layers. In Fig. 3 and 4, we see that the neural network approach produces accurate convergence results in all three examples. We then naturally raise a question: can a better choice of the hyper-parameters l , h , and $r = w/h$ further improve the accuracy of the neural network approach? To answer this question, we design experiments by changing these hyper-parameters and testing the performance in Example 3. We change the number of nodes in the periodic layer of the neural network in the left plot of Fig. 6. The results show that increasing the number of nodes beyond 10 in the first layer does not improve the L^2 -error. The results for tuning the number of layers and the number of nodes per layer for non-periodic layers are illustrated in the right plot of Fig. 6. We observe that adding more hidden layers to the fully-connected part of the network does not improve the accuracy of the periodic-NN solution. However, we do observe larger standard deviation for the choice of $r = 3$ for the width-depth ratio in comparison to the choices of $r = 5, 10, 20$. We also observe diminishing returns in further increasing the ratio to $r = 20$ as the cost of training increases due to the larger width of the network. Since we observe that the relative L^2 -error is not sensitive to these hyper-parameters, we choose hyper-parameters as stated at the beginning of §4.2. Analogous results hold for higher dimensions (we have run experiments in dimensions $d = 10$ and 20), but they are omitted for the sake of brevity.

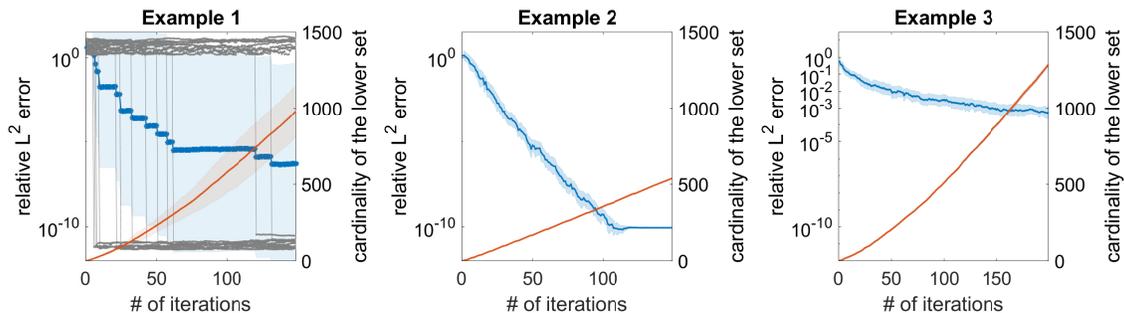


Figure 7: (Number of iterations) Relative L^2 -error (blue curve, left y-label) and cardinality (red curve, right y-label) of the lower set versus the number of iterations for the exact solutions defined in Eq. (27)-(29). Other parameters include dimension $d = 6$ and sample points $m = 3000$.

4.3 Adaptive Lower OMP Results

In this subsection, we represent numerical results for the adaptive lower OMP method (Algorithm 1), including the convergence results and a comparison with the non-adaptive Fourier compressive collocation method described in Wang and Brugiapaglia (2024).

4.3.1 NUMBER OF ITERATIONS

The lower OMP method is an adaptive method. On each iteration, the method adaptively changes the size of the lower index set and outputs an approximation to the PDE solution. As in the previous sections, we use the relative L^2 -error to measure accuracy. Here the size of the index set directly relates to the computational cost of the lower OMP method. Hence, we study the relative L^2 -error and the cardinality of the lower index set over the number of iterations to measure the performance of the method. From the numerical experiments in Fig. 7, we observe that for the different examples considered, the convergence rate and the increments in the size of the index set vary substantially.

In the case of Example 1, once the algorithm discovers the index for the complex Fourier basis function $\sin 4\pi x_1 \sin 2\pi x_2$, then the L^2 -error decays to less than 10^{-10} , otherwise, the approach does not converge. In this example, the visualization using sample geometric mean can not fully describe the convergence behavior, so we also plot each run's relative L^2 -error using black data points. As the number of iterations increases, we observe the algorithm is more likely to discover the correct Fourier basis function (i.e., more runs reach relative L^2 -error below 10^{-10}). However, even after 150 iterations, not all trials have converged to the exact solution as we see that some data points are still near the top of the figure. This suboptimal behaviour is not surprising since the multi-index set corresponding to the basis function $\sin 4\pi x_1 \sin 2\pi x_2$, i.e., $\{\pm 2\mathbf{e}_1 \pm \mathbf{e}_2\}$, is not lower. Hence, Example 1 is not ideally suited for the adaptive lower OMP setting.

The situation is different for the other two examples. Example 2 has only two activated variables, and the coefficients satisfy a lower structure. The adaptive lower OMP method successfully finds the appropriate two-dimensional lower set and converges to the exact

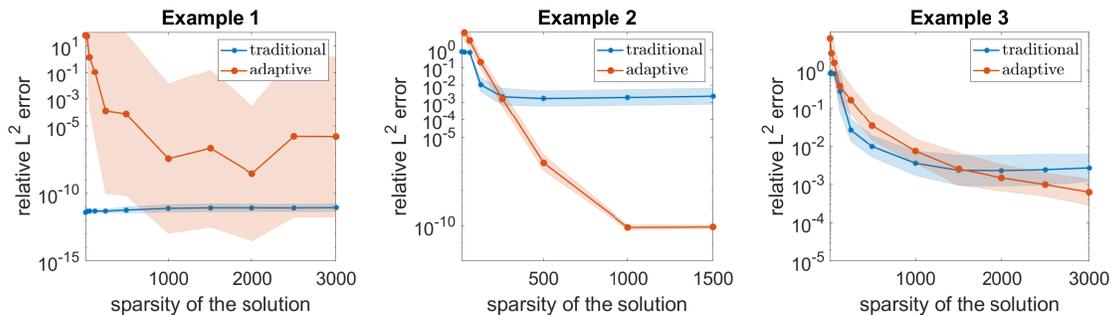


Figure 8: (Comparison OMP) Relative L^2 -error versus the sparsity of the OMP solution for the exact solutions defined in Eq. (28)-(29), Here dimension $d = 6$, and we choose the hyperbolic cross multi-index set $\Lambda_{d,n}^{\text{HC}}$ with $n = 18$ (cardinality = 3418) as ambient space for the traditional OMP.

solution after 100 iterations for all trials. The cardinality of the lower set increases at a slower rate in this example because the lower set only expands in the direction of x_1 and x_2 . For Example 3, the L^2 -error converges to approximately 10^{-3} within 200 iterations. In this example we also observe that the index set's cardinality significantly increases because the index set is being adaptively and anisotropically extended in all active dimensions.

4.3.2 COMPARISON WITH TRADITIONAL OMP METHOD

We now compare the performance for both the lower OMP and traditional OMP methods in moderate dimension $d = 6$. We note that the traditional OMP method is not applicable in higher dimensional problems due to the requirement of using an enormous index set. On the other hand, the lower OMP method can produce results in problems as high as $d = 30$ dimensions because it adaptively searches for the best index set (see §4.4 for more lower OMP results with $d = 30$). Fig.8 compares the L^2 -error of the traditional and the lower OMP methods with the same solution sparsity. Traditional OMP has a better performance when the Fourier expansion of the exact solution is fully captured by the traditional OMP ambient space (see, e.g., the left plot for Example 1). Otherwise, the lower OMP method adaptively extends the index set using the existing residual data and achieves better accuracy after more iterations (see, e.g., the middle and right plots for Examples 2 and 3).

4.4 Periodic PINNs vs. CFC

In this section we compare the performance of periodic PINNs and CFC with adaptive lower OMP recovery on high-dimensional test problems. To make a fair comparison, we use identical sets of sample points as the training data for both methods. For the neural network, we use 30000 epochs to examine the performance of the PINNs, which we empirically observe allows for the networks to saturate on the training data. In the adaptive lower OMP algorithm, the stopping criterion is set to be the size of the index set (which we recall is also an indicator of the computational cost). We choose this size to be greater than half of

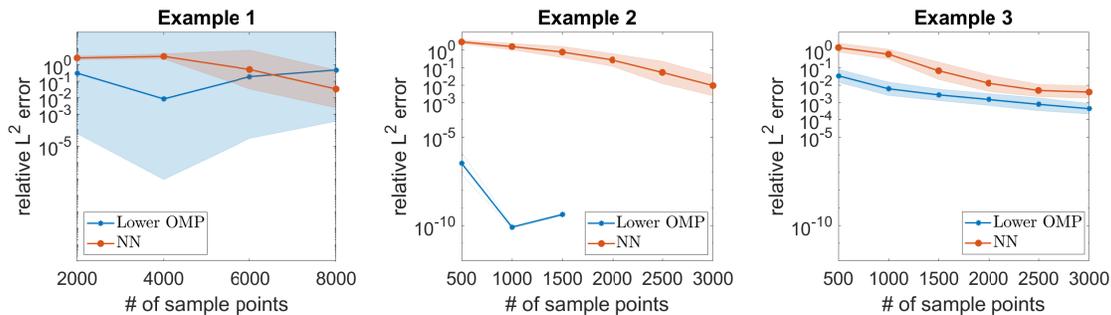


Figure 9: (Comparison) Relative L^2 -error versus the number of sample points m for the exact solutions defined in Eq. (28)-(29), Here dimension $d = 30$, number of sample points $m = 3000$. The neural networks have the same structure as in §4.2 and the relative error is measured after 30000 epochs. The cardinality of the index sets in the Lower adaptive OMP is constrained to be less than or equal to $m/2$

the number of sample points. As illustrated in Fig. 9, the adaptive lower OMP method has advantages in discovering the underlying anisotropy of the high-dimensional functions, see, e.g., the results in Examples 2 and 3 where the function is inactive in most variables or when the coefficients satisfy a lower structure. The adaptive lower OMP method captures the most important terms in Example 2 as the number of samples reaches 1000. In all three cases, the neural network accurately approximates the solution as the number of sample points increases and reaches error approximately 10^{-2} (Example 1) and 10^{-3} (Examples 2 and 3). In Example 3, the error of the periodic PINN reaches a plateau after approximately 2500 samples (see also §4.2). However, similar to Example 1 in Fig. 8, the lower OMP method struggles to converge for Example 1, while the periodic neural network provides slightly better average relative L^2 -error at 8000 sample points with tighter spread in the standard deviation over the trials. In summary, CFC can achieve much higher accuracy than periodic PINNs (gaining from 1 to 8 orders of magnitude), but its performance depends on the sparsity properties of u . On the other hand, periodic PINNs are able to achieve a consistent accuracy level (10^{-2} relative L^2 -error) on all examples. Note that this phenomenon is not explained by Theorem 1. While a possible cause for this behavior is the use of stochastic optimizers for PINNs' training, a rigorous understanding of it remains an interesting open problem.

5. Proof of the Practical Existence Theorem

In order to prove Theorem 1, we first extend the CFC Theorem of Wang and Brugiapaglia (2024) to handle periodic diffusion equations with a reaction term in §5.1. This requires first outlining the theory of bounded Riesz systems and forming a connection with the CFC matrix and its corresponding Gram matrix. After illustrating this CFC convergence result, we will show a proof of Theorem 1 in §5.2.

5.1 Compressive Fourier Collocation for Diffusion-Reaction Problems

In this subsection we extend the CFC convergence analysis of Wang and Brugiapaglia (2024) from diffusion to diffusion-reaction problems of the form (1). The general proof strategy is to present sufficient conditions for the PDE coefficients a and ρ such that the CFC matrix (25) is a random sampling matrix from a bounded Riesz system (Brugiapaglia et al., 2021a). This step is the most technical aspect of the proof. The rest uses the existing framework of sparse recovery in bounded Riesz systems whereby a sufficient condition on the sampling complexity is chosen such that the CFC satisfies the robust null space property with high-probability. From compressive sensing theory, we then obtain recovery guarantees for the CFC approximation.

5.1.1 SOME PRELIMINARY FACTS

We start by recalling the definition of the bounded Riesz system. We restrict our attention to Riesz systems in $L^2(\mathbb{T}^d)$, although the definition can be extended to general Hilbert spaces. Note in the following definition that $\ell^2(\Lambda; \mathbb{C})$ is the set of ℓ^2 -integrable complex sequences, indexed by Λ , i.e., $\ell^2(\Lambda; \mathbb{C}) = \{z = (z_\nu)_{\nu \in \Lambda} : \sum_{\nu \in \Lambda} |z_\nu|^2 < \infty\}$.

Definition 4 (Bounded Riesz System) *Let $\Lambda \subseteq \mathbb{Z}^d$, $0 < b_\Phi \leq B_\Phi < \infty$, and let $\ell^2(\Lambda; \mathbb{C})$ denote the space of sequences $z = (z_\nu)_{\nu \in \Lambda}$ with $\|z\|_2 < \infty$. A set of functions $\{\Phi_\nu\}_{\nu \in \Lambda} \subset L^2(\mathbb{T}^d)$ is a Riesz system with constants b_Φ and B_Φ if*

$$b_\Phi \|z\|_2^2 \leq \left\| \sum_{\nu \in \Lambda} z_\nu \Phi_\nu \right\|_{L^2}^2 \leq B_\Phi \|z\|_2^2, \quad \forall z = (z_\nu)_{\nu \in \Lambda} \in \ell^2(\Lambda; \mathbb{C}).$$

The constants b_Φ and B_Φ are called lower and upper Riesz constants, respectively. Moreover, the system $\{\Phi_\nu\}_{\nu \in \Lambda}$ is bounded if there exists a constant $0 < K_\Phi < \infty$ such that

$$\|\Phi_\nu\|_{L^\infty} \leq K_\Phi, \quad \forall \nu \in \Lambda.$$

Note that any L^2 -orthonormal system is a Riesz system with $b_\Phi = B_\Phi = 1$. In particular, the Fourier system $\{F_\nu\}_{\nu \in \mathbb{Z}^d}$ is a bounded Riesz system with $b_\Phi = B_\Phi = K_\Phi = 1$

Recalling the definition (3) of \mathcal{L} , we define

$$\Phi_\nu = \mathcal{L}[\Psi_\nu], \quad \forall \nu \in \Lambda, \quad (30)$$

where $\{\Psi_\nu\}_{\nu \in \mathbb{Z}^d}$ is the renormalized Fourier system given in (14). To show that the system $\{\Phi_\nu\}_{\nu \in \Lambda}$ defined in (30) is a Riesz system, it is convenient to consider its *Gram matrix* $G \in \mathbb{C}^{N \times N}$, where $N = |\Lambda|$, defined by

$$G_{\nu\mu} = \langle \Phi_\nu, \Phi_\mu \rangle, \quad \forall \nu, \mu \in \Lambda. \quad (31)$$

We note in passing that, thanks to the normalization factor $1/\sqrt{m}$ in (25), we have $\mathbb{E}[A^*A] = G$ (this follows from a direct computation and the fact that the random collocation points $\mathbf{y}_1, \dots, \mathbf{y}_m$ are independently and uniformly distributed over \mathbb{T}^d). The significance of the Gram matrix G relies on the fact that it yields the following norm equivalence:

$$\left\| \sum_{\nu \in \Lambda} c_\nu \Phi_\nu \right\|_{L^2}^2 = \left\langle \sum_{\nu \in \Lambda} c_\nu \Phi_\nu, \sum_{\nu \in \Lambda} c_\nu \Phi_\nu \right\rangle = \mathbf{c}^T G \mathbf{c}, \quad \forall \mathbf{c} \in \mathbb{C}^N. \quad (32)$$

Note that G is a Hermitian positive semidefinite matrix. Hence, it has only real nonnegative eigenvalues. The Courant–Fischer–Weyl min-max principle implies that, if $0 < b_\Phi \leq B_\Phi < \infty$ are such that

$$b_\Phi \leq \lambda_{\min}(G) \leq \lambda_{\max}(G) \leq B_\Phi, \quad (33)$$

then $\{\Phi_\nu\}_{\nu \in \Lambda}$ is a Riesz system with constants b_Φ and B_Φ . Hence, estimating the lower and upper Riesz constants of $\{\Phi_\nu\}_{\nu \in \Lambda}$ corresponds to finding two-sided spectral bounds for the Gram matrix G . To obtain this type of spectral bounds, we will employ Gershgorin’s circle theorem, see, e.g., (Horn and Johnson, 2012, Theorem 6.1.1), of which Lemma 6 is a direct consequence. Here we state a useful corollary of the Gershgorin circle theorem.

Lemma 5 (Gershgorin’s circle theorem for Hermitian matrices) *Let $A \in \mathbb{C}^{N \times N}$ be a Hermitian matrix. Then, all eigenvalues of A lie in the real interval*

$$\left[\min_{i \in [N]} \left\{ A_{ii} - \sum_{j \neq i} |A_{ij}| \right\}, \max_{i \in [N]} \left\{ A_{ii} + \sum_{j \neq i} |A_{ij}| \right\} \right].$$

Using the Fourier expansion (16) of a , it is possible to compute an explicit formula for the entries of the Gram matrix G .

Lemma 6 (Explicit formula for the Gram matrix) *Let $\Lambda \subseteq \mathbb{Z}^d$, $r_\nu \in \mathbb{C}$, with $\nu \in \Lambda$, be generic rescaling constants for the Fourier system such that $\Psi_\nu = r_\nu F_\nu$, and consider a diffusion coefficient $a \in C^1(\mathbb{T}^d)$ having Fourier expansion (16) with $T = \mathbb{Z}^d \setminus \{\mathbf{0}\}$. Then, the elements of G admit the following explicit formula in terms of the Fourier coefficients $(a_\tau)_{\tau \in \mathbb{Z}^d}$ of a and the reaction term $\rho \in \mathbb{R}$:*

$$\begin{aligned} G_{\nu\mu} = & r_\nu \bar{r}_\mu \left(16\pi^4 \sum_{\tau \in \mathbb{Z}^d} \sum_{\tau' \in \mathbb{Z}^d} (\tau \cdot \nu + \|\nu\|_2^2)(\tau' \cdot \mu + \|\mu\|_2^2) a_\tau \bar{a}_{\tau'} \delta_{\tau+\nu, \tau'+\mu} + \rho^2 \delta_{\nu, \mu} \right. \\ & \left. + 4\pi^2 \rho \sum_{\tau \in \mathbb{Z}^d} (\tau \cdot \nu + \|\nu\|_2^2) a_\tau \delta_{\tau+\nu, \mu} + 4\pi^2 \rho \sum_{\tau' \in \mathbb{Z}^d} (\tau' \cdot \mu + \|\mu\|_2^2) \bar{a}_{\tau'} \delta_{\nu, \tau'+\mu} \right), \end{aligned}$$

for each ν, μ in $\Lambda \setminus \{\mathbf{0}\}$ and where $\delta_{\nu, \mu}$ denotes the Kronecker delta. Moreover,

$$G_{\mathbf{0}\mathbf{0}} = r_{\mathbf{0}}^2 \rho^2 \quad \text{and} \quad G_{\nu\mathbf{0}} = G_{\mathbf{0}\nu} = 0, \quad \forall \nu \in \Lambda \setminus \{\mathbf{0}\}. \quad (34)$$

Proof Before proving the identity, we note that gradients and Laplacians of the Fourier basis functions defined in (11) can be easily computed as

$$\nabla F_\nu = (2\pi i \nu) F_\nu \quad \text{and} \quad \Delta F_\nu = -4\pi^2 \|\nu\|_2^2 F_\nu, \quad \forall \nu \in \mathbb{Z}^d. \quad (35)$$

Moreover,

$$F_\nu F_\mu = F_{\nu+\mu}, \quad \forall \nu, \mu \in \mathbb{Z}^d. \quad (36)$$

To prove the desired formula for $G_{\nu, \mu}$, we expand the inner product in (31). Using the above properties, the L^2 -orthonormality of the Fourier basis $\{F_\nu\}_{\nu \in \mathbb{Z}^d}$, and recalling the

expansion (16) of a , if ρ is a constant, we see that

$$\begin{aligned}
 G_{\nu\mu} &= \langle \nabla a \cdot \nabla \Psi_\nu + a \Delta \Psi_\nu - \rho \Psi_\nu, \nabla a \cdot \nabla \Psi_\mu + a \Delta \Psi_\mu - \rho \Psi_\mu \rangle \\
 &= \langle \nabla a \cdot \nabla \Psi_\nu + a \Delta \Psi_\nu, \nabla a \cdot \nabla \Psi_\mu + a \Delta \Psi_\mu \rangle + \rho^2 \langle \Psi_\nu, \Psi_\mu \rangle \\
 &\quad - \rho \langle \nabla a \cdot \nabla \Psi_\nu + a \Delta \Psi_\nu, \Psi_\mu \rangle - \rho \langle \Psi_\nu, \nabla a \cdot \nabla \Psi_\mu + a \Delta \Psi_\mu \rangle \\
 &= \sum_{\tau} \sum_{\tau'} \langle ((2i\pi\tau) \cdot (2i\pi\nu) - 4\pi^2 \|\nu\|_2^2) a_\tau F_\tau \Psi_\nu, ((2i\pi\tau') \cdot (2i\pi\mu) - 4\pi^2 \|\mu\|_2^2) a_{\tau'} F_{\tau'} \Psi_\mu \rangle \\
 &\quad + \rho^2 \langle \Psi_\nu, \Psi_\mu \rangle - \rho \sum_{\tau} \langle ((2i\pi\tau) \cdot (2i\pi\nu) - 4\pi^2 \|\nu\|_2^2) a_\tau F_\tau \Psi_\nu, \Psi_\mu \rangle \\
 &\quad - \rho \sum_{\tau'} \langle \Psi_\nu, ((2i\pi\tau') \cdot (2i\pi\mu) - 4\pi^2 \|\mu\|_2^2) a_{\tau'} F_{\tau'} \Psi_\mu \rangle \\
 &= 16\pi^4 r_\nu \bar{r}_\mu \sum_{\tau} \sum_{\tau'} (\tau \cdot \nu + \|\nu\|_2^2) (\tau' \cdot \mu + \|\mu\|_2^2) a_\tau \bar{a}_{\tau'} \langle F_\tau F_\nu, F_{\tau'} F_\mu \rangle + \rho^2 r_\nu \bar{r}_\mu \langle F_\nu, F_\mu \rangle \\
 &\quad + 4\pi^2 \rho r_\nu \bar{r}_\mu \sum_{\tau} (\tau \cdot \nu + \|\nu\|_2^2) a_\tau \langle F_\tau F_\nu, F_\mu \rangle + 4\pi^2 \rho r_\nu \bar{r}_\mu \sum_{\tau'} (\tau' \cdot \mu + \|\mu\|_2^2) \bar{a}_{\tau'} \langle F_\nu, F_{\tau'} F_\mu \rangle \\
 &= r_\nu \bar{r}_\mu \left(16\pi^4 \sum_{\tau} \sum_{\tau'} (\tau \cdot \nu + \|\nu\|_2^2) (\tau' \cdot \mu + \|\mu\|_2^2) a_\tau \bar{a}_{\tau'} \delta_{\tau+\nu, \tau'+\mu} + \rho^2 \delta_{\nu, \mu} \right. \\
 &\quad \left. + 4\pi^2 \rho \sum_{\tau} (\tau \cdot \nu + \|\nu\|_2^2) a_\tau \delta_{\tau+\nu, \mu} + 4\pi^2 \rho \sum_{\tau'} (\tau' \cdot \mu + \|\mu\|_2^2) \bar{a}_{\tau'} \delta_{\nu, \tau'+\mu} \right),
 \end{aligned}$$

where all the summations are over $\tau, \tau' \in \mathbb{Z}^d$, and both ν, μ not equal to $\mathbf{0}$. Moreover, for any $\nu \neq \mathbf{0}$, we have

$$\begin{aligned}
 G_{\nu\mathbf{0}} &= \langle \nabla a \cdot \nabla \Psi_\nu + a \Delta \Psi_\nu - \rho \Psi_\nu, -\rho \Psi_{\mathbf{0}} \rangle \\
 &= \sum_{\tau} \langle ((2i\pi\tau) \cdot (2i\pi\nu) - 4\pi^2 \|\nu\|_2^2) a_\tau F_\tau \Psi_\nu, -\rho \Psi_{\mathbf{0}} \rangle \\
 &= 4\pi^2 \rho \sum_{\tau} (\tau \cdot \nu + \|\nu\|_2^2) a_\tau \rho \langle F_\tau F_\nu, F_{\mathbf{0}} \rangle \\
 &= 4\pi^2 \rho \sum_{\tau} (\tau \cdot \nu + \|\nu\|_2^2) a_\tau \rho \delta_{\tau+\nu, \mathbf{0}} = 0.
 \end{aligned}$$

Similarly, one can show that $G_{\mathbf{0}\mu} = 0$, and $G_{\mathbf{0}\mathbf{0}} = \rho^2 \langle r_{\mathbf{0}} F_{\mathbf{0}}, r_{\mathbf{0}} F_{\mathbf{0}} \rangle = r_{\mathbf{0}}^2 \rho^2$. \blacksquare

We also need an auxiliary result about norm equivalencies in $H^2(\mathbb{T}^d)$ that will be necessary to derive the final error bound of Theorem 8. Letting $\|u\|^2 := \|u\|_{L^2}^2 + \|\Delta u\|_{L^2}^2$, the following lemma shows that this norm is equivalent to the canonical H^2 -norm.

Lemma 7 *The norms $\|\cdot\|$ and $\|\cdot\|_{H^2}$ are equivalent. Specifically, $\sqrt{2/3}\|u\|_{H^2} \leq \|u\| \leq \|u\|_{H^2}$ for every $u \in H^2(\mathbb{T}^d)$.*

Proof Let $u(\mathbf{x}) = \sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} c_{\boldsymbol{\nu}} \exp(2\pi i \boldsymbol{\nu} \cdot \mathbf{x})$ and recall that $\|u\|_{H^2}^2 = \|u\|_{L^2}^2 + \|\nabla u\|_{L^2}^2 + \|\nabla^2 u\|_{L^2}^2$. Using Parseval's identity, we have

$$\begin{aligned} \|\nabla u\|_{L^2}^2 &= (2\pi)^2 \int_{\mathbb{T}^d} \sum_{k=1}^d \left| \sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} c_{\boldsymbol{\nu}} \nu_k \exp(2\pi i \boldsymbol{\nu} \cdot \mathbf{x}) \right|^2 d\mathbf{x} = (2\pi)^2 \sum_{k=1}^d \sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} |c_{\boldsymbol{\nu}}|^2 |\nu_k|^2 \\ &= (2\pi)^2 \sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} |c_{\boldsymbol{\nu}}|^2 \|\boldsymbol{\nu}\|_2^2. \end{aligned}$$

Moreover, recalling that $\|\nabla^2 u\|_{L^2}^2 = \int_{\mathbb{T}^d} \|\nabla^2 u(\mathbf{x})\|_F^2 d\mathbf{x}$ and using Parseval's identity again, we obtain

$$\begin{aligned} \|\nabla^2 u\|_{L^2}^2 &= (2\pi)^4 \int_{\mathbb{T}^d} \sum_{k=1}^d \sum_{l=1}^d \left| \sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} c_{\boldsymbol{\nu}} \nu_k \nu_l \exp(2\pi i \boldsymbol{\nu} \cdot \mathbf{x}) \right|^2 d\mathbf{x} \\ &= (2\pi)^4 \sum_{k=1}^d \sum_{l=1}^d \sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} |c_{\boldsymbol{\nu}}|^2 (\nu_k \nu_l)^2 = (2\pi)^4 \sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} |c_{\boldsymbol{\nu}}|^2 \|\boldsymbol{\nu}\|_2^4. \end{aligned}$$

Now, observe that $\|u\|_{L^2} = \|\mathbf{c}\|_2$. Thus,

$$\|u\|_{H^2}^2 = \sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} |c_{\boldsymbol{\nu}}|^2 (1 + (2\pi)^2 \|\boldsymbol{\nu}\|_2^2 + (2\pi)^4 \|\boldsymbol{\nu}\|_2^4) = \mathbf{c}^* D_1 \mathbf{c},$$

with $D_1 = \text{diag}((1 + (2\pi)^2 \|\boldsymbol{\nu}\|_2^2 + (2\pi)^4 \|\boldsymbol{\nu}\|_2^4)_{\boldsymbol{\nu} \in \mathbb{Z}^d})$. Moreover,

$$\|\Delta u\|_{L^2}^2 = (2\pi)^4 \int_{\mathbb{T}^d} \left| \sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} c_{\boldsymbol{\nu}} \|\boldsymbol{\nu}\|_2^2 \exp(2\pi i \boldsymbol{\nu} \cdot \mathbf{x}) \right|^2 d\mathbf{x} = (2\pi)^4 \sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} |c_{\boldsymbol{\nu}}|^2 \|\boldsymbol{\nu}\|_2^4.$$

Therefore,

$$\|u\|^2 = \sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} |c_{\boldsymbol{\nu}}|^2 (1 + (2\pi)^4 \|\boldsymbol{\nu}\|_2^4) = \mathbf{c}^* D_2 \mathbf{c},$$

where $D_2 = \text{diag}((1 + (2\pi)^4 \|\boldsymbol{\nu}\|_2^4)_{\boldsymbol{\nu} \in \mathbb{Z}^d})$. Hence,

$$\frac{\|u\|_{H^2}^2}{\|u\|^2} = \frac{\mathbf{c}^* D_1 \mathbf{c}}{\mathbf{c}^* D_2 \mathbf{c}} = \frac{\mathbf{d}^* D_3 \mathbf{d}}{\mathbf{d}^* \mathbf{d}} = \frac{\sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} (D_3)_{\boldsymbol{\nu}, \boldsymbol{\nu}} |d_{\boldsymbol{\nu}}|^2}{\sum_{\boldsymbol{\nu} \in \mathbb{Z}^d} |d_{\boldsymbol{\nu}}|^2} \leq \sup_{\boldsymbol{\nu} \in \mathbb{Z}^d} ((D_3)_{\boldsymbol{\nu}, \boldsymbol{\nu}}),$$

where we made the change of variable $\mathbf{d} = D_2^{\frac{1}{2}} \mathbf{c}$ and where

$$D_3 = D_2^{-\frac{1}{2}} D_1 D_2^{-\frac{1}{2}} = \text{diag} \left(\left(\frac{1 + (2\pi)^2 \|\boldsymbol{\nu}\|_2^2 + (2\pi)^4 \|\boldsymbol{\nu}\|_2^4}{1 + (2\pi)^4 \|\boldsymbol{\nu}\|_2^4} \right)_{\boldsymbol{\nu} \in \mathbb{Z}^d} \right).$$

In particular,

$$\sup_{\boldsymbol{\nu} \in \mathbb{Z}^d} ((D_3)_{\boldsymbol{\nu}, \boldsymbol{\nu}}) = \sup_{\boldsymbol{\nu} \in \mathbb{Z}^d} \frac{1 + (2\pi)^2 \|\boldsymbol{\nu}\|_2^2 + (2\pi)^4 \|\boldsymbol{\nu}\|_2^4}{1 + (2\pi)^4 \|\boldsymbol{\nu}\|_2^4} \leq \frac{3}{2}$$

The last bound is obtained by maximizing $x \mapsto (1 + x^2 + x^4)/(1 + x^4)$ over \mathbb{R} . Similarly, we obtain the second inequality in the statement of equivalence from

$$\frac{\|u\|^2}{\|u\|_{H^2}^2} \leq \sup_{\nu \in \mathbb{Z}^d} \frac{1 + (2\pi)^4 \|\nu\|_2^4}{1 + (2\pi)^2 \|\nu\|_2^2 + (2\pi)^4 \|\nu\|_2^4} \leq 1.$$

This concludes the proof. ■

5.1.2 CFC CONVERGENCE

We are now ready to illustrate our CFC convergence result. Its proof is based on the same tools employed in Wang and Brugiapaglia (2024) and it is therefore an extension of the analysis presented therein. Keeping the CFC setup of §3 in mind, in order to recover a compressible solution \hat{c} to the linear system (24), we utilize the *Square-Root LASSO* (in short, SR-LASSO), see (Adcock et al., 2019; Belloni et al., 2011) and references therein, a modified version of the original LASSO (Least Absolute Shrinkage and Selection Operator), which lacks a power of 2 on the data fidelity term. The SR-LASSO is defined as

$$\min_{z \in \mathbb{C}^N} \{\|Az - \mathbf{b}\|_2 + \lambda \|z\|_1\}. \quad (37)$$

The main benefit of this formulation is that the optimal choice of the tuning parameter $\lambda > 0$ is independent of the noise level, see (Adcock et al., 2022b, §6.6.2).

Theorem 8 (Convergence of CFC for diffusion-reaction problems) *Given a dimension $d \in \mathbb{N}$, target sparsity $s \in \mathbb{N}$, hyperbolic cross order $n \in \mathbb{N}$ and probability of failure $\varepsilon \in (0, 1)$, let $\Lambda = \Lambda_{d,n}^{\text{HC}} \subset \mathbb{Z}^d$ and suppose $a \in C^1(\mathbb{T}^d)$ and $\rho \in \mathbb{R}$ satisfy (2) and (16)–(17). Then, the system $\{\Phi_\nu\}_{\nu \in \Lambda}$ defined in (30) is a bounded Riesz system in the sense of Definition 4 with constants*

$$b_\Phi = a_0^2 - \left(2a_0 + \frac{\rho}{2\pi^2}\right) \beta - \beta^2 > 0, \quad (38)$$

$$B_\Phi = \|a\|_{H^1}^2 + \frac{\rho^2}{16\pi^4} + \frac{a_0\rho}{2\pi^2} + \left(2a_0 + \frac{\rho}{2\pi^2}\right) \beta + \beta^2, \quad (39)$$

where $\beta = \sqrt{|T|} \|a - a_0\|_{H^1}$, and $K_\Phi = a_0 + \beta + \frac{\rho}{4\pi^2}$. Moreover, the CFC solution \hat{u} in (26) approximating the high-dimensional periodic diffusion-reaction equation (1), whose coefficients \hat{c} are computed by solving the SR-LASSO problem (37) with A and \mathbf{b} defined as in (25) and with tuning parameter $d_{a,\rho}^{(1)} \sqrt{B_\Phi/s} < \lambda \leq d_{a,\rho}^{(2)} \sqrt{B_\Phi/s}$, satisfies the following with probability at least $1 - \varepsilon$: if

$$m \geq c_{a,\rho}^{(3)} \cdot s \cdot \log^2 \left(c_{a,\rho}^{(4)} \cdot s \right) \cdot \left(\min\{\log(n) + d, \log(2n) \log(2d)\} + \log(\varepsilon^{-1}) \right), \quad (40)$$

then

$$\|u - \hat{u}\|_{L^2} + \|(\Delta - \rho)(u - \hat{u})\|_{L^2} \leq C_{a,\rho}^{(1)} \cdot \frac{\sigma_s(\mathbf{c}_\Lambda)_1}{\sqrt{s}} + C_{a,d,\rho}^{(2)} \cdot \left(\frac{\|u - u_\Lambda\|_{W^{2,\infty}}}{\sqrt{s}} + \|u - u_\Lambda\|_{H^2} \right). \quad (41)$$

Moreover, if $\rho < 1$, we also have

$$\|u - \hat{u}\|_{H^2} \leq C_{a,\rho}^{(3)} \cdot \frac{\sigma_s(\mathbf{c}_\Lambda)_1}{\sqrt{s}} + C_{a,d,\rho}^{(4)} \cdot \left(\frac{\|u - u_\Lambda\|_{W^{2,\infty}}}{\sqrt{s}} + \|u - u_\Lambda\|_{H^2} \right). \quad (42)$$

Here each constant depends only on the subscripted parameters. Moreover, the dependence of each constant on d is at most linear (when present).

Remark 9 The above theorem holds for OMP recovery as well, under a sufficient condition on the ratio b_Φ/B_Φ , see (Wang and Brugiapaglia, 2024, Theorem 3).

Remark 10 There is a gap between Theorem 8 and the corresponding result in the diffusion setting (Wang and Brugiapaglia, 2024, Theorem 3.5). Namely, in the latter the diffusion coefficient a can be nonsparse. Extending Theorem 8 to allow for more general diffusion coefficients is an open problem.

Proof Given sparse diffusion and reaction terms a and ρ as in (2) and (16)–(17), we determine that $\{\Phi_\nu\}_{\nu \in \Lambda}$ is a bounded Riesz system using the explicit form of the Gram matrix in Lemma 6.

Step 1: Riesz property. We find lower and upper Riesz constants b_Φ and B_Φ by establishing a two-sided spectral bound for the Gram matrix G , recall equation (31). First, observing that the first row (and column) of G only has one nonzero entry (namely, $G_{\mathbf{0}\mathbf{0}}$), one eigenvalue of G is $G_{\mathbf{0}\mathbf{0}} = a_0^2$. To estimate the remaining eigenvalues, we use a special version of Gershgorin's circle theorem for Hermitian matrices, presented in Lemma 5, see, e.g., (Horn and Johnson, 2012, Theorem 6.1.1) for the general statement. Using Lemma 6 with the generic normalization, the diagonal entries $G_{\nu\nu}$ with $\nu \in \Lambda \setminus \{\mathbf{0}\}$ are given by

$$\begin{aligned} G_{\nu\nu} &= 16\pi^4 |r_\nu|^2 \sum_{\tau \in T \cup \{\mathbf{0}\}} (\tau \cdot \nu + \|\nu\|_2^2)^2 |a_\tau|^2 + \rho^2 |r_\nu|^2 + 8\pi^2 |r_\nu|^2 \|\nu\|_2^2 \rho a_0 \\ &= 16\pi^4 \|\nu\|_2^4 |r_\nu|^2 \left[\left(a_0 + \frac{\rho}{4\pi^2 \|\nu\|_2^2} \right)^2 + \sum_{\tau \in T} \left(\frac{\tau \cdot \nu + \|\nu\|_2^2}{\|\nu\|_2^2} \right)^2 |a_\tau|^2 \right]. \end{aligned}$$

Using the Cauchy-Schwarz inequality and inputting the actual normalization used in (14), we see that, for all $\nu \in \Lambda \setminus \{\mathbf{0}\}$,

$$|r_\nu|^2 \left(a_0 + \frac{\rho}{4\pi^2 \|\nu\|_2^2} \right)^2 \leq \frac{|G_{\nu\nu}|}{16\pi^4 \|\nu\|_2^4} \leq |r_\nu|^2 \left[a_0^2 + \frac{\rho^2}{16\pi^4} + \frac{|a_0|\rho}{2\pi^2} + \sum_{\tau \in T} (\|\tau\|_2 + 1)^2 |a_\tau|^2 \right]$$

which, in turn, implies

$$a_0^2 \leq |G_{\nu\nu}| \leq \|a\|_{H^1}^2 + \frac{\rho^2}{16\pi^4} + \frac{|a_0|\rho}{2\pi^2}. \quad (43)$$

The inequality defining the upper bound can be proved as follows. Using the definition of H^1 -norm, the differentiation properties (35), and the fact that the Fourier system $\{F_\nu\}_{\nu \in \mathbb{Z}^d}$

is L^2 -orthonormal, we obtain

$$\begin{aligned}
 \|a\|_{H^1}^2 &= \|a\|_{L^2}^2 + \sum_{l=1}^d \left\| \frac{\partial a}{\partial x_l} \right\|_{L^2}^2 = \sum_{\tau \in T \cup \{\mathbf{0}\}} |a_\tau|^2 + \sum_{l=1}^d \sum_{\tau \in T \cup \{\mathbf{0}\}} |2\pi a_\tau \tau_l|^2 \\
 &= |a_{\mathbf{0}}|^2 + \sum_{\tau \in T} \left(1 + (2\pi)^2 \|\tau\|_2^2 \right) |a_\tau|^2 \geq |a_{\mathbf{0}}|^2 + \sum_{\tau \in T} (\|\tau\|_2 + 1)^2 |a_\tau|^2, \quad (44)
 \end{aligned}$$

which proves (43). To apply Gershgorin's circle theorem, we now bound the sum of all off-diagonal entries in the ν -th row of G . Using Lemma 6 again, the definition of the Kronecker delta, (34) and the fact that $|r_\nu| \leq 1/(4\pi^2 \|\nu\|_2^2) \leq 1/4\pi^2$ for every $\nu \neq \mathbf{0}$, we obtain

$$\begin{aligned}
 \sum_{\mu \in \Lambda \setminus \{\nu\}} |G_{\nu\mu}| &= \sum_{\mu \in \Lambda \setminus \{\nu\}} \left| r_\nu \bar{r}_\mu \left(16\pi^4 \sum_{\tau \in \mathbb{Z}^d} (\tau \cdot \nu + \|\nu\|_2^2) ((\tau + \nu - \mu) \cdot \mu + \|\mu\|_2^2) a_\tau \bar{a}_{\tau + \nu - \mu} \right. \right. \\
 &\quad \left. \left. + 4\pi^2 \rho ((\mu - \nu) \cdot \nu + \|\nu\|_2^2) a_{\mu - \nu} + 4\pi^2 \rho ((\nu - \mu) \cdot \mu + \|\mu\|_2^2) \bar{a}_{\nu - \mu} \right) \right| \\
 &\leq \sum_{\tau \in \mathbb{Z}^d} \left(\frac{|\tau \cdot \nu|}{\|\nu\|_2^2} + 1 \right) |a_\tau| \sum_{\mu \in \Lambda \setminus \{\nu\}} \left(\frac{|(\tau + \nu - \mu) \cdot \mu|}{\|\mu\|_2^2} + 1 \right) |a_{\tau + \nu - \mu}| \\
 &\quad + \frac{\rho}{4\pi^2} \sum_{\mu \in \Lambda \setminus \{\nu\}} \left(\frac{|(\mu - \nu) \cdot \nu|}{\|\nu\|_2^2} + 1 \right) |a_{\mu - \nu}| \\
 &\quad + \frac{\rho}{4\pi^2} \sum_{\mu \in \Lambda \setminus \{\nu\}} \left(\frac{|(\nu - \mu) \cdot \mu|}{\|\mu\|_2^2} + 1 \right) |\bar{a}_{\nu - \mu}|.
 \end{aligned}$$

Using the Cauchy-Schwarz inequality $|\tau \cdot \nu| \leq \|\tau\|_2 \cdot \|\nu\|_2$ and the fact that $\|\mu\|_2 \geq 1$ and $\|\nu\|_2 \geq 1$ for all $\mu, \nu \neq \mathbf{0}$, we see that, for all $\nu \in \Lambda \setminus \{\mathbf{0}\}$,

$$\begin{aligned}
 \sum_{\mu \in \Lambda \setminus \{\nu\}} |G_{\nu\mu}| &\leq \sum_{\tau \in T \cup \{\mathbf{0}\}} (\|\tau\|_2 + 1) |a_\tau| \sum_{\mu \in \Lambda \setminus \{\nu\}} (\|\tau + \nu - \mu\|_2 + 1) |a_{\tau + \nu - \mu}| \\
 &\quad + \frac{\rho}{4\pi^2} \sum_{\tau \in T} (\|\tau\|_2 + 1) |a_\tau| + \frac{\rho}{4\pi^2} \sum_{\tau \in T} (\|\tau\|_2 + 1) |\bar{a}_\tau|.
 \end{aligned}$$

Substituting $\boldsymbol{\tau}' = \boldsymbol{\tau} + \boldsymbol{\nu} - \boldsymbol{\mu}$ (which implies $\boldsymbol{\tau}' \neq \boldsymbol{\tau}$), recalling that $(a_{\boldsymbol{\nu}})_{\boldsymbol{\nu} \in \mathbb{Z}^d}$ is supported on $T \cup \{\mathbf{0}\}$, and separating the $a_{\mathbf{0}}$ term, we obtain

$$\begin{aligned}
 \sum_{\boldsymbol{\mu} \in \Lambda \setminus \{\boldsymbol{\nu}\}} |G_{\boldsymbol{\nu}\boldsymbol{\mu}}| &\leq \sum_{\boldsymbol{\tau} \in T \cup \{\mathbf{0}\}} (\|\boldsymbol{\tau}\|_2 + 1) |a_{\boldsymbol{\tau}}| \sum_{\boldsymbol{\tau}' \in T \cup \{\mathbf{0}\} \setminus \{\boldsymbol{\tau}\}} (\|\boldsymbol{\tau}'\|_2 + 1) |a_{\boldsymbol{\tau}'}| + \frac{\rho}{2\pi^2} \sum_{\boldsymbol{\tau} \in T} (\|\boldsymbol{\tau}\|_2 + 1) |a_{\boldsymbol{\tau}}| \\
 &= |a_{\mathbf{0}}| \sum_{\boldsymbol{\tau}' \in T} (\|\boldsymbol{\tau}'\|_2 + 1) |a_{\boldsymbol{\tau}'}| + \sum_{\boldsymbol{\tau} \in T} (\|\boldsymbol{\tau}\|_2 + 1) |a_{\boldsymbol{\tau}}| \sum_{\boldsymbol{\tau}' \in T \cup \{\mathbf{0}\} \setminus \{\boldsymbol{\tau}\}} (\|\boldsymbol{\tau}'\|_2 + 1) |a_{\boldsymbol{\tau}'}| \\
 &\quad + \frac{\rho}{2\pi^2} \sum_{\boldsymbol{\tau} \in T} (\|\boldsymbol{\tau}\|_2 + 1) |a_{\boldsymbol{\tau}}| \\
 &\leq 2|a_{\mathbf{0}}| \sum_{\boldsymbol{\tau}' \in T} (\|\boldsymbol{\tau}'\|_2 + 1) |a_{\boldsymbol{\tau}'}| + \sum_{\boldsymbol{\tau} \in T} (\|\boldsymbol{\tau}\|_2 + 1) |a_{\boldsymbol{\tau}}| \sum_{\boldsymbol{\tau}' \in T} (\|\boldsymbol{\tau}'\|_2 + 1) |a_{\boldsymbol{\tau}'}| \\
 &\quad + \frac{\rho}{2\pi^2} \sum_{\boldsymbol{\tau} \in T} (\|\boldsymbol{\tau}\|_2 + 1) |a_{\boldsymbol{\tau}}| \\
 &= \left(2a_{\mathbf{0}} + \frac{\rho}{2\pi^2}\right) \sum_{\boldsymbol{\tau}' \in T} (\|\boldsymbol{\tau}'\|_2 + 1) |a_{\boldsymbol{\tau}'}| + \left(\sum_{\boldsymbol{\tau} \in T} (\|\boldsymbol{\tau}\|_2 + 1) |a_{\boldsymbol{\tau}}|\right)^2.
 \end{aligned}$$

Applying the Cauchy–Schwarz inequality, denoting $t = |T|$, and separating the $a_{\mathbf{0}}$ term, we obtain the bound

$$\sum_{\boldsymbol{\tau} \in T} (\|\boldsymbol{\tau}\|_2 + 1) |a_{\boldsymbol{\tau}}| \leq \sqrt{t} \left(\sum_{\boldsymbol{\tau} \in T} (\|\boldsymbol{\tau}\|_2 + 1)^2 |a_{\boldsymbol{\tau}}|^2 \right)^{1/2} \leq \sqrt{t} \left(\|a\|_{H^1}^2 - a_{\mathbf{0}}^2 \right)^{1/2} := \beta. \quad (45)$$

Combining the above inequalities yields

$$\sum_{\boldsymbol{\mu} \in \Lambda \setminus \{\boldsymbol{\nu}\}} |G_{\boldsymbol{\nu}\boldsymbol{\mu}}| \leq \left(2a_{\mathbf{0}} + \frac{\rho}{2\pi^2}\right) \beta + \beta^2. \quad (46)$$

Finally, applying the Gershgorin circle theorem on G combining (43) and (46) and recalling that ρ^2 is an eigenvalue of G , we obtain the Riesz constants

$$b_{\Phi} = a_{\mathbf{0}}^2 - \left(2a_{\mathbf{0}} + \frac{\rho}{2\pi^2}\right) \beta - \beta^2 > 0, \quad (47)$$

$$B_{\Phi} = \|a\|_{H^1}^2 + \frac{\rho^2}{16\pi^4} + \frac{a_{\mathbf{0}}\rho}{2\pi^2} + \left(2a_{\mathbf{0}} + \frac{\rho}{2\pi^2}\right) \beta + \beta^2. \quad (48)$$

We note that the positivity of b_{Φ} is required for equation (40) to be well-defined and this is assured by the sufficient condition (17).

Step 2: Boundedness. To bound the essential supremum of the Riesz system, we use properties (35) and (36)

$$\begin{aligned}
\|\Phi_{\nu}\|_{L^\infty} &= \|\mathcal{L}[\Psi_{\nu}]\|_{L^\infty} = \|\nabla \cdot (a \nabla \Psi_{\nu}) + \rho \Psi_{\nu}\|_{L^\infty} \\
&\leq \|\nabla a \cdot \nabla \Psi_{\nu}\|_{L^\infty} + \|a \Delta \Psi_{\nu}\|_{L^\infty} + \|\rho \Psi_{\nu}\|_{L^\infty} \\
&\leq \left\| \sum_{\tau \in T \cup \{0\}} 2\pi i \tau a_{\tau} F_{\tau} \right\|_{L^\infty} \cdot \left\| \frac{2\pi i \nu}{4\pi^2 \|\nu\|_2^2} F_{\nu} \right\|_{L^\infty} + \left\| \sum_{\tau \in T \cup \{0\}} a_{\tau} F_{\tau} \right\|_{L^\infty} + \frac{\rho}{4\pi^2 \|\nu\|_2^2} \\
&\leq \sum_{\tau \in T \cup \{0\}} \frac{|\tau \cdot \nu|}{\|\nu\|_2^2} |a_{\tau}| + \sum_{\tau \in T \cup \{0\}} |a_{\tau}| + \frac{\rho}{4\pi^2 \|\nu\|_2^2} \\
&= \sum_{\tau \in T \cup \{0\}} \left(\frac{|\tau \cdot \nu|}{\|\nu\|_2^2} + 1 \right) |a_{\tau}| + \frac{\rho}{4\pi^2 \|\nu\|_2^2}
\end{aligned}$$

Finally, applying the Cauchy inequality, equation (44) and the fact that $\|\nu\|_2 \geq 1$, we obtain

$$\begin{aligned}
\|\Phi_{\nu}\|_{L^\infty} &\leq \sum_{\tau \in T \cup \{0\}} \left(\frac{|\tau \cdot \nu|}{\|\nu\|_2^2} + 1 \right) |a_{\tau}| + \frac{\rho}{4\pi^2 \|\nu\|_2^2} \\
&\leq a_0 + \sum_{\tau \in T} (\|\tau\| + 1) |a_{\tau}| + \frac{\rho}{4\pi^2} \leq a_0 + \beta + \frac{\rho}{4\pi^2}. \tag{49}
\end{aligned}$$

This proves that $\{\Phi_{\nu}\}_{\nu \in \Lambda}$ is a bounded Riesz system (Definition 4). Now, it remains to show accurate and stable recovery guarantees for the problem (37). The machinery for this is enabled by recent advances in sparse recovery for bounded Riesz systems (Brugiapaglia et al., 2021a).

Step 3: Bounded Riesz property \implies error bound. The strategy is to pick a lower bound on the sample complexity, such that A satisfies the robust null-space property with high probability, and then appeal to the recovery bounds already supplied by the literature surrounding SR-LASSO. We begin with a simple re-normalization of the compressive Fourier collocation matrix, by letting $\tilde{A} = A/\sqrt{B_{\Phi}}$. We thus consider a rescaled version of the SR-LASSO problem

$$\hat{c} \in \arg \min_{z \in \mathbb{C}^N} \|Az - \mathbf{b}\|_2 + \lambda \|z\|_1, \tag{50}$$

defined as

$$\hat{c} \in \arg \min_{\tilde{z} \in \mathbb{C}^N} \|\tilde{A}\tilde{z} - \mathbf{b}\|_2 + \tilde{\lambda} \|\tilde{z}\|_1, \tag{51}$$

where $\tilde{z} = \sqrt{B_{\Phi}}z$ and $\tilde{\lambda} = \lambda/\sqrt{B_{\Phi}}$. The two minimizers are such that $\sqrt{B_{\Phi}}\hat{c} = \hat{\tilde{c}}$. An inspection of the proof of (Brugiapaglia et al., 2021a, Theorem 2.6) reveals that condition

$$m \geq c_0 \left(\frac{\max\{1, B_{\Phi}\}}{b_{\Phi}} \right)^2 K_{\Phi}^2 s \log^2 \left(s K_{\Phi}^2 \frac{\max\{1, B_{\Phi}\}}{b_{\Phi}} \right) \log(eN), \tag{52}$$

where $c_0 > 0$ is a universal constant, is sufficient for \tilde{A} to satisfy the rNSP with constants $\rho = 1/2$ and $\gamma = 2B_\Phi/b_\Phi$ and with probability at least $1 - \varepsilon/2$.¹

Note that condition (52) is implied by (40) with

$$c_{a,\rho}^{(3)} = c_0 \left(\frac{\max\{1, B_\Phi\}}{b_\Phi} \right)^2 K_\Phi^2 \quad \text{and} \quad c_{a,\rho}^{(4)} = K_\Phi^2 \frac{\max\{1, B_\Phi\}}{b_\Phi},$$

and thanks to the upper bound (13) on the cardinality of the hyperbolic cross $N = |\Lambda|$.

Hence, we can apply (Adcock et al., 2022b, Theorem 6.29) with $\mathbf{w} = 1$ and $\mathbf{b} = \tilde{A}\tilde{\mathbf{c}}_\Lambda + \mathbf{e} = A\mathbf{c}_\Lambda + \mathbf{e}$, where

$$\mathbf{e} = \frac{1}{\sqrt{m}} (\mathcal{L}[u - u_\Lambda](\mathbf{x}_i))_{i \in [m]} \in \mathbb{C}^m,$$

yielding the following coefficient recovery guarantee:

$$\|\hat{\mathbf{c}} - \tilde{\mathbf{c}}_\Lambda\|_2 \leq c_1 \frac{\sigma_s(\tilde{\mathbf{c}}_\Lambda)_1}{\sqrt{s}} + \frac{1}{2} \left(\frac{c_1}{\sqrt{s\lambda}} + c_2 \right) \|\mathbf{e}\|_2, \quad (53)$$

for constants $c_1 = 2(1 + \rho)^2/(1 - \rho) = 9$ and $c_2 = 2(3 + \rho)\gamma/(1 - \rho) = 28B_\Phi/b_\Phi$, see also Eq. (6.15) from (Adcock et al., 2022b). Note also that Eq. (6.47) (Adcock et al., 2022b) allows us to pick a specific range for λ and, consequently, for λ :

$$\sqrt{B_\Phi} \frac{d_{a,\rho}^{(1)}}{\sqrt{s}} \leq \lambda \leq \sqrt{B_\Phi} \frac{d_{a,\rho}^{(2)}}{\sqrt{s}}, \quad (54)$$

where $d_{a,\rho}^{(1)}$ and $d_{a,\rho}^{(2)}$ are such that $0 < d_{a,\rho}^{(1)} \leq d_{a,\rho}^{(2)} \leq \frac{1+\rho}{(3+\rho)\gamma} = \frac{3b_\Phi}{14B_\Phi}$. Finally, changing back to the original variables, and inserting values for constants, we obtain

$$\begin{aligned} \|\hat{\mathbf{c}} - \mathbf{c}_\Lambda\|_2 &\leq \frac{c_1 \sigma_s(\mathbf{c}_\Lambda)_1}{\sqrt{s}} + \frac{1}{2\sqrt{B_\Phi}} \left(\frac{c_1 \sqrt{B_\Phi}}{\sqrt{s\lambda}} + c_2 \right) \|\mathbf{e}\|_2 \\ &= \frac{9\sigma_s(\mathbf{c}_\Lambda)_1}{\sqrt{s}} + \frac{1}{2} \left(\frac{9}{\sqrt{s\lambda}} + \frac{28\sqrt{B_\Phi}}{b_\Phi} \right) \|\mathbf{e}\|_2 \\ &\leq \frac{9\sigma_s(\mathbf{c}_\Lambda)_1}{\sqrt{s}} + \frac{1}{2} \left(\frac{9}{\sqrt{B_\Phi}d_1} + \frac{28\sqrt{B_\Phi}}{b_\Phi} \right) \|\mathbf{e}\|_2. \end{aligned} \quad (55)$$

Thanks to the rescaling used in (14), we have $(-\Delta + \rho)\Psi_\nu = F_\nu$ for all $\nu \in \mathbb{Z}^d$. Hence, for any $v \in H^2(\mathbb{T}^d)$ such that $v = \sum_{\nu \in \mathbb{Z}^d} d_\nu \Psi_\nu$, using Parseval's identity, we obtain

$$\|(\Delta - \rho)v\|_{L^2} = \|\mathbf{d}\|_2 \quad \text{and} \quad \|v\|_{L^2} = \|\mathbf{Rd}\|_2 \leq (a_0/\rho)\|\mathbf{d}\|_2, \quad (56)$$

where $R = \text{diag}((r_\nu)_{\nu \in \mathbb{Z}^d})$ with $r_\nu = 1/(4\pi^2\|\nu\|_2^2 + \rho/a_0)$ being the rescaling factors, and where we used the fact that $r_\nu \leq a_0/\rho$ for every $\nu \in \mathbb{Z}^d$. Using the triangle inequality, we see that

$$\begin{aligned} \|(\Delta - \rho)(u - \hat{u})\|_{L^2} &\leq \|(\Delta - \rho)(u - u_\Lambda)\|_{L^2} + \|(\Delta - \rho)(u_\Lambda - \hat{u})\|_{L^2} \\ \|u - \hat{u}\|_{L^2} &\leq \|u - u_\Lambda\|_{L^2} + \|u_\Lambda - \hat{u}\|_{L^2}. \end{aligned}$$

1. Observe that the probability of failure (Brugiapaglia et al., 2021a, Theorem 2.6) is $4 \exp -c'_0(b_\Phi/\max\{1, B_\Phi\})^2 m/(sK_\Phi^2)$, for some constant $c'_0 > 0$, and which is bounded by ε thanks to (40).

These, combined with (55) and (56) (with $v = \hat{u} - u_\Lambda$ and $\mathbf{d} = \hat{\mathbf{c}} - \mathbf{c}_\Lambda$), imply the following error bounds:

$$\begin{aligned} \|(\Delta - \rho)(u - \hat{u})\|_{L^2} &\leq \|(\Delta - \rho)(u - u_\Lambda)\|_{L^2} + \frac{9}{\sqrt{s}}\sigma_s(\mathbf{c}_\Lambda)_1 + \frac{1}{2} \left(\frac{9}{\sqrt{B_\Phi}d_1} + \frac{28\sqrt{B_\Phi}}{b_\Phi} \right) \|\mathbf{e}\|_2 \\ \|u - \hat{u}\|_{L^2} &\leq \|u - u_\Lambda\|_{L^2} + \frac{a_0}{\rho} \left(\frac{9}{\sqrt{s}}\sigma_s(\mathbf{c}_\Lambda)_1 + \frac{1}{2} \left(\frac{9}{\sqrt{B_\Phi}d_1} + \frac{28\sqrt{B_\Phi}}{b_\Phi} \right) \|\mathbf{e}\|_2 \right). \end{aligned}$$

Summing these two inequalities and observing that, thanks to Lemma 7, for any $v \in H^2(\mathbb{T}^d)$,

$$\|(\Delta - \rho)v\|_{L^2} + \|v\|_{L^2} \leq \|\Delta u\|_{L^2} + (1 + \rho)\|v\|_{L^2} \leq \sqrt{2}(1 + \rho)\|v\| \leq \sqrt{2}(1 + \rho)\|v\|_{H^2},$$

we see that

$$\begin{aligned} \|u - \hat{u}\|_{L^2} + \|(\Delta - \rho)(u - \hat{u})\|_{L^2} &\leq \\ &\sqrt{2}(1 + \rho)\|u - u_\Lambda\|_{H^2} + \left(1 + \frac{a_0}{\rho}\right) \left(\frac{9}{\sqrt{s}}\sigma_s(\mathbf{c}_\Lambda)_1 + \frac{1}{2} \left(\frac{9}{\sqrt{B_\Phi}d_1} + \frac{28\sqrt{B_\Phi}}{b_\Phi} \right) \|\mathbf{e}\|_2 \right). \end{aligned} \quad (57)$$

Step 4: Bounding the truncation error $\|\mathbf{e}\|_2$ and further simplifications. Following the discussion of (Wang and Brugiapaglia, 2024, page 9) and the notation of (Adcock et al., 2022b, §7.6.1), we define

$$\|\mathbf{e}\|_2 = \sqrt{\frac{1}{m} \sum_{i=1}^m |\mathcal{L}[u - u_\Lambda](y_i)|^2} =: E_{\Lambda, \text{disc}}(\mathcal{L}[u]).$$

Now, (40) implies $m \geq 2s \log(4/\varepsilon)$, giving the following bound for $E_{\Lambda, \text{disc}}(\mathcal{L}[u])$ with probability $1 - \varepsilon/2$:

$$E_{\Lambda, \text{disc}}(\mathcal{L}[u]) \leq \sqrt{2} \left(\frac{E_{\Lambda, \infty}(\mathcal{L}[u])}{\sqrt{s}} + E_{\Lambda, 2}(\mathcal{L}[u]) \right), \quad (58)$$

thanks to (Adcock et al., 2022b, Lemma 7.11), where $E_{\Lambda, \infty}(\mathcal{L}[u]) = \|\mathcal{L}[u - u_\Lambda]\|_{L^\infty}$ and $E_{\Lambda, 2}(\mathcal{L}[u]) = \|\mathcal{L}[u - u_\Lambda]\|_{L^2}$. Hence, (40) is sufficient so that both the rNSP and (58) hold with probability $1 - \varepsilon$ and we obtain

$$\|\mathbf{e}\|_2 \leq \sqrt{2} \left(\frac{\|\mathcal{L}[u - u_\Lambda]\|_{L^\infty}}{\sqrt{s}} + \|\mathcal{L}[u - u_\Lambda]\|_{L^2} \right), \quad (59)$$

Now, we estimate the two error terms on the right-hand side. Note that, for $1 \leq p \leq \infty$, and any $v \in H^2(\mathbb{T}^d)$,

$$\begin{aligned} \|\mathcal{L}[v]\|_{L^p} &= \|-\nabla \cdot (a\nabla v) + \rho v\|_{L^p} \\ &= \|-\nabla a \cdot \nabla v - a\Delta v + \rho v\|_{L^p} \\ &\leq \|\nabla a \cdot \nabla v\|_{L^p} + \|a\Delta v\|_{L^p} + \|\rho v\|_{L^p}. \end{aligned} \quad (60)$$

Using the inequality $\|\mathbf{v}\|_1 \leq d^{1-\frac{1}{p}}\|\mathbf{v}\|_p$, for all $\mathbf{v} \in \mathbb{R}^d$, we begin by analyzing the first two terms. First, we see that

$$\begin{aligned} \|\nabla a \cdot \nabla v\|_{L^p} &= \left\| \sum_{k=1}^d \frac{\partial a}{\partial x_k} \frac{\partial v}{\partial x_k} \right\|_{L^p} \leq \sum_{k=1}^d \left\| \frac{\partial a}{\partial x_k} \frac{\partial v}{\partial x_k} \right\|_{L^p} \leq \max_{k \in [d]} \left\| \frac{\partial a}{\partial x_k} \right\|_{L^\infty} \sum_{k=1}^d \left\| \frac{\partial v}{\partial x_k} \right\|_{L^p} \\ &= \max_{k \in [d]} \left\| \frac{\partial a}{\partial x_k} \right\|_{L^\infty} \left\| \left(\left\| \frac{\partial v}{\partial x_k} \right\|_{L^p} \right)_{k=1}^d \right\|_1 \leq \|a\|_{W^{1,\infty}} \cdot d^{1-\frac{1}{p}} \left\| \left(\left\| \frac{\partial v}{\partial x_k} \right\|_{L^p} \right)_{k=1}^d \right\|_p \\ &\leq d^{1-\frac{1}{p}} \|a\|_{W^{1,\infty}} \|v\|_{W^{2,p}}. \end{aligned}$$

Second, we have

$$\begin{aligned} \|a\Delta v\|_{L^p} &\leq \|a\|_{L^\infty} \left\| \sum_{k=1}^d \frac{\partial^2 v}{\partial x_k^2} \right\|_{L^p} \leq \|a\|_{L^\infty} \sum_{k=1}^d \left\| \frac{\partial^2 v}{\partial x_k^2} \right\|_{L^p} = \|a\|_{L^\infty} \left\| \left(\left\| \frac{\partial^2 v}{\partial x_k^2} \right\|_{L^p} \right)_{k=1}^d \right\|_1 \\ &\leq \|a\|_{L^\infty} \cdot d^{1-\frac{1}{p}} \left\| \left(\left\| \frac{\partial^2 v}{\partial x_k^2} \right\|_{L^p} \right)_{k=1}^d \right\|_p \leq d^{1-\frac{1}{p}} \|a\|_{L^\infty} \|v\|_{W^{2,p}}. \end{aligned}$$

The third term is more easily estimated as $\|\rho v\|_{L^p} \leq \|\rho\|_{L^\infty} \|v\|_{L^p} \leq \|\rho\|_{L^\infty} \|v\|_{W^{2,p}}$. Altogether, when $p = \infty$, we have

$$\|\mathcal{L}[v]\|_{L^\infty} \leq \|\nabla a \cdot \nabla v\|_{L^\infty} + \|a\Delta v\|_{L^\infty} + \|\rho v\|_{L^\infty} \leq (2d\|a\|_{W^{1,\infty}} + \|\rho\|_{L^\infty}) \|v\|_{W^{2,\infty}}, \quad (61)$$

and, for $p = 2$,

$$\|\mathcal{L}[v]\|_{L^2} \leq \|\nabla a \cdot \nabla v\|_{L^2} + \|a\Delta v\|_{L^2} + \|\rho v\|_{L^2} \leq \left(2\sqrt{d}\|a\|_{W^{1,\infty}} + \|\rho\|_{L^\infty}\right) \|v\|_{H^2}. \quad (62)$$

Letting $v = u - u_\Lambda$ and combining (58) with the above estimates yields

$$\|\mathbf{e}\|_2 \leq \sqrt{2} (\|a\|_{W^{1,\infty}} + \|\rho\|_{L^\infty}) \left(\frac{\|u - u_\Lambda\|_{W^{2,\infty}}}{\sqrt{s}} + \|u - u_\Lambda\|_{H^2} \right).$$

Plugging the above bound into (57) leads to the recovery guarantee (41), where

$$\begin{aligned} C_{a,\rho}^{(1)} &= 9 \left(1 + \frac{a_0}{\rho} \right) \\ C_{a,d,\rho}^{(2)} &= \frac{1}{2} \left(1 + \frac{a_0}{\rho} \right) \left(\frac{9}{\sqrt{B_\Phi} d_1} + \frac{28\sqrt{B_\Phi}}{b_\Phi} \right) \sqrt{2} (\|a\|_{W^{1,\infty}} + \|\rho\|_{L^\infty} + 1) + \sqrt{2}(1 + \rho). \end{aligned}$$

When $\rho < 1$, the bound (42) can be obtained by invoking Lemma 7 again and observing that, for every $v \in H^2(\mathbb{T}^d)$,

$$\|(\Delta - \rho)v\|_{L^2} + \|v\|_{L^2} \geq \|\Delta v\|_{L^2} + (1 - \rho)\|v\|_{L^2} \geq (1 - \rho)\|v\| \geq (1 - \rho)\sqrt{\frac{2}{3}}\|v\|_{H^2}.$$

Combining this inequality with (57), we obtain (42), where

$$C_{a,\rho}^{(3)} = \frac{C_{a,\rho}^{(1)}}{\sqrt{2/3}(1 - \rho)} \quad \text{and} \quad C_{a,d,\rho}^{(4)} = \frac{C_{a,d,\rho}^{(2)}}{\sqrt{2/3}(1 - \rho)}.$$

This concludes the proof. ■

5.2 Proof of Theorem 1

We proceed with a proof of the main result. For a summary of the general proof strategy, we refer to the proof sketch after the statement of Theorem 1. Most of the technical efforts will be devoted to constructing the class of periodic PINNs, \mathcal{N} and bounding the depth and width of its neural networks. This will allow us to leverage the CFC convergence theory (Theorem 8) to derive the desired error bounds (22)–(23).

Construction of the network class \mathcal{N} . As explained in the proof sketch, we will construct networks $\psi \in \mathcal{N}$ able to produce linear combinations of the form (15). Letting $N = |\Lambda|$, this class is of the form

$$\mathcal{N} = \left\{ \psi : \mathbb{R}^d \rightarrow \mathbb{C} : \psi(\mathbf{x}) = \mathbf{z}^\top \psi_\Lambda(\mathbf{x}), \mathbf{z} \in \mathbb{C}^N \right\}, \quad (63)$$

where $\psi_\Lambda(\mathbf{x}) = (\psi_\nu(\mathbf{x}))_{\nu \in \Lambda}$ and ψ_ν is a network replicating the ν th (rescaled) Fourier function, i.e., $\psi_\nu \equiv \Psi_\nu$, with Ψ_ν defined as in (14). The main task is therefore to explicitly construct the networks ψ_ν . We will do so by composing different network modules, aimed at producing more and more complex functions of the input variables $\mathbf{x} = (x_j)_{j \in [d]}$, as follows:

1. First, using a periodic layer $\mathbf{v}^{(2)} \circ \mathbf{q}^{(1)}$ with $\mathbf{v}^{(2)}(\mathbf{x}) = \mathbf{x}$ (see §2.1), we will generate univariate trigonometric functions $\sin(2\pi x_j)$ and $\cos(2\pi x_j)$, for $j \in [d]$.
2. Then, through a network module $\mathbf{v}_{\text{trig} \rightarrow \text{multi-freq}}$, we will compute univariate trigonometric functions corresponding to all the frequencies we need to generate functions in the system $\{\Psi_\nu\}_{\nu \in \Lambda}$. These are given by $\cos(2\pi\nu x_j)$ and $\sin(2\pi\nu x_j)$, for $\nu \in \{-\nu_{\max}, \dots, \nu_{\max}\}$ and $j \in [d]$, where

$$\nu_{\max} := \max\{|\nu_j| : \boldsymbol{\nu} = (\nu_j)_{j=1}^d \in \Lambda\} = n - 1,$$

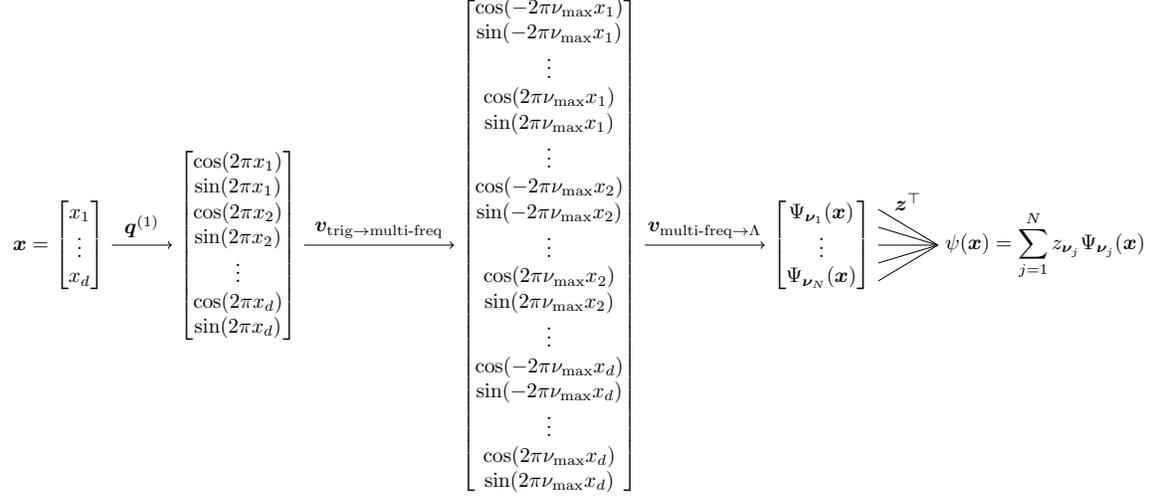
is the maximum absolute frequency of multi-indices in the hyperbolic cross $\Lambda = \Lambda_{d,n}^{\text{HC}}$.

3. Finally, using a module $\mathbf{v}_{\text{multi-freq} \rightarrow \Lambda}$, we will generate the rescaled Fourier functions $\Psi_\nu(\mathbf{x})$ through tensorization.

In summary, we have $\psi_\Lambda = \mathbf{v}_{\text{multi-freq} \rightarrow \Lambda} \circ \mathbf{v}_{\text{trig} \rightarrow \text{multi-freq}} \circ \mathbf{q}^{(1)}$. After constructing ψ_Λ , we add a last linear layer with weights $\mathbf{z}^\top \in \mathbb{R}^{1 \times N}$ to generate linear combinations of the form (15) (with $c_\nu = z_\nu$). Fig. 10 summarizes our network construction.

A key ingredient of the proof is the fact that neural networks with RePU activation can exactly replicate products of real numbers. Specifically, for any $\ell, k \in \mathbb{N}$ with $\ell \geq 2$, (Adcock et al., 2025, Lemma 7.3), which, in turn relies on techniques from (Opschoor et al., 2022; Schwab and Zech, 2019), establishes the existence (along with an explicit construction) of a feedforward neural network $P^{(k)} : \mathbb{R}^k \rightarrow \mathbb{R}$ with RePU_ℓ activation able to exactly reproduce the product of k numbers, i.e.,

$$P^{(k)}(x_1, \dots, x_k) = \prod_{j=1}^k x_j, \quad \forall \mathbf{x} = (x_j)_{j=1}^k \in \mathbb{R}^k.$$


 Figure 10: Summary of a generic network $\psi \in \mathcal{N}$.

Moreover, the width and depth of this network are such that

$$\text{width}(P^{(k)}) \leq p_\ell^{(1)} \cdot k \quad \text{and} \quad \text{depth}(P^{(k)}) \leq p^{(2)} \cdot \log_2(k), \quad (64)$$

where $p_\ell^{(1)} > 0$ depends on ℓ only and $p^{(2)} > 0$ is a universal constant. Note that we will not track the dependence of $P^{(k)}$ on ℓ since this parameter is assumed to be fixed throughout the proof. Before showing in detail how to construct the various modules composing ψ_Λ , we establish a connection between the training of periodic PINNs in \mathcal{N} and CFC approximation, which will yield the desired recovery guarantees.

CFC convergence theory \Rightarrow error bounds (22)–(23). To apply Theorem 8, we leverage two important facts: (i) the spectral basis (14) is exactly replicated by the network ψ_Λ and (ii) the weights of the last layer $\mathbf{z}^\top \in \mathbb{R}^{1 \times N}$, corresponding to the coefficients of linear combinations of the form (15), are the only trainable parameters of ψ . In this setting, training the network ψ by minimizing the regularized RMSE loss in (21) with regularization term

$$\mathcal{R}(\psi) = \|\mathbf{z}\|_1, \quad (65)$$

is equivalent to solving the SR-LASSO problem (37) with A and \mathbf{b} as in (25). Hence, we can simply apply Theorem 8 to obtain the desired conclusion.

The rest of the proof is devoted to illustrating the construction of the periodic layer $\mathbf{q}^{(1)} \circ \mathbf{v}^{(2)}$ and the network modules $\mathbf{v}_{\text{trig} \rightarrow \text{multi-freq}}$ and $\mathbf{v}_{\text{multi-freq} \rightarrow \Lambda}$ in full detail.

Construction of ψ_Λ (Step 1): the periodic layer. Recalling the notation introduced in (4), we let $l = 2$, with $\phi_{i1} = 0$ and $\phi_{i2} = -\pi/2$, for $i \in [d]$. This leads to

$$\mathbf{x} \in \mathbb{R}^d \mapsto \mathbf{q}^{(1)}(\mathbf{x}) = \begin{bmatrix} \cos(2\pi x_1) \\ \cos(2\pi x_1 - \pi/2) \\ \vdots \\ \cos(2\pi x_d) \\ \cos(2\pi x_d - \pi/2) \end{bmatrix} = \begin{bmatrix} \cos(2\pi x_1) \\ \sin(2\pi x_1) \\ \vdots \\ \cos(2\pi x_d) \\ \sin(2\pi x_d) \end{bmatrix} =: \begin{bmatrix} c_1 \\ s_1 \\ \vdots \\ c_d \\ s_d \end{bmatrix} \in \mathbb{R}^{2d},$$

where we have used the short-hand notation $c_j = \cos(2\pi x_j)$ and $s_j = \sin(2\pi x_j)$. The second layer within the periodic layer is a linear (identity) layer, i.e. $\mathbf{v}^{(2)}(\mathbf{x}) = \mathbf{x}$. Hence, in summary,

$$\text{depth}(\mathbf{v}^{(2)} \circ \mathbf{q}^{(1)}) = 3 \quad \text{and} \quad \text{width}(\mathbf{v}^{(2)} \circ \mathbf{q}^{(1)}) = 2d. \quad (66)$$

Construction of ψ_Λ (Step 2): the module $\mathbf{v}_{\text{trig} \rightarrow \text{multi-freq}}$. To generate sine and cosine functions at multiple frequencies given c_i, s_i , we first generate powers of c_i via an inner network submodule module $\mathbf{v}_{\text{trig} \rightarrow \text{pow}_i}$ and then apply multi-angle trigonometric formulas (see, e.g., Weisstein (2024)).

We start by presenting the inner module $\mathbf{v}_{\text{trig} \rightarrow \text{pow}_i}$, which constructs the powers of cosine for a fixed dimension i . This is given by

$$\begin{bmatrix} c_1 \\ s_1 \\ \vdots \\ c_d \\ s_d \end{bmatrix} \xrightarrow[\text{restriction}]{\text{linear}} \begin{bmatrix} c_i \\ s_i \end{bmatrix} \xrightarrow[\mathbf{x} \mapsto W\mathbf{x} + \mathbf{b}]{\text{affine layer}} \begin{bmatrix} s_i \\ 1 \\ c_i \\ c_i \\ \vdots \\ c_i \end{bmatrix} \xrightarrow{\text{powers}} \begin{bmatrix} s_i \\ 1 \\ c_i \\ c_i^2 \\ \vdots \\ c_i^{\nu_{\max}} \end{bmatrix} := \text{pow}_i, \quad (67)$$

where the weights and biases in $W \in \mathbb{R}^{(\nu_{\max}+2) \times 2}$ and $\mathbf{b} \in \mathbb{R}^{\nu_{\max}+2}$ and the network corresponding to $\xrightarrow{\text{powers}}$ are illustrated in Fig. 11. Note that this network contains both RePU and linear activations, in accordance with (6). Recalling (64) and observing that consecutive affine layers can be combined together without altering depth (since the composition of affine maps is an affine map), the depth and width of this module are bounded by

$$\begin{aligned} \text{width}(\mathbf{v}_{\text{trig} \rightarrow \text{pow}_i}) &\leq 3 + \sum_{k=2}^{\nu_{\max}} \text{width}(P^{(k)}) \leq C_1 \cdot \nu_{\max}^2 \cdot p_\ell^{(1)}, \\ \text{depth}(\mathbf{v}_{\text{trig} \rightarrow \text{pow}_i}) &\leq \text{depth}(P^{(\nu_{\max})}) \leq p^{(2)} \cdot \log_2(\nu_{\max}), \end{aligned}$$

where $C_1 > 0$ is a universal constant (recall that $p_\ell^{(1)}$ depends on the degree ℓ of the RePU activation, whereas $p^{(2)}$ does not).

Now, the output pow_i of $\mathbf{v}_{\text{trig} \rightarrow \text{pow}_i}(\mathbf{x})$ is transformed using multi-angle formulae (see, e.g., Weisstein (2024)) to obtain univariate sine and cosine functions at multiple frequencies.

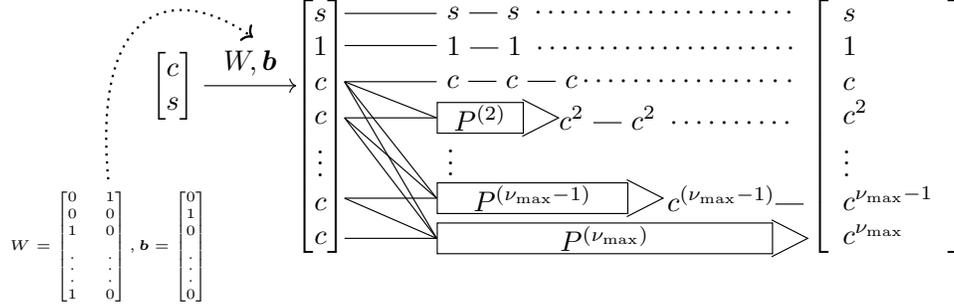


Figure 11: Part of the network submodule $\mathbf{v}_{\text{trig} \rightarrow \text{pow}_i}$ defined in (67), used to generate powers of cosine at dimension i .

This corresponds to a linear layer

$$\text{pow}_i = \begin{bmatrix} s_i \\ 1 \\ c_i \\ \vdots \\ c_i^{\nu_{\max}} \end{bmatrix} \xrightarrow[\text{formulae}]{\text{multi-angle}} \begin{bmatrix} \cos(-\nu_{\max} x_i) \\ \sin(-\nu_{\max} x_i) \\ \vdots \\ \cos(\nu_{\max} x_i) \\ \sin(\nu_{\max} x_i) \end{bmatrix}.$$

For the sake of completeness, we quickly review how to apply the multi-angle formulae in this context. These formulae allow us to generate functions of the form $\cos(2\pi\nu x)$ and $\sin(2\pi\nu x)$ from linear combinations of $\sin(2\pi x)$ and powers of $\cos(2\pi x)$. In fact, for any $\nu \in \{-\nu_{\max}, \dots, \nu_{\max}\}$ and $x \in \mathbb{R}$, we have

$$\cos(2\pi\nu x) = \sum_{i=0}^{\lfloor \nu/2 \rfloor} \sum_{j=0}^i (-1)^{i-j} \binom{\nu}{2i} \binom{i}{j} [\cos(2\pi x)]^{\nu-2(i-j)} = \sum_{k=\nu-\lfloor \nu/2 \rfloor}^{\nu} a_k^{(\nu)} \cos^k(2\pi x), \quad (68)$$

for suitable values of $a_k^{(\nu)}$ and where we have used the reindexing $k = \nu - 2(i - j)$ in the second inequality and $\lfloor \cdot \rfloor$ denotes the floor function. We can use the following analogous formula for the construction of sine functions at multiple frequencies:

$$\sin(2\nu\pi x) = \sin(2\pi x) \cdot \sum_{i=0}^{\lfloor (\nu+1)/2 \rfloor} \sum_{j=0}^i (-1)^{i-j} \binom{\nu}{2i+1} \binom{i}{j} [\cos(2\pi x)]^{\nu-2(i-j)-1}.$$

Note that the multi-angle formula for sines does not need powers of sines but only powers of cosines. The above formula can be replicated by combining a linear layer with a network $P^{(2)}$ performing the product.

Taking into account the linear layer corresponding to the multi-angle formulae and the fact that the generation of functions $\cos(2\pi\nu x_i)$ and $\sin(2\pi\nu x_i)$ must be repeated for each

dimension $i \in [d]$, we obtain

$$\begin{aligned} \text{width}(\mathbf{v}_{\text{trig} \rightarrow \text{multi-freq}}) &= d \cdot \text{width}(\mathbf{v}_{\text{trig} \rightarrow \text{pow}_i}) \leq C_1 \cdot p_\ell^{(1)} \cdot d \cdot \nu_{\max}^2, \\ \text{depth}(\mathbf{v}_{\text{trig} \rightarrow \text{multi-freq}}) &= \text{depth}(\mathbf{v}_{\text{trig} \rightarrow \text{pow}_i}) + \text{depth}(P^{(2)}) \leq C_2 \cdot p^{(2)} \cdot \log_2(\nu_{\max}) \end{aligned}$$

for a universal constant $C_2 > 0$.

Construction of ψ_Λ (Step 3): the module $\mathbf{v}_{\text{multi-freq} \rightarrow \Lambda}$. The final module's objective is to generate the rescaled Fourier functions $\Psi_\nu(\mathbf{x})$ in (14) given the output of $\mathbf{v}_{\text{trig} \rightarrow \text{multi-freq}}(\mathbf{x})$. Recall that each Fourier function is a product of complex numbers:

$$F_\nu(\mathbf{x}) = \exp(2\pi i \nu \cdot \mathbf{x}) = \prod_{j \in \text{supp}(\nu)} \exp(2\pi i \nu_j x_j) = \prod_{j \in \text{supp}(\nu)} (\cos(2\pi i \nu_j x_j) + i \sin(2\pi i \nu_j x_j)).$$

In particular, the generation of $F_\nu(\mathbf{x})$ involves the product of $\|\nu\|_0$ complex numbers whose real and imaginary parts are given by the outputs of $\mathbf{v}_{\text{trig} \rightarrow \text{multi-freq}}(\mathbf{x})$. Moreover, since $\nu \in \Lambda = \Lambda_{d,n}^{\text{HC}}$, we have

$$n \geq \prod_{j=1}^d (|\nu_j| + 1) = \prod_{j \in \text{supp}(\nu)} (|\nu_j| + 1) \geq 2^{\|\nu\|_0} \implies \|\nu\|_0 \leq \log_2 n. \quad (69)$$

Therefore, computing each $F_\nu(\mathbf{x})$ requires at most $\min\{d, \log_2 n\}$ products.

Now, we describe how to compute products of complex numbers using RePU networks. We would like to compute $\prod_{j=1}^k z_j$ for $z_j = x_j + iy_j \in \mathbb{C}$ for generic values $x_j, y_j \in \mathbb{R}$. To do so, we first find explicit formulas for the real and imaginary parts of this product:

$$\begin{aligned} \prod_{j=1}^k z_j &= \prod_{j=1}^k (x_j + iy_j) = \sum_{j \in \{0,1\}^k} \prod_{t=1}^k x_t^{j_t} (iy_t)^{1-j_t} = \sum_{j \in \{0,1\}^k} i^{k - \|j\|_1} \prod_{t=1}^k x_t^{j_t} y_t^{1-j_t} \\ &= \underbrace{\left(\sum_{\substack{j \in \{0,1\}^k \\ k - \|j\|_1 \text{ even}}} (-1)^{\frac{k - \|j\|_1}{2}} \prod_{t=1}^k x_t^{j_t} y_t^{1-j_t} \right)}_{=\text{Re}(\prod_{j=1}^k z_j)} + i \cdot \underbrace{\left(\sum_{\substack{j \in \{0,1\}^k \\ k - \|j\|_1 \text{ odd}}} (-1)^{\frac{k - \|j\|_1 - 1}{2}} \prod_{t=1}^k x_t^{j_t} y_t^{1-j_t} \right)}_{=\text{Im}(\prod_{j=1}^k z_j)}. \end{aligned}$$

This shows that $\text{Re}(\prod_{j=1}^k z_j)$ and $\text{Im}(\prod_{j=1}^k z_j)$ can be computed as linear combinations of 2^k products of k real numbers. These correspond to the terms $\prod_{t=1}^k x_t^{j_t} y_t^{1-j_t}$, where we observe that

$$x_t^{j_t} y_t^{1-j_t} = \begin{cases} x_t & \text{if } j_t = 1, \\ y_t & \text{if } j_t = 0. \end{cases}$$

Hence, we can compute products of k complex numbers z_1, \dots, z_k via a RePU network $P_{\mathbb{C}}^{(k)} : \mathbb{R}^{2k} \rightarrow \mathbb{C}$ that takes the real and imaginary parts of the complex numbers z_j as

inputs and constructed by stacking 2^k copies of $P^{(k)}$ atop, adding a linear layer realizing the linear combinations above, and a final linear layer with complex weights $[1 \ i]$:

$$\begin{bmatrix} x_1 \\ y_1 \\ \vdots \\ x_k \\ y_k \end{bmatrix} \xrightarrow[\text{of } P^{(k)}]{2^k \text{ copies}} \left[\prod_{t=1}^k x_t^{j_t} y_t^{1-j_t} \right]_{j \in \{0,1\}^k} \xrightarrow[\text{layer}]{\text{linear}} \begin{bmatrix} \text{Re} \left(\prod_{j=1}^k z_j \right) \\ \text{Im} \left(\prod_{j=1}^k z_j \right) \end{bmatrix} \xrightarrow{[1 \ i]} \text{Re} \left(\prod_{j=1}^k z_j \right) + i \text{Im} \left(\prod_{j=1}^k z_j \right).$$

Therefore, recalling (64), we obtain

$$\begin{aligned} \text{width}(P_{\mathbb{C}}^{(k)}) &= 2^k \cdot \text{width}(P^{(k)}) \leq C_3 \cdot p_\ell^{(1)} \cdot 2^k \cdot k \\ \text{depth}(P_{\mathbb{C}}^{(k)}) &= \text{depth}(P^{(k)}) + 1 \leq C_4 \cdot p^{(2)} \cdot \log_2 k, \end{aligned}$$

for some universal constants $C_3, C_4 > 0$. To conclude, the module $\mathbf{v}_{\text{multi-freq} \rightarrow \Lambda}$ is such that

$$\begin{aligned} \mathbf{v}_{\text{multi-freq} \rightarrow \Lambda} \left(\begin{bmatrix} \cos(2\pi\nu x_j) \\ \sin(2\pi\nu x_j) \end{bmatrix}_{j \in [d], \nu \in \{-\nu_{\max}, \dots, \nu_{\max}\}} \right) &= \left[r_\nu \cdot P_{\mathbb{C}}^{(\|\nu\|_0)} \left(\begin{bmatrix} \cos(2\pi\nu_j x_j) \\ \sin(2\pi\nu_j x_j) \end{bmatrix}_{j \in \text{supp}(\nu)} \right) \right]_{\nu \in \Lambda} \\ &= r_\nu F_\nu(\mathbf{x}), \end{aligned}$$

where $r_\nu = 1/(4\pi^2 \|\nu\|_2^2 + \rho/a_0)$ as desired. Note that the networks $P_{\mathbb{C}}^{(\|\nu\|_0)}$, with $\nu \in \Lambda$, can be stacked using a construction analogous to that of Fig. 11. Also observe that for $\nu = \mathbf{0}$ we simply let $P_{\mathbb{C}}^{(0)} \equiv 1$. In addition, the multiplication by the rescaling factor r_ν can be realized by modifying weights in the last layer (we don't need to add a $P^{(2)}$ module because r_ν is independent of \mathbf{x}). In summary, recalling that $N = |\Lambda|$ and using (69) and the depth and width bounds for $P_{\mathbb{C}}^{(n)}$, the architecture bounds for the whole network module are

$$\begin{aligned} \text{width}(\mathbf{v}_{\text{multi-freq} \rightarrow \Lambda}) &= \sum_{\nu \in \Lambda} \text{width}(P_{\mathbb{C}}^{(\|\nu\|_0)}) \leq N \cdot \text{width}(P_{\mathbb{C}}^{(\min\{d, \log_2 n\})}) \\ &\leq p_\ell^{(1)} \cdot C_3 \cdot N \cdot \min\{2^d, n\} \cdot \min\{d, \log_2 n\} \\ \text{depth}(\mathbf{v}_{\text{multi-freq} \rightarrow \Lambda}) &= \max_{\nu \in \Lambda} \text{depth}(P_{\mathbb{C}}^{(\|\nu\|_0)}) = \text{depth}(P_{\mathbb{C}}^{(\min\{d, 2^n\})}) \leq C_4 \cdot p^{(2)} \cdot \log_2(\min\{d, 2^n\}). \end{aligned}$$

Final architecture bounds for ψ . Combining the architecture bounds obtained for each network module in Steps 1, 2 and 3, and recalling that $\psi(\mathbf{x}) = \mathbf{z}^\top \psi_\Lambda(\mathbf{x})$, we obtain

$$\begin{aligned} \text{width}(\psi) &= \max \left\{ \text{width}(\mathbf{v}^{(2)} \circ \mathbf{q}^{(1)}), \text{width}(\mathbf{v}_{\text{trig} \rightarrow \text{multi-freq}}), \text{width}(\mathbf{v}_{\text{multi-freq} \rightarrow \Lambda}) \right\} \\ &\leq \max \left\{ 2d, C_1 \cdot p_\ell^{(1)} \cdot d \cdot \nu_{\max}^2, p_\ell^{(1)} \cdot C_3 \cdot N \cdot \min\{2^d, n\} \cdot \min\{d, \log_2 n\} \right\} \\ &\leq \min \left\{ 4n^5 16^d, e^2 n^{2+\log_2 d} \right\} \cdot c_{(\ell)}^{(1)} \cdot d. \end{aligned}$$

Using the hyperbolic cross cardinality bound (13), we see that

$$\text{width}(\psi) \leq C_5 \cdot p_\ell^{(1)} \cdot \min \left\{ 4n^5 16^d, e^2 n^{2+\log_2 d} \right\} \cdot d \cdot \min\{2^d, n\},$$

for some universal constant $C_5 > 0$. Moreover,

$$\begin{aligned} \text{depth}(\psi) &= \text{depth}(\mathbf{v}^{(2)} \circ \mathbf{q}^{(1)}) + \text{depth}(\mathbf{v}_{\text{trig} \rightarrow \text{multi-freq}}) + \text{depth}(\mathbf{v}_{\text{multi-freq} \rightarrow \Lambda}) + 1 \\ &\leq C_6 \cdot p^{(2)} \cdot (\log_2(n) + \min\{\log_2 d, n\}), \end{aligned}$$

where C_6 is a universal constant. Letting $c_\ell^{(1)} = C_5 \cdot p_\ell^{(1)}$ and $c^{(2)} = C_6 \cdot p^{(2)}$ yields (18) and (19) and concludes the proof. \blacksquare

6. Conclusions and Open Problems

We have shown a new convergence result for PINNs (Theorem 1) in the form of a practical existence theorem for the numerical solution of high-dimensional, periodic diffusion-reaction problems. This result establishes the existence of a class of periodic PINNs able to achieve the same accuracy as a sparse approximation method (namely, CFC) and using a number of training samples that scales only logarithmically or, at worst, linearly with the PDE domain's dimension d . The mild scaling of the sample complexity with respect to d is numerically confirmed through experimentation (see Fig. 5). Our practical existence theorem relies on a new CFC convergence result for diffusion-reaction problems (Theorem 8) and an explicit construction of periodic PINNs able to replicate the Fourier basis. We have also experimentally confirmed the robustness of periodic PINNs to solve high-dimensional PDEs with respect to the network hyperparameters. Finally, we compared periodic PINNs with CFC, showing that the latter could achieve much higher (but, sometimes, worse) accuracy than the former depending on the sparsity properties of the PDE solution. On the other hand, the performance of periodic PINNs is numerically observed to be consistent across the three examples considered (see Fig. 9).

We conclude by mentioning some gaps between theory and practice and open problems for future research. First, Theorem 1 relies on RePU or linear activations and not on more standard ones such as ReLU or tanh. The main reason to work with RePU activations is that they allow to exactly replicate products and, hence, using the construction illustrated in §5.2, Fourier functions. The argument of Theorem 1 could be generalized to more general activations, such as ReLU or tanh. In that case, however, products (and, hence, Fourier basis functions) could only be approximated and not exactly replicated. This issue could be handled, see (Adcock et al., 2024) and references therein, but it will introduce nontrivial technical difficulties due to the fact that one would have to deal with an approximate CFC matrix, whose error with respect to the true CFC matrix should be carefully controlled in the analysis.

Another important aspect is the presence of sufficient conditions (16) and (17) on the PDE coefficients a and ρ in Theorem 1. These are inherited by the CFC convergence analysis (Theorem 8) and are an artifact of its proof. It was shown in the diffusion equation case that these conditions are sufficient but far from being necessary, see (Wang and Brugiapaglia, 2024, §4.4). This can also be seen by observing that, for $\rho \neq 0$, problem (1) with coefficients (a, ρ, f) is equivalent to the same PDE with coefficients $(\rho^{-1}a, 1, \rho^{-1}f)$. This invariance to rescaling by ρ is not inherited by condition (17). This is likely to be the case in the diffusion-reaction case as well since the convergence analysis is based on the same argument.

The optimal network $\hat{\psi}$ of Theorem 1 is assumed to be trained by *exactly* minimizing a regularized RMSE loss. Of course, this is not what happens in practice, where the loss is

usually only approximately minimized using a stochastic gradient descent method. Taking into account the error introduced by the training algorithm in the analysis is an important open question. Moreover, only the last layer of $\hat{\psi}$ is trained, whereas the previous layers are explicitly constructed. This gap between theory and practice is an intrinsic limitation of the argument that Theorem 1 relies on. Other types of training loss could also be considered. The regularized loss (10) considered in this paper is of the square-root LASSO type, which has the advantage of having an optimal choice of tuning parameter λ independent of the noise corrupting the samples. However, finding this optimal value in practice through e.g., cross-validation could be more challenging than using a more standard LASSO-type loss (obtained by replacing the RMSE data-fidelity term with an MSE one in (10))—see Adcock et al. (2019); Berk et al. (2024). The theory could be extended to a LASSO-type loss, but the corresponding restrictions on λ would be depending on u and Λ since the noise corrupting the samples contains the truncation error (recall Step 4 in the proof of Theorem 8).

Hence, differences exist between the theoretical setting of Theorem 1 and the practical implementation of periodical PINNs in §4, for which we have followed a setup closer to what is commonly employed in the literature. These differences include the choice of activation, the network’s width and depth (which do not saturate the bounds (18)–(19) of Theorem 1), the use of complex-valued weights and the presence of a regularization term in the loss function. It would be interesting to study whether a numerical implementation that follows the setting of Theorem 1 more closely would give any practical benefit over the setup considered in §4—see Appendix A for some initial results in this direction. On this note, it is also worth observing that the optimal network $\hat{\psi}$ of Theorem 1 is sparsely connected and its last layer is approximately sparse due to the presence of ℓ^1 regularization. Understanding the potential practical benefits of sparsely-connected PINNs is also an interesting avenue of future work.

An important direction of future investigation aimed at bridging the gap with real-world applications is the study of time-dependent PDEs. In order to solve a time-dependent diffusion-reaction equation, a natural option would be to employ an implicit Euler discretization in time. This would, in turn, lead to solving a steady-state diffusion equation of the form (1) at each time step. In the case of CFC, this scheme could be studied using tools from spectral methods analysis (see, e.g., Canuto et al. (2006)). Yet, building a rigorous connection between this approach and a PINN-based approximation of a time-dependent diffusion-reaction equation where the loss involves both space and time variables (as it is typically done) does not appear to be a straightforward task and deserves further research. Moving beyond reaction-diffusion problems is also a possible direction of future work. One could consider, e.g., more general advection-diffusion-reaction operators of the form $\mathcal{L}[u](\mathbf{x}) = -\nabla \cdot (a(\mathbf{x})\nabla u(\mathbf{x})) + \mathbf{b}(\mathbf{x}) \cdot \nabla u(\mathbf{x}) + \rho(\mathbf{x})u(\mathbf{x})$, with suitable restrictions on $a(\mathbf{x})$, $\mathbf{b}(\mathbf{x})$, and $\rho(\mathbf{x})$ to ensure ellipticity. The main task required to extend the theory to this case would be the estimation of the smallest and largest eigenvalues of the Gram matrix G (see Step 1 in the proof of Theorem 8). While this can be done using the same tools employed in this paper, this is likely to lead to substantially more technically involved computations and to the introduction of extra assumptions on $a(\mathbf{x})$, $\mathbf{b}(\mathbf{x})$, and $\rho(\mathbf{x})$ for the Riesz property to hold.

An additional limitation of our setting is that it only covers periodic boundary conditions. From a practical viewpoint, PINNs can handle different types of boundary conditions

in at least two ways: (i) through the addition of a penalty term in the loss function aimed at enforcing the specific boundary conditions considered; (ii) in the case of homogeneous Dirichlet boundary conditions, by considering modified network ansatzes of the form, e.g., $\psi(\mathbf{x}) = N(\mathbf{x}) \cdot \prod_{i=1}^d x_i(1 - x_i)$, where $N(\mathbf{x})$ is a trainable network and the extra factor enforces zero boundary conditions on $[0, 1]^d$. Compressive spectral collocation methods can also be implemented for non-periodic boundary conditions—see Brugiapaglia (2020) for numerical and theoretical results on steady-state diffusion equations over $[0, 1]^d$. Note, however, that the approach in Brugiapaglia (2020) is heavily affected by the curse of dimensionality. Similarly to the PINNs case, one could add a penalty term to enforce boundary conditions or modify the spectral basis functions via multiplication by an extra term. Although this is feasible in practice, the corresponding theoretical analysis is expected to be considerably harder than the Fourier case due to the lack of nice algebraic properties of the spectral basis. In summary, extending CFC to non-periodic high dimensional problems and establishing a rigorous link between CFC and PINNs in that setting would be a natural open question stemming from this work.

In summary, there are still important gaps between theory and practice and the convergence theory of PINNs remains a key area of active research. Nonetheless, Theorem 1 builds a first important bridge between PINNs’ convergence analysis and sparse approximation theory. We hope it will enable further research advances in the theoretical analysis and numerical implementation of physics-informed deep learning in the coming years.

Acknowledgments

SB acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) through grant RGPIN-2020-06766, the Fonds de Recherche du Québec Nature et Technologies (FRQNT) through grants 313276 and 359708, the Faculty of Arts and Science of Concordia University, and Applied Math Lab of the Centre de Recherches Mathématiques (CRM). ND acknowledges the support of Florida State University through the CRC 2022-2023 FYAP grant program. WW acknowledges the support of Concordia University through the Horizon Postdoctoral Fellowship program (2021-2023). SB thanks Prof. Alexander Shnirelman (Concordia University) for helpful discussions about norm equivalences in $H^2(\mathbb{T}^d)$. The authors acknowledge the Digital Research Alliance of Canada for providing the computational resources needed to run the numerical experiments. The authors would also like to acknowledge the anonymous reviewers, whose feedback substantially improved the paper’s quality.

Appendix A. Additional Numerical Experiments

This appendix contains additional numerical experiments aimed at testing the numerical impact of some of the constructions adopted in Theorem 1. In particular, we investigate the impact of explicit sparse regularization on the last layer during training and of the use of Fourier features (as opposed to a trainable periodic layer).

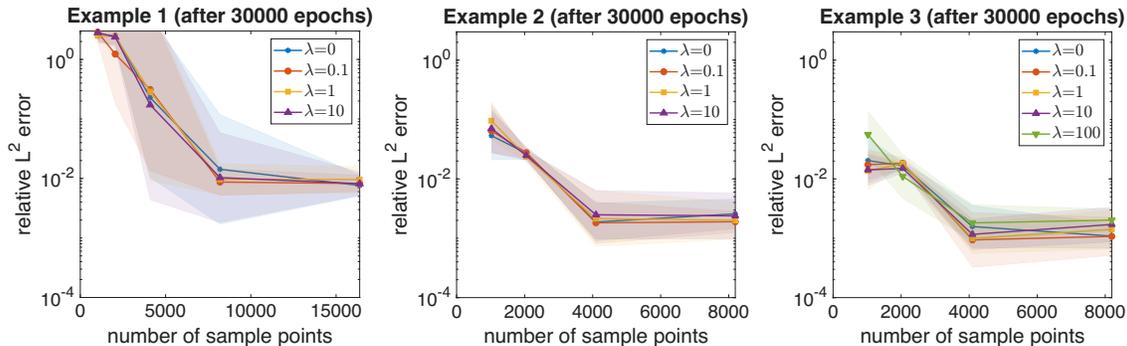


Figure 12: (Impact of sparse regularization) Relative L^2 -error versus training sample count for PINNs trained with with last layer ℓ^1 -regularization and tuning parameter of the form $\lambda = 10^k$ or without regularization ($\lambda = 0$). Plots from left to right refer to Examples 1–3 ($d = 10$), respectively, defined in (27)–(29).

The impact of sparse regularization. Theorem 1 shows the existence of a neural network with an approximately sparse final layer due to the regularization term $\mathcal{R}(\psi)$ employed during training (recall (65)). On the other hand, in §4, we trained periodic PINNs without enforcing any sparsity and observed that this tends to produce solutions with relative L^2 -error between 10^{-2} and 10^{-3} for all the examples considered (assuming to use enough collocation points). In the experiment of Fig. 12, we test whether explicitly promoting sparsity can further improve accuracy.

We take the same periodic PINN architecture and training procedure described in §4 and compare (i) a baseline model with no sparsity penalty, and (ii) the same model with an additional ℓ^1 penalty term applied to the weights of the last layer. We consider values of the tuning parameter spanning several orders of magnitude, of the form $\lambda = 10^k$ for different integer values of k . For each λ , and for each number m of collocation points, we train for a fixed budget of 30000 epochs and show the corresponding relative L^2 -error.

We observe two consistent trends. In the low-sample regime (smaller m), adding an ℓ_1 penalty on the last layer yields accuracy comparable to unregularized baseline when the penalty weight λ is appropriately tuned, suggesting that sparse regularization does not harm (and could help) data efficiency when the number of sample points is small. As m increases, both the unregularized and ℓ_1 -regularized models converge toward a similar error floor, so the effect of regularization becomes negligible. Overall, across all tested settings the ℓ_1 penalty with optimal choice of tuning parameter is never worse than the baseline for Examples 1-3.

Trainable periodic layer vs. Fourier features. Theorem 1 constructs a periodic PINN whose first few layers map the input $\mathbf{x} \in \mathbb{R}^d$ to fixed trigonometric features of the form $(\cos(2\pi\nu x_i), \sin(2\pi\nu x_i))$, for $i \in [d]$ and $\nu \in \{-\nu_{\max}, \dots, \nu_{\max}\}$ (see Figure 1). This construction is the first step towards showing the existence of a class \mathcal{N} of trainable networks with controlled width and depth that can accurately approximate the PDE solution from limited samples. A natural question is whether these Fourier features are only a theoretical device or whether

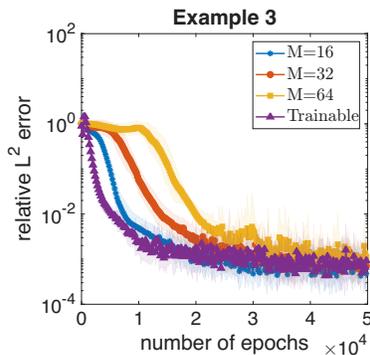


Figure 13: (Trainable vs. non-trainable periodic layer) Comparison of the standard trainable periodic layer versus non-trainable first Fourier features layer $(\cos(2\pi\nu x_i), \sin(2\pi\nu x_i))$ with $i \in [d]$ and $\nu \in [M]$ and frequency cap $M = \{16, 32, 64\}$. The plot shows the relative L^2 -error as a function of the number of training epochs.

it can also be numerically effective when implemented. To address this, we use a frozen set of Fourier features in the first layer instead of learning the periodic layer’s phase parameters as in the PINN tests in §4, where the periodic layer in (4) was considered.

The results are shown in Fig. 13. We fix the first layer to a predetermined set of trigonometric features of the form $(\cos(2\pi\nu x_i), \sin(2\pi\nu x_i))$ with integer frequencies up to a chosen cutoff M , i.e., $\nu \in [M]$. The subsequent layers and training are the same as in §4. We then compare this “frozen periodic layer” model with the standard model where the first periodic layer is trainable by plotting the relative L^2 error versus training epochs. All runs use the same settings (exact solution from Example 3, $d = 6$, number of sample points $m = 10000$).

For a sufficiently large number of epochs, the non-trainable periodic layer model reaches essentially the same final relative L^2 error as the PINN with the trainable first periodic layer. We observe that architectures with larger M are harder to train. This is likely due to the fact that the first layer has width equal to $2M$. The final error is slightly better than the baseline for $M = 16, 32$ and slightly worse than the baseline for $M = 64$ (although the relative L^2 -error seems to still be decreasing). This demonstrates that the explicit Fourier basis used in Theorem 1 is numerically effective and its performance is comparable to that of PINNs with a trainable periodic layer.

References

- B. Adcock and N. Dexter. The gap between theory and practice in function approximation with deep neural networks. *SIAM J. Math. Data Sci.*, 3(2):624–655, 2021.
- B. Adcock, A. Bao, and S. Brugiapaglia. Correcting for unknown errors in sparse high-dimensional function approximation. *Numer. Math.*, 142(3):667–711, 2019.

- B. Adcock, S. Brugiapaglia, N. Dexter, and S. Moraga. Deep neural networks are effective at learning high-dimensional Hilbert-valued functions from limited data. In J. Bruna, J. Hesthaven, and L. Zdeborova, editors, *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference*, volume 145 of *Proc. Mach. Learn. Res.*, pages 1–36. PMLR, 16–19 Aug 2022a.
- B. Adcock, S. Brugiapaglia, and C. G. Webster. *Sparse Polynomial Approximation of High-Dimensional Functions*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2022b.
- B. Adcock, S. Brugiapaglia, N. Dexter, and S. Moraga. Learning smooth functions in high dimensions: From sparse polynomials to deep neural networks. In S. Mishra and A. Townsend, editors, *Numerical Analysis meets Machine Learning*, volume 25 of *Handb. Numer. Anal.*, pages 1–52. Elsevier, 2024.
- B. Adcock, S. Brugiapaglia, N. Dexter, and S. Moraga. Near-optimal learning of Banach-valued, high-dimensional functions via deep neural networks. *Neural Netw.*, 181:106761, 2025.
- M. Bachmayr and W. Dahmen. Adaptive near-optimal rank tensor approximation for high-dimensional operator equations. *Found. Comput. Math.*, 15(4):839–898, 2015.
- M. Bachmayr, R. Schneider, and A. Uschmajew. Tensor networks and hierarchical tensors for the solution of high-dimensional partial differential equations. *Found. Comput. Math.*, 16(6):1423–1472, 2016.
- N. Baker, F. Alexander, T. Bremer, A. Hagberg, Y. Kevrekidis, H. Najm, M. Parashar, A. Patra, J. Sethian, S. Wild, K. Willcox, and S. Lee. Workshop report on basic research needs for scientific machine learning: Core technologies for artificial intelligence. Technical report, USDOE Office of Science (SC), Washington, DC (United States), 2019.
- R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.
- R. Bellman. *Adaptive Control Processes: A Guided Tour*. Princeton University Press, Princeton, NJ, 1961.
- A. Belloni, V. Chernozhukov, and L. Wang. Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4):791–806, 2011.
- A. Berk, S. Brugiapaglia, and T. Hoheisel. Square root LASSO: Well-posedness, Lipschitz stability, and the tuning trade-off. *SIAM J. Optim.*, 34(3):2609–2637, 2024.
- S. Brugiapaglia. A compressive spectral collocation method for the diffusion equation under the restricted isometry property. In M. D’Elia, M. Gunzburger, and G. Rozza, editors, *Quantification of Uncertainty: Improving Efficiency and Technology*, volume 137 of *Lect. Notes Comput. Sci. Eng.*, pages 15–40. Springer, 2020.
- S. Brugiapaglia, S. Micheletti, and S. Perotto. Compressed solving: A numerical approximation technique for elliptic PDEs based on Compressed Sensing. *Comput. Math. Appl.*, 70(6):1306–1335, 2015.

- S. Brugiapaglia, L. Tamellini, and M. Tani. Compressive isogeometric analysis. *Comput. Math. Appl.*, 80(12):3137–3155, 2020.
- S. Brugiapaglia, S. Dirksen, H. C. Jung, and H. Rauhut. Sparse recovery in bounded Riesz systems with applications to numerical methods for PDEs. *Appl. Comput. Harmon. Anal.*, 53:231–269, 2021a.
- S. Brugiapaglia, S. Micheletti, F. Nobile, and S. Perotto. Wavelet–Fourier CORSING techniques for multidimensional advection–diffusion–reaction equations. *IMA J. Numer. Anal.*, 41(4):2744–2781, 2021b.
- C. Canuto, M. Y. Hussaini, A. Quarteroni, and T. A. Zang. *Spectral Methods: Fundamentals in Single Domains*. Springer, 2006.
- J. Chen, X. Chi, W. E, and Z. Yang. Bridging traditional and machine learning-based algorithms for solving PDEs: The random feature method. *J. Mach. Learn.*, 1(3):268–298, 2022.
- A. Chernov and D. Dũng. New explicit-in-dimension estimates for the cardinality of high-dimensional hyperbolic crosses and approximation of functions having mixed smoothness. *J. Complexity*, 32(1):92–121, 2016.
- A. Cohen and G. Migliorati. Multivariate approximation in downward closed polynomial spaces. In J. Dick, F. Y. Kuo, and H. Woźniakowski, editors, *Contemporary Computational Mathematics-A celebration of the 80th birthday of Ian Sloan*, pages 233–282. Springer Cham, 2018.
- A. Cohen, W. Dahmen, and R. DeVore. Compressed sensing and best k -term approximation. *J. Amer. Math. Soc.*, 22(1):211–231, 2009.
- W. Dahmen, R. Devore, L. Grasedyck, and E. Süli. Tensor-sparsity of solutions to high-dimensional elliptic partial differential equations. *Found. Comput. Math.*, 16(4):813–874, 2016.
- I. Daubechies, O. Runborg, and J. Zou. A sparse spectral method for homogenization multiscale problems. *Multiscale Model. Simul.*, 6(3):711–740, 2007.
- J. Daws and C. Webster. Analysis of deep neural networks with quasi-optimal polynomial approximation rates. *arXiv:1912.02302*, 2019.
- T. De Ryck and S. Mishra. Error analysis for physics-informed neural networks (PINNs) approximating Kolmogorov PDEs. *Adv. Comput. Math.*, 48(6):79, 2022.
- T. De Ryck, S. Lanthaler, and S. Mishra. On the approximation of functions by tanh neural networks. *Neural Netw.*, 143:732–750, 2021.
- S. Dong and N. Ni. A method for representing periodic functions and enforcing exactly periodic boundary conditions with deep neural networks. *J. Comput. Phys.*, 435:110242, 2021.

- A. Doostan and H. Owhadi. A non-adapted sparse approximation of PDEs with stochastic inputs. *J. Comput. Phys.*, 230(8):3015–3034, 2011.
- N. Doumèche, G. Biau, and C. Boyer. On the convergence of PINNs. *Bernoulli*, 31(3):2127–2151, 2025.
- D. Dũng, V. Temlyakov, and T. Ullrich. *Hyperbolic Cross Approximation*. Springer, 2018.
- W. E, J. Han, and A. Jentzen. Algorithms for solving high dimensional PDEs: From nonlinear Monte Carlo to machine learning. *Nonlinearity*, 35(1):278, 2021.
- D. Elbrächter, D. Perekrestenko, P. Grohs, and H. Bölcskei. Deep neural network approximation theory. *IEEE Trans. Inf. Theory*, 67(5):2581–2623, 2021.
- L. C. Evans. *Partial Differential Equations*, volume 19. American Mathematical Society, 2010.
- S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Appl. Numer. Harmon. Anal. Birkhäuser, New York, NY, 2013.
- N. R. Franco and S. Brugiapaglia. A practical existence theorem for reduced order models based on convolutional autoencoders. *Foundations of Data Science*, 7(1):72–98, 2025.
- T. Gerstner and M. Griebel. Dimension–adaptive tensor–product quadrature. *Computing*, 71:65–87, 2003.
- C. Gross and M. Iwen. Sparse spectral methods for solving high-dimensional and multiscale elliptic PDEs. *Found. Comput. Math.*, 25(3):765–811, 2025.
- J. Han and A. Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Commun. Math. Stat.*, 5(4):349–380, 2017.
- J. Han, A. Jentzen, and W. E. Solving high-dimensional partial differential equations using deep learning. *Proc. Nat. Acad. Sci.*, 115(34):8505–8510, 2018.
- R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 2012.
- Z. Hu, K. Shukla, G. E. Karniadakis, and K. Kawaguchi. Tackling the curse of dimensionality with physics-informed neural networks. *Neural Netw.*, 176:106369, 2024.
- S. Jokar, V. Mehrmann, M. E. Pfetsch, and H. Yserentant. Sparse approximate solution of partial differential equations. *Appl. Numer. Math.*, 60(4):452–472, 2010.
- G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2017.

- T. Kühn, W. Sickel, and T. Ullrich. Approximation of mixed order Sobolev functions on the d -torus: asymptotics, preasymptotics, and d -dependence. *Constr. Approx.*, 42(3): 353–398, 2015.
- F. Kupka. Sparse grid spectral methods for the numerical solution of partial differential equations with periodic boundary conditions. *PhD dissertation (Universität Wien)*, 1997.
- I. Lagaris, A. Likas, and D. Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Trans. Neural Netw. Learn. Syst.*, 9(5):987–1000, 1998.
- H. Lee and I. Kang. Neural algorithm for solving differential equations. *J. Comput. Phys.*, 91(1):110–131, 1990.
- G. Migliorati. Adaptive polynomial approximation by means of random discrete least squares. In *Numerical Mathematics and Advanced Applications-ENUMATH 2013: Proceedings of ENUMATH 2013, the 10th European Conference on Numerical Mathematics and Advanced Applications, Lausanne, August 2013*, pages 547–554. Springer, 2014.
- G. Migliorati. Adaptive approximation by optimal weighted least-squares methods. *SIAM J. Numer. Anal.*, 57(5):2217–2245, 2019.
- J. A. A. Opschoor, C. Schwab, and J. Zech. Exponential ReLU DNN expression of holomorphic maps in high dimension. *Constr. Approx.*, 55:537–582, 2022.
- E. Putri, M. Shahab, M. Iqbal, I. Mukhlash, A. Hakam, L. Mardianto, and H. Susanto. A deep-genetic algorithm (deep-GA) approach for high-dimensional nonlinear parabolic partial differential equations. *Comput. Math. Appl.*, 154:120–127, 2024.
- M. Raissi, P. Perdikaris, and G. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.*, 378:686–707, 2019. ISSN 0021-9991.
- H. Rauhut and R. Ward. Sparse Legendre expansions via ℓ_1 -minimization. *J. Approx. Theory*, 164(5):517–533, 2012.
- C. Schwab and J. Zech. Deep learning in high dimension: Neural network expression rates for generalized polynomial chaos expansions in uq. *Anal. Appl.*, 17(01):19–55, 2019.
- J. Shen and H. Yu. Efficient spectral sparse grid methods and applications to high-dimensional elliptic problems. *SIAM J. Sci. Comput.*, 32(6):3228–3250, 2010.
- J. Shen and H. Yu. Efficient spectral sparse grid methods and applications to high-dimensional elliptic equations II. Unbounded domains. *SIAM J. Sci. Comput.*, 34(2): A1141–A1164, 2012.
- Y. Shin, J. Darbon, and G. Karniadakis. On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type PDEs. *Commun. Comput. Phys.*, 28(5), 2020.

- Y. Shin, Z. Zhang, and G. Karniadakis. Error estimates of residual minimization using neural networks for linear PDEs. *J. Mach. Learn. Model. Comput.*, 4(4), 2023.
- J. Sirignano and K. Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *J. Comput. Phys.*, 375:1339–1364, 2018.
- H. Son, J. W. Jang, W. J. Han, and H. J. Hwang. Sobolev training for physics informed neural networks. *arXiv:2101.08932*, 2021.
- V. Temlyakov. *Multivariate Approximation*, volume 32. Cambridge University Press, 2018.
- C. Wang, S. Li, D. He, and L. Wang. Is L^2 physics informed loss always suitable for training physics informed neural network? *Adv. Neural Inf. Process. Syst.*, 35:8278–8290, 2022.
- W. Wang and S. Brugiapaglia. Compressive Fourier collocation methods for high-dimensional diffusion equations with periodic boundary conditions. *IMA J. Numer. Anal.*, 44(6):3780–3814, 2024.
- E. W. Weisstein. Multiple-angle formulas. <https://mathworld.wolfram.com/Multiple-AngleFormulas.html>, 2024. Accessed: 2024-02-21.
- S. Zeng, Z. Zhang, and Q. Zou. Adaptive deep neural networks methods for high-dimensional partial differential equations. *J. Comput. Phys.*, 463:111232, 2022.