

On universal approximation and error bounds for Fourier Neural Operators

N. Kovachki and S. Lanthaler and S. Mishra

Research Report No. 2021-23
July 2021

Seminar für Angewandte Mathematik
Eidgenössische Technische Hochschule
CH-8092 Zürich
Switzerland

ON UNIVERSAL APPROXIMATION AND ERROR BOUNDS FOR FOURIER NEURAL OPERATORS

NIKOLA KOVACHKI, SAMUEL LANTHALER, AND SIDDHARTHA MISHRA

ABSTRACT. Fourier neural operators (FNOs) have recently been proposed as an effective framework for learning operators that map between infinite-dimensional spaces. We prove that FNOs are universal, in the sense that they can approximate any continuous operator to desired accuracy. Moreover, we suggest a mechanism by which FNOs can approximate operators associated with PDEs efficiently. Explicit error bounds are derived to show that the size of the FNO, approximating operators associated with a Darcy type elliptic PDE and with the incompressible Navier-Stokes equations of fluid dynamics, only increases sub (log)-linearly in terms of the reciprocal of the error. Thus, FNOs are shown to efficiently approximate operators arising in a large class of PDEs.

1. INTRODUCTION

Deep neural networks have been extremely successful in diverse fields of science and engineering including image classification, speech recognition, natural language understanding, autonomous systems, game intelligence and protein folding, [16] and references therein. Moreover, deep neural networks are being increasingly used successfully in scientific computing, particular in simulating physical and engineering systems modeled by partial differential equations (PDEs). Examples include the use of physics informed neural networks [29, 30, 27, 28] for solving forward and inverse problems for PDEs and supervised learning algorithms for high-dimensional parabolic PDEs [11] and parametric elliptic [14, 32] and hyperbolic [22, 23] PDEs, among others.

The success of deep neural networks at a wide variety of learning tasks can be attributed to a confluence of several factors such as the availability of massive labeled data sets, the design of novel architectures and training algorithms as well as the abundance of high-end computing platforms such as GPUs [12]. Still, it is fair to surmise that this edifice of success partly rests on the foundation of *universal approximation* [1, 9, 13], i.e., the ability of neural networks to approximate any continuous (even measurable) function, mapping a finite-dimensional input space into another finite-dimensional output space, to arbitrary accuracy.

However, many interesting learning tasks entail learning *operators* i.e., mappings between an infinite-dimensional input Banach space and (possibly) an infinite-dimensional output space. A prototypical example in scientific computing is provided by nonlinear operators that map the initial datum into the (time series of) solution of a nonlinear time-dependent PDE such as the Navier-Stokes equations of fluid dynamics. A priori, it is unclear if neural networks can be successfully employed for learning such operators from data, given that their universality only pertains to finite-dimensional functions.

The first successful use of neural networks in the context of such *operator learning* was provided in [8], where the authors proposed a novel neural network based learning architecture, which they termed as *operator networks* and proved that these operator networks possess a surprising universal approximation property for

infinite-dimensional nonlinear operators. Operator networks are based on two different neural networks, a *branch net* and a *trunk net*, which are trained concurrently to learn from data. More recently, the authors of [21] have proposed using deep, instead of shallow, neural networks in both the trunk and branch net and have christened the resulting architecture as a *DeepOnet*. In a recent article [15], the universal approximation property of DeepOnets was extended, making it completely analogous to universal approximation results for finite-dimensional functions by neural networks. The authors of [15] were also able to show that DeepOnets can break the curse of dimensionality for a large variety of PDE learning tasks. Hence, in spite of the underlying infinite-dimensional setting, DeepOnets are capable of approximating a large variety of nonlinear operators efficiently. This is further validated by the success of DeepOnets in many interesting examples in scientific computing [26, 6, 20] and references therein.

An alternative operator learning framework is provided by the concept of *neural operators*, first proposed in [18]. Just as canonical artificial neural networks are a concatenated composition of multiple hidden layers, with each hidden layer composing an affine function with a scalar nonlinear activation function, neural operators also compose multiple hidden layers, with each hidden layer composing an affine *operator* with a local, scalar nonlinear activation operator. The infinite-dimensional setup is reflected in the fact that the affine operator can be significantly more general than in the finite-dimensional case, where it is represented by a weight matrix and bias vector. On the other hand, for neural operators, one can even use *non-local* linear operators, such as those defined in terms of an integral kernel. The evaluation of such integral kernels can be performed either with graph kernel networks [18] or with multipole expansions [17].

More recently, the authors of [19] have proposed using convolution-based integral kernels within neural operators. Such kernels can be efficiently evaluated in the Fourier space, leading to the resulting neural operators being termed as *Fourier Neural Operators* (FNOs). In [19], the authors discuss the advantages, in terms of computational efficiency, of FNOs over the other neural operators mentioned above. Moreover, they present several convincing numerical experiments to demonstrate that FNOs can very efficiently approximate a variety of operators that arise in simulating PDEs.

However, the theoretical basis for neural operators has not yet been properly investigated. In particular, it is unclear if neural operators such as FNOs are *universal* i.e., if they can approximate a large class of nonlinear infinite-dimensional operators. Moreover in this infinite-dimensional setting, universality does not suffice to indicate computational viability or efficiency as the size of the underlying neural networks might grow exponentially with respect to increasing accuracy, see discussion in [15] on this issue. Hence in addition to universality, it is natural to ask if neural operators can *efficiently* approximate a large class of operators, such as those arising in the simulation of parametric PDEs.

The investigation of these questions is the main rationale for the current paper. We focus our attention here on FNOs as they appear to be the most promising of the neural operator based operator learning frameworks. Our main result in this paper is to show that FNOs are *universal* in possessing the ability to approximate a very large class of continuous nonlinear operators. This result highlights the potential of FNOs for operator learning.

As argued before, a universality result is only a first step and by itself, does not constitute evidence for efficient approximation by FNOs. In fact, we show that in the worst case, the network size might grow exponentially with respect to accuracy, when approximating general operators. Hence, there is a need to derive

explicit bounds on the network size in terms of the desired error tolerance. In this context, we consider a concrete computational realization of FNOs, that we term as *pseudospectral FNO* or Ψ -FNO (for short). In addition to proving universality for Ψ -FNOs, we will suggest a mechanism through which Ψ -FNOs can approximate operators arising from PDEs, efficiently. We also derive explicit error bounds for this architecture in approximating PDEs, for two widely used prototypical examples of PDEs i.e, a Darcy type elliptic equation and the incompressible Navier-Stokes equations of fluid dynamics. In particular, we prove that the size of Ψ -FNOs in approximating the underlying operators for both these PDEs, under suitable hypotheses, only scales polynomially (log-linearly) in the error. Thus, FNOs can approximate these operators efficiently and these results validate some of the computational findings of [19]. Together, these results constitute the first theoretical justification for the use of FNOs.

The rest of the paper is organized as follows: in section 2, we introduce FNOs and state the universality result. We also introduce Ψ -FNOs in this section. In section 3, we show that Ψ -FNOs can efficiently approximate operators, stemming from the Darcy-type elliptic equation as well as the incompressible Navier-Stokes equations. In section 4, we compare FNOs with DeepOnets and the results of the article are discussed in section 5. The mathematical notation, used in this paper, is summarized in Appendix A and we present all the technical details and proofs in other appendices.

2. APPROXIMATION BY FOURIER NEURAL OPERATORS

In this section, we present Fourier Neural Operators (FNOs) and discuss their approximation of a class of nonlinear operators specified below:

2.1. Setting for Operator Learning.

Setting 2.1. We fix a spatial dimension $d \in \mathbb{N}$, and denote by $D \subset \mathbb{R}^d$ a domain in \mathbb{R}^d . We consider the approximation of operators $\mathcal{G} : \mathcal{A}(D; \mathbb{R}^{d_a}) \rightarrow \mathcal{U}(D; \mathbb{R}^{d_u})$, $a \mapsto u := \mathcal{G}(a)$, where the input $a \in \mathcal{A}(D; \mathbb{R}^{d_a})$, $d_a \in \mathbb{N}$, is a function $a : D \rightarrow \mathbb{R}^{d_a}$ with d_a components, and the output $u \in \mathcal{U}(D; \mathbb{R}^{d_u})$, $d_u \in \mathbb{N}$, is a function $u : D \rightarrow \mathbb{R}^{d_u}$ with d_u components. Here $\mathcal{A}(D; \mathbb{R}^{d_a})$ and $\mathcal{U}(D; \mathbb{R}^{d_u})$ are Banach spaces (or suitable subsets of Banach spaces). Typical examples of \mathcal{A} and \mathcal{U} include the space of continuous functions $C(D; \mathbb{R}^{d_u})$, or Sobolev spaces $H^s(D; \mathbb{R}^{d_u})$ of order $s \geq 0$ (see Appendix B for definitions.).

Concrete examples for operators \mathcal{G} , involving solution operators of PDEs, are given in section 3.

2.2. Neural Operators. With the above setting 2.1 and as defined in [18], a neural operator $\mathcal{N} : \mathcal{A}(D; \mathbb{R}^{d_a}) \rightarrow \mathcal{U}(D; \mathbb{R}^{d_u})$, $a \mapsto \mathcal{N}(a)$ is a mapping of the form

$$\mathcal{N}(a) = \mathcal{Q} \circ \mathcal{L}_L \circ \mathcal{L}_{L-1} \circ \cdots \circ \mathcal{L}_1 \circ \mathcal{R}(a),$$

for a given depth $L \in \mathbb{N}$, where $\mathcal{R} : \mathcal{A}(D; \mathbb{R}^{d_a}) \rightarrow \mathcal{U}(D; \mathbb{R}^{d_v})$, $d_v \geq d_u$, is a *lifting* operator (acting locally), of the form

$$\mathcal{R}(a)(x) = Ra(x), \quad R \in \mathbb{R}^{d_v \times d_a}, \quad (2.1)$$

and $\mathcal{Q} : \mathcal{U}(D; \mathbb{R}^{d_v}) \rightarrow \mathcal{U}(D; \mathbb{R}^{d_u})$ is a local *projection* operator, of the form

$$\mathcal{Q}(v)(x) = Qv(x), \quad Q \in \mathbb{R}^{d_u \times d_v}. \quad (2.2)$$

Remark 2.2. In practice, it has been found that improved results can be obtained if the simple *linear* lifting and projection operators \mathcal{R} (2.1) and \mathcal{Q} (2.2) are replaced instead by *non-linear* mappings of the form

$$\widehat{\mathcal{R}}(a)(x) = \widehat{R}(a(x), x), \quad \widehat{\mathcal{Q}}(v)(x) = \widehat{Q}(v(x), x),$$

where $\widehat{R} : \mathbb{R}^{d_a} \times D \rightarrow \mathbb{R}^{d_v}$ and $\widehat{Q} : \mathbb{R}^{d_v} \times D \rightarrow \mathbb{R}^{d_u}$ are neural networks with activation function σ . Our error estimates will rely on the (more restrictive) linear choice of lifting and projection operators, given by (2.1), (2.2). The linear choice has the theoretical benefit of ensuring *compositionality*, i.e. that a composition of neural operators can again be represented by a neural operator (cf. Lemma D.4). Despite this technical distinction, we emphasize that all of our error and complexity estimates continue to hold also for neural operators with non-linear lifting and projections, since linear operators can always be approximated by non-linear ones (cp. Lemma C.1). In fact, in the non-linear case, our results imply that \widehat{Q}, \widehat{R} can be chosen to be *shallow* networks.

In analogy with canonical finite-dimensional neural networks, the layers $\mathcal{L}_1, \dots, \mathcal{L}_L$ are non-linear operator layers, $\mathcal{L}_\ell : \mathcal{U}(D; \mathbb{R}^{d_v}) \rightarrow \mathcal{U}(D; \mathbb{R}^{d_v})$, $v \mapsto \mathcal{L}_\ell(v)$, which we assume to be of the form

$$\mathcal{L}_\ell(v)(x) = \sigma \left(W_\ell v(x) + b_\ell(x) + (\mathcal{K}(a; \theta_\ell)v)(x) \right), \quad \forall x \in D.$$

Here, the weight matrix $W_\ell \in \mathbb{R}^{d_v \times d_v}$ and bias $b_\ell(x) \in \mathcal{U}(D; \mathbb{R}^{d_v})$ define an affine pointwise mapping $W_\ell v(x) + b_\ell(x)$. The richness of linear operators in the infinite-dimensional setting can partly be realized by defining the following *non-local* linear operator,

$$\mathcal{K} : \mathcal{A} \times \Theta \rightarrow L(\mathcal{U}(D; \mathbb{R}^{d_v}), \mathcal{U}(D; \mathbb{R}^{d_v})),$$

that maps the input field a and a parameter $\theta \in \Theta$ in the parameter-set Θ to a bounded linear operator $\mathcal{K}(a, \theta) : \mathcal{U}(D; \mathbb{R}^{d_v}) \rightarrow \mathcal{U}(D; \mathbb{R}^{d_v})$, and the non-linear activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is applied component-wise. As proposed in [18], the linear operators $\mathcal{K}(a, \theta)$ are integral operators of the form

$$(\mathcal{K}(a; \theta)v)(x) = \int_D \kappa_\theta(x, y; a(x), a(y))v(y) dy, \quad \forall x \in D. \quad (2.3)$$

Here, the integral kernel $\kappa_\theta : \mathbb{R}^{2(d+d_a)} \rightarrow \mathbb{R}^{d_v \times d_v}$ is a neural network parametrized by $\theta \in \Theta$. Specific examples of the integral kernel (2.3) include those evaluated with a graph kernel network as in [18] or with a multipole expansion [17].

2.3. Fourier Neural Operators. As defined in [19], Fourier Neural operators (FNOs) are special cases of general neural operators (2.3), in which the kernel $\kappa_\theta(x, y; a(x), a(y))$ is of the form $\kappa_\theta = \kappa_\theta(x - y)$. In this case, (2.3) can be written as a convolution

$$(\mathcal{K}(\theta)v)(x) = \int_D \kappa_\theta(x - y)v(y) dy, \quad \forall x \in D. \quad (2.4)$$

For concreteness, we consider the periodic domain $D = \mathbb{T}^d$ (which we identify with the standard torus $\mathbb{T}^d = [0, 2\pi]^d$), although non-periodic, rectangular domains D can also be handled in a straightforward manner.

Given this periodic framework, the convolution operator in (2.4) can be computed using the Fourier transform \mathcal{F} and the inverse Fourier transform \mathcal{F}^{-1} (see Appendix B (B.1) and (B.2) for notation and definitions), resulting in the following equivalent representation of the kernel (2.3),

$$(\mathcal{K}(\theta)v)(x) = \mathcal{F}^{-1} \left(P_\theta(k) \cdot \mathcal{F}(v)(k) \right)(x), \quad \forall x \in \mathbb{T}^d. \quad (2.5)$$

Here, $P_\theta(k) \in \mathbb{C}^{d_v \times d_v}$ is a full matrix indexed by $k \in \mathbb{Z}^d$, and is related to the integral kernel $\kappa_\theta(x)$ in (2.4) via the Fourier transform, $P_\theta(k) = \mathcal{F}(\kappa_\theta)(k)$. Note that we must impose that $P_\theta(-k) = P_\theta(k)^\dagger$ coincides with the Hermitian transpose for all $k \in \mathbb{Z}^d$, to ensure that the image function $(\mathcal{K}(\theta)v)(x)$ is a real-valued function for real-valued $v(x)$. Consequently, the form of Fourier neural operators (FNOs) for the periodic domain \mathbb{T}^d is that of a mapping $\mathcal{N} : \mathcal{A}(D; \mathbb{R}^{d_a}) \rightarrow \mathcal{U}(D; \mathbb{R}^{d_u})$, of the form

$$\mathcal{N}(a) := \mathcal{Q} \circ \mathcal{L}_L \circ \mathcal{L}_{L-1} \circ \cdots \circ \mathcal{L}_1 \circ \mathcal{R}(a), \quad (2.6)$$

where the lifting and projection operators \mathcal{R} and \mathcal{Q} are given by (2.1) and (2.2), respectively, and where the non-linear layers \mathcal{L}_ℓ are of the form

$$\mathcal{L}_\ell(v)(x) = \sigma \left(W_\ell v(x) + b_\ell(x) + \mathcal{F}^{-1} \left(P_\ell(k) \cdot \mathcal{F}(v)(k) \right) (x) \right). \quad (2.7)$$

Here, $W_\ell \in \mathbb{R}^{d_v \times d_v}$ and $b_\ell(x)$ define a pointwise affine mapping (corresponding to weights and biases), and $P_\ell : \mathbb{Z}^d \rightarrow \mathbb{C}^{d_v \times d_v}$ defines the coefficients of a non-local, linear mapping via the Fourier transform.

Remark 2.3. The simplest example for a FNO, as defined by (2.6),(2.4) is as follows; let $\widehat{\mathcal{N}} : \mathbb{R}^{d_a} \rightarrow \mathbb{R}^{d_u}$ be a canonical finite-dimensional neural network with activation function σ . We can associate to $\widehat{\mathcal{N}}$ the mapping $\mathcal{N} : L^2(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{d_u})$, given by $a(x) \mapsto \widehat{\mathcal{N}}(a(x))$. We easily observe that \mathcal{N} is a FNO as we can write it in the form,

$$\widehat{\mathcal{N}} = \widehat{\mathcal{Q}} \circ \widehat{\mathcal{L}}_L \circ \cdots \circ \widehat{\mathcal{L}}_1 \circ \widehat{\mathcal{R}},$$

where $\widehat{\mathcal{R}}(y) = Ry$ with $R \in \mathbb{R}^{d_u \times d_a}$, and each layer $\widehat{\mathcal{L}}_\ell$ is of the form $\widehat{\mathcal{L}}_\ell(y) = \sigma(W_\ell y + b_\ell)$ for some $W_\ell \in \mathbb{R}^{d_v \times d_v}$, $b_\ell \in \mathbb{R}^{d_v}$, with $\widehat{\mathcal{Q}}$ being an affine output layer of the form $\widehat{\mathcal{Q}}(y) = Qy + q$ with $Q \in \mathbb{R}^{d_u \times d_v}$, $q \in \mathbb{R}^{d_u}$. Replacing the input y by a function $v(x)$, these layers clearly are a special case of the FNO lifting layer (2.1), the non-linear layers (2.7) (with $P_\ell \equiv 0$ and constant bias $b_\ell(x) \equiv b_\ell$), and the projection layer (2.2). Thus, any finite-dimensional neural network can be identified with a FNO as defined above.

For the remainder of this work, we make the following assumption,

Assumption 2.4 (Activation function). Unless explicitly stated otherwise, the activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ in (2.7) is assumed to be non-polynomial, (globally) Lipschitz continuous and $\sigma \in C^3$.

2.4. Universal Approximation by FNOs. Next, we will show that FNOs (2.6) are *universal* i.e., given a large class of operators, as defined in setting 2.1, one can find an FNO that approximates it to desired accuracy. To be more precise, we have the following theorem.

Theorem 2.5 (Universal approximation). Let $s, s' \geq 0$. Let $\mathcal{G} : H^s(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow H^{s'}(\mathbb{T}^d; \mathbb{R}^{d_u})$ be a continuous operator. Let $K \subset H^s(\mathbb{T}^d; \mathbb{R}^{d_a})$ be a compact subset. Then for any $\epsilon > 0$, there exists a FNO $\mathcal{N} : H^s(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow H^{s'}(\mathbb{T}^d; \mathbb{R}^{d_u})$, of the form (2.6), continuous as an operator $H^s \rightarrow H^{s'}$, such that

$$\sup_{a \in K} \|\mathcal{G}(a) - \mathcal{N}(a)\|_{H^{s'}} \leq \epsilon.$$

Sketch of proof. The detailed proof of this universal approximation theorem is provided in Appendix D.4 and we outline it here. For notational simplicity, we set $d_a = d_u = 1$, and first observe the following lemma, proved in Appendix D.1:

Lemma 2.6. Assume that the universal approximation Theorem 2.5 holds for $s' = 0$. Then it holds for arbitrary $s' \geq 0$.

The main objective is thus to prove Theorem 2.5 for the special case $s' = 0$; i.e. given a continuous operator $\mathcal{G} : H^s(\mathbb{T}^d) \rightarrow L^2(\mathbb{T}^d)$, $K \subset H^s(\mathbb{T}^d)$ compact, and $\epsilon > 0$, we wish to construct a FNO $\mathcal{N} : H^s(\mathbb{T}^d) \rightarrow L^2(\mathbb{T}^d)$, such that $\sup_{a \in K} \|\mathcal{G}(a) - \mathcal{N}(a)\|_{L^2} \leq \epsilon$.

To this end, we start by defining the following operator,

$$\mathcal{G}_N : H^s(\mathbb{T}^d) \rightarrow L^2(\mathbb{T}^d), \quad \mathcal{G}_N(a) := P_N \mathcal{G}(P_N a), \quad (2.8)$$

with P_N being the orthogonal Fourier projection operator, defined in Appendix B (B.6). Thus, \mathcal{G}_N can be thought of loosely as the *Fourier projection* of the continuous operator \mathcal{G} .

Next, we can show that for any given $\epsilon > 0$, there exists $N \in \mathbb{N}$, such that

$$\|\mathcal{G}(a) - \mathcal{G}_N(a)\|_{L^2} \leq \epsilon, \quad \forall a \in K. \quad (2.9)$$

Thus, the proof boils down to finding a FNO (2.6) that can approximate the operator \mathcal{G}_N to any desired accuracy.

To this end, we introduce a set of Fourier wavenumbers $k \in \mathcal{K}_N$, by

$$\mathcal{K}_N := \{k \in \mathbb{Z}^d \mid |k|_\infty \leq N\}, \quad (2.10)$$

and define a *Fourier conjugate* or *Fourier dual* operator of the form $\widehat{\mathcal{G}}_N : \mathbb{C}^{\mathcal{K}_N} \rightarrow \mathbb{C}^{\mathcal{K}_N}$,

$$\widehat{\mathcal{G}}_N(\widehat{a}_k) := \mathcal{F}_N \left(\mathcal{G}_N \left(\operatorname{Re} \left(\mathcal{F}_N^{-1}(\widehat{a}_k) \right) \right) \right), \quad (2.11)$$

such that the identity

$$\mathcal{G}_N(a) = \mathcal{F}_N^{-1} \circ \widehat{\mathcal{G}}_N \circ \mathcal{F}_N(P_N a), \quad (2.12)$$

holds for all *real-valued* $a \in L^2(\mathbb{T}^d)$. Here, \mathcal{F}_N is the discrete Fourier transform and \mathcal{F}_N^{-1} is the discrete inverse Fourier transform, with both being defined in Appendix B (B.13) and (B.14), respectively.

The next steps in the proof are to leverage the natural decomposition of the projection \mathcal{G}_N in (2.12) in terms of the discrete Fourier transform $\mathcal{F}_N \circ P_N$, the discrete inverse Fourier transform \mathcal{F}_N^{-1} and the Fourier conjugate operator $\widehat{\mathcal{G}}_N$ and approximate each of these operators by Fourier neural operators.

We start by denoting,

$$\mathbb{R}^{2\mathcal{K}_N} = (\mathbb{R}^2)^{\mathcal{K}_N} (\simeq \mathbb{C}^{\mathcal{K}_N}), \quad (2.13)$$

as the set consisting of coefficients $\{(v_{1,k}, v_{2,k})\}_{k \in \mathcal{K}_N}$, where $v_{\ell,k} \in \mathbb{R}$ are indexed by a tuple (ℓ, k) , $\ell \in \{1, 2\}$, $k \in \mathcal{K}_N$, and interpreting the operator $\mathcal{F}_N \circ P_N$ as a mapping $\mathcal{F}_N \circ P_N : a \mapsto \{(\operatorname{Re}(\widehat{a}_k), \operatorname{Im}(\widehat{a}_k))\}_{|k| \leq N}$, with input $a \in L^2(\mathbb{T}^d)$ and the output $\{\operatorname{Re}(\widehat{a}_k), \operatorname{Im}(\widehat{a}_k)\}_{|k| \leq N} \in \mathbb{R}^{2\mathcal{K}_N}$ is viewed as a *constant* function in $L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$. The approximation of this operator is a straightforward consequence of the following Lemma, proved in Appendix D.2,

Lemma 2.7. Let $B > 0$ and $N \in \mathbb{N}$ be given. For all $\epsilon > 0$, there exists a FNO $\mathcal{N} : L^2(\mathbb{T}^d) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$, $v \mapsto \{\mathcal{N}(v)_{\ell,k}\}$, with *constant* output functions (constant as a function of $x \in \mathbb{T}^d$), and such that

$$\left. \begin{aligned} \|\operatorname{Re}(\widehat{v}_k) - \mathcal{N}(v)_{1,k}\|_{L^\infty} &\leq \epsilon \\ \|\operatorname{Im}(\widehat{v}_k) - \mathcal{N}(v)_{2,k}\|_{L^\infty} &\leq \epsilon \end{aligned} \right\} \quad \forall k \in \mathbb{Z}^d, |k|_\infty \leq N,$$

for all $\|v\|_{L^2} \leq B$, and where $\widehat{v}_k \in \mathbb{C}$ denotes the k -th Fourier coefficient of v .

In the next step, we approximate the (discrete) inverse Fourier transform \mathcal{F}_N^{-1} by an FNO. We recall that FNOs act on functions rather than on constants. Therefore, to connect \mathcal{F}_N^{-1} and FNOs, we are going to interpret the mapping

$$\mathcal{F}_N^{-1} : [-R, R]^{2\mathcal{K}_N} \subset \mathbb{R}^{2\mathcal{K}_N} \rightarrow L^2(\mathbb{T}^d),$$

as a mapping

$$\mathcal{F}_N^{-1} : \begin{cases} L^2(\mathbb{T}^d; [-R, R]^{2\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d), \\ \{\text{Re}(\widehat{v}_k), \text{Im}(\widehat{v}_k)\}_{|k| \leq N} \mapsto v(x), \end{cases}$$

where the input $\{\text{Re}(\widehat{v}_k), \text{Im}(\widehat{v}_k)\}_{|k| \leq N} \in [-R, R]^{2\mathcal{K}_N}$ is identified with a *constant* function in $L^2(\mathbb{T}^d; [-R, R]^{2\mathcal{K}_N})$. The existence of a FNO of the form (2.6) that can approximate (D.17) to desired accuracy is a consequence of the following lemma, proved in Appendix D.3,

Lemma 2.8. Let $B > 0$ and $N \in \mathbb{N}$ be given. For all $\epsilon > 0$, there exists a FNO $\mathcal{N} : L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d)$, such that for any $v \in L^2_N(\mathbb{T}^d)$ with $\|v\|_{L^2} \leq B$, we have

$$\|v - \mathcal{N}(w)\|_{L^2} \leq \epsilon,$$

where $w(x) := \{(\text{Re}(\widehat{v}_k), \text{Im}(\widehat{v}_k))\}_{k \in \mathcal{K}_N}$, i.e. $w \in L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$ is a constant function collecting the real and imaginary parts of the Fourier coefficients \widehat{v}_k of v .

Finally, by setting $\widehat{K} := \mathcal{F}_N(P_N K) \subset \mathbb{C}^{\mathcal{K}_N}$ as the (compact) image of K under the continuous mapping $\mathcal{F}_N \circ P_N : L^2(\mathbb{T}^d) \rightarrow \mathbb{C}^{\mathcal{K}_N}$ and identifying $\mathbb{C}^{\mathcal{K}_N} \simeq \mathbb{R}^{2\mathcal{K}_N}$, where $\widehat{v}_{1,k} := \text{Re}(\widehat{v}_k)$ and $\widehat{v}_{2,k} := \text{Im}(\widehat{v}_k)$ for $k \in \mathcal{K}_N$, we can view $\widehat{\mathcal{G}}_N$ as a continuous mapping

$$\widehat{\mathcal{G}}_N : \widehat{K} \subset \mathbb{R}^{2\mathcal{K}_N} \rightarrow \mathbb{R}^{2\mathcal{K}_N},$$

on a compact subset. Hence, by the universal approximation theorem for finite-dimensional neural networks [1, 13], one can readily show that there exists an FNO, with only *local* weights (see remark 2.3), which will approximate this continuous mapping $\widehat{\mathcal{G}}_N$ on compact subsets to desired accuracy.

Hence, each of the component operators of the decomposition (2.12) can be approximated to desired accuracy by FNOs and the universal approximation theorem follows by composing these FNOs and estimating the resulting error, with details provided in appendix D.4. \square

In the following theorem, we will show that the universal approximation theorem 2.5 can be extended to include operators defined on function spaces with Lipschitz domains. In fact, the Lipschitz condition can be relaxed to include all locally uniform domains using ideas from [31]; we will, however, not pursue this for simplicity of the exposition. We show that one can construct a period extension of the input function and a FNO so that the restriction of the FNO's periodic output to the domain of interest gives a suitable approximation to any continuous operator. Similar ideas have been pursued in the design of numerical algorithms for solving PDEs and usually go by the name of *Fourier continuations* [5, 24]. A major challenge for these methods is designing a suitable periodic function whose restriction gives the solution of interest. We show that FNOs can learn the output representation automatically.

Theorem 2.9. Let $s, s' \geq 0$ and $\Omega \subset [0, 2\pi]^d$ be a domain with Lipschitz boundary. Let $\mathcal{G} : H^s(\Omega; \mathbb{R}^{d_a}) \rightarrow H^{s'}(\Omega; \mathbb{R}^{d_u})$ be a continuous operator. Let $K \subset H^s(\Omega; \mathbb{R}^{d_a})$ be a compact subset. Then there exists a continuous, linear operator

$\mathcal{E} : H^s(\Omega; \mathbb{R}^{d_a}) \rightarrow H^s(\mathbb{T}^d; \mathbb{R}^{d_a})$ such that $\mathcal{E}(a)|_\Omega = a$ for all $a \in H^s(\Omega; \mathbb{R}^{d_a})$. Furthermore, for any $\epsilon > 0$, there exists a FNO $\mathcal{N} : H^s(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow H^{s'}(\mathbb{T}^d; \mathbb{R}^{d_u})$ of the form (2.6), such that

$$\sup_{a \in K} \|\mathcal{G}(a) - \mathcal{N} \circ \mathcal{E}(a)|_\Omega\|_{H^{s'}} \leq \epsilon.$$

Proof. Since Ω is open we have that $\text{dist}(\Omega, \partial[0, 2\pi]^d) > 0$ hence the conclusion of Lemma B.3 in Appendix B follows with the hypercube $B = [0, 2\pi]^d$, in particular, there exists a continuous, linear operator $\mathcal{E} : H^s(\Omega; \mathbb{R}^{d_a}) \rightarrow H^s([0, 2\pi]^d; \mathbb{R}^{d_a})$ such that $\mathcal{E}(a)|_\Omega = a$ and $\mathcal{E}(a)$ is periodic on $[0, 2\pi]^d$ for all $a \in H^s(\Omega; \mathbb{R}^{d_a})$. Therefore $\mathcal{E} : H^s(\Omega; \mathbb{R}^{d_a}) \rightarrow H^s(\mathbb{T}^d; \mathbb{R}^{d_a})$. Similarly, we can construct an extension operator $\mathcal{E}' : H^{s'}(\Omega; \mathbb{R}^{d_u}) \rightarrow H^{s'}(\mathbb{T}^d; \mathbb{R}^{d_u})$.

We can then associate to $\mathcal{G} : H^s(\Omega; \mathbb{R}^{d_a}) \rightarrow H^{s'}(\Omega; \mathbb{R}^{d_u})$ another continuous operator $\bar{\mathcal{G}} : H^s(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow H^{s'}(\mathbb{T}^d; \mathbb{R}^{d_u})$, by defining $\bar{\mathcal{G}}(a) := \mathcal{E}' \circ \mathcal{G} \circ \mathcal{R}(a)$. Here $\mathcal{R}(a) := a|_\Omega$ denotes the restriction to Ω which is clearly linear and continuous. By the continuity of \mathcal{E} , we have that $K' := \mathcal{E}(K)$ is compact in $H^s(\mathbb{T}^d; \mathbb{R}^{d_a})$. By the universal approximation theorem 2.5, for any $\epsilon > 0$, there exists a FNO $\bar{\mathcal{N}} : H^s(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow H^{s'}(\mathbb{T}^d; \mathbb{R}^{d_u})$, such that

$$\sup_{a' \in K'} \|\bar{\mathcal{G}}(a') - \bar{\mathcal{N}}(a')\|_{H^{s'}} \leq \epsilon.$$

But then, using the fact that $\mathcal{R} \circ \mathcal{E} = \text{Id}$, $\mathcal{R} \circ \mathcal{E}' = \text{Id}$, the mapping $\mathcal{N} : H^s(\Omega; \mathbb{R}^{d_a}) \rightarrow H^{s'}(\Omega; \mathbb{R}^{d_u})$, given by $\mathcal{N} := \mathcal{R} \circ \bar{\mathcal{N}} \circ \mathcal{E}$, satisfies

$$\begin{aligned} \sup_{a \in K} \|\mathcal{G}(a) - \mathcal{N}(a)\|_{H^{s'}} &= \sup_{a \in K} \|\mathcal{R} \circ \mathcal{E}' \circ \mathcal{G} \circ \mathcal{R} \circ \mathcal{E}(a) - \mathcal{R} \circ \bar{\mathcal{N}} \circ \mathcal{E}(a)\|_{H^{s'}} \\ &= \sup_{a \in K} \|\mathcal{R} \circ \bar{\mathcal{G}} \circ \mathcal{E}(a) - \mathcal{R} \circ \bar{\mathcal{N}} \circ \mathcal{E}(a)\|_{H^{s'}} \\ &\leq \sup_{a \in K} \|\bar{\mathcal{G}} \circ \mathcal{E}(a) - \bar{\mathcal{N}} \circ \mathcal{E}(a)\|_{H^{s'}} \\ &= \sup_{a' \in K'} \|\bar{\mathcal{G}}(a') - \bar{\mathcal{N}}(a')\|_{H^{s'}} \\ &\leq \epsilon. \end{aligned}$$

□

Remark 2.10. The form of the universal approximation theorem 2.5 stated above, shows that any continuous operator $\mathcal{G} : H^s \rightarrow H^{s'}$ can be approximated to arbitrary accuracy by a FNO, on a given compact subset $K \subset H^s$. The restriction to compact subsets may not always be very natural. For example, to train FNOs in practice, it might be more convenient to draw training samples from a measure μ such as the law of a Gaussian random field, which does not have compact support. Furthermore, the operator \mathcal{G} may not always be continuous. To address these issues, one can follow the recent paper [15], where the authors prove a more general version for the universal approximation of operators for DeepOnets; for any input measure μ , and a Borel measurable operator \mathcal{G} , such that $\int \|\mathcal{G}(a)\|_{L^2}^2 d\mu(a) < \infty$, it is shown that for any $\epsilon > 0$, there exists a DeepOnet $\mathcal{N}(a) \approx \mathcal{G}(a)$ such that

$$\int \|\mathcal{G}(a) - \mathcal{N}(a)\|_{L^2}^2 d\mu(a) < \epsilon.$$

In particular, there are no restrictions on the topological support of μ . The result of [15] was for the alternative operator learning framework of DeepOnets, but the ideas and the proof can be analogously extended to FNOs.

2.5. Ψ -Fourier neural operators. In practice, one needs to compute the FNO, of form (2.6), both during training as well as for the evaluation of the neural operator. Thus, given any input function a , one should be able to readily calculate the FNO $\mathcal{N}(a)$, requiring the efficient computation of the Fourier transform \mathcal{F} (B.1) and the inverse Fourier transform \mathcal{F}^{-1} (B.2). In general, this is not possible as evaluating the Fourier transform (B.1) entails computing an integral exactly. Therefore, approximations are necessary to realize the action of FNOs on functions. Following [19], one can efficiently approximate the Fourier transform and its inverse by the discrete Fourier transform (B.13) and the discrete inverse Fourier transform (B.14), respectively. This amounts to performing a pseudo(Ψ)-spectral Fourier projection between successive layers of the FNO and leading to the following precise definition,

Definition 2.11 (Ψ -FNO). A Ψ -FNO (or Ψ -spectral FNO) is a mapping

$$\mathcal{N} : \mathcal{A}(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow \mathcal{U}(\mathbb{T}^d; \mathbb{R}^{d_u}), \quad a \mapsto \mathcal{N}(a),$$

of the form

$$\mathcal{N}(a) = \mathcal{Q} \circ \mathcal{I}_N \circ \mathcal{L}_L \circ \mathcal{I}_N \circ \cdots \circ \mathcal{L}_1 \circ \mathcal{I}_N \circ \mathcal{R}(a), \quad (2.14)$$

where \mathcal{I}_N denotes the pseudo-spectral Fourier projection onto trigonometric polynomials of degree $N \in \mathbb{N}$ (B.10), the lifting operator $\mathcal{R} : \mathcal{A}(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow \mathcal{U}(\mathbb{T}^d; \mathbb{R}^{d_v})$, the projection $\mathcal{Q} : \mathcal{U}(\mathbb{T}^d; \mathbb{R}^{d_v}) \rightarrow \mathcal{U}(\mathbb{T}^d; \mathbb{R}^{d_u})$ are defined as in (2.1), (2.2), and the non-linear layers \mathcal{L}_ℓ , for $\ell = 1, \dots, N$, are of the form

$$\mathcal{L}_\ell(v)(x) = \sigma \left(W_\ell v(x) + b_\ell(x) + \mathcal{F}^{-1} \left(P_\ell(k) \cdot \mathcal{F}(v)(k) \right) (x) \right).$$

Here, $W_\ell \in \mathbb{R}^{d_v \times d_v}$ and $b_\ell(x) \in \mathcal{U}(\mathbb{T}^d; \mathbb{R}^{d_v})$ define a pointwise affine mapping $v \mapsto W_\ell v(x) + b_\ell(x)$, and the coefficients $P_\ell(k) \in \mathbb{R}^{d_v \times d_v}$ ($k \in \mathcal{K}_N$) define a (non-local) convolution operator via the Fourier transform.

Note that a Ψ -FNO \mathcal{N} is uniquely defined, as an operator, by its restriction to the finite-dimensional subspace $L_N^2(\mathbb{T}^d; \mathbb{R}^{d_a}) \subset \mathcal{A}(\mathbb{T}^d; \mathbb{R}^{d_a})$ (see Appendix B for the definition of L_N^2). Furthermore, we have that the image $\text{Im}(\mathcal{N}) \subset L_N^2(\mathbb{T}^d; \mathbb{R}^{d_u})$. To indicate that a Ψ -FNO is of the form 2.14, for some $N \in \mathbb{N}$, we shall thus more simply say that “ $\mathcal{N} : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_u})$ is a Ψ -FNO”.

At the level of numerical implementation, a Ψ -FNO can be naturally identified with a *finite-dimensional* mapping

$$\widehat{\mathcal{N}} : \mathbb{R}^{d_a \times \mathcal{J}_N} \rightarrow \mathbb{R}^{d_u \times \mathcal{J}_N}, \quad \mathbf{a} \mapsto \widehat{\mathcal{N}}(\mathbf{a}),$$

with input $\mathbf{a} = \{a_j\}_{j \in \mathcal{J}_N} \in \mathbb{R}^{d_a \times \mathcal{J}_N}$ corresponding to the point-values $a_j = a(x_j)$ on the grid $\{x_j\}_{j \in \mathcal{J}_N}$, and $\mathcal{J}_N := \{0, \dots, 2N\}^d$. Here, $\widehat{\mathcal{N}}$ is of the form

$$\widehat{\mathcal{N}}(\mathbf{a}) = \widehat{\mathcal{Q}} \circ \widehat{\mathcal{L}}_L \circ \widehat{\mathcal{L}}_{L-1} \circ \cdots \circ \widehat{\mathcal{L}}_1 \circ \widehat{\mathcal{R}}(\mathbf{a}),$$

where the lifting operator $\widehat{\mathcal{R}} : \mathbb{R}^{d_a \times \mathcal{J}_N} \rightarrow \mathbb{R}^{d_v \times \mathcal{J}_N}$, $\mathbf{a} \mapsto \widehat{\mathcal{R}}(\mathbf{a})$, the projection $\widehat{\mathcal{Q}} : \mathbb{R}^{d_v \times \mathcal{J}_N} \rightarrow \mathbb{R}^{d_u \times \mathcal{J}_N}$, $\mathbf{v} \mapsto \widehat{\mathcal{Q}}(\mathbf{v})$, are given by

$$\begin{aligned} \widehat{\mathcal{R}}(\mathbf{a}) &= \{R a_j\}_{j \in \mathcal{J}_N}, \quad (R \in \mathbb{R}^{d_v \times d_a}), \\ \widehat{\mathcal{Q}}(\mathbf{v}) &= \{Q v_j\}_{j \in \mathcal{J}_N}, \quad (Q \in \mathbb{R}^{d_u \times d_v}), \end{aligned}$$

and the non-linear layers $\widehat{\mathcal{L}}_\ell$, for $\ell = 1, \dots, N$, are of the form

$$\widehat{\mathcal{L}}_\ell(\mathbf{v})_j = \sigma \left(W_\ell v_j + b_{\ell,j} + \mathcal{F}_N^{-1} \left(P_\ell(k) \cdot \mathcal{F}_N(\mathbf{v})(k) \right)_j \right) \quad (2.15)$$

for $j \in \mathcal{J}_N$. Here, $W_\ell \in \mathbb{R}^{d_v \times d_v}$, $b_{\ell,j} = b_\ell(x_j) \in \mathbb{R}^{d_v \times \mathcal{J}_N}$ defines a pointwise affine mapping $W_\ell v_j + b_{\ell,j}$, the coefficients $P_\ell(k) \in \mathbb{C}^{d_v \times d_v}$ ($k \in \mathcal{K}_N$) satisfy the Hermitian conjugacy condition $P_\ell(-k) = P_\ell(k)^\dagger$ and define a (non-local) convolution operator via the discrete Fourier transform, and the non-linear activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is extended componentwise to a function $\mathbb{R}^{d_v \times \mathcal{J}_N} \rightarrow \mathbb{R}^{d_v \times \mathcal{J}_N}$. Comparing \mathcal{N} with the corresponding discretization $\widehat{\mathcal{N}}$, it is easy to see that

$$\widehat{\mathcal{N}}(\{a(x_j)\}_{j \in \mathcal{J}_N})_j = \mathcal{N}(a)(x_j), \quad \forall j \in \mathcal{J}_N.$$

In particular, this implies that $\mathcal{N}(a)(x)$ can in practice be computed for any $x \in \mathbb{T}^d$ via the Fourier interpolation of the grid values $\widehat{\mathcal{N}}(\{a(x_j)\}_{j \in \mathcal{J}_N})$. In contrast to general FNOs, Ψ -FNOs therefore allow for efficient numerical implementation. Furthermore, the discrete (inverse) Fourier transforms in each hidden layer in (2.15) can be very efficiently computed using the fast Fourier transform (FFT).

The above discussion also leads to a very natural definition of the size of a Ψ -FNO below:

Definition 2.12 (Depth, width, lift and size). The **depth** and **width** of a Ψ -FNO \mathcal{N} (cp. Definition 2.11), are defined by

$$\text{depth}(\mathcal{N}) := L, \quad \text{width}(\mathcal{N}) := d_v |\mathcal{J}_N| = d_v |\mathcal{K}_N| = (2N + 1)^d d_v.$$

We refer to the dimension d_v , as the **lift** of \mathcal{N} , i.e. we set

$$\text{lift}(\mathcal{N}) := d_v.$$

The **size** of a Ψ -FNO \mathcal{N} is defined as the total number of degrees of freedom in a Ψ -FNO. A simple calculation shows that

$$\text{size}(\mathcal{N}) = \underbrace{d_u d_v}_{\text{size}(\mathcal{Q})} + L \underbrace{(d_v^2 + d_v |\mathcal{J}_N| + d_v^2 |\mathcal{J}_N|)}_{\text{size}(\mathcal{L}_\ell)} + \underbrace{d_a d_v}_{\text{size}(\mathcal{R})}.$$

The precise size of a Ψ -FNO will not be of any particular relevance for our asymptotic complexity estimates. Instead, we will usually content ourselves with the simple estimate

$$\text{size}(\mathcal{N}) \lesssim \text{depth}(\mathcal{N}) \text{width}(\mathcal{N}) \text{lift}(\mathcal{N}),$$

where we assume that $\max(d_a, d_u) \leq d_v$; under this condition, the above estimate follows from the fact that $\text{size}(\mathcal{N}) \sim L d_v^2 |\mathcal{J}_N|$.

Given our discussion, it is natural to ask whether any FNO $\widehat{\mathcal{N}} = \mathcal{Q} \circ \mathcal{L}_L \circ \mathcal{L}_{L-1} \circ \dots \circ \mathcal{L}_1 \circ \mathcal{R}$ can be approximated to arbitrary accuracy by an associated Ψ -FNO $\mathcal{N} : L_N^2 \rightarrow L_N^2$,

$$\mathcal{N} = \mathcal{Q} \circ \mathcal{I}_N \circ \mathcal{L}_L \circ \mathcal{I}_N \circ \dots \circ \mathcal{L}_1 \circ \mathcal{I}_N \circ \mathcal{R},$$

for sufficiently large $N \in \mathbb{N}$? An affirmative answer can be given for a natural class of FNOs of *finite width*, defined as follows.

Definition 2.13. A FNO $\widehat{\mathcal{N}} : \mathcal{A}(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow \mathcal{U}(\mathbb{T}^d; \mathbb{R}^{d_u})$ is said to be **of finite width**, if $\widehat{\mathcal{N}}$ is a composition $\widehat{\mathcal{N}} = \mathcal{Q} \circ \mathcal{L}_L \circ \dots \circ \mathcal{L}_1 \circ \mathcal{R}$, with layers \mathcal{L}_ℓ of the form (2.7), and for which there exists a “width” $W \in \mathbb{N}$, such that the Fourier multiplier $P_\ell(k) \equiv 0$, for $|k|_\infty > W$.

We can now state the following theorem, which shows that Ψ -FNOs \mathcal{N} provide an arbitrarily close approximation of a given FNO $\widehat{\mathcal{N}}$:

Theorem 2.14. Assume that the activation function $\sigma \in C^m$ is globally Lipschitz continuous. Let $\widehat{\mathcal{N}} : H^s(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{d_u})$ be a FNO of finite width, with $s > d/2$. and assume that $m > s$. Then for any $\epsilon, B > 0$, there exists $N \in \mathbb{N}$ and a Ψ -FNO $\mathcal{N} : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_u})$, such that

$$\sup_{\|a\|_{H^s} \leq B} \|\widehat{\mathcal{N}}(a) - \mathcal{N}(a)\|_{L^2} \leq \epsilon.$$

For the proof, we refer to Appendix D.5. In particular, the last theorem implies an extension of the universal approximation theorem 2.5 to Ψ -FNOs, provided that the input functions have sufficient regularity for the pseudo-spectral projection (B.10) to be well-defined:

Theorem 2.15 (Universal approximation for Ψ -FNOs). Let $s > d/2$, and let $s' \geq 0$. Let $\mathcal{G} : H^s(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow H^{s'}(\mathbb{T}^d; \mathbb{R}^{d_u})$ be a continuous operator. And let $K \subset H^s(\mathbb{T}^d; \mathbb{R}^{d_a})$ be a compact subset. Then for any $\epsilon > 0$, there exists $N \in \mathbb{N}$ and a Ψ -FNO $\mathcal{N} : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_u})$, such that

$$\sup_{a \in K} \|\mathcal{G}(a) - \mathcal{N}(a)\|_{H^{s'}} \leq \epsilon.$$

Proof. Similar to the proof of the universal approximation theorem for FNOs, we again note that the general case $s' \geq 0$ can be deduced from the statement of Theorem 2.15 for the special case $s' = 0$. This is the content of the following lemma, whose proof is provided in Appendix D.6:

Lemma 2.16. Assume that Theorem 2.15 holds for $s' = 0$. Then it holds for arbitrary $s' \geq 0$.

The special case $s' = 0$ follows immediately from Theorem 2.14 and the observation that the FNO approximation constructed in the proof of the universal approximation theorem for FNOs, Theorem 2.5, has finite width. \square

2.5.1. *Structure and properties of Ψ -FNOs.* We conclude this section by pointing out some aspects of the structure of Ψ -FNOs (2.14) that will be relevant in the following. To start with, we can simplify Ψ -FNOs by viewing them in terms of two types of layers. which we will refer to as σ - and \mathcal{F} -layers, respectively. A **σ -layer** $\mathcal{L} = \mathcal{L}_\sigma$ of a Ψ -FNO is a **local, non-linear** layer of the form $\mathcal{L}_\sigma(v)(x) = \mathcal{I}_N \sigma(A \mathcal{I}_N v(x) + b)$, or, in the numerical implementation (cp. (2.15))

$$\mathcal{L}_\sigma(v)_j = \sigma(Av_j + b_j), \quad \forall j \in \mathcal{J}_N,$$

with $A \in \mathbb{R}^{d_v \times d_v}$, and $b_j \in \mathbb{R}^{\mathcal{J}_N \times d_v}$ defining an affine mapping. A **\mathcal{F} -layer** $\mathcal{L} = \mathcal{L}_\mathcal{F}$ of a Ψ -FNO is a **non-local, linear** layer of the form $\mathcal{L}_\mathcal{F}(v)(x) = \mathcal{F}^{-1}(P(k) \cdot \mathcal{F}(\mathcal{I}_N v)(k))(x)$, which in a practical implementation corresponds to

$$\mathcal{L}_\mathcal{F}(v)_j = \mathcal{F}_N^{-1} \left(P(k) \cdot \mathcal{F}_N(v)(k) \right)_j, \quad \forall j \in \mathcal{J}_N,$$

where $P : \mathcal{K}_N \rightarrow \mathbb{C}^{d_v \times d_v}$ is a collection of complex weights, with $P(-k) = P(k)^\dagger$ the Hermitian transpose of $P(k)$, and \mathcal{F}_N (\mathcal{F}_N^{-1}) denotes the discrete (inverse) Fourier transform.

The main point of these definitions is that each Ψ -FNO can be decomposed into a finite number of σ -layers and \mathcal{F} -layers, and that the converse is also true; i.e. any composition of σ -layers and \mathcal{F} -layers can be represented by a Ψ -FNO. These statements are made precise in a series of technical Lemmas, which are stated and proved in Appendix D.7.

3. APPROXIMATION OF PDES BY Ψ -FNOs

We have shown in the previous section that FNOs (2.6) as well as their computational realizations (Ψ -FNOs (2.14)) are universal i.e., they approximate any continuous operator, defined in the setting 2.1, to desired accuracy. However, as repeatedly discussed in the introduction, universality alone does not suffice to claim that FNOs can approximate operators efficiently. In particular, it could happen that the size of the FNO is unfeasibly large to ensure a given accuracy of the approximation. That this is indeed the case is made precise in the following remark.

Remark 3.1. We observe from the proof of Theorem 2.5 that the desired FNO, approximating the operator \mathcal{G} , is constructed as $\mathcal{N}_{\text{IFT}} \circ \widehat{\mathcal{N}} \circ \mathcal{N}_{\text{FT}}$, with $\mathcal{N}_{\text{FT}}, \mathcal{N}_{\text{IFT}}$ approximating the Fourier and Inverse Fourier transforms, respectively, whereas $\widehat{\mathcal{N}} : \mathbb{R}^{2\mathcal{K}_N} \rightarrow \mathbb{R}^{2\mathcal{K}_N}$ is a canonical finite-dimensional neural network approximation of the “Fourier conjugate operator” (2.11): $\widehat{\mathcal{G}}_N : \mathbb{R}^{2\mathcal{K}_N} \rightarrow \mathbb{R}^{2\mathcal{K}_N}$. We note that N herein has to be chosen sufficiently large in order to yield the desired error tolerance of ϵ . By Theorem B.2, this depends on the smoothness of the input space, i.e., if the input $a \in K \subset H^s$, for some $s > 0$, then we need to choose N such that $N^{-s} \sim \epsilon$. Further assuming that the mapping \mathcal{G} is Lipschitz continuous, implies that the Fourier conjugate operator $\widehat{\mathcal{G}}$ is also Lipschitz continuous as a mapping from $\mathbb{R}^{2\mathcal{K}_N}$ to $\mathbb{R}^{2\mathcal{K}_N}$. Hence, neural network approximation results, such as those of [34] for ReLU activations or [10] for tanh activations, yield that the width of the approximating neural network $\widehat{\mathcal{N}}$ scales as $\text{width}(\widehat{\mathcal{N}}) \gtrsim \epsilon^{-D}$, where D is the dimension of the domain of $\widehat{\mathcal{G}}_N$. In the present case, we have $D = |\mathcal{K}_N| \sim N^d \sim \epsilon^{-d/s}$, yielding that

$$\text{width}(\widehat{\mathcal{N}}) \gtrsim \epsilon^{-\epsilon^{-d/s}}. \quad (3.1)$$

This scaling represents a *super-exponential* growth in the size of the FNO \mathcal{N} , with respect to the error ϵ , incurred in approximating the underlying operator \mathcal{G} .

Given the above remark, we infer that in the worst case, a FNO approximating a generic Lipschitz continuous operator \mathcal{G} , can require extremely large sizes to achieve the desired accuracy, making it unfeasible in practice. The same holds for Ψ -FNOs of the form (2.14). This super-exponential growth appears as a form of *curse of dimensionality* i.e., exponential growth of complexity (measured here in the size of the FNO), with respect to the error.

Hence, it is reasonable to ask how these extremely pessimistic complexity bounds on FNOs (Ψ -FNOs), can be reconciled to their robust numerical performance for approximating PDEs, as reported in [19]. The rest of the section investigates this fundamental question.

The starting point of our explanation for the robustness of FNOs in approximating PDEs is the observation that operators which arise in the context of PDEs have a special structure and are not merely generic continuous operators mapping one infinite-dimensional function space to another. To see this, we point out that many time-dependent PDEs arising in physics can be written in the general abstract form,

$$\partial_t u + \nabla \cdot F(u, \nabla u) = 0, \quad (3.2)$$

where for any $(t, x) \in [0, T] \times D \subset \mathbb{R}^d$, $u(t, x) \in \mathbb{R}^{d_u}$ is a vector of physical quantities, describing e.g. density, velocity or temperature of a fluid or other material at a given point $x \in D$ in the domain D and at time $t \in [0, T]$. Equation (3.2) describes the general form of a conservation law for the physical quantities u with a flux function $F(u, \nabla u)$, which is typically *non-linear*, and can e.g. represent advection

or diffusion terms. The flux function $F(u, \nabla u)$ may also depend on u in a *non-local* manner. For example, for the incompressible Navier-Stokes equations in \mathbb{R}^d ,

$$u(x, t) = (u_1(x, t), \dots, u_d(x, t)) \in \mathbb{R}^d,$$

represent the fluid velocity at (x, t) , and the flux is defined by

$$F(u, \nabla u) = -u \otimes u - p + \nu \nabla u,$$

where $p = p(u)$ depends on u in a non-local manner:

$$p = \mathcal{R} : (u \otimes u), \quad \mathcal{R} := (-\Delta)^{-1}(\nabla \otimes \nabla),$$

where \mathcal{R} is a (non-local) operator closely related to the Riesz transform.

A popular numerical method for time-dependent PDEs, of the form (3.2), particularly on periodic domains $D = \mathbb{T}^d$, is the *pseudo-spectral method* [7], wherein (3.2) is discretized as,

$$\partial_t u_N + \nabla \cdot \mathcal{I}_N F(u_N, \nabla u_N) = 0, \quad (3.3)$$

where $u_N \in L_N^2$ is a trigonometric polynomial of degree $\leq N$.

The resulting system of ODEs (3.3) can be further discretized in time using a time-marching scheme. For simplicity, the forward Euler discretization with time step τ leads to,

$$u_N^{n+1} = u_N^n - \tau \nabla \cdot \mathcal{I}_N F(u_N^n, \nabla u_N^n). \quad (3.4)$$

One might prove that the system (3.4) provides a convergent approximation for the underlying time-dependent PDE (3.2) for many different choices of the flux F . In order to connect the approximation (3.4) with FNOs, we decompose the right hand side of (3.4) as,

$$\begin{aligned} u_N^n &\xrightarrow{\mathcal{R}} \begin{Bmatrix} u_N^n \\ u_N^n \end{Bmatrix} \xrightarrow{\mathcal{F}} \begin{Bmatrix} u_N^n \\ \nabla u_N^n \end{Bmatrix} \\ &\xrightarrow{\sigma} \dots \xrightarrow{\sigma} \begin{Bmatrix} u_N^n \\ F_N^\sigma(u_N^n, \nabla u_N^n) \end{Bmatrix} \\ &\xrightarrow{\mathcal{F}} \begin{Bmatrix} u_N^n \\ \nabla \cdot F_N^\sigma(u_N^n, \nabla u_N^n) \end{Bmatrix} \mapsto u_N^n - \tau \nabla \cdot F_N^\sigma(u_N^n, \nabla u_N^n). \end{aligned}$$

Here, \mathcal{R} is the lifting operator and σ , \mathcal{F} are the σ - and \mathcal{F} -layers, respectively, of a Ψ -FNO, that are defined in section 2.5.1. The above representation suggests that the Fourier \mathcal{F} -layers of a Ψ -FNO allow us to take *exact* derivatives, and a composition of σ -layers of a Ψ -FNO allows us to approximate continuous functions to any desired accuracy (via the standard universal approximation theorem for finite-dimensional neural networks); in particular, a composition of σ -layers can provide an approximation

$$(u, \nabla u) \mapsto F_N^\sigma(u, \nabla u) \approx \mathcal{I}_N F(u, \nabla u).$$

Thus, by a suitable composition of σ - and \mathcal{F} -layers, Ψ -FNOs can *emulate* pseudo-spectral methods, providing a mechanism by which such neural operators can approximate solution operators for a large class of PDEs efficiently.

We will make this intuition precise in the following. However, instead of considering a generic abstract form of PDEs, we focus on two PDEs, often encountered in physics, which serve as prototypes for a wide variety of PDEs. We start with an elliptic PDE below.

3.1. Stationary Darcy Flow. We consider the elliptic PDE of the form,

$$-\nabla \cdot (a \nabla u) = f. \quad (3.5)$$

Here, $u \in H^1(\mathbb{T}^d)$, can correspond to the steady-state pressure for a fluid, flowing according to the Darcy's law, in a porous medium with the positive coefficient $a \in L^\infty(\mathbb{T}^d)$, denoting the rock permeability. Another model for (3.5) is that of a diffusion equation, with u modeling the temperature and a the thermal conductivity of the medium.

For simplicity, we assume periodic boundary conditions on \mathbb{T}^d , and we impose that $\int_{\mathbb{T}^d} f dx = \int_{\mathbb{T}^d} u dx = 0$. Employing a suitable rescaling, we will furthermore assume that a can be written in the form:

$$a = 1 + \tilde{a}, \quad \tilde{a} \in H^s(\mathbb{T}^d), \quad s > d/2.$$

We note that the assumption $s > d/2$ ensures that $\|\tilde{a}\|_{L^\infty} < \infty$, via the Sobolev embedding $H^s(\mathbb{T}^d) \hookrightarrow L^\infty(\mathbb{T}^d)$ (cp. Theorem B.1). To ensure that (3.5) is well-posed, we assume the following coercivity condition: There exists $\lambda > 0$, such that $1 + \tilde{a} \geq \lambda$. In fact, we shall assume the slightly stronger condition that

$$\|\tilde{a}\|_{L^\infty} \leq C \|\tilde{a}\|_{H^s} \leq 1 - \lambda, \quad (3.6)$$

where C is the norm of the embedding $H^s(\mathbb{T}^d) \hookrightarrow L^\infty(\mathbb{T}^d)$. The condition (3.6) clearly implies the λ -coercivity of (3.5).

The underlying operator $\mathcal{G} : L^\infty(\mathbb{T}^d) \rightarrow H^1(\mathbb{T}^d)$, maps a coefficient $a \in L^\infty(\mathbb{T}^d)$ to the solution $u \in H^1(\mathbb{T}^d)$ of (3.5). Our aim is to *learn* this operator efficiently using a Ψ -FNO. To this end, we will follow the program discussed above and first present a pseudo-spectral method that approximates the Darcy flow PDE (3.5) accurately. Then, this pseudo-spectral method will be emulated by a suitable Ψ -FNO.

3.1.1. A Fourier-Galerkin approximation of (3.5). We fix $N \in \mathbb{N}$ and assume that the coefficient field and right-hand side $a = 1 + \tilde{a}, f \in H^s(\mathbb{T}^d)$, for some $s > d/2$, such that the pseudo-spectral projections $\mathcal{I}_N \tilde{a}$ and $\mathcal{I}_N f$ are well-defined. We can now define a Fourier-Galerkin approximation of (3.5) as the (unique) solution $u_N \in L_N^2(\mathbb{T}^d)$, with $\int u_N dx = 0$, and such that

$$-\dot{P}_N \nabla \cdot (a_N \nabla u_N) = f_N, \quad (3.7)$$

where we set $a_N := 1 + \tilde{a}_N$ and

$$\begin{cases} \tilde{a}_N := \dot{P}_N \mathcal{I}_{2N} \tilde{a}, \\ f_N := \dot{P}_N \mathcal{I}_{2N} f. \end{cases} \quad (3.8)$$

See notation for P_N, \dot{P}_N and \mathcal{I}_N in appendix B. We observe from (3.8) that \tilde{a}_N and f_N are obtained by first carrying out a pseudo-spectral projection of \tilde{a}, f on the regular grid $\{x_j\}_{j \in \mathcal{J}_{2N}}$ (with $2N$ grid points in each direction), yielding a representation of the form

$$\mathcal{I}_{2N} \tilde{a} = \sum_{|k| \leq 2N} \hat{a}_k e^{i\langle k, x \rangle}, \quad \mathcal{I}_{2N} f = \sum_{|k| \leq 2N} \hat{f}_k e^{i\langle k, x \rangle},$$

for coefficients $\hat{a}_k, \hat{f}_k \in \mathbb{C}$, and then projecting these expressions onto $\dot{L}_N^2(\mathbb{T}^d)$:

$$\dot{P}_N \mathcal{I}_{2N} \tilde{a} = \sum_{0 < |k| \leq N} \hat{a}_k e^{i\langle k, x \rangle}, \quad \dot{P}_N \mathcal{I}_{2N} f = \sum_{0 < |k| \leq N} \hat{f}_k e^{i\langle k, x \rangle}.$$

The reason for this particular choice of the projection (3.8) lies in the fact that the mapping

$$L_N^2(\mathbb{T}^d) \times L_N^2(\mathbb{T}^d) \rightarrow L_N^2(\mathbb{T}^d), \quad (a_N, u_N) \mapsto -P_N \nabla \cdot (a_N \nabla u_N),$$

can be exactly represented by a pseudo-spectral method on $\{x_j\}_{j \in \mathcal{J}_{2N}}$, i.e., on a denser grid of $2N$ grid points in each direction. Indeed, it is easy to see that if $a_N, u_N \in L_N^2$, then $\nabla u_N \in L_N^2$, and the values $\nabla u_N(x_j)$ can be computed exactly via the discrete Fourier transform from the values $u_N(x_j)$, $j \in \mathcal{J}_{2N}$. Since the product $a_N \nabla u_N \in L_{2N}^2$ is a trigonometric polynomial of degree $\leq 2N$, the Fourier coefficients of $a_N \nabla u_N$ can also be recovered from knowledge of the point-values $a_N(x_j) \nabla u_N(x_j)$ at the grid points $x_j \in \mathcal{J}_{2N}$. Finally, since the divergence and projection P_N can be evaluated exactly via discrete Fourier transforms, we conclude that the mapping $(a_N, u_N) \mapsto -P_N \nabla \cdot (a_N \nabla u_N)$ can be computed based on knowledge of the grid values $a_N(x_j)$, and $u_N(x_j)$. The above procedure of computing a product of two trigonometric polynomials of degree $\leq N$ *exactly*, based on the point-values on a finer grid of size $2N + 1$ in each direction is well-known in the context of pseudo-spectral methods, and is usually referred in the numerical analysis literature as *de-aliasing* (cf. section 3.4.2 of the textbook [7]).

In order to computationally realize the Fourier-Galerkin method (3.7), we are going to recast it as a fixed point problem $u_N = F_N(u_N; a_N, f_N)$, where $F_N(\cdot; a_N, f_N) : L_N^2(\mathbb{T}^d) \rightarrow L_N^2(\mathbb{T}^d)$ is defined by

$$F_N(u_N; a_N, f_N) = \dot{P}_N(-\Delta)^{-1} \nabla \cdot (\tilde{a}_N \nabla u_N) + (-\Delta)^{-1} f_N. \quad (3.9)$$

In Lemma E.1 Appendix E, we show that the map (3.9) is a *contraction* and can be efficiently realized by a Picard type fixed-point iteration scheme. This leads to the following algorithm for realizing the Fourier-Galerkin method (3.7) computationally,

Algorithm 3.2 (Approximate solution of Darcy equation).

Input: $N \in \mathbb{N}$, $a \in H^s(\mathbb{T}^d)$, $f \in \dot{H}^{k-1}(\mathbb{T}^d)$, with $s > d/2 + k > d$, $k \in \mathbb{N}$, $k \geq 1$.

Output: $u_N \in H_N^1(\mathbb{T}^d)$, such that $u_N \approx u$, where $u \in H^1(\mathbb{T}^d)$ solves (3.5) with coefficient field a and right-hand side f .

- (1) Compute pseudo-spectral projections $\tilde{a}_N, f_N \in L_N^2(\mathbb{T}^d)$, defined via the values on “de-aliased” grid $\{x_j\}_{j \in \mathcal{J}_{2N}}$ (cp. (3.8)):

$$\left\{ \begin{array}{l} \tilde{a}_N := \dot{P}_N \mathcal{I}_{2N}(a - 1), \\ f_N := \dot{P}_N \mathcal{I}_{2N} f \end{array} \right\} \in L_N^2(\mathbb{T}^d).$$

- (2) Set

$$K := \left\lceil \frac{\log(\lambda^{-1} N^{-k})}{\log(1 - \lambda/2)} \right\rceil.$$

- (3) Set $u_N^0 := 0 \in L_N^2(\mathbb{T}^d)$.

- (4) For $k = 1, 2, \dots, K$: compute

$$u_N^k := \dot{P}_N(-\Delta)^{-1} \nabla \cdot (\tilde{a}_N \nabla u_N^{k-1}) + (-\Delta)^{-1} f_N.$$

- (5) Set $u_N := u_N^K$.

We have the following theorem on the convergence of the algorithm 3.2,

Theorem 3.3. Let u be the unique solution of (3.5). Let $k \in \mathbb{N}$ be given, with $k > d/2 + 1$, and assume that the coefficient $a \in H^s(\mathbb{T}^d)$ for $s > d/2 + k$ satisfies the coercivity condition (3.6). Assume furthermore that $f \in H^{k-1}(\mathbb{T}^d)$. Then there exists $N_0 = N_0(s, d, \|a\|_{H^s}, \lambda) \in \mathbb{N}$, such that for any integer $N \geq N_0$, there exists a unique solution of the discretized elliptic equation (3.7), and there exists a constant $C = C(s, d, \lambda, \|a\|_{H^s}, \|f\|_{H^{k-1}}) > 0$, such that

$$\|u - u_N\|_{H^1} \leq CN^{-k},$$

where u_N is the output of Algorithm 3.2.

The proof of Theorem 3.3 relies only on standard techniques of numerical analysis and is presented in appendix E.1.

3.1.2. Ψ -FNO approximation of the Darcy equations. We will emulate the pseudo-spectral fixed point algorithm 3.2 by a Ψ -FNO, allowing us to derive approximation bounds. We consider the following setting,

Setting 3.4. Let $s \geq d/2 + k + \delta$ for some $k \in \mathbb{N}$, $\delta \in (0, 1)$ be a given Sobolev regularity parameter, and let $\lambda \in (0, 1)$ be a (coercivity) constant. Denote by $\mathcal{A}_\lambda^s(\mathbb{T}^d) \subset H^s(\mathbb{T}^d)$ the set of λ -coercive coefficients $a \in H^s(\mathbb{T}^d)$ of the form $a = 1 + \tilde{a}$, and such that

$$\|a\|_{H^s} \leq \lambda^{-1}, \quad \|\tilde{a}\|_{L^\infty} \leq C \|\tilde{a}\|_{H^{d/2+\delta}} \leq 1 - \lambda.$$

Here $C > 0$ denotes the norm of the embedding $H^{d/2+\delta}(\mathbb{T}^d) \hookrightarrow L^\infty(\mathbb{T}^d)$. We consider the operator $\mathcal{G} : \mathcal{A}_\lambda^s(\mathbb{T}^d) \rightarrow \dot{H}^1(\mathbb{T}^d)$, $a \mapsto u$, where u solves the Darcy equation

$$-\nabla \cdot (a \nabla u) = f, \quad \int_{\mathbb{T}^d} u(x) dx = 0,$$

on the periodic torus \mathbb{T}^d , with right-hand side $f \in \dot{H}^{k-1}$.

With this setting in place, we can now state our main FNO approximation theorem for the solution operator $\mathcal{G} : \mathcal{A}_\lambda^s(\mathbb{T}^d) \rightarrow H^1(\mathbb{T}^d)$ of the Darcy problem:

Theorem 3.5. Assume Setting 3.4, and assume that the activation function σ is non-linear and $\sigma \in C^3(\mathbb{R})$. Then there exists a constant $C = C(s, \lambda, d) > 0$, such that for any $N \in \mathbb{N}$, there exists a Ψ -FNO $\mathcal{N} : \mathcal{A}_\lambda^s(\mathbb{T}^d) \rightarrow H^1(\mathbb{T}^d)$, such that

$$\sup_{a \in \mathcal{A}_\lambda^s} \|\mathcal{G}(a) - \mathcal{N}(a)\|_{H^1(\mathbb{T}^d)} \leq CN^{-k}$$

and

$$\text{width}(\mathcal{N}) \leq CN^d, \quad \text{depth}(\mathcal{N}) \leq C \log(N), \quad \text{lift}(\mathcal{N}) \leq C.$$

In particular, we have

$$\text{size}(\mathcal{N}) \lesssim N^d \log(N).$$

The proof of this theorem, presented in Appendix E.3, relies crucially on the efficient approximation of quadratic nonlinearities by shallow neural networks with smooth activation functions, see Lemma E.3 in Appendix E for the precise statement and proof.

Remark 3.6. To achieve a FNO approximation error of order ϵ for the Darcy problem, Theorem 3.5 shows that a Ψ -FNO \mathcal{N} with

$$\text{size}(\mathcal{N}) \lesssim \left(\frac{1}{\epsilon}\right)^{d/k} \log(\epsilon^{-1}), \quad \text{depth}(\mathcal{N}) \lesssim \log(\epsilon^{-1}), \quad (3.10)$$

is sufficient. Furthermore, the lifting dimension d_v can be kept uniformly bounded, independently of ϵ . In particular, for $k > d$, the required total size of the Ψ -FNO scales sub-log-linearly in the approximation accuracy ϵ , indicating that Ψ -FNOs may provide a very efficient approximation to the solution operator of the Darcy problem, in this case.

As a concrete example for illustrating the approximation of Darcy equations by FNOs, we consider the following example.

Example 3.7. A possible model for coefficients a with a typical length scale $\ell > 0$, is to assume an expansion of $a = a(x; Y)$ in terms of random variables $Y = (Y_1, Y_2, \dots) \in [-1, 1]^{\mathbb{N}}$ (not necessarily i.i.d.), of the following form, similar to the ansatz in [32]:

$$a(x; Y) = 1 + \sum_{k \in \mathbb{Z}^d \setminus \{0\}} b_k Y_k e^{i\langle k, x \rangle},$$

where we assume that the coefficients b_k satisfy a decay condition of the form $|b_k| \leq C_b \exp(-\ell|k|)$ for a constant $C_b > 0$. We also assume that $a(x; Y)$ satisfy the coercivity condition

$$\|a(\cdot; Y) - 1\|_{L^\infty} \leq C \|a(\cdot; Y) - 1\|_{H^{d/2+\delta}(\mathbb{T}^d)} \leq 1 - \lambda,$$

for some $\delta, \lambda \in (0, 1)$, and that the source term $f \in C^\infty(\mathbb{T}^d)$. Under these assumptions, we have $a \in H^s$ for any $s > 0$, and hence it follows from Theorem 3.5 that for any $r \in \mathbb{N}$, there exists a constant $C = C(r, \ell, \|f\|_{C^k}, \delta)$, with the following property (cp. also Remark 3.6): For any $\epsilon > 0$, there exists a Ψ -FNO \mathcal{N} , such that

$$\sup_{Y \in [-1, 1]^{\mathbb{N}}} \|\mathcal{G}(a(\cdot, Y)) - \mathcal{N}(a(\cdot, Y))\|_{H^1} < \epsilon,$$

and

$$\text{size}(\mathcal{N}) \leq C\epsilon^{-1/r}, \quad \text{depth}(\mathcal{N}) \leq C \log(\epsilon^{-1}). \quad (3.11)$$

The complexity bounds (3.10) and (3.11) suggest that the size of a Ψ -FNO approximating the operator \mathcal{G} stemming from the Darcy equations, scales only sub-linearly (or even better) in the desired accuracy ϵ . This should be contrasted with the fact that the size of a Ψ -FNO for a generic Lipschitz operator grows super-exponentially in the desired error (3.1). Thus, we are able to show that a Ψ -FNO can approximate this PDE solution operator much more efficiently than it might a generic infinite-dimensional operator.

3.2. Incompressible Navier-Stokes equations. The motion of a viscous, incompressible Newtonian fluid is modeled by the incompressible Navier-Stokes equations,

$$\begin{cases} \partial_t u + u \cdot \nabla u + \nabla p = \nu \Delta u, \\ \text{div}(u) = 0, \quad u(t=0) = u_0, \end{cases} \quad (3.12)$$

Here, $u \in \mathbb{R}^d$ is the fluid velocity and $p \in \mathbb{R}$ is the pressure of the fluid, acting as a Lagrange multiplier to enforce the divergence-free constraint $\text{div}(u) = 0$. The initial fluid velocity is denoted by u_0 . For simplicity, we assume periodic boundary conditions in the domain \mathbb{T}^d . The viscosity is denoted by $\nu \geq 0$ and we would like to state that the subsequent analysis also applies for $\nu = 0$, where (3.12) reduces to the incompressible Euler equations modeling an ideal fluid. For definiteness, we recall the following well-known theorem for the well-posedness of the Navier-Stokes equations (3.12),

Theorem 3.8 (see e.g. [25, Thm 3.4]). Let $r > d/2 + 2$. For any $u_0 \in H^r$, there exists $T > 0$ and a unique classical solution of the Navier-Stokes equations (3.12), such that $u \in C([0, T]; H^r) \cap C^1([0, T]; H^{r-2})$ with $u(t=0) = u_0$.

It is well known that in two space dimensions $d = 2$, the time interval for existence of solutions $[0, T]$ can be extended to any finite T as long as $u_0 \in H^r$, whereas the corresponding finite-time well-posedness result for three space-dimensions is an outstanding open problem [25].

We recall that if the initial data u_0 of (3.12) belongs to $\dot{L}^2(\mathbb{T}^d; \mathbb{R}^d)$, i.e. if

$$\int_{\mathbb{T}^d} u_0(x) dx = 0,$$

then we also have that the corresponding solution $u(x, t) \in \dot{L}^2(\mathbb{T}^d; \mathbb{R}^d)$ (reflecting momentum conservation). Next, we introduce the **Leray-projection operator** $\mathbb{P} : L^2(\mathbb{T}^d; \mathbb{R}^d) \rightarrow \dot{L}^2(\mathbb{T}^d; \text{div})$, as the L^2 -orthogonal projection onto the subspace $\dot{L}^2(\mathbb{T}^d; \text{div}) \subset \dot{L}^2(\mathbb{T}^d; \mathbb{R}^d)$, consisting of divergence-free vector fields; i.e. we have $u \in \dot{L}^2(\mathbb{T}^d; \text{div})$ if, and only if, $u \in \dot{L}^2(\mathbb{T}^d; \mathbb{R}^d)$ and

$$\int_{\mathbb{T}^d} u(x) \cdot \nabla \phi(x) dx = 0, \quad \forall \phi \in C^\infty(\mathbb{T}^d).$$

In terms of Fourier series, the Leray projection $\mathbb{P} : \dot{L}^2(\mathbb{T}^d; \mathbb{R}^d) \rightarrow \dot{L}^2(\mathbb{T}^d; \text{div})$ is explicitly given by

$$\mathbb{P} \left(\sum_{k \in \mathbb{Z}^d} \hat{u}_k e^{i(k, x)} \right) = \sum_{k \in \mathbb{Z}^d \setminus \{0\}} \left(1 - \frac{k \otimes k}{|k|^2} \right) \hat{u}_k e^{i(k, x)}. \quad (3.13)$$

In terms of the Leray projection \mathbb{P} , we can now equivalently write the incompressible Navier-Stokes equations (3.12) as the following equation on the Hilbert space $\dot{L}^2(\mathbb{T}^d; \text{div})$ as,

$$\begin{cases} \partial_t u = -\mathbb{P}(u \cdot \nabla u) + \nu \Delta u, \\ u(t=0) = u_0. \end{cases} \quad (3.14)$$

Given this background, our main objective in this section is to construct a Ψ -FNO that will approximate the operator \mathcal{G} which maps the initial data u_0 to the solution $u(\cdot, T)$ (at the final time T) of the incompressible Navier-Stokes equations (3.12), (3.14). To this end, we will follow the general program outlined at the beginning of this section and introduce a suitable pseudo-spectral method for approximating the Navier-Stokes equations. Then, we construct a Ψ -FNO that can efficiently emulate this pseudo-spectral method.

3.2.1. *A fully-discrete Ψ -spectral approximation of the Navier-Stokes equations (3.12).* The form of the Leray-projected Navier-Stokes equations (3.14) naturally suggests the following fully-discrete approximation of (3.12):

$$\begin{cases} \frac{u_N^{n+1} - u_N^n}{\tau} + \mathbb{P}_N(u_N^n \cdot \nabla u_N^{n+1}) = \nu \Delta u_N^{n+1}, \\ u_N^0 = \mathcal{I}_N u(t=0). \end{cases} \quad (3.15)$$

Here, we fix $N \in \mathbb{N}$ and introduce the space, $\dot{L}_N^2(\mathbb{T}^d; \text{div}) := \dot{L}^2(\mathbb{T}^d; \text{div}) \cap \dot{L}_N^2(\mathbb{T}^d; \mathbb{R}^d)$. We fix a time-step $\tau > 0$ and let $u_N^n \in \dot{L}_N^2(\mathbb{T}^d; \text{div})$, for all $n = 0, \dots, n_T$, with n_T such that $\tau n_T = T$. Moreover, we use the following *finite-dimensional* Leray-Fourier projection operator $\mathbb{P}_N : \dot{L}^2(\mathbb{T}^d; \mathbb{R}^d) \rightarrow \dot{L}_N^2(\mathbb{T}^d; \text{div})$ in analogy with (3.13):

$$\mathbb{P}_N \left(\sum_{k \in \mathbb{Z}^d} \hat{u}_k e^{i(k, x)} \right) := \sum_{0 < |k|_\infty \leq N} \left(1 - \frac{k \otimes k}{|k|^2} \right) \hat{u}_k e^{i(k, x)}, \quad (3.16)$$

to complete the description of the scheme (3.15).

We observe that the scheme (3.15) is *implicit* i.e., at each time step n , one has to solve an operator equation to compute the velocity field u_N^{n+1} at the next time step. Thus, one needs to show the solvability of this operator equation in order to

ensure that the scheme (3.15) is *well-defined*. Under the following CFL condition for choosing a small enough time step τ ,

$$\tau \|u_N^n\|_{L^\infty} N \leq \frac{1}{2}, \quad (3.17)$$

we prove in Appendix F.1 that the scheme (3.15) is well-defined.

Next, in practice, one has to numerically approximate the solutions of the implicit equation (3.15) for evaluating the velocity field u_N^{n+1} , at the next time-step. We choose to do so by recasting the solution of the implicit equation (3.15) to finding a *fixed point* for the mapping,

$$w_N \mapsto F(w_N) := (1 - \nu\tau\Delta)^{-1}u_N^n - \tau(1 - \nu\tau\Delta)^{-1}\mathbb{P}_N(u_N^n \cdot \nabla w_N). \quad (3.18)$$

In Appendix F.1 Lemma F.4, we show that a standard Picard-type iteration converges to a fixed point for the map (3.18). This suggests the following numerical algorithm for approximating strong solutions of the incompressible Navier-Stokes equations (3.12),

Algorithm 3.9 (Pseudo-spectral approximation of the Navier-Stokes equations (3.12)).

Input: $U > 0$, $N \in \mathbb{N}$, $T > 0$, a time-step $\tau > 0$, such that $n_T = T/\tau \in \mathbb{N}$, and $\tau UN^{d/2+1} \leq \frac{1}{2e}$, initial data $u_N^0 \in L_N^2(\mathbb{T}^d; \text{div})$, such that $\|u_N^0\|_{L^2} \leq U$.
Output: $u_N^{n_T} \in L_N^2(\mathbb{T}^d; \text{div})$ an approximation of the solution $u_N^{n_T} \approx u(t = T)$ of (3.12) at time $t = T$.

(1) Set

$$\kappa_0 := \left\lceil \frac{\log(T^2/\tau^2)}{\log(2)} \right\rceil \in \mathbb{N}.$$

(2) For $n = 0, \dots, n_T - 1$:

- (a) Set $w_N^{n,0} := 0$,
(b) For $k = 1, \dots, \kappa_0$: Compute

$$w_N^{n,k} := (1 - \nu\tau\Delta)^{-1}u_N^n - \tau(1 - \nu\tau\Delta)^{-1}\mathbb{P}_N(u_N^n \cdot \nabla w_N^{n,k-1}),$$

- (c) Set $u_N^{n+1} := w_N^{n,\kappa_0}$,

The convergence of the algorithm 3.9, together with a convergence rate, to the strong solution of the Navier-Stokes equations is summarized in the following theorem,

Theorem 3.10. Let $U, T > 0$. Consider the Navier-Stokes equations on \mathbb{T}^d , for $d \geq 2$. Assume that $r \geq d/2 + 2$, and let $u \in C([0, T]; H^r) \cap C^1([0, T]; H^{r-2})$ be a solution of the Navier-Stokes equations (3.12), such that $\|u\|_{L^2} \leq U$. Choose a time-step τ , such that $\tau UN^{d/2+1} \leq (2e)^{-1}$. There exists a constant

$$C = C(T, d, r, \|u\|_{C_t(H_x^r)}, \|u\|_{C_t^1(H_x^{r-2})}) > 0,$$

such that with $u_N^0 := \mathcal{I}_N u(0)$, and for the sequence $u_N^1, \dots, u_N^{n_T} \in L_N^2(\mathbb{T}^d; \text{div})$ generated by Algorithm 3.9, we have

$$\max_{n=0, \dots, n_T} \|u_N^n - u(t^n)\|_{L^2} \leq C(\tau + N^{-r}),$$

where $n_T \tau = T$. In particular, choosing $\tau \sim N^{-r}$, we have

$$\max_{n=0, \dots, n_T} \|u_N^n - u(t^n)\|_{L^2} \leq CN^{-r},$$

with $n_T \sim N^r$ (and enlarging the constant $C > 0$ by a constant factor).

The proof of this theorem relies on several techniques from numerical analysis and is presented in detail in Appendix F.2.

3.2.2. Approximation of algorithm 3.9 by Ψ -FNOs. Next, we are going to construct a Ψ -FNO of the form (2.14), which can efficiently emulate the pseudo-spectral algorithm 3.9. To this end, we have the following result (proved in Appendix F.3) on the efficient approximation of the non-linear term in the Navier-Stokes equations by FNOs,

Lemma 3.11. Assume that the activation function $\sigma \in C^3$ is three times continuously differentiable and non-linear. There exists a constant $C > 0$, such that for any $N \in \mathbb{N}$, and for any $\epsilon, B > 0$, there exists a Ψ -FNO $\mathcal{N} : L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d) \times L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d) \rightarrow L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d)$, with

$$\text{depth}(\mathcal{N}), \text{lift}(\mathcal{N}) \leq C, \quad \text{width}(\mathcal{N}) \leq CN^d,$$

such that we have

$$\|\mathbb{P}_N(u_N \cdot \nabla w_N) - \mathcal{N}(u_N, w_N)\|_{L_N^2} \leq \epsilon,$$

for all trigonometric polynomials $u_N, w_N \in L_N^2(\mathbb{T}^d; \mathbb{R}^d) \subset L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d)$ of degree $|k|_\infty \leq N$, satisfying the bound $\|u_N\|_{L^2}, \|w_N\|_{L^2} \leq B$.

Thus, from the preceding Lemma, we have that the nonlinearities in algorithm 3.9 can be efficiently approximated by Ψ -FNOs. This paves the way for the following theorem on the emulation of the pseudo-spectral algorithm 3.9 by Ψ -FNOs,

Theorem 3.12. Let $U, T > 0$ and viscosity $\nu \geq 0$. Consider the Navier-Stokes equations on \mathbb{T}^d , for $d \geq 2$. Assume that $r \geq d/2 + 2$, and let $\mathcal{V} \subset C([0, T]; H^r) \cap C^1([0, T]; H^{r-2})$ be a set of solutions of the Navier-Stokes equations (3.12), such that $\sup_{u \in \mathcal{V}} \|u\|_{L^2} \leq U$, and

$$\bar{U} := \sup_{u \in \mathcal{V}} \left\{ \|u\|_{C_t(H_x^r)} + \|u\|_{C_t^1(H_x^{r-2})} \right\} < \infty.$$

For $t \in [0, T]$, denote $\mathcal{V}_t := \{u(t) \mid u \in \mathcal{V}\}$. Let $\mathcal{G} : \mathcal{V}_0 \rightarrow \mathcal{V}_T$ denote the solution operator of (3.12), mapping initial data $u_0 = u(t=0)$, to the solution $u(T)$ at $t = T$ of the incompressible Navier-Stokes equations. There exists a constant

$$C = C(d, r, U, \bar{U}, T) > 0,$$

such that for $N \in \mathbb{N}$ there exists a Ψ -FNO $\mathcal{N} : L_N^2(\mathbb{T}^d; \mathbb{R}^d) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^d)$, such that

$$\sup_{u \in \mathcal{V}_0} \|\mathcal{G}(u) - \mathcal{N}(u)\|_{L^2} \leq CN^{-r},$$

and such that

$$\text{width}(\mathcal{N}) \leq CN^d, \quad \text{depth}(\mathcal{N}) \leq CN^r \log(N), \quad \text{lift}(\mathcal{N}) \leq C.$$

The proof of this theorem is provided in Appendix F.3.

Remark 3.13. It is straightforward to observe from Theorem 3.12 that the size of a Ψ -FNO to achieve a desired error tolerance of $\epsilon > 0$, scales (neglecting log-terms) as

$$\text{size}(\mathcal{N}) \leq C\epsilon^{-(1+\frac{d}{r})}, \quad (3.19)$$

Given that we need $r \geq d/2 + 2$, we observe from (3.19) that the size of the Ψ -FNO, approximating the initial data to solution operator \mathcal{G} , for the Navier-Stokes equations (3.12), scales at most *sub-quadratically* with respect to the error tolerance ϵ for the physically relevant values $d = 2, 3$. This polynomial scaling should be compared with the super-exponential growth (see Remark 3.1) of the size

of FNOs in approximating a generic Lipschitz-continuous operator. Thus, we are able to demonstrate that Ψ -FNOs can approximate the solutions of Navier-Stokes equations far more efficiently than what the universal approximation theorem 2.15 suggests.

Remark 3.14. From the convergence theorem 3.10, we observe that the underlying scheme (3.15) is *first-order* in time. This low accuracy of the scheme necessitates a large number of time steps and affects the overall complexity. We describe a second-order accurate time discretized version of the pseudo-spectral method for approximating the Navier-Stokes equations (3.12) in Appendix F.4 and in complete analogy with Theorem 3.12 (see Theorem F.8 in Appendix F.4), we can construct a Ψ -FNO to emulate this second-order in time pseudo-spectral scheme, resulting in a Ψ -FNO of

$$\text{size}(\mathcal{N}) \leq C\epsilon^{-\left(\frac{1}{2} + \frac{d}{r}\right)}, \quad (3.20)$$

to obtain a desired accuracy of ϵ . Thus, we can obtain a more efficient approximation of the underlying operator than Ψ -FNO emulating the first-order time scheme (3.15). In particular for $r \geq 2d$, we obtain that the size of a Ψ -FNO only grows sub-linearly in terms of the desired accuracy, making this FNO approximation comparable in complexity to the FNO approximation of the Darcy equation (see (3.10)).

4. COMPARISON OF FNOs WITH DEEPONETS

In this section, we will compare Ψ -FNOs with another operator learning framework, namely DeepOnets of [8, 21], defined as,

Definition 4.1. Fix $m, p \in \mathbb{N}$. A **DeepOnet** \mathcal{N} with output dimension p and **sensor points** $x_1, \dots, x_m \in \mathbb{T}^d$ is a mapping $\mathcal{N} : C(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow C(\mathbb{T}^d; \mathbb{R}^{d_u})$ of the form

$$\mathcal{N}(a)(x) = \sum_{k=1}^p \beta_k(a(x_1), \dots, a(x_p)) \tau_k(x),$$

where $\beta : \mathbb{R}^{m \times d_a} \rightarrow \mathbb{R}^{p \times d_u}$, $\alpha \mapsto \beta(\alpha) = (\beta_1(\alpha), \dots, \beta_p(\alpha))$, and $\tau : \mathbb{R}^d \mapsto \mathbb{R}^p$, $x \mapsto \tau(x) = (\tau_1(x), \dots, \tau_p(x))$ are (ordinary) neural networks. We refer to β and τ as the **branch** and **trunk** nets, respectively.

In the following theorem (proved in Appendix G), we show that a Ψ -FNO can naturally be viewed as a DeepOnet with a specific choice of the branch and trunk-nets; where the trunk net is fixed to represent a trigonometric basis, and the branch net is constrained by a specific choice of hidden layer architecture, which provides a more parsimonious representation compared to a canonical DeepOnet implementation based on dense layers.

Theorem 4.2 (DeepOnet approximation of Ψ -FNOs). Let $\mathcal{N} : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_u})$ be a Ψ -FNO, and fix $B > 0$. For any $\epsilon > 0$, there exists $p \in \mathbb{N}$ and a DeepOnet (β, τ) , with branch net β , trunk net τ , and sensor points $\{x_j\}_{j \in \mathcal{J}_N}$, such that

$$\sup_{\|a\|_{L^\infty} \leq B} \sup_{y \in \mathbb{T}^d} \left| \mathcal{N}(a)(y) - \sum_{k=1}^p \beta_k(a) \tau_k(y) \right| \leq \epsilon,$$

where the first supremum is taken over all $a \in C(\mathbb{T}^d)$, such that $\|a\|_{L^\infty} \leq B$. Furthermore, we have

$$\text{width}(\beta) = \text{width}(\mathcal{N}), \quad \text{depth}(\beta) = \text{depth}(\mathcal{N}),$$

and the trunk net τ defines a mapping $\tau : \mathbb{R}^d \rightarrow \mathbb{R}^{\mathcal{K}_N}$, which approximates an (arbitrary) orthonormal trigonometric basis $\{\mathbf{e}_k\}_{k \in \mathcal{K}_N}$ with $\text{span}\{\mathbf{e}_k\}_{k \in \mathcal{K}_N} = \text{span}\{e^{i\langle k, x \rangle}\}_{k \in \mathcal{K}_N}$, such that

$$\max_{k \in \mathcal{K}_N} \|\mathbf{e}_k - \tau_k\|_{L^\infty} \leq \epsilon/\bar{B},$$

where

$$\bar{B} := (2N + 1)^d \left(\sup_{\|a\|_{L^\infty} \leq B} \|\mathcal{N}(a)\|_{L^2} \right).$$

Remark 4.3. We note that due to the particular architecture of Ψ -Fourier neural operators, the total size of a Ψ -FNO \mathcal{N} is upper bounded by

$$\text{size}(\mathcal{N}) \lesssim \text{depth}(\mathcal{N})\text{width}(\mathcal{N})\text{lift}(\mathcal{N}).$$

In contrast, since β and τ in the DeepOnet approximation of Theorem 4.2 are conventional deep neural networks, the total number of degrees of freedom (i.e. the total number of weights and biases) is given, in general, by

$$\text{size}(\beta) \sim \text{depth}(\beta)\text{width}(\beta)^2 = \text{depth}(\mathcal{N})\text{width}(\mathcal{N})^2,$$

and $\text{size}(\tau) \sim \text{depth}(\tau)\text{width}(\tau)^2$. So that in particular, comparing the sizes of the Ψ -FNO and the corresponding (fully connected) DeepOnet approximation, we have

$$\frac{\text{size}(\mathcal{N})}{\text{size}(\beta, \tau)} \lesssim \frac{\text{lift}(\mathcal{N})}{\text{width}(\mathcal{N})}.$$

Based on our explicit complexity estimates (3.10) and (3.19), we can expect that $\text{lift}(\mathcal{N}) \ll \text{width}(\mathcal{N})$ for a wide range of problems which can be efficiently approximated by FNOs. In particular, this indicates that the Ψ -FNO provides a more parsimonious approximation of the underlying operator than the DeepOnet “emulator” constructed in Theorem 4.2, with $\text{size}(\mathcal{N}) \ll \text{size}(\beta, \tau)$.

5. SUMMARY AND DISCUSSION.

Many learning tasks, particularly, but not exclusively, in scientific computing, are naturally formulated as learning operators mapping one infinite-dimensional space to another. Neural operators have recently been proposed as a framework for operator learning. A particular form, the so-called Fourier Neural Operators (FNOs) (2.6), have been shown to be efficient in approximating a wide variety of operators that arise in PDEs [19]. Our main aim in this paper was to analyze FNOs and Ψ -FNOs (2.14), which is a concrete computational realization of FNOs. To this end, we have presented the following results,

- We showed in Theorem 2.5 and Theorem 2.15 that FNOs (resp. Ψ -FNOs) are *universal* i.e., they can approximate any continuous operator to desired accuracy. Our proof relies heavily on the ability of FNOs to approximate the Fourier transform and its inverse, together with the neural network approximation of the finite-dimensional Fourier conjugate operator (2.11). Thus, FNOs have the same universal approximation property as canonical neural networks for finite-dimensional functions and DeepOnets for operators [15]. This universality result paves the way for the widespread use of FNOs in the context of operator learning.
- However as stated in remark 3.1, in the worst case, the size of a FNO can grow *super-exponentially* in terms of the desired error for approximating a general Lipschitz continuous operator. This might inhibit the use of FNOs. On the other hand, we argue in the beginning of section 3 that Ψ -FNOs, which are a concrete computational realization of FNOs, can approximate the nonlinearities and differential operators that define PDEs,

very efficiently. Hence, one can think of Ψ -FNOs as a new form of pseudo-spectral methods for PDEs, which in practice are adapted to, and optimized based on the given training data. Thus, one can expect that Ψ -FNOs can approximate PDEs efficiently.

- We consider two widely used prototypical PDEs, namely the elliptic PDE (3.5) that arises in a stationary Darcy flow and the well-known incompressible Navier-Stokes equations for fluid dynamics. For both these PDEs, we prove rigorously that there exists a Ψ -FNO which can approximate the underlying nonlinear operators efficiently, as we can show that the size of the Ψ -FNO only needs to grow polynomially in terms of the error. In fact, we show that the size grows sub-linearly in terms of the error. Thus, FNOs can approximate these widely used PDEs efficiently, corroborating the empirical results presented in [19].

Hence, our analysis provides very strong theoretical evidence that FNOs are an effective framework for operator learning. Moreover, we also compare FNOs to an alternative operator learning framework, that of DeepOnets [21] and show that Ψ -FNOs can be thought of a special case of DeepOnets with a trunk net approximating trigonometric functions and sensor points being equi-distributed Cartesian grid points. Given its special architecture, we argue that a Ψ -FNO can allow for a more parsimonious representation of operators than a DeepOnet, enabling a cheaper approximation of certain operators.

The comparison with DeepOnets also brings out some obvious limitations of FNOs. In particular, FNOs are efficient on rectangular domains as the Ψ -FNO can be evaluated efficiently with FFT. Although one can use FNOs for operators defined on arbitrary domains using suitable extension and restriction operators (see Theorem 2.9), it is unclear if these operators can be realized computationally in an efficient manner. Moreover, FNOs fix the trigonometric basis as the trunk net in the underlying DeepOnet (see Theorem 4.2). On the other hand, a general DeepOnet can learn trunk nets from the data during training, allowing the possibility of learning a more suitable representation from data. These considerations call for a more thorough computational comparison between DeepOnets and FNOs, and possibly other operator learning frameworks such as the one from [2].

Similarly, extending the analysis of this article to other neural operators can be readily envisaged. The use of FNOs for more general operators, particularly those arising in non-scientific computing settings, such as images, text and speech, also needs to be investigated.

REFERENCES

- [1] A. R. BARRON, *Universal approximation bounds for superpositions of a sigmoidal function*, IEEE Trans. Inform. Theory., 39 (1993), pp. 930–945.
- [2] K. BHATTACHARYA, B. HOSSEINI, N. B. KOVACHKI, AND A. M. STUART, *Model reduction and neural networks for parametric PDEs*, The SMAI journal of computational mathematics, 7 (2021), pp. 121–157.
- [3] J. P. BOYD, *Fourier embedded domain methods: Extending a function defined on an irregular region to a rectangle so that the extension is spatially periodic and C infinity*, Appl. Math. Comput., 161 (2005), p. 591–597.
- [4] H. BREZIS AND P. MIRONESCU, *Composition in fractional Sobolev spaces*, Discrete & Continuous Dynamical Systems-A, 7 (2001), p. 241.
- [5] O. P. BRUNO AND M. LYON, *High-order unconditionally stable FC-AD solvers for general smooth domains I. Basic elements*, Journal of Computational Physics, 229 (2010), pp. 2009–2033.
- [6] S. CAI, Z. WANG, L. LU, T. A. ZAKI, AND G. E. KARNIADAKIS, *Deepm&mnnet: Inferring the electroconvection multiphysics fields based on operator approximation by neural networks*, Journal of Computational Physics, 436 (2021), p. 110296.

- [7] C. CANUTO, M. Y. HUSSAINI, A. QUARTERONI, AND T. A. ZANG, *Spectral methods: fundamentals in single domains*, Springer Science & Business Media, 2007.
- [8] T. CHEN AND H. CHEN, *Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems*, IEEE Transactions on Neural Networks, 6 (1995), pp. 911–917.
- [9] G. CYBENKO, *Approximations by superpositions of sigmoidal functions*, Approximation theory and its applications, 9 (1989), pp. 17–28.
- [10] T. DERUYCK, S. LANTHALER, AND S. MISHRA, *On the approximation of functions by tanh neural networks*. Preprint, available from arXiv:2104:08938v1, 2021.
- [11] W. E, J. HAN, AND A. JENTZEN, *Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations*, Communications in Mathematics and Statistics, 5 (2017), pp. 349–380.
- [12] I. GOODFELLOW, Y. BENGIO, AND A. COURVILLE, *Deep learning*, MIT press, 2016.
- [13] K. HORNIK, M. STINCHCOMBE, AND H. WHITE, *Multilayer feedforward networks are universal approximators*, Neural networks, 2 (1989), pp. 359–366.
- [14] G. KUTYNIOK, P. PETERSEN, M. RASLAN, AND R. SCHNEIDER, *A theoretical analysis of deep neural networks and parametric pdes*, Constructive Approximation, (2021), pp. 1–53.
- [15] S. LANTHALER, S. MISHRA, AND G. E. KARNIADAKIS, *Error estimates for DeepOnets: A deep learning framework in infinite dimensions*, 2021.
- [16] Y. LECUN, Y. BENGIO, AND G. HINTON, *Deep learning*, Nature, 521 (2015), pp. 436–444.
- [17] Z. LI, N. KOVACHKI, K. AZIZZADENESHELI, B. LIU, K. BHATTACHARYA, A. STUART, AND A. ANANDKUMAR, *Multipole graph neural operator for parametric partial differential equations*, 2020.
- [18] ———, *Neural operator: Graph kernel network for partial differential equations*, 2020.
- [19] Z. LI, N. B. KOVACHKI, K. AZIZZADENESHELI, B. LIU, K. BHATTACHARYA, A. STUART, AND A. ANANDKUMAR, *Fourier neural operator for parametric partial differential equations*, in International Conference on Learning Representations, 2021.
- [20] C. LIN, Z. LI, L. LU, S. CAI, M. MAXEY, AND G. E. KARNIADAKIS, *Operator learning for predicting multiscale bubble growth dynamics*, The Journal of Chemical Physics, 154 (2021), p. 104118.
- [21] L. LU, P. JIN, AND G. E. KARNIADAKIS, *DeepONet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators*, arXiv preprint arXiv:1910.03193, (2019).
- [22] K. O. LYE, S. MISHRA, AND D. RAY, *Deep learning observables in computational fluid dynamics*, Journal of Computational Physics, (2020), p. 109339.
- [23] K. O. LYE, S. MISHRA, D. RAY, AND P. CHANDRASHEKAR, *Iterative surrogate model optimization (ismo): An active learning algorithm for pde constrained optimization with deep neural networks*, Computer Methods in Applied Mechanics and Engineering, 374 (2021), p. 113575.
- [24] M. LYON AND O. P. BRUNO, *High-order unconditionally stable FC-AD solvers for general smooth domains II. Elliptic, parabolic and hyperbolic PDEs; theoretical considerations*, Journal of Computational Physics, 229 (2010), pp. 3358–3381.
- [25] A. J. MAJDA AND A. L. BERTOZZI, *Vorticity and Incompressible Flow*, Cambridge Texts in Applied Mathematics, Cambridge University Press, 2001.
- [26] Z. MAO, L. LU, O. MARXEN, T. ZAKI, AND G. E. KARNIADAKIS, *DeepMandMnet for hyperionics: Predicting the coupled flow and finite-rate chemistry behind a normal shock using neural-network approximation of operators*. Preprint, available from arXiv:2011.03349v1, 2020.
- [27] S. MISHRA AND R. MOLINARO, *Estimates on the generalization error of physics informed neural networks (PINNs) for approximating PDEs*. Preprint, available from arXiv:2006:16144v1, 2020.
- [28] ———, *Estimates on the generalization error of physics informed neural networks (PINNs) for approximating PDEs II: A class of inverse problems*. Preprint, available from arXiv:2007:01138v1, 2020.
- [29] M. RAISSI AND G. E. KARNIADAKIS, *Hidden physics models: Machine learning of nonlinear partial differential equations*, Journal of Computational Physics, 357 (2018), pp. 125–141.
- [30] M. RAISSI, P. PERDIKARIS, AND G. E. KARNIADAKIS, *Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations*, Journal of Computational Physics, 378 (2019), pp. 686–707.
- [31] L. G. ROGERS, *Degree-independent Sobolev extension on locally uniform domains*, Journal of Functional Analysis, 235 (2006), pp. 619–665.
- [32] C. SCHWAB AND J. ZECH, *Deep learning in high dimension: Neural network expression rates for generalized polynomial chaos expansions in UQ*, Analysis and Applications, 17 (2019), pp. 19–55.

- [33] E. M. STEIN, *Singular Integrals and Differentiability Properties of Functions*, Princeton University Press, 1970.
- [34] D. YAROTSKY, *Error bounds for approximations with deep ReLU networks*, Neural Networks, 94 (2017), pp. 103–114.

APPENDIX A. GLOSSARY OF MATHEMATICAL NOTATION

Symbol	Description	Page
σ	activation function	
\mathbb{T}^d	periodic torus, identified with $[0, 2\pi]^d$	
d	spatial dimension of domain	p. 3
d_a, d_u, d_v	number of components of input, output and lifting	p. 3
$\mathcal{A}(D; \mathbb{R}^{d_a})$	input function space	p. 3
$\mathcal{U}(D; \mathbb{R}^{d_u})$	output function space	p. 3
$\mathcal{F}, \mathcal{F}^{-1}$	Fourier transform and inverse Fourier transform	p. 25
$\mathcal{F}_N, \mathcal{F}_N^{-1}$	discrete Fourier transform and inverse	p. 27
$\{x_j\}_{j \in \mathcal{J}_N}$	regular periodic grid, $x_j = 2\pi j / (2N + 1)$	p. 27
\mathcal{J}_N	grid point indices, $\mathcal{J}_N = \{0, \dots, 2N\}^d$	p. 27
\mathcal{K}_N	Fourier wavenumbers $\mathcal{K}_N = \{k \in \mathbb{Z}^d \mid k _\infty \leq N\}$	
\mathcal{R}	lifting operator	p. 3
\mathcal{L}_ℓ	neural operator layer	p. 5
\mathcal{Q}	projection operator	p. 3
\mathcal{F} -layer	linear, non-local layer; $v(x) \mapsto \mathcal{F}^{-1}(P\mathcal{F}v)(x)$	p. 11
σ -layer	non-linear, local layer; $v(x) \mapsto \sigma(Wv(x) + b(x))$	p. 11
Banach spaces		
L^2	Space of square-integrable functions	
\dot{L}^2	$\dot{L}^2 \subset L^2$ square-integrable functions with zero mean	p. 26
L_N^2	$L_N^2 \subset L^2$ trigonometric polynomials of degree $\leq N$	p. 26
\dot{L}_N^2	$\dot{L}_N^2 = \dot{L}^2 \cap L_N^2$ trigonometric polynomials with zero mean	p. 26
H^s	Sobolev space of smoothness s , with norm $\ \cdot\ _{H^s}$	p. 26
\dot{H}^s	Sobolev space with zero mean, with norm $\ \cdot\ _{\dot{H}^s}$	p. 26
Projection operators		
P_N	L^2 -orthogonal Fourier projection $P_N : L^2 \rightarrow L_N^2$	p. 26
\dot{P}_N	Fourier projection $\dot{P}_N : L^2 \rightarrow \dot{L}_N^2$ with zero mean	p. 26
\mathcal{I}_N	Pseudo-spectral Fourier projection, i.e. trigonometric interpolation on regular grid $\{x_j\}_{j \in \mathcal{J}_N}$	p. 27
\mathbb{P}	Leray projection onto divergence-free vector fields	p. 18
\mathbb{P}_N	Leray projection followed by projection P_N ; $\mathbb{P}_N = P_N \circ \mathbb{P}$	p. 18

APPENDIX B. NOTATION AND TECHNICAL PRELIMINARIES.

In this section, we introduce frequently used notation in the article main text and recall some essential facts about Fourier analysis.

In the main text, our focus is functions defined on the periodic torus \mathbb{T}^d , identified as $\mathbb{T}^d = [0, 2\pi]^d$. Following standard practice, we denote by $L^2(\mathbb{T}^d)$ the space of square-integrable functions. For any such function $v \in L^2(\mathbb{T}^d)$, we can define the *Fourier Transform* as,

$$\mathcal{F}(v)(k) := \frac{1}{(2\pi)^d} \int_{\mathbb{T}^d} v(x) e^{-i(k,x)} dx, \quad \forall k \in \mathbb{Z}^d. \quad (\text{B.1})$$

For any $k \in \mathbb{Z}^d$, the k -th Fourier coefficient of v is denoted by $\hat{v}_k = \mathcal{F}(v)(k)$.

Given a set of Fourier coefficients $\{\widehat{v}_k\}_{k \in \mathbb{Z}^d}$, the *inverse Fourier Transform* is defined as,

$$\mathcal{F}^{-1}(\widehat{v})(x) := \sum_{k \in \mathbb{Z}^d} \widehat{v}_k e^{i\langle k, x \rangle}, \quad \forall x \in \mathbb{T}^d. \quad (\text{B.2})$$

Using the Fourier transform (B.1) and for $s \geq 0$, one can denote by $H^s(\mathbb{T}^d)$ the **Sobolev space** of functions $v \in L^2(\mathbb{T}^d)$, with Fourier coefficients $\{\widehat{v}_k\}_{k \in \mathbb{Z}^d}$, having a finite H^s -norm:

$$\|v\|_{H^s}^2 := \frac{(2\pi)^d}{2} \sum_{k \in \mathbb{Z}^d} (1 + |k|^{2s}) |\widehat{v}_k|^2 < \infty. \quad (\text{B.3})$$

Note that with this definition, we have from Parseval's identity, that $\|v\|_{H^0} = \|v\|_{L^2}$, so that $H^0(\mathbb{T}^d) = L^2(\mathbb{T}^d)$.

We also introduce the corresponding **homogeneous Sobolev spaces** $\dot{H}^s(\mathbb{T}^d)$ (and $\dot{L}^2(\mathbb{T}^d) := \dot{H}^0(\mathbb{T}^d)$), consisting of functions $v(x) \in H^s(\mathbb{T}^d)$ and with zero mean $\int_{\mathbb{T}^d} v(x) dx = \widehat{v}_0 = 0$, and with norm

$$\|v\|_{\dot{H}^s} := \left((2\pi)^d \sum_{k \in \mathbb{Z}^d \setminus \{0\}} |k|^{2s} |\widehat{v}_k|^2 \right)^{1/2}. \quad (\text{B.4})$$

Given $N \in \mathbb{N}$, throughout this work, we will denote by $L_N^2(\mathbb{T}^d)$, the space of trigonometric polynomials $v_N : \mathbb{T}^d \rightarrow \mathbb{R}$, of the form

$$v_N(x) = \sum_{|k|_\infty \leq N} c_k e^{i\langle x, k \rangle}, \quad (\text{B.5})$$

where the summation is over all $k = (k_1, \dots, k_d) \in \mathbb{Z}^d$ such that

$$|k|_\infty := \max_{i=1, \dots, d} |k_i| \leq N.$$

The space $L_N^2(\mathbb{T}^d)$ is viewed as a normed vector space with norm $\|\cdot\|_{L^2}$. Similarly, for $s \geq 0$, we denote by $H_N^s(\mathbb{T}^d)$ the normed vector space of trigonometric polynomials v_N of degree $\leq N$, with norm $\|\cdot\|_{H^s}$.

We note that in order to ensure that $v_N(x) \in \mathbb{R}$ is real-valued for all $x \in \mathbb{T}^d$, the coefficients $c_k \in \mathbb{C}$ must satisfy the relations $c_{-k} = \overline{c_k}$ for all $|k|_\infty \leq N$, and where $\overline{c_k}$ denotes the complex conjugate of c_k .

We denote by

$$P_N : L^2(\mathbb{T}^d) \rightarrow L_N^2(\mathbb{T}^d), \quad v \mapsto P_N v, \quad (\text{B.6})$$

the **L^2 -orthogonal projection onto $L_N^2(\mathbb{T}^d)$** ; or more explicitly,

$$P_N \left(\sum_{k \in \mathbb{Z}^d} c_k e^{i\langle k, x \rangle} \right) = \sum_{|k|_\infty \leq N} c_k e^{i\langle k, x \rangle}, \quad \forall (c_k)_{k \in \mathbb{Z}^d} \in \ell^2(\mathbb{Z}^d).$$

In fact, the mapping P_N defines a projection $H^s(\mathbb{T}^d) \rightarrow H_N^s(\mathbb{T}^d)$ for any $s \geq 0$. We have the following spectral approximation estimate: Let $s > 0$ be given. There exists a constant $C = C(s, d) > 0$, such that for any $v \in H^s(\mathbb{T}^d)$, we have

$$\|v - P_N v\|_{H^s} \leq C N^{-(s-\varsigma)} \|v\|_{H^s}, \quad \text{for any } \varsigma \in [0, s]. \quad (\text{B.7})$$

We also define a natural projection

$$\dot{P}_N : L^2(\mathbb{T}^d) \rightarrow \dot{L}_N^2(\mathbb{T}^d), \quad (\text{B.8})$$

by removing the mean, i.e. $\dot{P}_N v = P_N v - \int_{\mathbb{T}^d} v(x) dx$, or equivalently:

$$\dot{P}_N \left(\sum_{k \in \mathbb{Z}^d} c_k e^{i\langle k, x \rangle} \right) = \sum_{0 < |k|_\infty \leq N} c_k e^{i\langle k, x \rangle}, \quad \forall (c_k)_{k \in \mathbb{Z}^d} \in \ell^2(\mathbb{Z}^d).$$

Furthermore, we denote by

$$\mathcal{I}_N : C(\mathbb{T}^d) \mapsto L_N^2(\mathbb{T}^d), \quad u \mapsto \mathcal{I}_N u, \quad (\text{B.9})$$

the **pseudo-spectral projection** onto $L_N^2(\mathbb{T}^d)$; we recall that the pseudo-spectral projection $\mathcal{I}_N v$ of a continuous function v is defined as the unique trigonometric polynomial $\mathcal{I}_N v \in L_N^2(\mathbb{T}^d)$, such that

$$\mathcal{I}_N v(x_j) = v(x_j), \quad \forall j \in \mathcal{J}_N, \quad (\text{B.10})$$

where $\{x_j\}_{j \in \mathcal{J}_N}$ denotes the set of all regular grid points $x_j \in \mathbb{Z}^d$ of the form $x_j = 2\pi j / (2N + 1) \in \mathbb{T}^d$, $j \in \mathbb{Z}^d$ (cp. equation (B.11)).

We also recall the following embedding theorem for the Sobolev spaces $H^s(\mathbb{T}^d)$:

Theorem B.1 (Sobolev embedding). Let $d \in \mathbb{N}$. For any $s > d/2$, we have a compact embedding $H^s(\mathbb{T}^d) \hookrightarrow C(\mathbb{T}^d)$ into the space of continuous functions. In particular, there exists a constant $C = C(s, d) > 0$, such that

$$\|v\|_{L^\infty} \leq C \|v\|_{H^s}, \quad \forall v \in H^s(\mathbb{T}^d).$$

The Sobolev embedding theorem implies in particular that the pseudo-spectral projection \mathcal{I}_N is well-defined as an operator $\mathcal{I}_N : H^s(\mathbb{T}^d) \rightarrow L_N^2(\mathbb{T}^d)$ for $s > d/2$. In the following theorem, we recall a well-known approximation error estimate for the pseudo-spectral projection \mathcal{I}_N :

Theorem B.2 (Pseudo-spectral approximation estimate). Let $d \in \mathbb{N}$. For any $s > d/2$ and $N \in \mathbb{N}$, the spectral interpolation operator $\mathcal{I}_N : H^s(\mathbb{T}^d) \rightarrow L_N^2(\mathbb{T}^d)$ is well-defined. Furthermore, there exists a constant $C = C(s, d) > 0$, such that the following approximation error estimate holds

$$\|(1 - \mathcal{I}_N)v\|_{H^\varsigma} \leq C N^{-(s-\varsigma)} \|v\|_{H^s}, \quad \forall v \in H^s(\mathbb{T}^d),$$

for any $\varsigma \in [0, s]$.

Finally, for $N \in \mathbb{N}$, we fix a regular grid $\{x_j\}_{j \in \mathcal{J}_N}$ of values

$$x_j = \frac{2\pi j}{2N + 1}, \quad (\text{B.11})$$

where the index $j \in \mathcal{J}_N$ belong to the index set

$$\mathcal{J}_N := \{0, \dots, 2N\}^d. \quad (\text{B.12})$$

Recall the set of Fourier wave numbers (2.10) and we define the **discrete Fourier transform** $\mathcal{F}_N : \mathbb{R}^{\mathcal{J}_N} \rightarrow \mathbb{C}^{\mathcal{K}_N}$ by

$$\mathcal{F}_N(v)(k) := \frac{1}{(2N + 1)^d} \sum_{j \in \mathcal{J}_N} v_j e^{-2\pi i \langle j, k \rangle / N}, \quad (\text{B.13})$$

with inverse $\mathcal{F}_N^{-1} : \mathbb{C}^{\mathcal{K}_N} \rightarrow \mathbb{R}^{\mathcal{J}_N}$,

$$\mathcal{F}_N^{-1}(\hat{v})(j) := \sum_{k \in \mathcal{K}_N} \hat{v}_k e^{2\pi i \langle j, k \rangle / N}. \quad (\text{B.14})$$

Lemma B.3 (Periodic extension operator). Let $\Omega \subset \mathbb{R}^d$ be a bounded, Lipschitz domain. There exists a continuous, linear operator $\mathcal{E} : W^{m,p}(\Omega) \rightarrow W^{m,p}(B)$ for any $m \geq 0$ and $1 \leq p \leq \infty$ where $B \Subset \mathbb{R}^d$ is a hypercube with $\Omega \subset B$ such that, for any $u \in W^{m,p}(\Omega)$,

- (1) $\mathcal{E}(u)|_\Omega = u$,
- (2) $\mathcal{E}(u)$ is periodic on B .

Proof. By [33, Chapter 6, Theorem 5], there exists a continuous, linear operator $\tilde{\mathcal{E}} : W^{m,p}(\Omega) \rightarrow W^{m,p}(\mathbb{R}^d)$ such that $\tilde{\mathcal{E}}(u)|_\Omega = u$ for any $u \in W^{m,p}(\Omega)$. Let $\Phi \in C^\infty(\mathbb{R}^d)$ be a mapping whose zeroth level set defines a curve $\partial\Omega'$ such that $\Omega \subset \Omega'$ and $\text{dist}(\partial\Omega', \partial\Omega) > 0$ where Ω' is a bounded domain with boundary $\partial\Omega'$. Furthermore, suppose that the first level set of Φ defines a curve $\partial\Lambda$ such that $\Omega' \subset \Lambda$ and $\text{dist}(\partial\Lambda, \partial\Omega') > 0$ where, again, Λ is a bounded domain with boundary $\partial\Lambda$. For example, we may take $\Phi(x) = x_1^2 + \dots + x_d^2 - r$ for some $r > 0$ large enough then Ω' is d -ball of radius \sqrt{r} enclosing Ω , and Λ is a d -ball of radius $\sqrt{1+r}$ clearly enclosing Ω' . We will now follow the approach of [3] and construct a windowing function $\rho \in C_c^\infty(\mathbb{R}^d)$. In particular, we define

$$\rho(x) = H(1 - 2\Phi(x)), \quad \forall x \in \mathbb{R}^d$$

where

$$H(t) = \frac{1}{2}(1 + S(t)), \quad \forall t \in \mathbb{R}$$

and

$$S(t) = \begin{cases} -1, & t < -1 \\ \text{erf}\left(\frac{t}{\sqrt{1-t^2}}\right), & -1 \leq t \leq 1, \\ 1, & t > 1 \end{cases} \quad \forall t \in \mathbb{R}$$

with erf, the Gauss error function,

$$\text{erf}(t) = \frac{2}{\sqrt{\pi}} \int_0^t e^{-z^2} dz, \quad \forall t \in \mathbb{R}.$$

Define $B \Subset \mathbb{R}^d$ to be a large enough d -cube such that $\Lambda \subset B$ and $\text{dist}(\partial\Lambda, \partial B) > 0$. It is easy to check that $\rho \in C_c^\infty(\mathbb{R}^d)$ and we have that $\rho|_\Omega = \rho|_{\Omega'} = 1$ and $\rho|_{B \setminus \bar{\Lambda}} = 0$ with ρ and all of its derivatives vanishing on $\partial\Lambda$. Furthermore,

$$\|\rho\|_{C^m(B)} < \infty$$

for any $m \geq 0$. We can thus define $\mathcal{E} : W^{m,p}(\Omega) \rightarrow W^{m,p}(B)$ by

$$\mathcal{E}(u) = \rho \tilde{\mathcal{E}}(u), \quad \forall u \in W^{m,p}(\Omega).$$

By Leibniz's rule and the generalized triangle inequality, there exists a constant $C_1 = C_1(d, m, p) > 0$ (independent of p in the case $p = \infty$) such that

$$\|\mathcal{E}(u)\|_{W^{m,p}(B)} \leq C_1 \|\rho\|_{C^m(B)} \|\tilde{\mathcal{E}}(u)\|_{W^{m,p}(B)}.$$

Since $\tilde{\mathcal{E}}$ is bounded, there is a constant $C_2 > 0$ such that

$$\begin{aligned} \|\mathcal{E}(u)\|_{W^{m,p}(B)} &\leq C_1 \|\rho\|_{C^m(B)} \|\tilde{\mathcal{E}}(u)\|_{W^{m,p}(B)} \\ &\leq C_1 \|\rho\|_{C^m(B)} \|\tilde{\mathcal{E}}(u)\|_{W^{m,p}(\mathbb{R}^d)} \\ &\leq C_1 C_2 \|\rho\|_{C^m(B)} \|u\|_{W^{m,p}(\Omega)} \end{aligned}$$

hence \mathcal{E} is a continuous, linear operator. Since $\rho|_\Omega = 1$, we immediately have $\mathcal{E}(u)|_\Omega = \tilde{\mathcal{E}}(u)|_\Omega = u$. Since $\text{supp}(\partial^\alpha \rho) \subseteq \bar{\Lambda}$ for any multi-index $0 \leq |\alpha| \leq m$, we conclude that $\partial^\alpha \mathcal{E}(u)|_{B \setminus \bar{\Lambda}} = 0$ hence $\mathcal{E}(u)$ is periodic on B as desired. \square

APPENDIX C. NON-LINEAR LIFTING AND PROJECTION OPERATORS

Lemma C.1. Assume that the activation function $\sigma \in C^2$ is non-linear and that $\|\sigma\|_{C^2(\mathbb{R})} < \infty$. Let $\mathcal{N} : L^2(D; \mathbb{R}^{d_a}) \rightarrow L^2(D; \mathbb{R}^{d_u})$ be a neural operator of the form $\mathcal{N} = \mathcal{Q} \circ \mathcal{L}_L \circ \dots \circ \mathcal{L}_1 \circ \mathcal{R}$, with Lipschitz continuous layers $\mathcal{L}_\ell : L^2(D; \mathbb{R}^{d_v}) \rightarrow L^2(D; \mathbb{R}^{d_v})$, and with linear lifting/projection operators \mathcal{R}, \mathcal{Q} of the form (2.1) and (2.2), respectively. For any $K \subset L^2(D; \mathbb{R}^{d_v})$ compact and $\epsilon > 0$, there exist neural networks $\widehat{R} : \mathbb{R}^{d_a} \times D \rightarrow \mathbb{R}^{d_v}$, $\widehat{Q} : \mathbb{R}^{d_v} \times D \rightarrow \mathbb{R}^{d_u}$ with a single hidden layer and width(\widehat{R}), width(\widehat{Q}) $\leq 2d_v$, such that the mapping $\widehat{\mathcal{N}} : \mathcal{A}(D; \mathbb{R}^{d_a}) \rightarrow \mathcal{U}(D; \mathbb{R}^{d_u})$ defined by

$$\widehat{\mathcal{N}}(a) := \widehat{Q} \circ \mathcal{L}_L \circ \dots \circ \mathcal{L}_1 \circ \widehat{R}(a),$$

where $\mathcal{R}(a)(x) := \widehat{R}(a(x), x)$, $\mathcal{Q}(v) := \widehat{Q}(v(x), x)$, approximates \mathcal{N} to order ϵ :

$$\sup_{a \in K} \|\mathcal{N}(a) - \widehat{\mathcal{N}}(a)\|_{L^2} \leq \epsilon.$$

Proof. We recall that, by our definition of a neural operator, we have $\mathcal{R}(a)(x) = Ra(x)$, $\mathcal{Q}(v)(x) = Qv(x)$, where $R \in \mathbb{R}^{d_v \times d_a}$, $Q \in \mathbb{R}^{d_u \times d_v}$. By assumption on the non-linearity of σ , there exists $z_0 \in \mathbb{R}$, such that $\sigma'(z_0) \neq 0$. For any $h > 0$, we now define

$$\begin{cases} \widehat{Q}_h(v, x) := \frac{\sigma(z_0 + hQv) - \sigma(z_0 - hQv)}{2h\sigma'(z_0)}, \\ \widehat{R}_h(a, x) := \frac{\sigma(z_0 + hRa) - \sigma(z_0 - hRa)}{2h\sigma'(z_0)}. \end{cases}$$

We remark that in the above expressions, the addition of $z_0 \in \mathbb{R}$ and $hQv \in \mathbb{R}^{d_u}$, $hRa \in \mathbb{R}^{d_v}$ is carried out componentwise, as is the evaluation of σ .

Since the interior layers $\mathcal{L}_1, \dots, \mathcal{L}_L$ are fixed, we introduce the short-hand notation $\mathcal{L} := \mathcal{L}_L \circ \dots \circ \mathcal{L}_1$. By assumption on the activation function and the layers, we have that $\mathcal{L} : L^2(D; \mathbb{R}^{d_v}) \rightarrow L^2(D; \mathbb{R}^{d_v})$ is a Lipschitz continuous mapping. Our goal is to show that for any $K \subset L^2$ compact and $\epsilon > 0$, there exists a sufficiently small $h > 0$, such that

$$\sup_{a \in K} \|\mathcal{N}(a) - \widehat{\mathcal{N}}_h(a)\|_{L^2} \leq \epsilon,$$

where $\widehat{\mathcal{N}}_h(a) := \widehat{Q}_h \circ \mathcal{L} \circ \widehat{R}_h(a)$. To this end, we note that for any $a \in K$, we have

$$\begin{aligned} \|\mathcal{N}(a) - \widehat{\mathcal{N}}_h(a)\|_{L^2} &= \|\mathcal{Q} \circ \mathcal{L} \circ \mathcal{R}(a) - \widehat{Q}_h \circ \mathcal{L} \circ \widehat{R}_h(a)\|_{L^2} \\ &\leq \|\mathcal{Q} \circ \mathcal{L} \circ \mathcal{R}(a) - \widehat{Q}_h \circ \mathcal{L} \circ \mathcal{R}(a)\|_{L^2} \\ &\quad + \|\widehat{Q}_h \circ \mathcal{L} \circ \mathcal{R}(a) - \widehat{Q}_h \circ \mathcal{L} \circ \widehat{R}_h(a)\|_{L^2}. \end{aligned}$$

Introducing $\widetilde{K} := \mathcal{L} \circ \mathcal{R}(K)$, we note that $\widetilde{K} \subset L^2$ is compact, and

$$\begin{aligned} \sup_{a \in K} \|\mathcal{N}(a) - \widehat{\mathcal{N}}_h(a)\|_{L^2} &\leq \sup_{v \in \widetilde{K}} \|\mathcal{Q}(v) - \widehat{Q}_h(v)\|_{L^2} \\ &\quad + \text{Lip}(\widehat{Q}_h) \text{Lip}(\mathcal{L}) \sup_{a \in K} \|\mathcal{R}(a) - \widehat{R}_h(a)\|_{L^2}. \end{aligned}$$

The proof of this lemma thus follows from the following three claims:

Claim 1: There exists $C > 0$, independent of h , such that $\text{Lip}(\widehat{Q}_h) \leq C$ for all $h > 0$.

Claim 2: For any $\epsilon > 0$ and $\widetilde{K} \subset L^2$ compact, there exists $h > 0$, such that

$$\sup_{v \in \widetilde{K}} \|\widehat{Q}_h(v) - \mathcal{Q}(v)\|_{L^2} \leq \epsilon.$$

Claim 3: For any $\epsilon > 0$ and $K \subset L^2$ compact, there exists $h > 0$, such that

$$\sup_{a \in K} \|\widehat{\mathcal{R}}_h(a) - \mathcal{R}(a)\|_{L^2} \leq \epsilon.$$

Clearly, the only difference between Claims 2 and 3 is notational, hence it suffices show Claims 1 and 2 to conclude the proof of the present lemma.

Proof of Claim 1: We note that for any $v(x), v'(x) \in \mathbb{R}^{d_v}$, we have

$$|\widehat{Q}_h(v(x), x) - \widehat{Q}_h(v'(x), x)| \leq \frac{1}{2h|\sigma'(z_0)|} 2\text{Lip}(\sigma) |hQ(v(x) - v'(x))| \quad (\text{C.1})$$

$$\leq \frac{\text{Lip}(\sigma)\|Q\|}{|\sigma'(z_0)|} |v(x) - v'(x)|, \quad (\text{C.2})$$

where $|\cdot|$ denotes the Euclidean norm and $\|Q\|$ denotes the operator norm of Q . Hence, it follows that for any $v, v' \in L^2(D; \mathbb{R}^{d_v})$, we have

$$\|\widehat{Q}_h(v) - \widehat{Q}_h(v')\|_{L^2} \leq \frac{\text{Lip}(\sigma)\|Q\|}{|\sigma'(z_0)|} \|v - v'\|_{L^2},$$

and thus,

$$\text{Lip}(\widehat{Q}_h) \leq \frac{\text{Lip}(\sigma)\|Q\|}{|\sigma'(z_0)|},$$

is bounded independently of $h > 0$.

Proof of Claim 2: It follows from Taylor expansion that $\widehat{Q}_h(v(x), x) = Qv(x) + hR(h; Qv(x))$, where the remainder R satisfies a uniform bound

$$|R(h; Qv(x))| \leq \frac{C\|\sigma\|_{C^2}}{|\sigma'(z_0)|} |Qv(x)|^2, \quad \forall h \in (0, 1], v(x) \in \mathbb{R}^{d_v}.$$

In particular, considering the L^2 -norm of the left-hand side and assuming $x \mapsto v(x)$ to be *bounded*, we conclude that there exists a constant $C = C(\sigma, \|Q\|) > 0$, depending only on the activation function σ and the (Euclidean) operator norm $\|Q\|$ of $Q : (\mathbb{R}^{d_v}, |\cdot|) \rightarrow (\mathbb{R}^{d_u}, |\cdot|)$, such that

$$\|\widehat{Q}_h(v) - Q(v)\|_{L^2} = \|R(h; Qv)\|_{L^2} \leq C\|v\|_{L^\infty} \|v\|_{L^2}, \quad \forall v \in L^2 \cap L^\infty. \quad (\text{C.3})$$

Next, for $M > 0$ we introduce a cut-off operator $p_M : L^2 \rightarrow L^2 \cap L^\infty$, by

$$p_M(v(x)) := \begin{cases} v(x), & (|v(x)| \leq M), \\ 0, & (|v(x)| > M). \end{cases}$$

Since $\widetilde{K} \subset L^2$ is compact, it follows that

$$\lim_{M \rightarrow \infty} \sup_{v \in \widetilde{K}} \|v - p_M(v)\|_{L^2} = 0.$$

We can thus choose $M > 0$ sufficiently large, such that

$$\sup_{v \in \widetilde{K}} \|v - p_M(v)\|_{L^2} \leq \frac{\epsilon}{2 \left(\sup_{h \in (0,1]} \text{Lip}(\widehat{Q}_h) + \text{Lip}(Q) \right)},$$

where we note that $\sup_{h \in (0,1]} \text{Lip}(\widehat{Q}_h) < \infty$, by Claim 1. Furthermore, and again by compactness of $\widetilde{K} \subset L^2$, there exists a constant $\widetilde{M} > 0$, such that

$$\sup_{v \in \widetilde{K}} \|v\|_{L^2} \leq \widetilde{M}.$$

With this \widetilde{M} , our choice of M and the estimate (C.3) on the remainder R , it follows that for any $v \in \widetilde{K}$:

$$\begin{aligned} \|\widehat{\mathcal{Q}}_h(v) - \mathcal{Q}(v)\|_{L^2} &\leq \|\widehat{\mathcal{Q}}_h(v) - \widehat{\mathcal{Q}}_h(p_M(v))\|_{L^2} \\ &\quad + \|\widehat{\mathcal{Q}}_h(p_M(v)) - \mathcal{Q}(p_M(v))\|_{L^2} \\ &\quad + \|\mathcal{Q}(p_M(v)) - \mathcal{Q}(v)\|_{L^2} \\ &\leq \left(\text{Lip}(\widehat{\mathcal{Q}}_h) + \text{Lip}(\mathcal{Q})\right) \|v - p_M(v)\|_{L^2} \\ &\quad + Ch\|p_M(v)\|_{L^\infty}\|p_M(v)\|_{L^2} \\ &\leq \epsilon/2 + CM\widetilde{M}h. \end{aligned}$$

Finally, choosing $h = \epsilon/(2CM\widetilde{M})$ implies that

$$\sup_{v \in \widetilde{K}} \|\widehat{\mathcal{Q}}_h(v) - \mathcal{Q}(v)\|_{L^2} \leq \epsilon.$$

This concludes the proof. \square

APPENDIX D. PROOFS AND TECHNICAL DETAILS FOR SECTION 2

D.1. Proof of Lemma 2.6. The proof of Lemma 2.6 will rely on the following technical lemma:

Lemma D.1. Let $s' \geq 0$, $N \in \mathbb{N}$ be given. Let $K \subset H^{s'}$ be compact, and assume that $\sigma \in C^m$, where $m > s'$ is integer. Then for any $\epsilon > 0$, there exists a single-layer FNO $\mathcal{L} : H^{s'} \rightarrow H^{s'}$, such that

$$\sup_{v \in K} \|P_N v - \mathcal{L}(v)\|_{H^{s'}} \leq \epsilon.$$

Proof of Lemma D.1. First, we note that the Fourier projection $P_N : H^{s'} \rightarrow H^{s'}$ is a continuous operator, and hence the image $P_N K \subset H^{s'}$ is compact. Furthermore, P_N maps into a finite-dimensional subspace of $H^{s'}$. Due to the norm-equivalence on finite-dimensional spaces, there thus exists $C_0 = C_0(N, K) > 0$, such that

$$\sup_{v \in K} \|P_N v\|_{L^\infty} \leq C_0, \quad \sup_{v \in K} \|P_N v\|_{H^m} \leq C_0. \quad (\text{D.1})$$

Let $x_0 \in \mathbb{R}$ be such that $\sigma'(x_0) \neq 0$. We define, for $h > 0$,

$$\psi_h(x) := \frac{\sigma(x_0 + hx) - \sigma(x_0 - hx)}{2h\sigma'(x_0)}. \quad (\text{D.2})$$

One readily shows that $\psi_h \in C^m$, and that there exists a constant $C_1 = C_1(\sigma, C_0) > 0$, such that

$$\|\psi_h\|_{C^m([-C_0, C_0])} \leq C_1, \quad \forall h \in (0, 1]. \quad (\text{D.3})$$

Furthermore, by Taylor expansion, we have

$$|\psi_h(x) - x| \leq Ch, \quad \forall x \in [-C_0, C_0], \forall h \in (0, 1]. \quad (\text{D.4})$$

By the composition rule for Sobolev functions, we have that $\psi_h \circ P_N a \in H^m$, for $P_N a \in H^m$, and there exists a constant $C_2 = C_2(C_1, C_0) > 0$, such that

$$\|\psi_h(P_N v)\|_{H^m} \leq C_2, \quad \forall v \in K. \quad (\text{D.5})$$

We finally observe that the mapping $v \mapsto \mathcal{L}_h(v) := \psi_h(P_N v)$ can be represented by a single-layer FNO, and by (D.5), we have

$$\|\mathcal{L}_h(v)\|_{H^m} \leq C_2, \quad \forall v \in K. \quad (\text{D.6})$$

From the interpolation inequality between Sobolev spaces, it follows that

$$\|\mathcal{L}_h(v) - P_N v\|_{H^{s'}} \leq \|\mathcal{L}_h(v) - P_N v\|_{L^2}^\theta \|\mathcal{L}_h(v) - P_N v\|_{H^m}^{1-\theta},$$

where $\theta = 1 - s'/m > 0$. By (D.6) and (D.1), the second factor can be bounded independently of $h > 0$. By (D.4) and (D.1), we have

$$\|\mathcal{L}_h(v) - P_N v\|_{L^2} = \|\psi_h(P_N v) - P_N v\|_{L^2} \leq Ch,$$

for a constant $C > 0$, independent of h and $v \in K$. We conclude that

$$\|\mathcal{L}_h(v) - P_N v\|_{H^{s'}} \leq Ch^\theta \rightarrow 0,$$

as $h \rightarrow 0$, for some constant $C > 0$, independent of h . This proves the claim. \square

We can now prove:

Proof of Lemma 2.6. Let $\mathcal{G} : H^s \rightarrow H^{s'}$ be a continuous operator, and let $K \subset H^s$ be compact. We assume that FNOs are universal approximators of operators $H^s \rightarrow L^2$, and we wish to show that for any $\epsilon > 0$, there exists a FNO approximation of $\mathcal{G} : H^s \rightarrow H^{s'}$ to accuracy ϵ . We first note that by the compactness of $\mathcal{G}(K) \subset H^{s'}$, there exists $N \in \mathbb{N}$, such that

$$\sup_{a \in K} \|\mathcal{G}(a) - P_N \mathcal{G}(a)\|_{H^{s'}} \leq \epsilon/2. \quad (\text{D.7})$$

Fix $\delta > 0$ for the moment. A suitable choice of δ will be specified at the end of this proof. By assumption on the universal approximation of operators $H^s \rightarrow L^2$, there exists a FNO $\tilde{\mathcal{N}} : H^s \rightarrow L^2$, continuous as an operator $H^s \rightarrow L^2$, such that

$$\sup_{a \in K} \|P_N \mathcal{G}(a) - \tilde{\mathcal{N}}(a)\|_{L^2} \leq \delta. \quad (\text{D.8})$$

One difficulty in the present construction is that there is no guarantee that $\tilde{\mathcal{N}}$ defines a mapping $H^s \rightarrow H^{s'}$, and indeed for $s' > s$ this is not generally the case. We circumvent this issue by composing with an additional FNO layer $\tilde{\mathcal{L}} : L^2 \rightarrow H^{s'}$. By Lemma D.1, there exists a single-layer FNO $v \mapsto \tilde{\mathcal{L}}(v)$, satisfying the identity

$$\tilde{\mathcal{L}}(v) = \tilde{\mathcal{L}}(P_N v), \quad (\text{D.9})$$

for all v , and defining a continuous operator $H^{s'} \rightarrow H^{s'}$, such that

$$\sup_{v \in K'} \|P_N v - \tilde{\mathcal{L}}(v)\|_{H^{s'}} \leq \delta, \quad (\text{D.10})$$

where $K' := P_N \tilde{\mathcal{N}}(K) \subset H^{s'}$ is a compact subset¹ of $H^{s'}$. Next, we define a new FNO by the composition $\mathcal{N} := \tilde{\mathcal{L}} \circ \tilde{\mathcal{N}} : H^s \rightarrow H^{s'}$. \mathcal{N} is a continuous operator $H^s \rightarrow H^{s'}$, since it can be written as the composition

$$H^s \xrightarrow{\tilde{\mathcal{N}}} L^2 \xrightarrow{P_N} H^{s'} \xrightarrow{\tilde{\mathcal{L}}} H^{s'},$$

of continuous operators. We observe that for any $a \in K$, we have the following bound:

$$\begin{aligned} \|P_N \mathcal{G}(a) - \mathcal{N}(a)\|_{H^{s'}} &\leq \|P_N \mathcal{G}(a) - P_N \tilde{\mathcal{N}}(a)\|_{H^{s'}} + \|P_N \tilde{\mathcal{N}}(a) - \mathcal{N}(a)\|_{H^{s'}} \\ &\leq CN^{s'} \|P_N \mathcal{G}(a) - \tilde{\mathcal{N}}(a)\|_{L^2} + \|P_N \tilde{\mathcal{N}}(a) - \tilde{\mathcal{L}}(P_N \tilde{\mathcal{N}}(a))\|_{H^{s'}}, \end{aligned}$$

having made use of the inequality $\|P_N v\|_{H^{s'}} \leq CN^{s'} \|P_N v\|_{L^2}$ for a constant $C = C(\mathbb{T}^d, s') > 0$ independent of N , and the fact that $\mathcal{N}(a) = \tilde{\mathcal{L}}(\tilde{\mathcal{N}}(a)) = \tilde{\mathcal{L}}(P_N \tilde{\mathcal{N}}(a))$ (cp. (D.9)). Using (D.8), we can estimate

$$CN^{s'} \|P_N \mathcal{G}(a) - \tilde{\mathcal{N}}(a)\|_{L^2} \leq CN^{s'} \delta.$$

¹We note that $\mathcal{N}(K) \subset L^2$ is compact as the continuous image of K , and that $P_N : L^2 \rightarrow H^{s'}$ defines a continuous mapping for fixed $N \in \mathbb{N}$.

The bound (D.10) implies that $\|P_N \tilde{\mathcal{N}}(a) - \tilde{\mathcal{L}}(P_N \tilde{\mathcal{N}}(a))\|_{H^{s'}} = \|P_N v - \tilde{\mathcal{L}}(v)\|_{H^{s'}} \stackrel{(D.10)}{\leq} \delta$, with $v := P_N \mathcal{N}(a) \in K'$. We thus obtain

$$\|P_N \mathcal{G}(a) - \mathcal{N}(a)\|_{H^{s'}} \leq (CN^{s'} + 1)\delta, \quad (\text{D.11})$$

where $C = C(\mathbb{T}^d, s') > 0$ is independent of δ . Since $\delta > 0$ was arbitrary, we can ensure that $(CN^{s'} + 1)\delta \leq \epsilon/2$. From this estimate, and the bound (D.7), we conclude that there exists a FNO $\mathcal{N} : H^s \rightarrow H^{s'}$, such that

$$\begin{aligned} \sup_{a \in K} \|\mathcal{G}(a) - \mathcal{N}(a)\|_{H^{s'}} &\leq \sup_{a \in K} \|\mathcal{G}(a) - P_N \mathcal{G}(a)\|_{H^{s'}} \\ &\quad + \sup_{a \in K} \|P_N \mathcal{G}(a) - \mathcal{N}(a)\|_{H^{s'}} \\ &\leq \epsilon, \end{aligned}$$

This concludes our proof. \square

D.2. Proof of Lemma 2.7.

Proof. Step 1: In this first step, for any $\epsilon > 0$, we will construct a FNO

$$\mathcal{N}_1 : L^2(\mathbb{T}^d; \mathbb{R}) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}),$$

such that

$$\begin{cases} \|\mathcal{N}_1(v)_{1,k} - P_N v(x) \cos(\langle k, x \rangle)\|_{L^\infty} < \epsilon, \\ \|\mathcal{N}_1(v)_{2,k} - P_N v(x) \sin(\langle k, x \rangle)\|_{L^\infty} < \epsilon, \end{cases} \quad \forall k \in \mathcal{K}_N. \quad (\text{D.12})$$

To see how to construct such \mathcal{N}_1 , we first define a lifting

$$\mathcal{R}_1 : L^2(\mathbb{T}^d; \mathbb{R}) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{4\mathcal{K}_N}), \quad v(x) \mapsto \hat{w}(x),$$

where $\hat{w}(x) := \{(v(x), 0, v(x), 0)\}_{k \in \mathcal{K}_N} \in \mathbb{R}^{4\mathcal{K}_N}$ for any $x \in \mathbb{T}^d$. In the following, we will identify $\mathbb{R}^{4\mathcal{K}_N} \simeq \mathbb{R}^{d_v}$, with $d_v = 4|\mathcal{K}_N|$. Next, we define the inner part of the first FNO layer (i.e. the matrix W , multiplier P , bias $b(x)$ in (2.7)), such that $P(k) \equiv \mathbf{1}_{|k| \leq N} \mathbf{1}_{d_v \times d_v}$ either vanishes (for $|k| > N$) or is the unit matrix (for $|k| \leq N$), $W \equiv 0$ is zero and the bias function $b(x) := \{(0, b_{2,k}(x), 0, b_{4,k}(x))\}_{k \in \mathcal{K}_N}$, $b_{2,k}(x) = \cos(\langle k, x \rangle)$, $b_{4,k}(x) = \sin(\langle k, x \rangle)$, such that

$$\begin{aligned} \hat{\mathcal{L}}_1(\hat{v})(x) &:= W \hat{w}(x) + b(x) + \mathcal{F}^{-1}(P \mathcal{F} \hat{w})(x) \\ &= \left\{ (P_N \hat{w}_{1,k}(x), \cos(\langle k, x \rangle), P_N \hat{w}_{3,k}(x), \sin(\langle k, x \rangle)) \right\}_{k \in \mathcal{K}_N}. \end{aligned}$$

Next, we recall that by assumption, we have $\|v\|_{L^2} \leq B$, and by construction, we have $\hat{w}_{1,k}(x) = \hat{w}_{3,k}(x) = v(x)$, implying that

$$\|P_N \hat{w}_{1,k}\|_{L^\infty} = \|P_N \hat{w}_{3,k}\|_{L^\infty} = \|P_N v\|_{L^\infty} \leq C \|P_N v\|_{L^2} \leq CB,$$

where $C = C(N) \propto N^{d/2}$ is a constant depending on N . By the universal approximation theorem for (ordinary) neural networks, there exists a neural network

$$\hat{\mathcal{N}} : [-CB, CB] \times [-1, 1] \times [-CB, CB] \times [-1, 1] \rightarrow \mathbb{R}^2$$

with activation function σ , such that $\max_{(a,b,c,d)} |\hat{\mathcal{N}}(a, b, c, d) - (ab, cd)| < \epsilon$, and where the maximum is taken over all $a, c \in [-CB, CB]$, $b, d \in [-1, 1]$. But then, the point-wise mapping

$$\mathcal{N}_1 = \hat{\mathcal{N}} \circ \hat{\mathcal{L}} \circ \mathcal{R}_1 : \begin{cases} L^2(\mathbb{T}^d; \mathbb{R}) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}), \\ v(x) \mapsto \tilde{w}(x) := \hat{\mathcal{L}}_1(\mathcal{R}_1(v))(x) \mapsto \hat{\mathcal{N}}(\tilde{w}(x)), \end{cases}$$

satisfies (D.12), and \mathcal{N}_1 can be represented by a FNO (cp. Remark 2.3). **Step 2:** By the last step, we have $\|\mathcal{N}_1(v)_{1,k} - P_N v(x) \cos(\langle k, x \rangle)\|_{L^\infty} < \epsilon$ and $\|\mathcal{N}_1(v)_{2,k} -$

$P_N v(x) \sin(\langle k, x \rangle) \|_{L^\infty} < \epsilon$ for all $v \in L^2(\mathbb{T}^d)$ with $\|v\|_{L^2} \leq B$. We next note that, since

$$P_N v(x) = \sum_{|k|_\infty \leq N} \widehat{v}_k e^{i\langle k, x \rangle},$$

and since $v(x)$ is real-valued, the Fourier coefficients \widehat{v}_k satisfy $\widehat{v}_{-k} = \overline{\widehat{v}_k}$, where $\overline{\widehat{v}_k}$ denotes the complex conjugate of \widehat{v}_k , and

$$\int_{\mathbb{T}^d} P_N v(x) \cos(\langle k, x \rangle) dx = \operatorname{Re}(\widehat{v}_k), \quad \int_{\mathbb{T}^d} P_N v(x) \sin(\langle k, x \rangle) dx = -\operatorname{Im}(\widehat{v}_k),$$

In particular, this implies that the 0-th Fourier modes of $P_N v(x) \cos(\langle k, x \rangle)$, and $P_N v(x) \sin(\langle k, x \rangle)$, respectively, are given by

$$\mathcal{F}\left[P_N v \cos(\langle k, \cdot \rangle)\right](0) = \operatorname{Re}(\widehat{v}_k), \quad \mathcal{F}\left[P_N v \sin(\langle k, \cdot \rangle)\right](0) = -\operatorname{Im}(\widehat{v}_k),$$

and, as a consequence, we have with the Fourier multiplier

$$\delta_0(k') = \begin{cases} 1, & (k' = 0) \\ 0, & (k' \neq 0), \end{cases} \quad \forall k' \in \mathbb{Z}^d,$$

and with $w(x) := \mathcal{N}_1(v)(x) \in L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$, written as $w(x) = \{(w_{1,k}(x), w_{2,k}(x))\}_{k \in \mathcal{K}_N}$, that

$$\begin{aligned} \left\| \mathcal{F}^{-1}\left(\delta_0(k') \mathcal{F}(w_{1,k})(k')\right)(x) - \operatorname{Re}(\widehat{v}_k) \right\|_{L^\infty} &< \epsilon, \\ \left\| \mathcal{F}^{-1}\left(\delta_0(k') \mathcal{F}(w_{2,k})(k')\right)(x) + \operatorname{Im}(\widehat{v}_k) \right\|_{L^\infty} &< \epsilon, \end{aligned}$$

for all $v \in L^2(\mathbb{T}^d)$ with $\|v\|_{L^2} \leq B$. We use this observation to define a suitable FNO layer; with local matrix $W = 0$, bias $b(x) = 0$ and a Fourier multiplier matrix $P : \mathbb{Z}^d \rightarrow \mathbb{C}^{d_v \times d_v}$, $k' \mapsto P(k')$ (where $d_v = 2|\mathcal{K}_N|$, and $\mathbb{C}^{d_v} \simeq \mathbb{C}^{2\mathcal{K}_N}$), with entries

$$[P(k')]_{(\ell,k),(\bar{\ell},\bar{k})} := \delta_0(k') \{\delta_{\ell=1}(\ell) - \delta_{\ell=2}(\ell)\} \mathbf{1}_{d_v \times d_v}.$$

With this definition of the Fourier multiplier P , we define $\widehat{\mathcal{L}}_2 : L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$, by

$$\widehat{\mathcal{L}}_2(w) := Ww(x) + b(x) + \mathcal{F}^{-1}\left(P\mathcal{F}w\right)(x).$$

Then, by construction, we have for any $v \in L^2(\mathbb{T}^d)$ with $\|v\|_{L^2} \leq B$ that the output $w := \widehat{\mathcal{L}}_2(\mathcal{N}_1(v))$, $w(x) = \{(w_{1,k}(x), w_{2,k}(x))\}_{k \in \mathcal{K}_N}$, satisfies

$$\begin{cases} \|w_{1,k}(x) - \operatorname{Re}(\widehat{v}_k)\|_{L^\infty} < \epsilon, \\ \|w_{2,k}(x) - \operatorname{Im}(\widehat{v}_k)\|_{L^\infty} < \epsilon. \end{cases}$$

We also note that $\widehat{\mathcal{L}}_2$ outputs only constant functions, by construction. This is almost the desired result, except for the fact that the composition $\widehat{\mathcal{L}}_2 \circ \mathcal{N}_1$ does not define a FNO, since the linear layer $\widehat{\mathcal{L}}_2$ is missing the non-linearity σ . This can be rectified by composition with a suitable (ordinary) σ -neural network $\widetilde{\mathcal{N}}$, which approximates the identity: Indeed, the last estimate implies in particular that $|w_{\ell,k}| \leq |\widehat{v}_k| + \epsilon \leq B + \epsilon$ for all ℓ, k . By the (ordinary) universal approximation theorem, there exists a neural network $\widetilde{\mathcal{N}} : \mathbb{R}^{2\mathcal{K}_N} \rightarrow \mathbb{R}^{2\mathcal{K}_N}$, such that

$$\|\widetilde{\mathcal{N}}(w) - w\|_{\ell^\infty} < \epsilon, \quad \forall w \in \mathbb{R}^{2\mathcal{K}_N}, \quad \|w\|_{\ell^\infty} \leq B + \epsilon.$$

Then the composition $\mathcal{N}_2 := \widetilde{\mathcal{N}} \circ \widehat{\mathcal{L}}_2 : L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$, given by $w(x) \mapsto \widetilde{\mathcal{N}}\left(\widehat{\mathcal{L}}_2(w)(x)\right)$, does define a FNO.

Step 3: We finally observe that since both \mathcal{N}_1 and \mathcal{N}_2 can be represented by FNOs, then also their composition

$$\mathcal{N} := \mathcal{N}_2 \circ \mathcal{N}_1 : L^2(\mathbb{T}^d) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}),$$

can be represented by a FNO with $\text{depth}(\mathcal{N}) = \text{depth}(\mathcal{N}_1) + \text{depth}(\mathcal{N}_2)$, $\text{width}(\mathcal{N}) = \max_{j=1,2} \text{width}(\mathcal{N}_j)$, $\text{lift}(\mathcal{N}) = \max_{j=1,2} \text{lift}(\mathcal{N}_j)$. Furthermore, the mapping $v \mapsto \mathcal{N}(v)$ maps to constant functions, and for any $v \in L^2(\mathbb{T}^d)$, $\|v\|_{L^2} \leq B$, we have

$$\begin{cases} \|\text{Re}(\widehat{v}_k) - \mathcal{N}(v)_{1,k}\|_{L^\infty} < 2\epsilon, \\ \|\text{Im}(\widehat{v}_k) - \mathcal{N}(v)_{2,k}\|_{L^\infty} < 2\epsilon. \end{cases}$$

Since $\epsilon > 0$ was arbitrary, the claim follows. \square

D.3. Proof of Lemma 2.8.

Proof. Step 1: We first construct a FNO $\mathcal{N}_1 : L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$, such that

$$\begin{cases} \|\mathcal{N}_1(w)_{1,k} - P_N w_{1,k}(x) \cos(\langle k, x \rangle)\|_{L^\infty} < \left(2|\mathcal{K}_N| |\mathbb{T}^d|^{1/2}\right)^{-1} \epsilon, \\ \|\mathcal{N}_1(w)_{2,k} - P_N w_{2,k}(x) \sin(\langle k, x \rangle)\|_{L^\infty} < \left(2|\mathcal{K}_N| |\mathbb{T}^d|^{1/2}\right)^{-1} \epsilon, \end{cases} \quad (\text{D.13})$$

for all $w(x) = (w_{1,k}(x), w_{2,k}(x))$, such that $\|w_{1,k}\|_{L^\infty}, \|w_{2,k}\|_{L^\infty} \leq B$, for all $k \in \mathcal{K}_N$. Here $|\mathbb{T}^d| = (2\pi)^d$ denotes the Lebesgue measure of \mathbb{T}^d and $|\mathcal{K}_N| = (2N+1)^d$ the number of elements of \mathcal{K}_N .

To construct such \mathcal{N}_1 , we first define a lifting $\mathcal{R} : L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{4\mathcal{K}_N})$, $\{(w_{1,k}(x), w_{2,k}(x))\}_{k \in \mathcal{K}_N} \mapsto \{(w_{1,k}(x), 0, w_{2,k}(x), 0)\}_{k \in \mathcal{K}_N}$, followed by a linear layer $\mathcal{L} : L^2(\mathbb{T}^d; \mathbb{R}^{4\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{4\mathcal{K}_N})$, which only introduces a bias (setting $W \equiv 0$, $b(x) = \{(0, \cos(\langle k, x \rangle), 0, \sin(\langle k, x \rangle))\}_{k \in \mathcal{K}_N}$, $P \equiv 0$), to yield

$$\mathcal{L} \circ \mathcal{R}(w) = \{(w_{1,k}(x), \cos(\langle k, x \rangle), w_{2,k}(x), \sin(\langle k, x \rangle))\}_{k \in \mathcal{K}_N},$$

for all $w \in L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$. There exists an ordinary neural network $\widehat{\mathcal{N}} : [-B, B] \times [-1, 1] \times [-B, B] \times [-1, 1] \rightarrow \mathbb{R}^2$, such that $\max_{a,b,c,d} |\widehat{\mathcal{N}}(a, b, c, d) - (ab, cd)| < (2|\mathcal{K}_N| |\mathbb{T}^d|^{1/2})^{-1} \epsilon$, where the maximum is taken over $a, c \in [-B, B]$ and $b, d \in [-1, 1]$. Since, by assumption, we have $\|w_{1,k}\|_{L^\infty}, \|w_{2,k}\|_{L^\infty} \leq B$ for the inputs of interest, we conclude that $\mathcal{N}_1(w) := \widehat{\mathcal{N}} \circ \mathcal{L} \circ \mathcal{R}(w)$ satisfies (D.13). Furthermore, by Lemma D.2 and Lemma D.4, \mathcal{N}_1 is represented by a FNO.

Step 2: We define a projection \mathcal{Q} (cp. the definition of neural operators (2.2)), mapping $w(x) \in L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$ to a scalar-valued function, by

$$\mathcal{Q} : L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d), \quad w(x) \mapsto \sum_{k \in \mathcal{K}_N} (w_{1,k}(x) - w_{2,k}(x)).$$

Then $\mathcal{N} := \mathcal{Q} \circ \mathcal{N}_1 : L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d)$ is an FNO, and for any $v \in L^2_{\mathbb{N}}(\mathbb{T}^d)$, and $w = \{(\text{Re}(\widehat{v}_k), \text{Im}(\widehat{v}_k))\}_{k \in \mathcal{K}_N}$ defined as in the statement of this lemma, we

have

$$\begin{aligned}
\|v - \mathcal{N}(w)\|_{L^2} &= \left\| v - \sum_{k \in \mathcal{K}_N} \left(\mathcal{N}_1(w)_{1,k} - \mathcal{N}_1(w)_{2,k} \right) \right\|_{L^2} \\
&\leq \sum_{k \in \mathcal{K}_N} \|\mathcal{N}_1(w)_{1,k} - \operatorname{Re}(\widehat{v}_k) \cos(\langle k, x \rangle)\|_{L^2} \\
&\quad + \sum_{k \in \mathcal{K}_N} \|\mathcal{N}_1(w)_{2,k} - \operatorname{Im}(\widehat{v}_k) \sin(\langle k, x \rangle)\|_{L^2} \\
&\leq \sum_{k \in \mathcal{K}_N} \|\mathcal{N}_1(w)_{1,k} - \operatorname{Re}(\widehat{v}_k) \cos(\langle k, x \rangle)\|_{L^2} \\
&\quad + \sum_{k \in \mathcal{K}_N} \|\mathcal{N}_1(w)_{2,k} - \operatorname{Im}(\widehat{v}_k) \sin(\langle k, x \rangle)\|_{L^2} \\
&\leq |\mathcal{K}_N| |\mathbb{T}^d|^{1/2} \max_{k \in \mathcal{K}_N} \|\mathcal{N}_1(w)_{1,k} - w_{1,k} \cos(\langle k, x \rangle)\|_{L^\infty} \\
&\quad + |\mathcal{K}_N| |\mathbb{T}^d|^{1/2} \max_{k \in \mathcal{K}_N} \|\mathcal{N}_1(w)_{2,k} - w_{2,k} \sin(\langle k, x \rangle)\|_{L^\infty} \\
&< \epsilon.
\end{aligned}$$

□

D.4. Proof of the Universal Approximation Theorem 2.5. As mentioned in the sketch of the proof of Theorem 2.5, we only need to consider the special case $s' = 0$, due to Lemma 2.6; i.e. given a continuous operator $\mathcal{G} : H^s(\mathbb{T}^d) \rightarrow L^2(\mathbb{T}^d)$, $K \subset H^s(\mathbb{T}^d)$ compact, and $\epsilon > 0$, we wish to construct a FNO $\mathcal{N} : H^s(\mathbb{T}^d) \rightarrow L^2(\mathbb{T}^d)$, such that $\sup_{a \in K} \|\mathcal{G}(a) - \mathcal{N}(a)\|_{L^2} \leq \epsilon$.

Throughout this proof, we set $d_a = d_u = 1$ for notational convenience. The general case with $d_a, d_u > 1$ follows analogously. For $N \in \mathbb{N}$, let P_N denote the orthogonal Fourier projection (B.6). First, we note that since $K \subset H^s(\mathbb{T}^d)$ is compact, one can prove by elementary arguments that the set \widetilde{K} given by

$$\widetilde{K} := K \cup \bigcup_{N \in \mathbb{N}} P_N K,$$

is compact. Since \mathcal{G} is continuous, its restriction to \widetilde{K} is uniformly continuous, i.e. there exists a modulus of continuity $\omega : [0, \infty) \rightarrow [0, \infty)$, such that

$$\|\mathcal{G}(a) - \mathcal{G}(a')\|_{L^2} \leq \omega(\|a - a'\|_{H^s}), \quad \forall a, a' \in \widetilde{K}.$$

From the definition of the projection \mathcal{G}_N (2.8), we have,

$$\begin{aligned}
\|\mathcal{G}(a) - \mathcal{G}_N(a)\|_{L^2} &\leq \|\mathcal{G}(a) - P_N \mathcal{G}(a)\|_{L^2} + \|P_N \mathcal{G}(a) - P_N \mathcal{G}(P_N a)\|_{L^2} \\
&\leq \|\mathcal{G}(a) - P_N \mathcal{G}(a)\|_{L^2} + \|\mathcal{G}(a) - \mathcal{G}(P_N a)\|_{L^2} \\
&\leq \sup_{v \in \mathcal{G}(\widetilde{K})} \|(1 - P_N)v\|_{L^2} + \omega \left(\sup_{a \in \widetilde{K}} \|(1 - P_N)a\|_{H^s} \right). \quad (\text{D.14})
\end{aligned}$$

Since \widetilde{K} is compact, also the image $\mathcal{G}(\widetilde{K})$ is compact leading to,

$$\limsup_{N \rightarrow \infty} \sup_{u \in \mathcal{G}(\widetilde{K})} \|(1 - P_N)u\|_{L^2} = 0 = \limsup_{N \rightarrow \infty} \sup_{a \in \widetilde{K}} \|(1 - P_N)a\|_{H^s}.$$

In particular, there exists $N \in \mathbb{N}$, such that

$$\|\mathcal{G}(a) - \mathcal{G}_N(a)\|_{L^2} \leq \epsilon, \quad \forall a \in K \subset \widetilde{K}. \quad (\text{D.15})$$

In the remainder of this proof, we will construct an FNO approximation of \mathcal{G}_N . In fact, we note that \mathcal{G}_N defines a continuous operator $\mathcal{G}_N : L^2(\mathbb{T}^d) \rightarrow L^2(\mathbb{T}^d)$,

via $a \mapsto P_N \mathcal{G}(P_N a)$, and the compact set K remains compact also as a subset of $L^2(\mathbb{T}^d)$. We will show that there exists a FNO $\mathcal{N} : L^2(\mathbb{T}^d) \rightarrow L^2(\mathbb{T}^d)$, such that

$$\sup_{a \in K} \|\mathcal{G}_N(a) - \mathcal{N}(a)\|_{L^2} < \epsilon.$$

Then the restriction of \mathcal{N} to $H^s(\mathbb{T}^d) \subset L^2(\mathbb{T}^d)$ provides an approximation of \mathcal{G} , such that

$$\sup_{a \in K} \|\mathcal{G}(a) - \mathcal{N}(a)\|_{L^2} < 2\epsilon.$$

Since $\epsilon > 0$ was arbitrary, the claim follows from this.

As outlined in the sketch, our proof for the existence of a FNO approximating \mathcal{G}_N relies on the decomposition (2.12), which in turn is defined via the Fourier conjugate operator $\widehat{\mathcal{G}}_N$ (2.11). Our aim is to show that each operator in the decomposition (2.12) can be approximated by FNOs to desired accuracy. To this end, let $\epsilon > 0$ be given, and choose $R_K, R_{\widehat{K}}, R_{\widehat{\mathcal{G}}} > 0$, such that

$$\begin{cases} K \subset B_{R_K}(0) := \{\|u\|_{L^2} \leq R_K\} \subset L^2(\mathbb{T}^d), \\ \mathcal{F}_N \circ P_N(B_{R_K}(0)) \subset \left[-\frac{R_{\widehat{K}}}{2}, \frac{R_{\widehat{K}}}{2}\right]^{2\mathcal{K}_N}, \\ \widehat{\mathcal{G}}_N([-R_{\widehat{K}}, R_{\widehat{K}}]^{2\mathcal{K}_N}) \subset \left[-\frac{R_{\widehat{\mathcal{G}}}}{2}, \frac{R_{\widehat{\mathcal{G}}}}{2}\right]^{2\mathcal{K}_N}. \end{cases} \quad (\text{D.16})$$

The reason for introducing $R_K, R_{\widehat{K}}, R_{\widehat{\mathcal{G}}}$ lies in the fact that in order to approximate \mathcal{G}_N by a composition of FNOs, we need to ensure that each FNO in this composition maps its own domain into the domain of the next FNO. The FNO approximations of the individual steps in the composition $\mathcal{G}_N = \mathcal{F}_N^{-1} \circ \widehat{\mathcal{G}}_N \circ (\mathcal{F}_N \circ P_N)$ are constructed below:

FNO approximation of \mathcal{F}_N^{-1} : We start with our construction of a FNO approximation of the last step in the composition. To this end, we are going to interpret the mapping

$$\mathcal{F}_N^{-1} : [-R, R]^{2\mathcal{K}_N} \subset \mathbb{R}^{2\mathcal{K}_N} \rightarrow L^2(\mathbb{T}^d),$$

as a mapping

$$\mathcal{F}_N^{-1} : \begin{cases} L^2(\mathbb{T}^d; [-R, R]^{2\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d), \\ \{\text{Re}(\widehat{v}_k), \text{Im}(\widehat{v}_k)\}_{|k| \leq N} \mapsto v(x), \end{cases} \quad (\text{D.17})$$

where the input $\{\text{Re}(\widehat{v}_k), \text{Im}(\widehat{v}_k)\}_{|k| \leq N} \in [-R, R]^{2\mathcal{K}_N}$ is identified with a *constant* function in $L^2(\mathbb{T}^d; [-R, R]^{2\mathcal{K}_N})$ (for non-constant input function $v(x)$, we apply \mathcal{F}_N^{-1} to the constant function $x \mapsto \int_{\mathbb{T}^d} v(\xi) d\xi$ to define the mapping (D.17) for general inputs). By Lemma 2.8, the mapping (D.17) can be approximated to any desired accuracy by an FNO $\mathcal{N}_{\text{IFT}} : L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d)$, such that

$$\|\mathcal{N}_{\text{IFT}}(\widehat{v}) - \mathcal{F}_N^{-1}(\widehat{v})\|_{L^2} \leq \epsilon/3, \quad (\text{D.18})$$

for all *constant* input functions $\widehat{v} \in L^2(\mathbb{T}^d; [-R, R]^{2\mathcal{K}_N})$.

FNO approximation of $\widehat{\mathcal{G}}_N$: We view the Fourier conjugate operator $\widehat{\mathcal{G}}_N$ (2.11) as a continuous mapping

$$\widehat{\mathcal{G}}_N : [-R_{\widehat{K}}, R_{\widehat{K}}]^{2\mathcal{K}_N} \subset \mathbb{R}^{2\mathcal{K}_N} \rightarrow \mathbb{R}^{2\mathcal{K}_N}.$$

Since $[-R_{\widehat{K}}, R_{\widehat{K}}]^{2\mathcal{K}_N}$ is compact, there exists a finite-dimensional canonical neural network $\widehat{\mathcal{N}} : \mathbb{R}^{2\mathcal{K}_N} \rightarrow \mathbb{R}^{2\mathcal{K}_N}$, such that

$$\sup_{\widehat{v} \in [-R_{\widehat{K}}, R_{\widehat{K}}]^{2\mathcal{K}_N}} \|\widehat{\mathcal{G}}_N(\widehat{v}) - \widehat{\mathcal{N}}(\widehat{v})\|_{\ell^2} \leq \epsilon/3. \quad (\text{D.19})$$

Furthermore, by (D.16), we have

$$\widehat{G}_N([-R_{\widehat{K}}, R_{\widehat{K}}]^{2\mathcal{K}_N}) \subset \left[-\frac{R_{\widehat{G}}}{2}, \frac{R_{\widehat{G}}}{2}\right]^{2\mathcal{K}_N}.$$

Therefore, by choosing a neural network approximation $\widehat{\mathcal{N}}$ with sufficiently high accuracy, we can also ensure that

$$\widehat{\mathcal{N}}([-R_{\widehat{K}}, R_{\widehat{K}}]^{2\mathcal{K}_N}) \subset [-R_{\widehat{G}}, R_{\widehat{G}}]^{2\mathcal{K}_N},$$

in addition to (D.19). Finally, we note that for $v \in L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$, the corresponding mapping

$$\widehat{\mathcal{N}} : L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}), \quad v(x) \mapsto \widehat{\mathcal{N}}(v(x)),$$

is in fact an FNO, with only *local* layers of the form

$$v_\ell(x) \mapsto \sigma(A_\ell v_\ell(x) + b_\ell), \quad (A_\ell \in \mathbb{R}^{d_v \times d_v}, b_\ell \in \mathbb{R}^{d_v}),$$

and where $d_v := 2|\mathcal{K}_N|$, i.e. a FNO for which all $P_\ell \equiv 0$ (cp. Remark 2.3). We shall thus identify $\widehat{\mathcal{N}}$ with this particular FNO, in the following.

FNO approximation of $\mathcal{F}_N \circ P_N$: Finally, we can similarly interpret

$$\mathcal{F}_N \circ P_N : B_{R_K}(0) \subset L^2(\mathbb{T}^d) \rightarrow \mathbb{R}^{2\mathcal{K}_N},$$

as a mapping

$$\mathcal{F}_N \circ P_N : \begin{cases} B_{R_K}(0) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N}), \\ v \mapsto \{\operatorname{Re}(\widehat{v}_k), \operatorname{Im}(\widehat{v}_k)\}_{|k| \leq N}, \end{cases} \quad (\text{D.20})$$

where the output $\{\operatorname{Re}(\widehat{v}_k), \operatorname{Im}(\widehat{v}_k)\}_{|k| \leq N} \in \mathbb{R}^{2\mathcal{K}_N}$ is a *constant* function in $L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$. By Lemma 2.7, the mapping (D.20) can be approximated to any desired accuracy by an FNO $\mathcal{N}_{\text{FT}} : B_{R_K}(0) \rightarrow L^2(\mathbb{T}^d; \mathbb{R}^{2\mathcal{K}_N})$ (with constant output functions). In particular, denoting $\operatorname{Lip}(\widehat{\mathcal{N}})$ the Lipschitz constant of the FNO constructed in the previous step, we can ensure that

$$\operatorname{Lip}(\widehat{\mathcal{N}}) \|\mathcal{F}_N P_N v - \mathcal{N}_{\text{FT}}(v)\|_{\ell^2} \leq \epsilon/3, \quad \forall v \in B_{R_K}(0), \quad (\text{D.21})$$

and furthermore, since by (D.16), we have

$$\mathcal{F}_N \circ P_N(B_{R_K}(0)) \subset \left[-\frac{R_{\widehat{K}}}{2}, \frac{R_{\widehat{K}}}{2}\right]^{2\mathcal{K}_N},$$

we can in addition ensure that

$$\mathcal{N}_{\text{FT}}(B_{R_K}(0)) \subset [-R_{\widehat{K}}, R_{\widehat{K}}]^{2\mathcal{K}_N}.$$

Error estimate for resulting FNO approximation: We now define a FNO $\mathcal{N}(a) := \mathcal{N}_{\text{FT}} \circ \widehat{\mathcal{N}} \circ \mathcal{N}_{\text{FT}}(a)$, where the right-hand side terms have been constructed

above. We note that

$$\begin{aligned}
& \sup_K \|\mathcal{G}_N - \mathcal{N}\|_{L^2} \\
& \leq \sup_{B_{R_K}(0)} \left\| \mathcal{F}_N^{-1} \circ \widehat{\mathcal{G}}_N \circ \mathcal{F}_N \circ P_N - \mathcal{N}_{\text{IFT}} \circ \widehat{\mathcal{N}} \circ \mathcal{N}_{\text{FT}} \right\|_{L^2} \\
& \leq \sup_{B_{R_K}(0)} \left\| \mathcal{F}_N^{-1} \circ \widehat{\mathcal{G}}_N \circ \mathcal{F}_N \circ P_N - \mathcal{F}_N^{-1} \circ \widehat{\mathcal{N}} \circ \mathcal{N}_{\text{FT}} \right\|_{L^2} \\
& \quad + \sup_{B_{R_K}(0)} \left\| \mathcal{F}_N^{-1} \circ \widehat{\mathcal{N}} \circ \mathcal{N}_{\text{FT}} - \mathcal{N}_{\text{IFT}} \circ \widehat{\mathcal{N}} \circ \mathcal{N}_{\text{FT}} \right\|_{L^2} \\
& \leq \sup_{B_{R_K}(0)} \left\| \widehat{\mathcal{G}}_N \circ \mathcal{F}_N \circ P_N - \widehat{\mathcal{N}} \circ \mathcal{N}_{\text{FT}} \right\|_{L^2} \\
& \quad + \sup_{\widehat{\mathcal{N}}(\mathcal{N}_{\text{FT}}(B_{R_K}(0)))} \left\| \mathcal{F}_N^{-1} - \mathcal{N}_{\text{IFT}} \right\|_{L^2} \\
& =: (I) + (II).
\end{aligned}$$

For the second term (II), we note that

$$\widehat{\mathcal{N}}(\mathcal{N}_{\text{FT}}(B_{R_K}(0))) \subset \widehat{\mathcal{N}}([-R_{\widehat{K}}, R_{\widehat{K}}]^{2\mathcal{K}_N}) \subset [-R_{\widehat{\mathcal{G}}}, R_{\widehat{\mathcal{G}}}]^{2\mathcal{K}_N},$$

and hence, by (D.18), we can bound

$$(II) \leq \sup_{[-R_{\widehat{\mathcal{G}}}, R_{\widehat{\mathcal{G}}}]^{2\mathcal{K}_N}} \left\| \mathcal{F}_N^{-1} - \mathcal{N}_{\text{IFT}} \right\|_{L^2} \leq \epsilon/3.$$

To estimate the first term (I), we note that

$$\begin{aligned}
(I) & = \sup_{B_{R_K}(0)} \left\| \widehat{\mathcal{G}}_N \circ \mathcal{F}_N \circ P_N - \widehat{\mathcal{N}} \circ \mathcal{N}_{\text{FT}} \right\|_{L^2} \\
& \leq \sup_{B_{R_K}(0)} \left\| \widehat{\mathcal{G}}_N \circ \mathcal{F}_N \circ P_N - \widehat{\mathcal{N}} \circ \mathcal{F}_N \circ P_N \right\|_{L^2} \\
& \quad + \sup_{B_{R_K}(0)} \left\| \widehat{\mathcal{N}} \circ \mathcal{F}_N \circ P_N - \widehat{\mathcal{N}} \circ \mathcal{N}_{\text{FT}} \right\|_{L^2} \\
& =: (Ia) + (Ib).
\end{aligned}$$

To estimate (Ia), we note that

$$\mathcal{F}_N(P_N(B_{R_K}(0))) \subset [-R_{\widehat{K}}, R_{\widehat{K}}]^{2\mathcal{K}_N},$$

and hence

$$(Ia) \leq \sup_{[-R_{\widehat{K}}, R_{\widehat{K}}]^{2\mathcal{K}_N}} \left\| \widehat{\mathcal{G}}_N - \widehat{\mathcal{N}} \right\|_{L^2} \leq \epsilon/3,$$

by (D.19). Finally, to estimate (Ib), we note that

$$(Ib) \leq \text{Lip}(\widehat{\mathcal{N}}) \sup_{B_{R_K}(0)} \left\| \mathcal{F}_N \circ P_N - \mathcal{N}_{\text{FT}} \right\|_{L^2} \leq \epsilon/3,$$

by (D.21). Combining these estimates, we conclude that

$$\sup_{a \in K} \|\mathcal{G}_N(a) - \mathcal{N}(a)\|_{L^2} \leq \epsilon,$$

This shows that the continuous operator \mathcal{G}_N can be approximated by a FNO \mathcal{N} to any desired accuracy $\epsilon > 0$, and together with (2.9) concludes our proof of the universal approximation theorem 2.5 for the special case $s' = 0$. The general case with $s' \geq 0$ now follows from Lemma 2.6.

D.5. Proof of Theorem 2.14.

Proof. The proof involves 4 steps.

Step 1: We may wlog assume that all biases $b_\ell(x) \in C^\infty$.

Indeed, it is easy to see that under the finite width assumption and for Lipschitz continuous σ , each layer \mathcal{L}_ℓ defines a Lipschitz continuous mapping $L^2 \rightarrow L^2$. Replacing $b_\ell(x)$ by its δ -mollification $b_\ell^\delta(x)$ we obtain a new layer \mathcal{L}_ℓ^δ , i.e.

$$\mathcal{L}_\ell^\delta(v) = \sigma(W_\ell v(x) + b_\ell^\delta(x) + \mathcal{F}^{-1}(P_\ell \mathcal{F}v)(x)).$$

Then we have

$$\|\mathcal{L}_\ell(v) - \mathcal{L}_\ell^\delta(v)\|_{L^2} \lesssim_{\text{Lip}(\sigma), \ell} \|b_\ell - b_\ell^\delta\|_{L^2} \rightarrow 0, \quad (\text{D.22})$$

and uniformly in $\delta > 0$, we have

$$\|\mathcal{L}_\ell^\delta(v) - \mathcal{L}_\ell^\delta(v')\|_{L^2} \lesssim_{\text{Lip}(\sigma), \ell} \|v - v'\|_{L^2}. \quad (\text{D.23})$$

In particular, this implies that $\text{Lip}(\mathcal{L}_\ell^\delta) \leq C(\text{Lip}(\sigma), \ell)$ is uniformly bounded in $\delta \geq 0$. Thus, there exists $\bar{M} \geq 1$, such that $\text{Lip}(\mathcal{Q}) \leq \bar{M}$ and $\text{Lip}(\mathcal{L}_\ell^\delta) \leq \bar{M}$, for all $\ell = 1, \dots, L$ and $\delta > 0$. Properties (D.22) and (D.23) are sufficient to show that with a sufficiently small choice of $\delta > 0$, we have

$$\sup_{\|a\|_{H^s} \leq B} \|\mathcal{N}(a) - \mathcal{N}^\delta(a)\|_{L^2} \leq \epsilon, \quad (\text{D.24})$$

where $\mathcal{N}^\delta := \mathcal{Q} \circ \mathcal{L}_L^\delta \circ \dots \circ \mathcal{L}_1^\delta \circ \mathcal{R}$: To see this, we introduce $\mathcal{N}_\ell^\delta := \mathcal{Q} \circ \mathcal{L}_L^\delta \circ \dots \circ \mathcal{L}_{\ell+1}^\delta \circ \mathcal{L}_\ell \circ \dots \circ \mathcal{L}_1 \circ \mathcal{R}$, and we note that for any $a \in L^2$, we have

$$\|\mathcal{N}_\ell^\delta(a) - \mathcal{N}_{\ell-1}^\delta(a)\|_{L^2} \leq \text{Lip}(\mathcal{Q} \circ \mathcal{L}_L^\delta \circ \dots \circ \mathcal{L}_{\ell+1}^\delta) \|\mathcal{L}_\ell(a') - \mathcal{L}_\ell^\delta(a')\|_{L^2},$$

where $a' := \mathcal{L}_{\ell-1} \circ \dots \circ \mathcal{L}_1 \circ \mathcal{R}(a)$. Using the uniform Lipschitz bound \bar{M} and the estimate (D.22), we obtain

$$\|\mathcal{N}_\ell^\delta(a) - \mathcal{N}_{\ell-1}^\delta(a)\|_{L^2} \leq \bar{M}^L \|b_\ell - b_\ell^\delta\|_{L^2}.$$

It now follows from the telescoping sum $\mathcal{N}(a) - \mathcal{N}^\delta(a) = \mathcal{N}_L^\delta(a) - \mathcal{N}_0^\delta(a) = \sum_{\ell=1}^L (\mathcal{N}_\ell^\delta(a) - \mathcal{N}_{\ell-1}^\delta(a))$ that

$$\begin{aligned} \|\mathcal{N}(a) - \mathcal{N}^\delta(a)\|_{L^2} &\leq \sum_{\ell=1}^L \|\mathcal{N}_\ell^\delta(a) - \mathcal{N}_{\ell-1}^\delta(a)\|_{L^2} \\ &\leq \bar{M}^L \sum_{\ell=1}^L \|b_\ell - b_\ell^\delta\|_{L^2}. \end{aligned}$$

Since the bound on the right-hand side is independent of a , and since $\lim_{\delta \rightarrow 0} \|b_\ell - b_\ell^\delta\|_{L^2} = 0$, choosing $\delta > 0$ sufficiently small, we obtain (D.24).

Step 2: Using Step 1, we will assume wlog that all biases are smooth, in the following. Given $a \in H^s$, let $a_\ell := \mathcal{L}_\ell \circ \dots \circ \mathcal{L}_1 \circ \mathcal{R}(a)$. We claim that there exists a constant $B' > 0$, depending only on B , the activation function σ and the biases $b_\ell(x)$, such that $\|a_\ell\|_{H^s} \leq B'$, provided that $\|a\|_{H^s} \leq B$ and $\ell = 1, \dots, L$.

Indeed, it is a well-known fact² that if $v \in H^s$, $s > d/2$, and if $\sigma \in C^{\lfloor s \rfloor + 1}$ (satisfied by assumption), then $\sigma \circ v \in H^s$. In fact, there exists a monotonically increasing function C_σ , such that $\|\sigma \circ v\|_{H^s} \leq C_\sigma(\|v\|_{H^s})$. The claimed existence of B' now follows immediately from this result on compositions with σ , and the observation that the linear part of each layer define a continuous (bounded, affine)

²This is well-known in the integer case $s \in \mathbb{N}$, but remains true also for fractional Sobolev spaces [4].

mapping $H^s \rightarrow H^s$, due to the finite width assumption, and the assumption that $b_\ell \in C^\infty \subset H^s$.

Step 3: Let ς_1, ς_2 be given such that $s \geq \varsigma_1 > \varsigma_2 > d/2$. Each layer defines a Hölder continuous mapping $\mathcal{L}_\ell : D_\ell \subset H^{\varsigma_1} \rightarrow H^{\varsigma_2}$, with Hölder exponent $\alpha = 1 - \varsigma_2/s$, and where D_ℓ denotes the image of the set $\{\|a\|_{H^s} \leq B\}$ under the previous $\ell - 1$ layers.

By Step 2, we have that $\|\mathcal{L}_\ell(v)\|_{H^s} \leq B'$ for all $v \in D_\ell$. Furthermore, as observed in Step 1, (D.23), $\mathcal{L}_\ell : L^2 \rightarrow L^2$ is Lipschitz continuous. It follows that for any $v \in D_\ell$, we have, by the interpolation inequality, that

$$\begin{aligned} \|\mathcal{L}_\ell(v) - \mathcal{L}_\ell(v')\|_{H^{\varsigma_2}} &\lesssim \|\mathcal{L}_\ell(v) - \mathcal{L}_\ell(v')\|_{L^2}^{1-\varsigma_2/s} \|\mathcal{L}_\ell(v) - \mathcal{L}_\ell(v')\|_{H^s}^{\varsigma_2/s} \\ &\lesssim_{\varsigma_2, B'} \|\mathcal{L}_\ell(v) - \mathcal{L}_\ell(v')\|_{L^2}^\alpha \\ &\lesssim_{\varsigma_2, B', \sigma, \ell} \|v - v'\|_{L^2}^\alpha \\ &\lesssim \|v - v'\|_{H^{\varsigma_1}}^\alpha. \end{aligned}$$

Step 4: We use Steps 1-3 to conclude that for a sufficiently large choice of $N \in \mathbb{N}$, the Ψ -FNO $\widehat{\mathcal{N}} : L_N^2 \rightarrow L_N^2$, given by

$$\widehat{\mathcal{N}}(a) := \mathcal{Q} \circ \mathcal{I}_N \circ \mathcal{L}_L \circ \mathcal{I}_N \circ \mathcal{L}_{L-1} \circ \mathcal{I}_N \circ \cdots \circ \mathcal{L}_1 \circ \mathcal{I}_N \circ \mathcal{R},$$

satisfies

$$\sup_{\|a\|_{H^s} \leq B} \|\mathcal{N}(a) - \widehat{\mathcal{N}}(a)\|_{L^2} \leq \epsilon.$$

In the following, we will denote by $\widehat{\mathcal{L}}_\ell := \mathcal{I}_N \circ \mathcal{L}_\ell \circ \mathcal{I}_N$ the ‘‘pseudo-spectral projection’’ of the layer \mathcal{L}_ℓ , and we observe that

$$\widehat{\mathcal{N}} = \mathcal{Q} \circ \widehat{\mathcal{L}}_L \circ \cdots \circ \widehat{\mathcal{L}}_1 \circ \mathcal{R}.$$

Fix a sequence $s = \varsigma_0 > \varsigma_1 > \varsigma_2 > \cdots > \varsigma_L > d/2$. By Step 3, we may view the ℓ -th layer \mathcal{L}_ℓ as a Hölder continuous mapping $\mathcal{L}_\ell : H^{\varsigma_{\ell-1}} \rightarrow H^{\varsigma_\ell}$. Furthermore, since $\varsigma_\ell > d/2$ for all ℓ , we also have that $\mathcal{I}_N : H^{\varsigma_\ell} \rightarrow H^{\varsigma_\ell}$ is a bounded linear operator, with operator norm that is uniformly bounded in N . In particular, this implies that there exists a constant $\overline{C} \geq 1$, independent of N , such that, for all $a, a' \in D_\ell$:

$$\begin{aligned} \|\widehat{\mathcal{L}}_\ell(a) - \widehat{\mathcal{L}}_\ell(a')\|_{H^{\varsigma_\ell}} &= \|\mathcal{I}_N \mathcal{L}_\ell(\mathcal{I}_N a) - \mathcal{I}_N \mathcal{L}_\ell(\mathcal{I}_N a')\|_{H^{\varsigma_\ell}} \\ &\leq \|\mathcal{I}_N\| \|\mathcal{L}_\ell\|_{C^{1-\varsigma_\ell/s}} \|\mathcal{I}_N\|^{1-\varsigma_\ell/s} \|a - a'\|_{H^{\varsigma_{\ell-1}}}^{1-\varsigma_\ell/s} \\ &\leq \overline{C} \|a - a'\|_{H^{\varsigma_{\ell-1}}}^{1-\varsigma_\ell/s}. \end{aligned}$$

We now introduce $\widehat{\mathcal{N}}_\ell := \mathcal{Q} \circ \widehat{\mathcal{L}}_L \circ \cdots \circ \widehat{\mathcal{L}}_{\ell+1} \circ \mathcal{L}_\ell \circ \cdots \circ \mathcal{L}_1 \circ \mathcal{R}$, for any $\ell = 0, \dots, L$. Then, $\mathcal{N}(a) - \widehat{\mathcal{N}}(a) = \sum_{\ell=1}^L (\widehat{\mathcal{N}}_\ell(a) - \widehat{\mathcal{N}}_{\ell-1}(a))$, and using the Hölder regularity of the $\widehat{\mathcal{L}}_\ell$, we obtain

$$\begin{aligned} \|\widehat{\mathcal{N}}_\ell(a) - \widehat{\mathcal{N}}_{\ell-1}(a)\|_{L^2} &\leq \|\widehat{\mathcal{N}}_\ell(a) - \widehat{\mathcal{N}}_{\ell-1}(a)\|_{H^{\varsigma_L}} \\ &\leq \text{Lip}(\mathcal{Q}) \overline{C}^L \|\mathcal{L}_\ell(a') - \widehat{\mathcal{L}}_\ell(a')\|_{H^{\varsigma_\ell}}^{\beta_\ell}, \end{aligned}$$

where $a' := \mathcal{L}_{\ell-1} \circ \cdots \circ \mathcal{L}_1 \circ \mathcal{R}(a)$ belongs to $D_\ell \subset \{a' \in H^s \mid \|a'\|_{H^s} \leq B'\}$ (cp. Step 3), and $\beta_\ell := \prod_{k=\ell+1}^L (1 - \frac{\varsigma_k}{s})$ is the Hölder exponent of the composition of the layers $\widehat{\mathcal{L}}_L \circ \cdots \circ \widehat{\mathcal{L}}_{\ell+1} : H^{\varsigma_{\ell+1}} \rightarrow H^{\varsigma_L}$. To estimate the last difference, we recall that if $\|a\|_{H^s} \leq B$, then $\|a'\|_{H^s} \leq B'$, by Step 2. In particular, it follows from the uniform

Hölder continuity of the layers \mathcal{L}_ℓ , established in Step 3 and the pseudo-spectral approximation estimate, that

$$\begin{aligned} \|\mathcal{L}_\ell(a') - \widehat{\mathcal{L}}_\ell(a')\|_{H^{s_\ell}} &= \|\mathcal{L}_\ell(a') - \mathcal{I}_N \mathcal{L}_\ell(\mathcal{I}_N a')\|_{H^{s_\ell}} \\ &\leq \|\mathcal{L}_\ell(a') - \mathcal{I}_N \mathcal{L}_\ell(a')\|_{H^{s_\ell}} + \|\mathcal{I}_N\| \|\mathcal{L}_\ell(a') - \mathcal{L}_\ell(\mathcal{I}_N a')\|_{H^{s_\ell}} \\ &\lesssim \frac{B'}{N^{-(s-s_\ell)}} + \|a' - \mathcal{I}_N a'\|_{H^{s_{\ell-1}}}^{1-s_\ell/s} \\ &\lesssim \frac{B'}{N^{-(s-s_\ell)}} + \left(\frac{B'}{N^{-(s-s_{\ell-1})}} \right)^{(1-s_\ell/s)}, \end{aligned}$$

where the implied constant is independent of N . In particular, we have

$$\lim_{N \rightarrow \infty} \sup_{a' \in D_\ell} \|\mathcal{L}_\ell(a') - \widehat{\mathcal{L}}_\ell(a')\|_{H^{s_\ell}} = \lim_{N \rightarrow \infty} \sup_{a' \in D_\ell} \|\mathcal{L}_\ell(a') - \mathcal{L}_\ell(\mathcal{I}_N a')\|_{H^{s_\ell}} = 0,$$

for all $\ell = 1, \dots, L$. Choosing N sufficiently large, we can thus ensure that

$$\begin{aligned} \sup_{\|a\|_{H^s} \leq B} \|\mathcal{N}(a) - \widehat{\mathcal{N}}(a)\|_{L^2} &\leq \sum_{\ell=1}^L \sup_{\|a\|_{H^s} \leq B} \|\widehat{\mathcal{N}}_\ell(a) - \widehat{\mathcal{N}}_{\ell-1}(a)\|_{L^2} \\ &\leq \text{Lip}(\mathcal{Q}) \overline{C}^L \sum_{\ell=1}^L \sup_{a' \in D_\ell} \|\mathcal{L}_\ell(a') - \widehat{\mathcal{L}}_\ell(a')\|_{H^{s_\ell}}^{\beta_\ell} \\ &\leq \epsilon. \end{aligned}$$

This concludes the proof. \square

D.6. Proof of Lemma 2.16.

Proof. Let $\epsilon > 0$ be given. Let $\mathcal{G} : H^s \rightarrow H^{s'}$ be continuous operator, and $K \subset H^s$ a compact subset. We wish to show that there exists $N_0 \in \mathbb{N}$, such that for any $N \geq N_0$, there exists a Ψ -FNO $\widehat{\mathcal{N}} : L_N^2 \rightarrow L_N^2$ such that

$$\sup_{a \in K} \|\mathcal{G}(a) - \widehat{\mathcal{N}}(a)\|_{H^{s'}} \leq \epsilon. \quad (\text{D.25})$$

Choose $M \in \mathbb{N}$, such that

$$\sup_{a \in K} \|P_M \mathcal{G}(a) - \mathcal{G}(a)\|_{H^{s'}} \leq \epsilon/2. \quad (\text{D.26})$$

Fix $\delta > 0$ for the moment. We will specify a suitable choice of $\delta = \delta(M, s', \mathbb{T}^d, \epsilon) > 0$ at the end of this proof. We emphasize that δ will depend only on parameters already introduced at this point of the proof. By assumption of the validity of the universal approximation theorem for Ψ -FNOs for $s' = 0$, we can choose $\tilde{N}_0 = \tilde{N}_0(\mathcal{G}, K, \delta) \in \mathbb{N}$, depending on \mathcal{G} , K and δ , such that for any $N \geq \tilde{N}_0$, there exists a Ψ -FNO $\mathcal{N} : L_N^2 \rightarrow L_N^2$, such that

$$\sup_{a \in K} \|\mathcal{G}(a) - \mathcal{N}(a)\|_{L^2} \leq \delta.$$

We define $N_0 := \max(\tilde{N}_0, M)$. By our choice of \tilde{N}_0 and M , the constant N_0 depends only on the underlying operator \mathcal{G} , the compact set K and the parameter $\delta > 0$. Let $N \geq N_0$. We claim that if $\delta = \delta(M, s', \mathbb{T}^d, \epsilon) > 0$ is been chosen sufficiently small, then for any $N \geq N_0$, there exists a Ψ -FNO $\widehat{\mathcal{N}} : L_N^2 \rightarrow L_N^2$ satisfying (D.25).

To see this, let $N \geq N_0$, and let $\mathcal{N} : L_N^2 \rightarrow L_N^2$ be a Ψ -FNO such that $\sup_{a \in K} \|\mathcal{G}(a) - \mathcal{N}(a)\|_{L^2} \leq \delta$. Then we have

$$\sup_{a \in K} \|P_M \mathcal{G}(a) - P_M \mathcal{N}(a)\|_{H^{s'}} \leq CN^{s'} \delta, \quad (\text{D.27})$$

where $C = C(\mathbb{T}^d, s') > 0$ is independent of N and δ . Let $m \in \mathbb{N}$ be the smallest natural number strictly bigger than s' and $d/2$, i.e. $m > \max(s', d/2)$. Let

$$K' := \mathcal{N}(K) \subset L_N^2, \quad (\text{D.28})$$

be the compact image of $K \subset H^s$ under the (continuous) mapping $\mathcal{N} : H^s \rightarrow L_N^2$. We note that since L_N^2 is a finite-dimensional space, the L^2 -norm is equivalent to the H^m -norm on L_N^2 . This implies that $K' \subset H^m$ is also compact when considered as a subset of H^m , with respect to norm $\|\cdot\|_{H^m}$. By Lemma D.1, there exists a single-layer FNO $\mathcal{L} : H^m \rightarrow H^m$, such that

$$\sup_{v \in K'} \|P_M v - \mathcal{L}(v)\|_{H^m} \leq \delta. \quad (\text{D.29})$$

We define a Ψ -FNO $\widehat{\mathcal{N}} : L_N^2 \rightarrow L_N^2$ by the following composition

$$\widehat{\mathcal{N}} := \mathcal{I}_N \circ \mathcal{L} \circ \mathcal{I}_N \circ \mathcal{N} = \mathcal{I}_N \circ \mathcal{L} \circ \mathcal{N},$$

where \mathcal{I}_N is the pseudo-spectral projection operator. We note that, since $N \geq N_0 \geq M$, we have $\mathcal{I}_N P_M = P_M$, and recall that for $m > d/2$, the mapping $\mathcal{I}_N : H^m \rightarrow H^m$ is continuous, with an operator norm that can be bounded independently of N . Hence, we can estimate, for any $a \in K$,

$$\begin{aligned} \|P_M \mathcal{N}(a) - \widehat{\mathcal{N}}(a)\|_{H^{s'}} &= \|\mathcal{I}_N \circ P_M \mathcal{N}(a) - \mathcal{I}_N \circ \mathcal{L} \circ \mathcal{N}(a)\|_{H^{s'}} \\ &\stackrel{(m > s')}{\leq} \|\mathcal{I}_N \circ P_M \mathcal{N}(a) - \mathcal{I}_N \circ \mathcal{L} \circ \mathcal{N}(a)\|_{H^m} \\ &\leq C \|P_M v - \mathcal{L}(v)\|_{H^m}, \end{aligned}$$

where $C = C(\mathbb{T}^d, m) > 0$ is independent of M and N , and where $v := \mathcal{N}(a) \in K'$, with $K' = \mathcal{N}(K)$ defined in (D.28). From (D.29), we conclude that

$$\sup_{a \in K} \|P_M \mathcal{N}(a) - \widehat{\mathcal{N}}(a)\|_{H^{s'}} \leq C\delta, \quad (\text{D.30})$$

where $C = C(\mathbb{T}^d, m) > 0$ is independent of M , N and δ . In fact, since m is defined as the smallest integer $> \max(s', d/2)$, the above constant only depends on s' and the dimension d of the domain, i.e. we have $C = C(\mathbb{T}^d, s')$. Combining (D.26), (D.27) and (D.30), we find that for $N \geq N_0$ and for any $\delta > 0$, there exists a Ψ -FNO $\widehat{\mathcal{N}} : L_N^2 \rightarrow L_N^2$, such that

$$\sup_{a \in K} \|\mathcal{G}(a) - \widehat{\mathcal{N}}(a)\|_{H^{s'}} \leq \epsilon/2 + C(1 + N^{s'})\delta,$$

where $C = C(\mathbb{T}^d, s') > 0$ is a constant independent of δ . Thus, if we choose $\delta = \delta(M, s', \mathbb{T}^d, \epsilon) > 0$ at the beginning of this proof sufficiently small to ensure that

$$C(1 + N^{s'})\delta \leq \epsilon/2,$$

then we conclude that for any $N \geq N_0$, there exists a Ψ -FNO $\widehat{\mathcal{N}} : L_N^2 \rightarrow L_N^2$, such that

$$\sup_{a \in K} \|\mathcal{G}(a) - \widehat{\mathcal{N}}(a)\|_{H^{s'}} \leq \epsilon.$$

This concludes the proof. \square

D.7. Technical Results on the structure of Ψ -FNOs. Recall section 2.5.1 where we introduced σ layers and \mathcal{F} -layers and claimed that Ψ -FNOs can be decomposed in terms of these layers. We have the following series of lemmas, which make this observation precise.

Lemma D.2. Let $\widehat{\mathcal{N}} : \mathbb{R}^{d_a} \rightarrow \mathbb{R}^{d_u}$ be a (ordinary) neural network with activation function σ . For any $N \in \mathbb{N}$, the mapping

$$\mathcal{N} : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_a}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_u}), \quad a(x) \mapsto \mathcal{N}(a)(x) := \mathcal{I}_N \widehat{\mathcal{N}}(a(x)),$$

can be represented by a Ψ -FNO, with

$$\text{width}(\mathcal{N}) = |\mathcal{J}_N| \text{width}(\widehat{\mathcal{N}}), \quad \text{depth}(\mathcal{N}) = \text{depth}(\widehat{\mathcal{N}}), \quad \text{lift}(\mathcal{N}) = \text{width}(\widehat{\mathcal{N}}).$$

The proof is straight-forward and completely analogous to the same statement for FNOs in Remark 2.3. We also note the following lemma, which allows us in practice to replace \mathcal{F} -layers by proper FNO layers:

Lemma D.3 (Linear approximation lemma). Assume that the activation function $\sigma \in C^2$ is twice continuously differentiable and non-constant. Let $\widehat{\mathcal{N}}$ be a linear Ψ -FNO layer of the form

$$\widehat{\mathcal{L}} : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_v}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_v}), \quad \widehat{\mathcal{L}}(v_N) := W v_N(x_j) + b_j + \mathcal{F}_N^{-1}(P \mathcal{F}_N v_N)_j.$$

Then there exists a constant $C > 0$, such that for any $\epsilon, B > 0$, there exists a Ψ -FNO $\mathcal{N} : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_v}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_v})$, such that

$$\sup_{\|v_N\|_{L^2} \leq B} \|\mathcal{N}(v_N) - \widehat{\mathcal{L}}(v_N)\|_{L^2} \leq \epsilon,$$

and

$$\text{width}(\mathcal{N}) \leq C N^d, \quad \text{depth}(\mathcal{N}) \leq C, \quad \text{lift}(\mathcal{N}) \leq C.$$

Proof. This follows from the observation that, by assumption, there exists $x_0 \in \mathbb{R}$, such that $\sigma'(x_0) \neq 0$. By Taylor expansion, we have

$$\frac{\sigma(x_0 + hy) - \sigma(x_0 - hy)}{h\sigma'(x_0)} = y + R(h, y),$$

where $|R(h, y)| \leq C \overline{B}^2 h$ for any $y \in [-\overline{B}, \overline{B}]$ and $0 < h \leq 1$. Replacing y by the output of the linear layer $\widehat{\mathcal{L}}$, and choosing \overline{B} sufficiently large, we obtain with

$$\mathcal{N}_h(v_n) := \frac{\sigma(x_0 + h\widehat{\mathcal{L}}(v_N)) - \sigma(x_0 - h\widehat{\mathcal{L}}(v_N))}{h\sigma'(x_0)},$$

that

$$\|\widehat{\mathcal{L}}(v_N) - \mathcal{N}_h(v_N)\|_{L^2} = O(h),$$

uniformly for all $v_N \in L_N^2(\mathbb{T}^d; \mathbb{R}^{d_v})$, such that $\|v_N\|_{L^2} \leq B$, and for $0 < h \leq 1$. In particular, choosing $h = h(\epsilon) > 0$ sufficiently small, we can ensure that

$$\|\widehat{\mathcal{L}}(v_N) - \mathcal{N}_h(v_N)\|_{L^2} \leq \epsilon,$$

for all $v_n \in L_N^2(\mathbb{T}^d; \mathbb{R}^{d_v})$, such that $\|v_N\|_{L^2} \leq B$. The proof is concluded by observing that the mapping $v_N \mapsto \mathcal{N}_h(v_N)$ defines a Ψ -FNO, and that the size (width, depth, lift) is uniformly bounded in h . \square

According to Lemma D.2, any composition of σ -layers can be identified with a Ψ -FNO. Lemma D.3 shows that we can approximate \mathcal{F} -layers to arbitrary accuracy with a Ψ -FNO of uniformly bounded size. In the following lemma, we record the simple fact that a composition of Ψ -FNOs is again representable by a Ψ -FNO:

Lemma D.4 (Composition Lemma). For $p \in \mathbb{N}$, let $\mathcal{N}_1, \dots, \mathcal{N}_p$ be Ψ -FNOs, defined with respect to the same grid $\{x_j\}_{j \in \mathcal{J}_N}$, and defining operators $\mathcal{N}_k : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_{k-1}}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_k})$ for $d_0, \dots, d_p \in \mathbb{N}$, such that the composition $\mathcal{N}_p \circ \dots \circ \mathcal{N}_1 : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_0}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_p})$ is well-defined. Then there exists a Ψ -FNO $\mathcal{N} : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_0}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_p})$, such that $\mathcal{N} = \mathcal{N}_p \circ \dots \circ \mathcal{N}_1$, and such that $\text{depth}(\mathcal{N}) \leq \sum_{k=1}^p \text{depth}(\mathcal{N}_k)$, $\text{width}(\mathcal{N}) \leq \max_{k=1, \dots, p} \text{width}(\mathcal{N}_k)$ and $\text{lift}(\mathcal{N}) \leq \max_{k=1, \dots, p} \text{lift}(\mathcal{N}_k)$.

The proof of the previous lemma is straight-forward, and only requires a padding of each layer with zeros to achieve a uniform lifting dimension d_v across all layers. The following lemma will be central for our approximation and complexity estimates:

Lemma D.5 (Replacement Lemma). Assume that σ is Lipschitz continuous. For $p \in \mathbb{N}$, let $\widehat{\mathcal{N}}_1, \dots, \widehat{\mathcal{N}}_p$ be continuous operators $\widehat{\mathcal{N}}_k : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_{k-1}}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_k})$ for $d_0, \dots, d_p \in \mathbb{N}$, such that the composition $\widehat{\mathcal{N}} := \widehat{\mathcal{N}}_p \circ \dots \circ \widehat{\mathcal{N}}_1 : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_0}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_p})$ is well-defined. Assume that there exist constants $D_k, W_k, L_k > 0$, $k = 1, \dots, p$, such that for any $\epsilon, M > 0$, there exists a Ψ -FNO $\mathcal{N}_k : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_{k-1}}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_k})$, such that

$$\sup_{\|u\|_{L^2} \leq M} \|\widehat{\mathcal{N}}_k(u) - \mathcal{N}_k(u)\|_{L_N^2} \leq \epsilon,$$

and $\text{depth}(\mathcal{N}_k) \leq D_k$, $\text{width}(\mathcal{N}_k) \leq W_k$, $\text{lift}(\mathcal{N}_k) \leq L_k$. Then for any $\epsilon > 0$ and $M > 0$, there exists a Ψ -FNO $\mathcal{N} : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_0}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_p})$, such that

$$\sup_{\|u\|_{L^2} \leq M} \|\widehat{\mathcal{N}}(u) - \mathcal{N}(u)\|_{L_N^2} \leq \epsilon,$$

and $\text{depth}(\mathcal{N}) \leq \sum_{k=1}^p D_k$, $\text{width}(\mathcal{N}) \leq \max_{k=1, \dots, p} W_k$, $\text{lift}(\mathcal{N}) \leq \max_{k=1, \dots, p} L_k$.

Proof. We prove the statement by induction on $p \in \mathbb{N}$. We start the induction at $p = 2$: Let $\epsilon > 0$ and $M > 0$ be given, and let $\widehat{\mathcal{N}}_1, \widehat{\mathcal{N}}_2$ satisfy the hypotheses of this lemma. Since $B_M(0) \subset L_N^2(\mathbb{T}^d; \mathbb{R}^{d_0})$, the closed ball of radius M , is compact in the finite-dimensional space $L_N^2(\mathbb{T}^d; \mathbb{R}^{d_0})$, and since $\widehat{\mathcal{N}}_1$ is continuous, it follows that the image $\widehat{\mathcal{N}}_1(B_M(0))$ is also compact; in particular, there exists $M_1 > 0$, such that $\widehat{\mathcal{N}}_1(B_M(0)) \subset B_{M_1}(0)$. By assumption, there exists a Ψ -FNO $\mathcal{N}_2 : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_1}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_2})$, such that

$$\sup_{\|u\|_{L^2} \leq 2M_1} \|\widehat{\mathcal{N}}_2(u) - \mathcal{N}_2(u)\|_{L_N^2} \leq \epsilon/2.$$

We note that (for Lipschitz continuous σ) $\widehat{\mathcal{N}}_2$ is Lipschitz continuous on $B_{2M_1}(0) \subset L_N^2$. Let $\text{Lip}(\widehat{\mathcal{N}}_2)$ denote the corresponding Lipschitz constant. By assumption on $\widehat{\mathcal{N}}_1$, there exists a Ψ -FNO $\mathcal{N}_1 : L_N^2(\mathbb{T}^d; \mathbb{R}^{d_0}) \rightarrow L_N^2(\mathbb{T}^d; \mathbb{R}^{d_1})$, such that

$$\sup_{\|u\|_{L^2} \leq M} \|\widehat{\mathcal{N}}_1(u) - \mathcal{N}_1(u)\|_{L_N^2} \leq \min\left(\frac{\epsilon}{2\text{Lip}(\widehat{\mathcal{N}}_2)}, M_1\right).$$

Note that this estimate implies in particular that

$$\|\mathcal{N}_1(u)\|_{L^2} \leq \|\widehat{\mathcal{N}}_1(u) - \mathcal{N}_1(u)\|_{L^2} + \|\widehat{\mathcal{N}}_1(u)\|_{L^2} \leq 2M_1.$$

Thus, we can estimate

$$\begin{aligned}
& \sup_{\|u\|_{L^2} \leq M} \|\widehat{\mathcal{N}}_2 \circ \widehat{\mathcal{N}}_1(u) - \mathcal{N}_2 \circ \mathcal{N}_1(u)\|_{L^2} \\
& \leq \sup_{\|u\|_{L^2} \leq M} \|\widehat{\mathcal{N}}_2 \circ \widehat{\mathcal{N}}_1(u) - \widehat{\mathcal{N}}_2 \circ \mathcal{N}_1(u)\|_{L^2} \\
& \quad + \sup_{\|u\|_{L^2} \leq M} \|\widehat{\mathcal{N}}_2 \circ \mathcal{N}_1(u) - \mathcal{N}_2 \circ \mathcal{N}_1(u)\|_{L^2} \\
& \leq \text{Lip}(\widehat{\mathcal{N}}_2) \sup_{\|u\|_{L^2} \leq M} \|\widehat{\mathcal{N}}_1(u) - \mathcal{N}_1(u)\|_{L^2} \\
& \quad + \sup_{\|v\| \leq 2M_1} \|\widehat{\mathcal{N}}_2(v) - \mathcal{N}_2(v)\|_{L^2} \\
& < \epsilon.
\end{aligned}$$

The estimate on the depth, width and lift is immediate (cp. Lemma D.4). This proves the base case $p = 2$ of the induction.

Let now $p > 2$, and assume that the claim holds for the composition of $p - 1$ operators. Given $\widehat{\mathcal{N}}_1, \dots, \widehat{\mathcal{N}}_p$ as in the statement of the lemma, we can apply the induction hypothesis for the composition of $p - 1$ operators, to see that the two continuous operators $\widetilde{\mathcal{N}}_1 := \widehat{\mathcal{N}}_1$ and $\widetilde{\mathcal{N}}_2 := (\widehat{\mathcal{N}}_p \circ \dots \circ \widehat{\mathcal{N}}_2)$ fulfill all assumptions of the lemma (with $p = 2$), and with depths, widths and lifts of the approximating FNOs given by $\widetilde{D}_1 = D_1$, $\widetilde{W}_1 = W_1$, $\widetilde{L}_1 = L_1$, and $\widetilde{D}_2 = \sum_{k=2}^p D_k$, $\widetilde{W}_2 = \max_{k=2, \dots, p} W_k$, $\widetilde{L}_2 = \max_{k=2, \dots, p} L_k$. The proof of the induction step now follows from the base case (with $p = 2$) already considered above. \square

APPENDIX E. TECHNICAL RESULTS AND PROOFS FOR SECTION 3.1

In this section, we collect some technical results and proofs for the material in section 3.1 of the main text. We start with the following Lemma on the contraction property of the map (3.9).

Lemma E.1. Let $k \in \mathbb{N}$, $k > d/2 + 1$, and let $s > d/2 + k$ be given. Assume that the coefficient $a \in H^s(\mathbb{T}^d)$ satisfies the λ -coercivity condition (3.6) for some $\lambda \in (0, 1)$, and that the right-hand side of the stationary Darcy equation (3.5) belongs to the Sobolev space $f \in H^{k-1}$. Then there exists $N_0 = N_0(s, d, \|a\|_{H^s}, \lambda) \in \mathbb{N}$, such that for any $N \geq N_0$, we have $a_N \geq \lambda/2$ and the mapping F_N defined by (3.9) is a contraction, with

$$\text{Lip}(F_N : \dot{H}^1 \rightarrow \dot{H}^1) \leq 1 - \frac{\lambda}{2}.$$

Proof. As $u_N \mapsto F_N(u_N)$ is an affine mapping, we can express the Lipschitz constant in terms of the following supremum

$$\text{Lip}(F_N : \dot{H}^1 \rightarrow \dot{H}^1) = \sup_{u_N \in \dot{H}^1 \setminus \{0\}} \frac{\left\| \dot{P}_N(-\Delta)^{-1} \nabla \cdot (\widetilde{a}_N \nabla u_N) \right\|_{\dot{H}^1}}{\|u_N\|_{\dot{H}^1}}.$$

We now observe that

$$\begin{aligned}
\left\| \dot{P}_N(-\Delta)^{-1} \nabla \cdot (\widetilde{a}_N \nabla u_N) \right\|_{\dot{H}^1} & \leq \|(-\Delta)^{-1} \nabla \cdot (\widetilde{a}_N \nabla u_N)\|_{\dot{H}^1} \\
& = \|\widetilde{a}_N \nabla u_N\|_{L^2} \\
& \leq \|\widetilde{a}_N\|_{L^\infty} \|u_N\|_{\dot{H}^1},
\end{aligned}$$

for all $u_N \in \dot{H}_N^1(\mathbb{T}^d)$. To finish the proof, we note that

$$\|\widetilde{a}_N\|_{L^\infty} \leq \|\widetilde{a}\|_{L^\infty} + \|\widetilde{a} - \widetilde{a}_N\|_{L^\infty},$$

where the first term is bounded by $1 - \lambda$ by the coercivity assumption. By our definition of \tilde{a}_N (cp. equation (3.8)), we have

$$\|\tilde{a} - \tilde{a}_N\|_{L^\infty} \leq \left\| \tilde{a} - \dot{P}_N \tilde{a} \right\|_{L^\infty} + \left\| \dot{P}_N \tilde{a} - \dot{P}_N \mathcal{I}_{2N} \tilde{a} \right\|_{L^\infty}.$$

But by the spectral/pseudo-spectral approximation estimates, we now have

$$\begin{aligned} \left\| \tilde{a} - \dot{P}_N \tilde{a} \right\|_{L^\infty} &\lesssim_{d,\delta} \left\| \tilde{a} - \dot{P}_N \tilde{a} \right\|_{H^{d/2+\delta}} \\ &\lesssim_{s,d,\delta} N^{-(s-d/2-\delta)} \|a\|_{H^s}, \end{aligned}$$

and

$$\begin{aligned} \left\| \dot{P}_N \tilde{a} - \dot{P}_N \mathcal{I}_{2N} \tilde{a} \right\|_{L^\infty} &\lesssim_{d,\delta} \left\| \dot{P}_N \tilde{a} - \dot{P}_N \mathcal{I}_{2N} \tilde{a} \right\|_{H^{d/2+\delta}} \\ &\leq \left\| \tilde{a} - \mathcal{I}_{2N} \tilde{a} \right\|_{H^{d/2+\delta}} \\ &\lesssim_{s,d,\delta} N^{-(s-d/2-\delta)} \|a\|_{H^s}. \end{aligned}$$

To be definite, let us choose $\delta := (s - d/2)/2$. In particular, it follows that there exists a constant $C = C(s, d)$, and $N_0 = N_0(s, d, \|a\|_{H^s}, \lambda) \in \mathbb{N}$, such that

$$\|\tilde{a} - \tilde{a}_N\|_{L^\infty} \leq \frac{\lambda}{2},$$

for $N \geq N_0$, and hence $a_N = 1 + \tilde{a}_N \geq 1 + \tilde{a} - \lambda/2 \geq \lambda/2$, and

$$\text{Lip}(F_N) \leq \|\tilde{a}_N\|_{L^\infty} \leq \|\tilde{a}\|_{L^\infty} + \|\tilde{a} - \tilde{a}_N\|_{L^\infty} \leq 1 - \frac{\lambda}{2},$$

for all $N \geq N_0$. The claim follows. \square

Next, we provide the detailed proof of Theorem 3.3 which guarantees convergence of the algorithm 3.2 to the solutions of the Darcy equation (3.5).

E.1. Proof of Theorem 3.3.

Proof. Choose $N_0 \in \mathbb{N}$ as in Lemma E.1.

Step 1: Well-posedness and error estimate of the FG approximation (3.7).

We consider the bilinear form $B : \dot{H}_N^1 \times \dot{H}_N^1 \rightarrow \mathbb{R}$, given by

$$B(u_N, w_N) := \int_{\mathbb{T}^d} a_N(x) \nabla u_N(x) \cdot \nabla w_N(x) dx.$$

We note that u_N solves (3.7), if and only if,

$$B(u_N, w_N) = \langle f_N, w_N \rangle_{L^2}, \quad (\text{E.1})$$

for all $w_N \in \dot{H}_N^1$. By Lemma E.1, we have $a_N \geq \lambda/2$ for all $N \geq N_0$, and hence

$$B(w_N, w_N) \geq \lambda/2 \|w_N\|_{\dot{H}^1}^2.$$

Thus, B is a coercive bilinear form. The existence and uniqueness of a solution $u_N^* \in \dot{H}_N^1$ of (E.1) follows from this for any right-hand side f_N . Furthermore, we have the estimate

$$\|u_N^*\|_{\dot{H}^1}^2 \leq 2 \|u_N^*\|_{\dot{H}^1}^2 \leq \frac{4}{\lambda} B(u_N^*, u_N^*) = \frac{4}{\lambda} \langle u_N^*, f_N \rangle_{L^2} \leq \frac{4}{\lambda} \|u_N^*\|_{H^1} \|f_N\|_{H^{-1}},$$

and thus, $\|u_N^*\|_{H^1} \leq 4\lambda^{-1} \|f_N\|_{H^{-1}}$.

It is also straightforward to show that by inductive differentiation of (3.5), one can also obtain higher-order estimates for $k \in \mathbb{N}$, $k \geq 1$ for the (non-discretized) elliptic equation (3.5), of the form:

$$\|u\|_{H^{k+1}} \leq C \|f\|_{H^{k-1}}, \quad (\text{E.2})$$

where $C = C(k, d, \|a\|_{C^k}) > 0$.

To prove the claimed error estimate, we note that if u solves (3.5), then $P_N u \in \dot{H}_N^1$ solves

$$\begin{aligned} -P_N \nabla \cdot (a_N \nabla P_N u) &= P_N f + P_N \nabla \cdot ((a - a_N) \nabla P_N u) \\ &\quad + P_N \nabla \cdot (a \nabla (1 - P_N) u). \end{aligned}$$

It thus follows that $w_N = P_N u - u_N^*$ solves

$$\begin{aligned} -P_N \nabla \cdot (a_N \nabla w_N) &= (P_N f - f_N) + P_N \nabla \cdot ((a - a_N) \nabla P_N u) \\ &\quad + P_N \nabla \cdot (a \nabla (1 - P_N) u) \\ &=: (I) + (II) + (III). \end{aligned}$$

The stability estimate then implies that $\|w_N\|_{H^1} \leq 4\lambda^{-1} \|(I) + (II) + (III)\|_{H^{-1}}$, can be bounded in terms of the H^{-1} -norm of the right-hand side. We can now estimate

$$\begin{aligned} \|(I)\|_{H^{-1}} &= \|(P_N f - f_N)\|_{H^{-1}} \\ &\leq \|(1 - P_N) f\|_{H^{-1}} + \|f - f_N\|_{H^{-1}} \lesssim_{d,s} N^{-k} \|f\|_{H^{k-1}}. \end{aligned}$$

The last inequality follows from the fact that $f \in H^{k-1}(\mathbb{T}^d)$, for $k-1 > d/2$, and the (pseudo-)spectral approximation estimate. For the second term, we fix a small $\delta > 0$, and obtain

$$\begin{aligned} \|(II)\|_{H^{-1}} &\leq \|(a - a_N) \nabla P_N u\|_{L^2} \\ &\lesssim_d \|a - a_N\|_{L^\infty} \|u\|_{H^1} \\ &\lesssim_{d,\delta,\lambda} \|a - a_N\|_{H^{d/2+\delta}} \|f\|_{H^{-1}}, \end{aligned}$$

where we have used the embedding $H^{d/2+\delta} \hookrightarrow L^\infty$ in the last step. Assuming that $\delta > 0$ is chosen sufficiently small, so that $s \geq d/2 + \delta$, we can then estimate

$$\|(II)\|_{H^{-1}} \lesssim_{s,d} \lambda^{-1} N^{-(s-d/2-\delta)} \|a\|_{H^s} \|f\|_{H^{-1}}.$$

Finally, the third term can be bounded as follows:

$$\begin{aligned} \|(III)\|_{H^{-1}} &\lesssim_d \|a \nabla (1 - P_N) u\|_{L^2} \\ &\lesssim_d \|a\|_{L^\infty} \|(1 - P_N) u\|_{H^1} \\ &\lesssim_{s,d} N^{-k} \|a\|_{H^s} \|u\|_{H^{k+1}} \\ &\lesssim_{s,d,\lambda,k,\|a\|_{C^k}} N^{-k} \|a\|_{H^s} \|f\|_{H^{k-1}}. \end{aligned}$$

In the last step, we used the higher regularity estimate (E.2). We note that for this step we assume that $a \in H^s$, $s > d/2 + k$, so that we have an embedding $H^s \hookrightarrow C^k$. Under this condition, we can further estimate (II) by

$$\begin{aligned} \|(II)\|_{H^{-1}} &\lesssim_{s,d} N^{-(s-d/2-\delta)} \lambda^{-1} \|a\|_{H^s} \|f\|_{H^{-1}} \\ &\lesssim_{s,d,\lambda,\|a\|_{H^s},\|f\|_{H^{-1}}} N^{-k}. \end{aligned}$$

Combining the above estimates for (I), (II) and (III), we conclude that for any $s \in \mathbb{R}$ and $k \in \mathbb{N}$ satisfying $s > d/2 + k$, we have

$$\|w_N\|_{H^1} = \|u_N^* - P_N u\|_{H^1} \leq C N^{-k},$$

provided that $a \in H^s$, $f \in H^{k-1}$, where $C = C(s, d, \lambda, k, \|a\|_{H^s}, \|f\|_{H^{k-1}}) > 0$ is independent of N .

The proof is now finished by observing that

$$\|u - u_N^*\|_{H^1}^2 = \|(1 - P_N) u\|_{H^1}^2 + \|P_N u - u_N^*\|_{H^1}^2,$$

and that, by the higher-regularity estimate (E.2), we have

$$\|(1 - P_N) u\|_{H^1} \lesssim_{k,d} N^{-k} \|u\|_{H^{k+1}} \lesssim_{s,k,d,\lambda,\|a\|_{H^s},\|f\|_{H^{k-1}}} N^{-k}.$$

Thus, it follows that there exists a constant $C = C(s, k, d, \lambda, \|a\|_{H^s}, \|f\|_{H^{k-1}}) > 0$, such that

$$\|u - u_N^*\|_{H^1} \leq CN^{-k}, \quad (\text{E.3})$$

where u_N^* is the solution of the Fourier-Galerkin discretization (3.7).

Step 2: Picard iteration estimate.

The output of Algorithm 3.2 is obtained by Picard iteration, i.e. by K -fold application of the mapping $F_N : \dot{H}_N^1 \rightarrow \dot{H}_N^1$, yielding a recursively defined sequence $u_N^0 := 0$, and $u_N^k = F_N(u_N^{k-1})$, for $k = 1, \dots, K$. Since F_N is a contraction with $\text{Lip}(F_N) \leq 1 - \lambda/2$ for $N \geq N_0$, by Lemma E.1, and since u_N^* is the unique fixed point of F_N , this implies that

$$\|u_N^K - u_N^*\|_{\dot{H}^1} \leq (1 - \lambda/2)^K \|u_N^*\|_{\dot{H}^1}.$$

Using the definition of H^1 , \dot{H}^1 , and the regularity estimate (E.2), we can further estimate

$$\|u_N^K - u_N^*\|_{H^1} \leq 2\|u_N^K - u_N^*\|_{\dot{H}^1} \leq 8\lambda^{-1}(1 - \lambda/2)^K \|f_N\|_{H^{-1}}.$$

By definition of K , we have

$$K \geq \frac{\log(\lambda N^{-k})}{\log(1 - \lambda/2)},$$

and hence $(1 - \lambda/2)^K \leq \lambda N^{-k}$. This yields

$$\|u_N^K - u_N^*\|_{H^1} \leq 8N^{-k} \|f_N\|_{H^{-1}}.$$

Combining Steps 1 and 2, implies that there exists a constant $C > 0$, depending on $s, k, d, \lambda, \|a\|_{H^s}$ and $\|f\|_{H^{k-1}} > 0$, such that

$$\|u - u_N\|_{H^1} \leq CN^{-k},$$

for all $N \geq N_0$, where $u_N := u_N^K$ is the output of Algorithm 3.2. This is the claimed estimate. \square

E.2. Neural network approximation of quadratic non-linearities. Our next aim is to prove Theorem 3.5. As stated in the main text, the proof relies crucially on the following Lemmas, which show that neural networks can efficiently approximate certain quadratic non-linearities. We start with the following result,

Lemma E.2. Let $\sigma \in C^3$ be a activation function. Let $d \in \mathbb{N}$. There exists a constant $C = C(d) > 0$, such that for any $\epsilon > 0$ and $B > 0$, there exists

- (1) a neural network $\widehat{\mathcal{N}}_1 : \mathbb{R}^2 \rightarrow \mathbb{R}$, such that

$$\sup_{|a|, |b| \leq B} |\widehat{\mathcal{N}}_1(a, b) - ab| \leq \epsilon,$$

$$\text{and width}(\widehat{\mathcal{N}}_1) \leq C, \text{ depth}(\widehat{\mathcal{N}}_1) \leq C,$$

- (2) a neural network $\widehat{\mathcal{N}}_2 : \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, such that

$$\sup_{|a|, \|v\|_{\ell^2} \leq B} |\widehat{\mathcal{N}}_2(a, v) - av| \leq \epsilon,$$

$$\text{and width}(\widehat{\mathcal{N}}_2) \leq C, \text{ depth}(\widehat{\mathcal{N}}_2) \leq C,$$

- (3) a neural network $\widehat{\mathcal{N}}_3 : \mathbb{R}^d \times \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^d$, such that

$$\sup_{\|v\|_{\ell^2}, \|U\|_{\ell^2 \rightarrow \ell^2} \leq B} |\widehat{\mathcal{N}}_3(v, U) - v \cdot U| \leq \epsilon,$$

$$\text{and width}(\widehat{\mathcal{N}}_3) \leq C, \text{ depth}(\widehat{\mathcal{N}}_3) \leq C.$$

Proof. Points (2) and (3) easily follow from (1), by parallelizing multiple networks $\widehat{\mathcal{N}}_1$. To see the first claim (1), we note that the quadratic function $y \mapsto y^2$ can be approximated for $y \in [-B, B]$, $B > 0$, to arbitrary precision by finite differences

$$y^2 = \underbrace{\frac{\sigma(x+hy) - 2\sigma(x) + \sigma(x-hy)}{h^2\sigma^{(2)}(x)}}_{\text{sq}_h(y)} + R(h; y),$$

where we assume that x is chosen so that the second derivative $\sigma^{(2)}(x) \neq 0$, and $|R(h, y)| \leq Ch$ for all $y \in [-B, B]$, and $C = C(B)$. Finally, following [34], we observe that the product ab of two numbers a, b can be expressed in the form

$$ab = \frac{1}{2} ((a+b)^2 - a^2 - b^2) = \frac{1}{2} (\text{sq}_h(a+b) - \text{sq}_h(a) - \text{sq}_h(b)) + \widetilde{R}(h; a, b),$$

where \widetilde{R} is related to R , and there exists a constant $C = C(B)$, such that $|\widetilde{R}(h; a, b)| \leq Ch$, for all $a, b \in [-B, B]$. Since sq_h is a neural network of finite width and depth (independent of h), and since the last expression

$$\widehat{\mathcal{N}}_h(a, b) := \frac{1}{2} (\text{sq}_h(a+b) - \text{sq}_h(a) - \text{sq}_h(b)),$$

is simply a linear combination of sq_h , we conclude that $\widehat{\mathcal{N}}_h$ is a neural network of uniformly bounded width and depth (uniform in h), and for sufficiently small $h > 0$, we have

$$\sup_{a, b \in [-B, B]} |\widehat{\mathcal{N}}_h(a, b) - ab| \leq \epsilon.$$

This concludes the proof. \square

Using the above lemma, one can prove the following result,

Lemma E.3. Assume that the activation function $\sigma \in C^3$ is three times continuously differentiable and non-linear. There exists a constant $C > 0$, such that for any $N \in \mathbb{N}$, and for any $\epsilon, B > 0$, there exists a Ψ -FNO $\mathcal{N} : L_{2N}^2(\mathbb{T}^d; \mathbb{R}) \times L_{2N}^2(\mathbb{T}^d; \mathbb{R}) \rightarrow L_{2N}^2(\mathbb{T}^d; \mathbb{R})$, with

$$\text{depth}(\mathcal{N}), \text{lift}(\mathcal{N}) \leq C, \quad \text{width}(\mathcal{N}) \leq CN^d,$$

such that we have

$$\|P_N(a_N \nabla u_N) - \mathcal{N}(a_N, u_N)\|_{L_{2N}^2} \leq \epsilon,$$

for all trigonometric polynomials $a_N, u_N \in L_{2N}^2(\mathbb{T}^d; \mathbb{R}) \subset L_{2N}^2(\mathbb{T}^d; \mathbb{R})$ of degree $|k|_\infty \leq N$, satisfying the bound $\|a_N\|_{L^2}, \|u_N\|_{L^2} \leq B$.

Proof. First, we observe that there exists a linear FNO layer $\mathcal{L} : L_{2N}^2 \rightarrow L_{2N}^2$, with a suitable choice of the Fourier multiplier matrix P , such that

$$\mathcal{L}(u_N) = \mathcal{F}_{2N}^{-1}(P\mathcal{F}(u_N)) = \mathcal{F}_{2N}^{-1} \left(\sum_{|k|_\infty \leq 2N} ik \widehat{u}_k e^{i\langle k, x \rangle} \right) = \nabla u_N,$$

is satisfied *exactly* for all $u_N \in L_{2N}^2(\mathbb{T}^d; \mathbb{R})$. We also note that if $\|u_N\|_{L^2} \leq B$, then $\|\nabla u_N\|_{L^2} \leq NB$. By the fact that all norms are equivalent on the finite-dimensional space $L_{2N}^2(\mathbb{T}^d; \mathbb{R}^{d \times d})$, there exists a constant $C' > 0$ (depending on N), such that

$$\sup_{x \in \mathbb{T}^d} \|\nabla u_N(x)\|_{\ell^2} \leq C' \|\nabla u_N\|_{L^2} \leq C' NB,$$

whenever $\|u_N\|_{L^2} \leq B$. We similarly see that by norm equivalence on $L_{2N}^2(\mathbb{T}^d; \mathbb{R})$, there exists a constant $C'' > 0$ (depending on N), such that we also have

$$\sup_{x \in \mathbb{T}^d} |a_N(x)| \leq C'' \|a_N\|_{L^2} \leq C'' B,$$

for any $\|a_N\|_{L^2} \leq B$. Let $\bar{B} := \max(C'NB, C''B)$. By Lemma E.2, there exists an ordinary neural network $\widehat{\mathcal{N}} : \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ with $\text{width}(\widehat{\mathcal{N}}), \text{depth}(\widehat{\mathcal{N}}) \leq C(d)$ (with $C = C(d)$ independent of N), such that

$$\sup_{|a|, \|v\|_{\ell^2} \leq \bar{B}} \|\widehat{\mathcal{N}}(a, v) - av\|_{\ell^2} \leq \epsilon.$$

By Lemmas D.2, D.3 and the composition lemma D.4, the composition

$$(a_N, u_N) \mapsto (a_N, \mathcal{L}(u_N)) = (a_N, \nabla u_N) \mapsto \widehat{\mathcal{N}}(a_N, \nabla u_N),$$

can be represented by a Ψ -FNO $\widetilde{\mathcal{N}} : L_{2N}^2(\mathbb{T}^d; \mathbb{R}^2) \rightarrow L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d)$, and by construction, we have

$$\sup_{x \in \mathbb{T}^d} \|a_N(x) \nabla u_N(x) - \widetilde{\mathcal{N}}(a_N, u_N)(x)\|_{\ell^2} \leq \epsilon, \quad (\text{E.4})$$

for all $a_N, u_N \in L_{2N}^2$, with $\|a_N\|_{L^2}, \|u_N\|_{L^2} \leq B$. Furthermore, since \mathcal{L} is a linear layer, and $\widehat{\mathcal{N}}$ is an ordinary neural network with $\text{width}(\widehat{\mathcal{N}}), \text{depth}(\widehat{\mathcal{N}}) \leq C = C(d)$, we in fact conclude that for some new constant $C = C(d) > 0$, we have

$$\text{width}(\widetilde{\mathcal{N}}) \leq CN^d, \quad \text{depth}(\widetilde{\mathcal{N}}) \leq C, \quad \text{lift}(\widetilde{\mathcal{N}}) \leq C.$$

Finally, we note that the projection $P_N : L_{2N}^2 \rightarrow L_{2N}^2$ onto Fourier modes with wavenumbers $|k|_\infty \leq N$ can again be represented exactly by a linear Ψ -FNO layer $\widehat{\mathcal{L}}$, and by Lemmas D.3, D.4, there exists a Ψ -FNO $\mathcal{N} : L_{2N}^2(\mathbb{T}^d; \mathbb{R}^2) \rightarrow L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d)$, such that $\|\mathcal{N}(a_N, u_N) - \widehat{\mathcal{L}} \circ \widetilde{\mathcal{N}}(a_N, u_N)\|_{L^2} \leq \epsilon$, for all $\|a_N\|_{L^2}, \|u_N\|_{L^2} \leq B$, and such that

$$\text{width}(\mathcal{N}) \leq CN^d, \quad \text{depth}(\mathcal{N}) \leq C, \quad \text{lift}(\mathcal{N}) \leq C,$$

where $C = C(d) > 0$ depends only on d . Combining this with (E.4), we conclude that \mathcal{N} satisfies

$$\begin{aligned} \|P_N(a_N \nabla u_N) - \mathcal{N}(a_N, u_N)\|_{L^2} &\leq \|\widehat{\mathcal{L}} \circ \widetilde{\mathcal{N}}(a_N, u_N) - \mathcal{N}(a_N, u_N)\|_{L^2} \\ &\quad + \|P_N(a_N \nabla u_N) - \widehat{\mathcal{L}} \circ \widetilde{\mathcal{N}}(a_N, u_N)\|_{L^2} \\ &= \|\widehat{\mathcal{L}} \circ \widetilde{\mathcal{N}}(a_N, u_N) - \mathcal{N}(a_N, u_N)\|_{L^2} \\ &\quad + \|P_N(a_N \nabla u_N) - P_N \widetilde{\mathcal{N}}(a_N, u_N)\|_{L^2} \\ &\leq \epsilon + \|a_N \nabla u_N - \widetilde{\mathcal{N}}(a_N, u_N)\|_{L^2} \\ &\leq \epsilon + |\mathbb{T}^d|^{1/2} \|a_N \nabla u_N - \widetilde{\mathcal{N}}(a_N, u_N)\|_{L^\infty} \\ &\leq \left(1 + (2\pi)^{d/2}\right) \epsilon, \end{aligned}$$

for all $\|a_N\|_{L^2}, \|u_N\|_{L^2} \leq B$. Since $\epsilon > 0$ was arbitrary, the claim follows. \square

E.3. Proof of Theorem 3.5. The stage is now set for the proof of Theorem 3.5 in the following,

Proof of Theorem 3.5. Since the claim is an asymptotic statement, it suffices to consider $N \geq N_0$, where N_0 is the constant of Lemma E.1. Indeed, the exceptional cases $N < N_0$ can be handled by suitably enlarging the constant C . We will thus assume that $N \geq N_0$, and $f \in \dot{H}^{k-1}$ are given. We fix $F_N := (-\Delta)f_N \in L_N^2(\mathbb{T}^d)$ for the rest of this proof, where $f_N := \dot{P}_N \mathcal{I}_{2N} f$ is defined as in Algorithm 3.2.

We define two operators $\widehat{\mathcal{N}}_1, \widehat{\mathcal{N}}_2$ as follows: We let

$$\widehat{\mathcal{N}}_1 : L_{2N}^2(\mathbb{T}^d; \mathbb{R}^2) \rightarrow L_{2N}^2(\mathbb{T}^d; \mathbb{R} \times \mathbb{R}^d), \quad \widehat{\mathcal{N}}_1(a, u) := (a, P_N(a \nabla u)),$$

and define $\widehat{\mathcal{N}}_2 : L_{2N}^2(\mathbb{T}^d; \mathbb{R} \times \mathbb{R}^d) \rightarrow L_{2N}^2(\mathbb{T}^d; \mathbb{R}^2)$ by

$$\widehat{\mathcal{N}}_2(a, U) := \left(a, \dot{P}_N(-\Delta)^{-1} \nabla \cdot U + F_N\right).$$

In terms of $\widehat{\mathcal{N}}_1, \widehat{\mathcal{N}}_2$, Algorithm 3.2, which defines a mapping $a \mapsto u_N = \widehat{\mathcal{N}}(a)$, can be written in the form

$$\widehat{\mathcal{N}}(a) = \widehat{\mathcal{Q}} \circ \underbrace{\widehat{\mathcal{N}}_2 \circ \widehat{\mathcal{N}}_1 \circ \cdots \circ \widehat{\mathcal{N}}_2 \circ \widehat{\mathcal{N}}_1}_{K\text{-fold composition}} \circ \widehat{\mathcal{R}}(a),$$

where $\widehat{\mathcal{R}}(a) := (a, 0)$, $\widehat{\mathcal{Q}}(a, u) := u$ and where $K \lesssim \log(N)$. By the composition lemma D.4, to prove the claim of this theorem, it therefore suffices to show the following

Claim: For any $B > 0$, there exists $C > 0$, such that for any $\epsilon > 0$, there exist Ψ -FNOs \mathcal{N}_1 and \mathcal{N}_2 , with $\text{width}(\mathcal{N}_1), \text{width}(\mathcal{N}_2) \leq CN^d$, $\text{depth}(\mathcal{N}_1), \text{depth}(\mathcal{N}_2)$, $\text{lift}(\mathcal{N}_1), \text{lift}(\mathcal{N}_2) \leq C$, and such that

$$\|\widehat{\mathcal{N}}_1(u) - \mathcal{N}_1(u)\|_{L^2}, \|\widehat{\mathcal{N}}_2(u) - \mathcal{N}_2(u)\|_{L^2} \leq \epsilon,$$

for all $u \in L^2_{2N}(\mathbb{T}^d)$ with bounded norm $\|u\|_{L^2} \leq B$.

For $\widehat{\mathcal{N}}_1$, the claim follows from Lemma E.3. For $\widehat{\mathcal{N}}_2$, we note that $\widehat{\mathcal{N}}_2$ can be represented exactly by a linear FNO layer with $W = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$, bias $b_j = (0, F_N(x_j))$, and with Fourier multiplier matrix $P(k) = \begin{pmatrix} 0 & 0 \\ 0 & \tilde{P}(k) \end{pmatrix}$, where $\tilde{P}(k) := 1_{\{|k|_\infty \leq N\}} \frac{ik^T}{|k|^2}$, so that

$$\widehat{\mathcal{N}}_2(v)_j = \widehat{\mathcal{N}}_2(v)(x_j) = Wv_j + b_j + \mathcal{F}_N^{-1}(P\mathcal{F}_N v)_j,$$

for $v_j = (a(x_j), U(x_j))$. The claim for $\widehat{\mathcal{N}}_2$ thus follows from the linear approximation lemma D.3. \square

APPENDIX F. TECHNICAL RESULTS AND PROOFS FROM SECTION 3.2

F.1. Properties of the Pseudo-spectral scheme (3.15). Our first aim is to show that the implicit operator equation that defines the scheme (3.15). To this end, we have following lemmas,

Lemma F.1. If $v \in L^2_N(\mathbb{T}^d; \mathbb{R}^d)$ and $\tau\|v\|_{L^\infty} N \leq \frac{1}{2}$, then we have

$$\|\tau\mathbb{P}_N(v \cdot \nabla w)\|_{L^2} \leq \frac{1}{2}\|w\|_{L^2},$$

for all $w \in L^2_N(\mathbb{T}^d; \mathbb{R}^d)$. In particular, this estimate holds provided that

$$\tau\|v\|_{L^2} N^{d/2+1} \leq \frac{1}{2}.$$

Proof. We first note that for any $v \in L^2_N$, we have

$$\tau\|v\|_{L^\infty} N \leq \tau\|v\|_{L^2} N^{d/2+1} \leq \frac{1}{2}.$$

The claim is now an immediate consequence of the fact that

$$\begin{aligned} \|\tau\mathbb{P}_N(v \cdot \nabla w)\|_{L^2} &\leq \|\tau v \cdot \nabla w\|_{L^2} \\ &\leq \tau\|v\|_{L^\infty} \|\nabla w\|_{L^2} \\ &\leq \tau\|v\|_{L^\infty} N \|w\|_{L^2}, \end{aligned}$$

for any $w \in L^2_N(\mathbb{T}^d; \mathbb{R}^d)$. \square

We can now state the following lemma on the well-posedness:

Lemma F.2. Let $U > 0$. If $\|u_N^n\|_{L^2} \leq U$ for $n = 0, \dots, n_T$, and if the CFL condition

$$\tau U N^{d/2+1} \leq \frac{1}{2}, \quad (\text{F.1})$$

is satisfied, then the recursion (3.15) is well-defined.

Proof. The recursion (3.15) can be written in the form

$$\mathcal{T}_n u_N^{n+1} = u^n, \quad (\text{F.2})$$

where the operator $\mathcal{T}_n : \dot{L}_N^2(\mathbb{T}^d; \text{div}) \rightarrow \dot{L}_N^2(\mathbb{T}^d; \text{div})$ is given by

$$\mathcal{T}_n w_N := w_N + \tau \mathbb{P}_N (u_N^n \cdot \nabla w_N) - \nu \tau \Delta w_N.$$

The claimed well-posedness of the recursion follows from the fact that, under the CFL assumption (F.1), the operator \mathcal{T}_n is invertible: Indeed, by Lemma F.1, this implies that for any $w_N \in L_N^2(\mathbb{T}^d; \text{div})$, we have

$$\|\tau \mathbb{P}_N (u_N^n \cdot \nabla w_N)\|_{L^2} \leq \frac{1}{2} \|w_N\|_{L^2}.$$

But then, we have for any $w_N \in L_N^2(\mathbb{T}^d; \text{div})$, that

$$\begin{aligned} \|\mathcal{T}_n w_N\|_{L^2} &\geq \|(1 - \tau \nu \Delta) w_N\|_{L^2} - \|\tau \mathbb{P}_N (u_N^n \cdot \nabla w_N)\|_{L_N^2} \\ &\geq \|w_N\|_{L^2} - \frac{1}{2} \|w_N\|_{L^2} \\ &= \frac{1}{2} \|w_N\|_{L^2}, \end{aligned}$$

where, in the first step, we have used the fact that $-\tau \nu \Delta$ is a non-negative operator. In particular, the estimate $\|\mathcal{T}_n w_N\|_{L^2} \geq \frac{1}{2} \|w_N\|_{L^2}$ implies that $\mathcal{T}_n : L_N^2(\mathbb{T}^d; \text{div}) \rightarrow L_N^2(\mathbb{T}^d; \text{div})$ is injective. Since $L_N^2(\mathbb{T}^d; \text{div})$ is finite-dimensional, we conclude that \mathcal{T}_n is actually invertible, and hence the recursion (3.15) is well-defined, i.e. it can be solved for u_N^{n+1} , given u_N^n . \square

By Lemma F.2, to prove the well-posedness of the scheme (3.15), it remains to be shown that with a suitable choice of the time-step τ , we have a uniform L^2 -energy bound of the form $\|u_N^n\|_{L^2} \leq U$, for some $U > 0$. This is the content of the following lemma:

Lemma F.3. Assume that $\|u_N^0\|_{L^2} \leq U$, then the first-order scheme (3.15) is well-defined for any time-step satisfying the CFL condition $\tau U N^{d/2+1} \leq \frac{1}{2}$, and we have $\|u_N^n\|_{L^2} \leq \|u_N^0\|_{L^2} \leq U$ for all $n = 0, \dots, n_T$.

Proof. We show inductively that if $\|u_N^n\|_{L^2} \leq U$, then $\|u_N^{n+1}\|_{L^2} \leq \|u_N^n\|_{L^2} \leq U$. The well-posedness then follows from Lemma F.2. To see that $\|u_N^{n+1}\|_{L^2} \leq \|u_N^n\|_{L^2}$, we integrate (3.15) against u_N^{n+1} , to find

$$\begin{aligned} \|u_N^{n+1}\|_{L^2}^2 &= \langle u_N^{n+1}, u_N^n \rangle - \tau \langle u_N^{n+1}, \mathbb{P}_N (u_N^n \cdot \nabla u_N^{n+1}) \rangle - \nu \tau \|\nabla u_N^{n+1}\|_{L^2}^2 \\ &= \langle u_N^{n+1}, u_N^n \rangle - \tau \underbrace{\langle u_N^{n+1}, u_N^n \cdot \nabla u_N^{n+1} \rangle}_{=0} - \nu \tau \|\nabla u_N^{n+1}\|_{L^2}^2 \\ &\leq |\langle u_N^{n+1}, u_N^n \rangle| \leq \|u_N^{n+1}\|_{L^2} \|u_N^n\|_{L^2}. \end{aligned}$$

This proves the claim. \square

Next, we have the following Lemma on the convergence of the iterations in algorithm 3.9,

Lemma F.4. Given $\|u_N^n\|_{L^2} \leq U$, assume that the CFL condition $\tau UN^{d/2+1} \leq \frac{1}{2}$ is satisfied. Define a recursive sequence $w_N^{n,k} \in \dot{L}_N^2(\mathbb{T}^d; \text{div})$, $k \in \mathbb{N}$, by $w_N^{n,0} := 0$, and $w_N^{n,k+1} := F(w_N^{n,k})$, where F is defined by (3.18). Then, we have $\|u_N^{n+1} - w_N^{n,k}\|_{L^2} \leq 2^{-k} \|u_N^n\|_{L^2}$, for $k \in \mathbb{N}$, and $\|w_N^{n,k}\|_{L^2} \leq (1 + 2^{-k}) \|u_N^n\|_{L^2}$.

Proof. By Lemma F.1, and the fact that $(1 - \nu\tau\Delta)^{-1}$ is a contraction, it follows that $\text{Lip}(F) \leq \frac{1}{2}$. By Picard iteration, and recalling that u_N^{n+1} is the unique fixed point of the recursion $w_N^{n,k}$ ($k \in \mathbb{N}$), it immediately follows that

$$\|u_N^{n+1} - w_N^{n,k}\|_{L^2} \leq \frac{1}{2^k} \|u_N^{n+1} - w_N^{n,0}\|_{L^2} \leq \frac{\|u_N^n\|_{L^2}}{2^k}.$$

The last step is a consequence of the a priori L^2 -bound $\|u_N^{n+1}\|_{L^2} \leq \|u_N^n\|_{L^2}$ proven in Lemma F.3. In particular, this estimate implies that

$$\|w_N^{n,k}\|_{L^2} \leq \|u_N^{n+1}\|_{L^2} + \|w_N^{n,k} - u_N^{n+1}\|_{L^2} \leq (1 + 2^{-k}) \|u_N^n\|_{L^2}.$$

□

Remark F.5. Note that as a consequence of Lemma F.4, we recursively find that

$$\begin{aligned} \|u_N^{n+1}\|_{L^2} &= \|w_N^{n,\kappa_0}\|_{L^2} \leq (1 + 2^{-\kappa_0}) \|u_N^n\|_{L^2} \leq \left(1 + \frac{\tau^2}{T^2}\right) \|u_N^n\|_{L^2} \\ &\leq \dots \leq \left(1 + \frac{\tau^2}{T^2}\right)^n \|u_N^0\|_{L^2} \leq \exp\left(\frac{\tau}{T}\right) U \leq \epsilon U, \end{aligned}$$

for $n = 0, \dots, n_T$. In particular, this ensures that the CFL condition (F.1) is satisfied for all u_N^n , generated by Algorithm 3.9.

Finally, we provide the proof of the convergence Theorem 3.10.

F.2. Proof of Theorem 3.10. In this appendix, we provide a detailed proof of the convergence estimate of Theorem 3.10, for the first-order scheme defined by Algorithm 3.9.

To this end, let $u(t)$ be an exact solution of (3.12), satisfying the assumptions of Theorem 3.10, and let $u_N^0, \dots, u_N^{n_T}$ denote the sequence generated by Algorithm 3.9. For any $n = 0, \dots, n_T - 1$, we denote by $u_N^{n+1,*}$ the solution of a single time-step with the semi-implicit scheme (3.15), starting from u_N^n , i.e. satisfying

$$\frac{u_N^{n+1,*} - u_N^n}{\tau} + \mathbb{P}_N \left(u_N^n \cdot \nabla u_N^{n+1,*} \right) = \nu \Delta u_N^{n+1,*}. \quad (\text{F.3})$$

We recall that by Lemma F.4 and Remark F.5, we have a uniform bound $\|u_N^n\|_{L^2} \leq \epsilon U$. Since by assumption, the time-step $\tau \epsilon UN^{d/2+1} \leq \frac{1}{2}$ satisfies the relevant CFL condition, it follows that a unique solution $u_N^{n+1,*}$ exists for all n . We also recall that by Lemma F.4, and by our choice of the number of iteration steps κ in Algorithm 3.9, we have

$$\|u_N^{n+1,*} - u_N^{n+1}\|_{L^2} \leq 2^{-\kappa} \epsilon U \leq C\tau^2,$$

where $C > 0$ depends only on U and the final time T .

Our first goal is to derive an estimate on the magnification of the approximation error due to a single timestep $u_N^n \mapsto u_N^{n+1}$. Let $u(t)$ be the exact solution of (3.12). Observing that

$$\begin{aligned} \|u_N^{n+1} - u(t^{n+1})\|_{L^2} &\leq \|u_N^{n+1} - u_N^{n+1,*}\|_{L^2} + \|u_N^{n+1,*} - u(t^{n+1})\|_{L^2} \\ &\leq C\tau^2 + \|u_N^{n+1,*} - u(t^{n+1})\|_{L^2}, \end{aligned} \quad (\text{F.4})$$

we only need to consider the error introduced by a single time-step $u_N^n \mapsto u_N^{n+1,*}$ of the semi-implicit scheme (3.15). To this end, we can write

$$\frac{u(t^{n+1}) - u(t^n)}{\tau} + \mathbb{P}_N (u(t^n) \cdot \nabla u(t^{n+1})) = \nu \Delta u(t^{n+1}) + \mathcal{E}^n,$$

where \mathcal{E}^n collects all error terms:

$$\mathcal{E}^n = \mathcal{E}_\tau^n + \mathcal{E}_{NL}^n + \mathcal{E}_P^n + \mathcal{E}_\nu^n,$$

where

$$\begin{aligned} \mathcal{E}_\tau^n &= \frac{u(t^{n+1}) - u(t^n)}{\tau} - \partial_t u(t^n), \\ \mathcal{E}_{NL}^n &= \mathbb{P} (u(t^n) \cdot \nabla (u(t^{n+1}) - u(t^n))), \\ \mathcal{E}_P^n &= (1 - \mathbb{P}_N) (u(t^n) \cdot \nabla u(t^n)), \\ \mathcal{E}_\nu^n &= -\nu \Delta (u(t^{n+1}) - u(t^n)). \end{aligned}$$

Subtracting (F.3), and introducing the short-hand notation $e^n := u(t^n) - u_N^n$, $e^{n+1} := u(t^{n+1}) - u_N^{n+1,*}$, we find

$$\frac{e^{n+1} - e^n}{\tau} = -\mathbb{P}_N (e^n \cdot \nabla u(t^{n+1})) - \mathbb{P}_N (u_N^n \cdot \nabla e^{n+1}) + \nu \Delta e^{n+1} + \mathcal{E}^n.$$

Next, integrate against $e_N^{n+1} := P_N e^{n+1}$ to find

$$\begin{aligned} &\frac{1}{2\tau} (\|e_N^{n+1}\|_{L^2}^2 + \|e_N^{n+1} - e_N^n\|_{L^2}^2 - \|e_N^n\|_{L^2}^2) \\ &\leq \|e^n\|_{L^2} \|\nabla u\|_{L_{t,x}^\infty} \|e_N^{n+1}\|_{L^2} - \langle \mathbb{P}_N (u_N^n \cdot \nabla e^{n+1}), e^{n+1} \rangle - \nu \|\nabla e^{n+1}\|_{L^2}^2 + \langle \mathcal{E}^n, e^{n+1} \rangle. \end{aligned}$$

We note that

$$\begin{aligned} \langle \mathbb{P}_N (u_N^n \cdot \nabla e^{n+1}), e^{n+1} \rangle &= \langle P_N (u_N^n \cdot \nabla e^{n+1}), e^{n+1} \rangle \\ &= \langle u_N^n \cdot \nabla e^{n+1}, e^{n+1} \rangle - \langle (1 - P_N) u_N^n \cdot \nabla e^{n+1}, e^{n+1} \rangle \\ &= 0 - \langle u_N^n \cdot \nabla e^{n+1}, (1 - P_N) e^{n+1} \rangle \\ &= -\langle u_N^n \cdot \nabla e^{n+1}, (1 - P_N) u(t^{n+1}) \rangle \\ &= \langle e^{n+1}, u_N^n \cdot \nabla (1 - P_N) u(t^{n+1}) \rangle \\ &= -\langle e^{n+1}, e^n \cdot \nabla (1 - P_N) u(t^{n+1}) \rangle \\ &\quad + \langle e^{n+1}, u(t^n) \cdot \nabla (1 - P_N) u(t^{n+1}) \rangle \end{aligned} \tag{F.5}$$

We now note that we can rewrite the last term as follows:

$$\langle e^{n+1}, u(t^n) \cdot \nabla (1 - P_N) u(t^{n+1}) \rangle = \langle (1 - P_N) (e^{n+1} \otimes u(t^n)), \nabla u(t^{n+1}) \rangle$$

Using the fact that

$$(1 - P_N) (e^{n+1} \otimes u(t^n)) = (1 - P_N) ((1 - P_{N/2}) e^{n+1} \otimes (1 - P_{N/2}) u(t^n)),$$

it then follows that we have

$$\begin{aligned} \langle P_N (u_N^n \cdot \nabla e^{n+1}), e^{n+1} \rangle &= -\langle e^{n+1}, e^n \cdot \nabla (1 - P_N) u(t^{n+1}) \rangle \\ &\quad - \langle (1 - P_N) (e^{n+1} \otimes u(t^n)), \nabla u(t^{n+1}) \rangle \\ &\leq \|e^{n+1}\|_{L^2} \|e^n\|_{L^2} \|(1 - P_N) \nabla u(t^{n+1})\|_{L^\infty} \\ &\quad + \|e^{n+1}\|_{L^2} \|(1 - P_{N/2}) u(t^n)\|_{L^2} \|(1 - P_N) \nabla u(t^{n+1})\|_{L^\infty}. \end{aligned}$$

For $r > d/2 + 1$, we have a continuous embedding $H^r \hookrightarrow W^{1,\infty}$, and an inequality of the form

$$\|(1 - P_{N/2}) u(t^n)\|_{L^2} \lesssim_{r,d} \frac{1}{N^r} \|u(t^n)\|_{H^r}.$$

Hence we can estimate

$$\begin{aligned} \langle P_N(u_N^n \cdot \nabla e^{n+1}), e^{n+1} \rangle &\lesssim_{r,d} \|e^{n+1}\|_{L^2} \|e^n\|_{L^2} \|u(t^n)\|_{H^r} \\ &\quad + \|e^{n+1}\|_{L^2} \frac{\|u(t^n)\|_{H^r}^2}{N^{2r}}. \end{aligned}$$

Estimating the products on the right-hand side using the inequality $ab \leq \epsilon a^2 + \frac{1}{4\epsilon} b^2$ with suitable $\epsilon > 0$, it follows that there exists a constant $C > 0$ (independent of $\nu > 0$, N and n), such that

$$\begin{aligned} \|e^{n+1}\|_{L^2}^2 &\leq \left(1 + C\tau \|u\|_{C_t(H_x^r)}^2\right) \|e^n\|_{L^2}^2 - \nu \|\nabla e^{n+1}\|_{L^2}^2 \\ &\quad + C\tau N^{-2r} \|u\|_{C_t(H_x^r)}^4 + C\tau |\langle \mathcal{E}^n, e^{n+1} \rangle|. \end{aligned} \quad (\text{F.6})$$

F.2.1. *Time-differencing error.* We note that

$$\begin{aligned} \langle \mathcal{E}_\tau^n, e^{n+1} \rangle &= \frac{1}{\tau} \int_{t^n}^{t^{n+1}} \int_{t^n}^t \langle \partial_t^2 u(s), e^{n+1} \rangle ds dt \\ &= \frac{1}{\tau} \int_{t^n}^{t^{n+1}} \int_{t^n}^t \langle -P \{ \partial_t u(s) \cdot \nabla u(s) + u(s) \cdot \nabla \partial_t u(s) \}, e^{n+1} \rangle ds dt \\ &\quad + \frac{1}{\tau} \int_{t^n}^{t^{n+1}} \int_{t^n}^t \langle \nu \Delta u(s), e^{n+1} \rangle ds dt \\ &=: (I)_\tau + (II)_\tau. \end{aligned}$$

The first term can be bounded from above by

$$\begin{aligned} (I)_\tau &\lesssim \tau \left(\|u\|_{C_t^1(L_x^2)} \|u\|_{C_t(W_x^{1,\infty})} + \|u\|_{C_t(L_x^\infty)} \|u\|_{C_t^1(H_x^1)} \right) \|e^{n+1}\|_{L^2} \\ &\lesssim_{r,d} \tau \|u\|_{C_t^1(H^{r-2})} \|u\|_{C_t(H^r)} \|e^{n+1}\|_{L^2}, \end{aligned}$$

provided that $r > d/2 + 2 \geq 3$ (the last bound is automatic for $d \geq 2$). For the second term, we derive the bound

$$(II)_\tau \lesssim \tau \nu \|u\|_{C_t(H_x^1)} \|\nabla e^{n+1}\|_{L^2}.$$

Thus, for any $\epsilon > 0$ (to be specified later), we have for some constant $C = C(r, d) > 0$:

$$\begin{aligned} |\langle \mathcal{E}_\tau^n, e^{n+1} \rangle| &\leq \epsilon (\|e^{n+1}\|_{L^2} + \nu \|\nabla e^{n+1}\|_{L^2}) \\ &\quad + \frac{C\tau^2}{\epsilon} \left(\|u\|_{C_t^1(H^{r-2})}^2 \|u\|_{C_t(H^r)}^2 + \nu \|u\|_{C_t(H_x^1)}^2 \right). \end{aligned} \quad (\text{F.7})$$

F.2.2. *Non-linear time-differencing error.* For the error associated with \mathcal{E}_{NL}^n , we simply estimate (for $\epsilon > 0$ to be determined later)

$$|\langle \mathcal{E}_{NL}^n, e^{n+1} \rangle| \leq \epsilon \|e^{n+1}\|_{L^2}^2 + \epsilon^{-1} \|\mathcal{E}_{NL}^n\|_{L^2}^2,$$

and we observe that

$$\|\mathcal{E}_{NL}^n\|_{L^2} \leq \tau \|u\|_{C_t(L_x^\infty)} \|u\|_{C_t^1(H_x^1)} \leq \tau \|u\|_{C_t(H_x^r)} \|u\|_{C_t^1(H_x^{r-2})},$$

assuming that $r > d/2 + 2 \geq 3$. This yields

$$|\langle \mathcal{E}_{NL}^n, e^{n+1} \rangle| \leq \epsilon \|e^{n+1}\|_{L^2}^2 + \frac{\tau^2}{\epsilon} \|u\|_{C_t(H_x^r)}^2 \|u\|_{C_t^1(H_x^{r-2})}^2 \quad (\text{F.8})$$

F.2.3. *Nonlinear projection error.* Again, we estimate the error \mathcal{E}_P^n using the simple estimate

$$|\langle \mathcal{E}_P^n, e^{n+1} \rangle| \leq \epsilon \|e^{n+1}\|_{L^2}^2 + \epsilon^{-1} \|\mathcal{E}_P^n\|_{L^2}^2,$$

with $\epsilon > 0$ to be specified later. We furthermore note that

$$(1 - P_N) \{u(t^n) \cdot \nabla u(t^{n+1})\} = (1 - P_N) \{(1 - P_{N/2})u(t^n) \cdot (1 - P_{N/2})\nabla u(t^{n+1})\},$$

which implies that

$$\|\mathcal{E}_P^n\|_{L^2} \leq \|(1 - P_{N/2})u(t^n)\|_{L^2} \|(1 - P_{N/2})\nabla u(t^{n+1})\|_{L^\infty} \lesssim_{r,d} N^{-r} \|u\|_{C_t(H_x^r)}^2.$$

In the last step, we have used the fact that $\|(1 - P_{N/2})v\|_{L^2} \lesssim_{r,d} N^{-r} \|v\|_{H^r}$ for $r \geq 0$, and that by Sobolev embedding

$$\|(1 - P_{N/2})\nabla u(t^{n+1})\|_{L^\infty} \lesssim_d \|(1 - P_{N/2})u(t^{n+1})\|_{H^r} \leq \|u\|_{C_t(H_x^r)},$$

for any $r > d/2 + 1$. Thus, there exists $C = C(d, r) > 0$, such that

$$|\langle \mathcal{E}_P^n, e^{n+1} \rangle| \leq \epsilon \|e^{n+1}\|_{L^2}^2 + \frac{C}{\epsilon N^{2r}} \|u\|_{C_t(H_x^r)}^2. \quad (\text{F.9})$$

F.2.4. *Viscosity error.* Finally, we note that

$$\begin{aligned} |\langle \mathcal{E}_\nu^n, e^{n+1} \rangle| &\leq \nu \|\nabla(u(t^{n+1}) - u(t^n))\|_{L^2} \|\nabla e^{n+1}\|_{L^2} \\ &\leq \nu \epsilon \|\nabla e^{n+1}\|_{L^2}^2 + \frac{\nu \tau^2}{\epsilon} \|u\|_{C_t^1(H_x^1)}^2 \\ &\leq \nu \epsilon \|\nabla e^{n+1}\|_{L^2}^2 + \frac{\nu \tau^2}{\epsilon} \|u\|_{C_t^1(H_x^{r-2})}^2, \end{aligned} \quad (\text{F.10})$$

for any $r > d/2 + 2 \geq 3$.

F.2.5. *The final stability estimate.* Choosing $\epsilon = 1/4$, it follows from (F.7), (F.8), (F.9) and (F.10), that the total error term $\langle \mathcal{E}^n, e^{n+1} \rangle$ can be estimated by

$$|\langle \mathcal{E}^n, e^{n+1} \rangle| \leq \|e^{n+1}\|_{L^2}^2 + \nu \|\nabla e^{n+1}\|_{L^2}^2 + C^* (\tau^2 + N^{-2r}). \quad (\text{F.11})$$

for some constant $C^* > 0$, depending only on $r > d/2 + 2$, the spatial dimension d and the norms $\|u\|_{C_t(H_x^r)}$, $\|u\|_{C_t^1(H_x^r)}$ of the exact solution u .

Substitution of the error estimate (F.11) in (F.6) finally yields

$$(1 - \tau) \|e^{n+1}\|_{L^2}^2 \leq (1 + C^* \tau) \|e^n\|_{L^2}^2 + \tau C^* (\tau^2 + N^{-2r}),$$

where the constant C^* depends only on $r \geq d/2 + 2$, the dimension d and the norms $\|u\|_{C_t(H_x^r)}$ and $\|u\|_{C_t^1(H_x^{r-2})}$. Assuming that $\tau \leq 1/2$, dividing by $(1 - \tau)$, and noting that

$$\frac{1 + C^* \tau}{1 - \tau} = 1 + \tau \left(C^* + \frac{(1 + C^* \tau)}{1 - \tau} \right) \leq 1 + \tau 2(C^* + 1),$$

and

$$\frac{C^*}{1 - \tau} \leq 2C^*,$$

we can clearly absorb the additional factor of $(1 - \tau)^{-1}$ by increasing the constant C^* , if necessary.

From this, we conclude that for a time-step $\tau \leq 1/2$ satisfying the CFL condition (F.1), there exists a constant $C^* = C^*(r, d, \|u\|_{C_t(H_x^r)}, \|u\|_{C_t^1(H_x^{r-2})}) > 0$, such that

$$\|u_N^{n+1,*} - u(t^{n+1})\|_{L^2}^2 \leq (1 + C^* \tau) \|u_N^n - u(t^n)\|_{L^2}^2 + \tau C^* (\tau^2 + N^{-2r}). \quad (\text{F.12})$$

In fact, recalling also that $\|u_N^{n+1} - u_N^{n+1,*}\|_{L^2} \leq C\tau^2$ by (F.4), we find that an inequality of the form (F.12) remains true with $u_N^{n+1,*}$ replaced by u_N^{n+1} . Indeed, we have

$$\begin{aligned} \|u_N^{n+1} - u(t^{n+1})\|_{L^2}^2 &\leq \|u_N^{n+1,*} - u(t^{n+1})\|_{L^2}^2 + 2C\tau^2 \|u_N^{n+1,*} - u(t^{n+1})\|_{L^2} + C^2\tau^4 \\ &\leq (1 + \epsilon) \|u_N^{n+1,*} - u(t^{n+1})\|_{L^2}^2 + C^2(1 + 4\epsilon^{-1})\tau^4 \\ &\stackrel{(\epsilon:=\tau)}{=} (1 + \tau) \|u_N^{n+1,*} - u(t^{n+1})\|_{L^2}^2 + C^2(1 + 4\tau^{-1})\tau^4 \\ &\leq (1 + C^*\tau) \|u_N^n - u(t^n)\|_{L^2}^2 + \tau C^* (\tau^2 + N^{-2r}), \end{aligned} \tag{F.13}$$

where the last estimate follows from (F.12), and C^* has been suitably enlarged (but still only depends on $r, d, \|u\|_{C_t(H_x^r)}, \|u\|_{C_t^1(H_x^{r-2})}$). In particular, denoting $E^n := \|u_N^n - u(t^n)\|_{L^2}^2$, we have obtained

$$E^{n+1} \leq (1 + C^*\tau)E^n + \tau C^* (\tau^2 + N^{-2r}),$$

and from Gronwall's inequality it now follows that

$$E^n \leq e^{C^*T} [E^0 + C^*T (\tau^2 + N^{-2r})].$$

We note that $E^0 = \|(1 - \dot{P}_N \mathcal{I}_{2N})u(t=0)\|_{L^2}^2 \leq N^{-2r} \|u\|_{H^r}^2$. And hence, we finally find, for $n = 0, \dots, n_T$, that

$$\|u_N^n - u(t^n)\|_{L^2} = \sqrt{E^n} \leq C (\tau + N^{-r}), \tag{F.14}$$

where $C > 0$ depends only on $T, r, d, \|u\|_{C_t(H_x^r)}$ and $\|u\|_{C_t^1(H_x^{r-2})}$.

F.3. Proof of Theorem 3.12. In this appendix, we provide a proof for Theorem 3.12. At this proof relies on Lemma 3.11, we prove this lemma below.

Proof of Lemma 3.11. The proof of Lemma 3.11 is almost identical to the proof of Lemma E.3; where in the present case, we replace $a_N \rightarrow u_N$ and refer to Lemma E.2 point (3), rather than point (2). The only main difference being that in the last step of the proof, the projection P_N is now replaced by the Leray projection \mathbb{P}_N . However, also for \mathbb{P}_N , we observe that $\mathbb{P}_N : L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d) \rightarrow L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d)$ can again be represented *exactly* by a linear Ψ -FNO layer. Thus, replacing the linear Ψ -FNO layer $\widehat{\mathcal{L}}$ which represents P_N in the proof of Lemma E.3 by a layer representing \mathbb{P}_N , an almost identical argument also applies in this case, and yields for any $\epsilon, B > 0$, a Ψ -FNO $\mathcal{N} : L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d) \rightarrow L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d)$, such that

$$\|\mathbb{P}_N(u_N \cdot \nabla u_N) - \mathcal{N}(u_N)\|_{L^2} \leq \epsilon,$$

for all $\|u_N\|_{L^2} \leq B$, with

$$\text{width}(\mathcal{N}) \leq CN^d, \quad \text{depth}(\mathcal{N}), \text{lift}(\mathcal{N}) \leq C,$$

where $C = C(d)$. \square

Proof of Theorem 3.12. Given $N \in \mathbb{N}$, choose $\tau \sim N^{-r}$, such that the CFL condition $\tau N^{d/2+1}U \leq \frac{1}{2}$ is satisfied. This is possible, since $r \geq d/2 + 2$, by assumption. It then follows that $n_T \sim N^r$, and we note that Algorithm 3.9 can be written as the composition of $O(n_T \log(n_T)) = O(N^r \log(N))$ mappings of the form

$$\widehat{\mathcal{N}}_1(u_N^n, w_N^{n,k}) := \begin{bmatrix} u_N^n \\ (1 - \nu\tau\Delta)^{-1}u_N^n - \tau(1 - \nu\tau\Delta)^{-1}\mathbb{P}_N(u_N^n \cdot \nabla w_N^{n,k}) \end{bmatrix},$$

where $\|u_N^n\|_{L^2}, \|w_N^{n,k}\|_{L^2} \leq 2U$ for all k, n . Applying the replacement lemma, Lemma D.5, the claim will thus follow if we can show that there exists a constant $C > 0$ independent of N , such that for any $\epsilon > 0$, there exists a Ψ -FNO

$\mathcal{N}_1 : L_{2N}^2(\mathbb{T}^d; \text{div}) \rightarrow L_{2N}^2(\mathbb{T}^d; \text{div})$, such that $\|\mathcal{N}_1(u_N, w_N) - \widehat{\mathcal{N}}_1(u_N, w_N)\|_{L^2} \leq \epsilon$, and

$$\text{width}(\mathcal{N}_1) \leq CN^d \sim C\epsilon^{-d/r}, \quad \text{depth}(\mathcal{N}_1) \leq C, \quad \text{lift}(\mathcal{N}_1) \leq C.$$

This is immediate for the approximation of the first component of $\widehat{\mathcal{N}}_1$. For the second component, we note that we can write it as a composition:

$$\begin{bmatrix} u_N^n \\ w_N^n \end{bmatrix} \mapsto \begin{bmatrix} u_N^n \\ \mathbb{P}_N \left(u_N^n \cdot \nabla w_N^{n,k} \right) \end{bmatrix} \mapsto (1 - \nu\tau\Delta)^{-1} \left\{ u_N^n - \tau\mathbb{P}_N \left(u_N^n \cdot \nabla w_N^{n,k} \right) \right\}.$$

To finish the proof, we note that the first mapping can be approximated to arbitrary accuracy $\epsilon > 0$ with a Ψ -FNO of width $\lesssim N^d$, and uniformly bounded depth and lift, by Lemma 3.11. The second mapping can be represented *exactly* by a linear FNO layer. Thus, \mathcal{N}_1 can be obtained as the composition of a Ψ -FNO approximating the quadratic non-linearity, and a linear FNO layer, implying the claimed complexity estimate. \square

F.4. A second-order in time accurate pseudo-spectral method for approximating the Navier-Stokes equation (3.12) and its emulation by Ψ -FNOs.

Our aim is to describe a second-order accurate (in time) version of the pseudo-spectral scheme (3.15). To this end, we propose the following scheme,

$$\frac{u_N^{n+1} - u_N^n}{\tau} + \mathbb{P}_N \left(\left[\frac{3}{2}u^n - \frac{1}{2}u^{n-1} \right] \cdot \nabla \frac{1}{2} [u^{n+1} + u^n] \right) = \nu\Delta \frac{1}{2} [u^{n+1} + u^n], \quad (\text{F.15})$$

In contrast to the first-order method (3.15) of the last section, to start the scheme (3.15), we now require two starting values $u_N^0 \approx u(t_0)$ and $u_N^1 \approx u(t_1)$. Given initial data $u_0 \in \dot{H}^r(\mathbb{T}^d; \text{div})$, $r \geq d/2$, we propose to define $u_N^0 := \mathcal{I}_N u_0$, where \mathcal{I}_N denotes the pseudo-spectral projection, and to generate u_N^1 by the first-order accurate Algorithm 3.9 applied on the time-interval $[0, \tau]$, with reduced time-step of size $\sim \tau^2$.

As in the case of (3.15), one needs to solve an implicit operator equation in the time update for (F.15). Analogously, we will use a fixed point iteration to approximate this implicit equation, resulting in the following algorithm,

Algorithm F.6 (Second-order in time approximation of (3.12)).

Input: $U > 0$, $N \in \mathbb{N}$, $T > 0$, $\nu \geq 0$, a time-step $\tau > 0$, such that $n_T = T/\tau \in \mathbb{N}$, and $\tau U N^{d/2+1} \leq \frac{1}{2e}$, initial data $u_N^0 \in L_N^2(\mathbb{T}^d; \text{div})$, such that $\|u_N^0\|_{L^2} \leq U$.

Output: $u_N^{n_T} \in L_N^2(\mathbb{T}^d; \text{div})$ an approximation of the solution $S_T(u_N^0)$ of (3.12) at time $t = T$.

(1) Set

$$\kappa := \left\lceil \frac{\log(T^3/\tau^3)}{\log(2)} \right\rceil \in \mathbb{N}.$$

(2) Compute $u_N^1 \approx u(t_1)$ by applying Algorithm 3.9 on the time-interval $[0, \tau]$, with n_T steps and with time-step $\tau' = \tau/n_T$.

(3) For $n = 1, \dots, n_T - 1$:

(a) Set $w_N^{n,0} := 0$,

(b) For $k = 1, \dots, \kappa_0$: Given the values u_N^{n-1}, u_N^n from the previous steps, compute

$$w_N^{n,k} := F_2^n(w_N^{n,k-1}),$$

(c) Set $u_N^{n+1} := w_N^{n,\kappa_0}$,

Next, we have the following convergence theorem for the algorithm F.6,

Theorem F.7. Let $U, T > 0$. Consider the Navier-Stokes equations on \mathbb{T}^d , for $d \geq 2$. Assume that $r \geq d+2$, and let $u \in C([0, T]; H^r) \cap C^1([0, T]; H^{r-2}) \cap C^2([0, T]; H^1)$ be a solution of the Navier-Stokes equations (3.12), such that $\|u\|_{L^2} \leq U$. Choose a time-step τ , such that $\tau U N^{d/2+1} \leq (2e)^{-1}$. There exists a constant

$$C = C(T, d, r, \|u\|_{C_t(H_x^r)}, \|u\|_{C_t^1(H_x^{r-2})}, \|u\|_{C_t^2(H_x^1)}) > 0,$$

such that with $u_N^0 := \mathcal{I}_N u(0)$, if $\|u(t_1) - u_N^1\|_{L^2} \leq \delta$, and for the sequence $u_N^2, \dots, u_N^{n_T} \in L_N^2(\mathbb{T}^d; \text{div})$ generated by Algorithm 3.9, we have

$$\max_{n=0, \dots, n_T} \|u_N^n - u(t^n)\|_{L^2} \leq C (\delta + \tau^2 + N^{-r}),$$

where $n_T \tau = T$. In particular, choosing a suitable time-step $\tau \sim N^{-r/2}$, and assuming that $\delta \leq CN^{-r}$, we have

$$\max_{n=0, \dots, n_T} \|u_N^n - u(t^n)\|_{L^2} \leq 3CN^{-r},$$

with $n_T \sim N^{r/2}$.

The proof uses very similar techniques as the proof of Theorem 3.10 and we omit it here.

Finally and in complete analogy with the proof of Theorem 3.12, one can prove the following Theorem for the approximation of scheme (F.15) with a Ψ -FNO,

Theorem F.8. Let $U, T > 0$ and viscosity $\nu \geq 0$. Consider the Navier-Stokes equations on \mathbb{T}^d , for $d \geq 2$. Assume that $r \geq d+2$, and let $\mathcal{V} \subset C([0, T]; H^r) \cap C^1([0, T]; H^{r-2}) \cap C^2([0, T]; H^1)$ be a set of solutions of the Navier-Stokes equations (3.12), such that $\sup_{u \in \mathcal{V}} \|u\|_{L^2} \leq U$, and

$$\bar{U} := \sup_{u \in \mathcal{V}} \left\{ \|u\|_{C_t(H_x^r)} + \|u\|_{C_t^1(H_x^{r-2})} + \|u\|_{C_t^2(H_x^1)} \right\}.$$

For $t \in [0, T]$, denote $\mathcal{V}_t := \{u(t) \mid u \in \mathcal{V}\}$. Let $\mathcal{G} : \mathcal{V}_0 \rightarrow \mathcal{V}_T$ denote the solution operator of (3.12), mapping initial data $u_0 = u(t=0)$, to the solution $u(T)$ at $t=T$ of the incompressible Navier-Stokes equations. Then there exists a constant

$$C = C(d, r, U, \bar{U}, T) > 0,$$

such that for any $N \in \mathbb{N}$ there exists a Ψ -FNO $\mathcal{N} : L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d) \rightarrow L_{2N}^2(\mathbb{T}^d; \mathbb{R}^d)$, such that

$$\sup_{u \in \mathcal{V}_0} \|\mathcal{G}(u) - \mathcal{N}(u)\|_{L^2} \leq CN^{-r},$$

and such that

$$\text{width}(\mathcal{N}) \leq CN^d, \quad \text{depth}(\mathcal{N}) \leq CN^{r/2} \log(N), \quad \text{lift}(\mathcal{N}) \leq C.$$

APPENDIX G. PROOF OF THEOREM 4.2

Proof. First, we note that each layer

$$v(x_j) \mapsto \sigma(Wv(x_j) + b_j + \mathcal{F}_N^{-1}(P\mathcal{F}_N v)),$$

is simply the composition of

- an affine mapping $\mathbb{R}^{d_v \times \mathcal{J}_N} \rightarrow \mathbb{R}^{d_v \times \mathcal{J}_N}$, and
- a componentwise application of the activation function σ .

In particular, the Ψ -FNO \mathcal{N} , interpreted as a mapping

$$\mathbb{R}^{\mathcal{J}_N \times d_a} \rightarrow \mathbb{R}^{\mathcal{J}_N \times d_v} \rightarrow \dots \rightarrow \mathbb{R}^{\mathcal{J}_N \times d_v} \rightarrow \mathbb{R}^{\mathcal{J}_N \times d_u},$$

can be represented by an ordinary neural network $\tilde{\beta} : \mathbb{R}^{\mathcal{J}_N \times d_a} \rightarrow \mathbb{R}^{\mathcal{J}_N \times d_u}$, with

$$\text{width}(\tilde{\beta}) = |\mathcal{J}_N| d_v = \text{width}(\mathcal{N}), \quad \text{depth}(\tilde{\beta}) = \text{depth}(\mathcal{N}).$$

In fact, by suitably modifying the linear output layer of $\tilde{\beta}$, we can map the grid values encoded in the output $\tilde{\beta}_j(a) = \mathcal{N}(a)(x_j)$ to the corresponding coefficients in a (real) trigonometric basis $\{\mathbf{e}_k\}_{k \in \mathcal{K}_N}$ with $\text{span}\{\mathbf{e}_k\}_{k \in \mathcal{K}_N} = \text{span}\{e^{i\langle k, x \rangle}\}_{k \in \mathcal{K}_N}$; i.e. by modifying the linear output layer of $\tilde{\beta}$ (and re-indexing the components of the output), we obtain another neural network $\beta : \mathbb{R}^{\mathcal{J}_N \times d_a} \rightarrow \mathbb{R}^{\mathcal{K}_N \times d_u}$, such that

$$\mathcal{N}(a)(x_j) = \sum_{k \in \mathcal{K}_N} \beta_k(a) \mathbf{e}_k(x_j), \quad \forall j \in \mathcal{J}_N. \quad (\text{G.1})$$

Since $\tilde{\beta}$ and β only differ in their output layers, we clearly have

$$\text{width}(\beta) = \text{width}(\tilde{\beta}) = \text{width}(\mathcal{N}), \quad \text{depth}(\beta) = \text{depth}(\tilde{\beta}) = \text{depth}(\mathcal{N}).$$

Since $\{\mathbf{e}_k\}_{k \in \mathcal{K}_N}$ has the same span as $\{e^{i\langle k, x \rangle}\}_{k \in \mathcal{K}_N}$, it follows from (G.1), that

$$\mathcal{N}(a)(x) = \sum_{k \in \mathcal{K}_N} \beta_k(a) \mathbf{e}_k(x), \quad \forall x \in \mathbb{T}^d.$$

To prove the claim, it thus suffices to observe that there exists a neural network

$$\tau : \mathbb{R}^d \rightarrow \mathbb{R}^{\mathcal{K}_N}, \quad x \mapsto \{\tau_k(x)\}_{k \in \mathcal{K}_N},$$

such that the DeepOnet defined by (β, τ) satisfies

$$\begin{aligned} & \sup_{\|a\|_{L^\infty} \leq B} \sup_{x \in \mathbb{T}^d} \left| \mathcal{N}(a)(x) - \sum_{k \in \mathcal{K}_N} \beta_k(a) \tau_k(x) \right| \\ &= \sup_{\|a\|_{L^\infty} \leq B} \sup_{x \in \mathbb{T}^d} \left| \sum_{k \in \mathcal{K}_N} \beta_k(a) [\mathbf{e}_k(x) - \tau_k(x)] \right| \\ &\leq (2N+1)^d \left(\sup_{\|a\|_{L^\infty} \leq B} \max_{k \in \mathcal{K}_N} |\beta_k(a)| \right) \sup_{x \in \mathbb{T}^d} \max_{k \in \mathcal{K}_N} |\mathbf{e}_k(x) - \tau_k(x)| \\ &\leq (2N+1)^d \left(\sup_{\|a\|_{L^\infty} \leq B} \|\mathcal{N}(a)\|_{L^2} \right) \sup_{x \in \mathbb{T}^d} \max_{k \in \mathcal{K}_N} |\mathbf{e}_k(x) - \tau_k(x)| \\ &= \bar{B} \sup_{x \in \mathbb{T}^d} \max_{k \in \mathcal{K}_N} |\mathbf{e}_k(x) - \tau_k(x)|. \end{aligned}$$

□

(Nikola Kovachki) COMPUTING AND MATHEMATICAL SCIENCES, CALIFORNIA INSTITUTE OF TECHNOLOGY, PASADENA, CALIFORNIA 91125, USA

(Samuel Lanthaler and S. Mishra) SEMINAR FOR APPLIED MATHEMATICS, ETH ZÜRICH, RÄMISTRASSE 101, 8092 ZÜRICH, SWITZERLAND