

Variationally correct neural residual regression for parametric PDEs: on the viability of controlled accuracy

MARKUS BACHMAYR*

Institut für Geometrie und Praktische Mathematik, RWTH Aachen University, Templergraben 55,
D-52062 Aachen, Germany

*Corresponding author: bachmayr@igpm.rwth-aachen.de

WOLFGANG DAHMEN

Department of Mathematics, University of South Carolina, 1523 Greene Street, Columbia, SC 29208,
USA

AND

MATHIAS OSTER

Institut für Geometrie und Praktische Mathematik, RWTH Aachen University, Templergraben 55,
D-52062 Aachen, Germany

[Received on 17 June 2024; revised on 31 January 2025]

This paper is about learning the parameter-to-solution map for systems of partial differential equations (PDEs) that depend on a potentially large number of parameters covering all PDE types for which a *stable variational formulation* (SVF) can be found. A central constituent is the notion of *variationally correct residual* loss function, meaning that its value is always uniformly proportional to the squared solution error in the norm determined by the SVF, hence facilitating rigorous *a posteriori* accuracy control. It is based on a single variational problem, associated with the family of parameter-dependent fibre problems, employing the notion of *direct integrals of Hilbert spaces*. Since in its original form the loss function is given as a dual test norm of the residual; a central objective is to develop equivalent computable expressions. The first critical role is played by *hybrid hypothesis classes*, whose elements are piecewise polynomial in (low-dimensional) spatio-temporal variables with parameter-dependent coefficients that can be represented, for example, by neural networks. Second, working with first-order SVFs we distinguish two scenarios: (i) the test space can be chosen as an L_2 -space (such as for elliptic or parabolic problems) so that residuals can be evaluated directly as elements of L_2 ; (ii) when trial and test spaces for the fibre problems depend on the parameters (as for transport equations) we use *ultra-weak* formulations. In combination with *discontinuous Petrov–Galerkin* concepts the hybrid format is then instrumental to arrive at variationally correct computable residual loss functions. Our findings are illustrated by numerical experiments representing (i) and (ii), namely elliptic boundary value problems with piecewise constant diffusion coefficients and pure transport equations with parameter-dependent convection fields.

Keywords: parametric partial differential equations; solution manifolds; physics-informed learning; deep neural networks; stable variational formulations; least-squares variational formulations; discontinuous Petrov–Galerkin methods.

1. Introduction

The problem of efficient computational exploration of families of solutions of parameter-dependent partial differential equations (PDEs), also called *solution manifolds*, arises in many different applications.

Elements of solution manifolds usually describe the viable states of a physical system, and one is often interested in the selection of parameters that best explain given measurements or observational data. Constructing an efficient approximation of the map taking observational data to the states in the manifold is therefore at the heart of many forward and inverse simulation tasks. Approximations of this type can be obtained, for example, by means of projection-based model order reduction methods or by operator learning methodologies based on deep neural networks (DNNs).

To explain this in more detail, we now assume as our starting point a family of PDEs, given in general form by some operator $\mathcal{F}$, depending on parameters p from a domain P . Given $p \in P$, we then seek a solution u in some Banach space $\mathbb{U}$, called the *trial space*, such that

$$\mathcal{F}(u; p) = 0. \quad (1.1)$$

Thus, u is not only a function of *spatio-temporal variables* in the domain on which the PDE is posed, but also of the *parametric variables* $p \in P$. The choice of function space $\mathbb{U}$ that accommodates solutions u depends on $\mathcal{F}$ and determines the metric in which we measure accuracy of approximations to u . For instance, for stationary PDE problems, $\mathbb{U}$ is typically a function space (such as a Sobolev space) on a spatial domain Ω . For time-dependent problems Ω would typically be a space-time cylinder.

In general, the parameter dependence may reflect incompleteness of the model that underlies (1.1). For instance, such models may be based on constitutive laws that are not precisely known or depend on missing data. Assuming that all possibilities of interest are covered by parameters in P we can assume that (1.1) determines a *unique* solution $u = u(p) \in \mathbb{U}$ for each parameter p (for an alternative regularization concept for incomplete models, see Dahmen *et al.* (2021, 2023)). The *solution operator* $\mathcal{S} : p \mapsto u(p)$ is then well defined, and its range

$$\mathcal{M} = \mathcal{M}(P) := \{u(p) : p \in P\}$$

is commonly referred to as the solution manifold associated with the parametric problem.

We postpone a brief discussion of established methods for approximating $\mathcal{S}$ (and hence $\mathcal{M}$) to Section 1.1, but we remark here that the universality of DNNs and their prominent role in modern machine learning offer promising perspectives for learning-based approximations of solution manifolds. An efficient exploration of $\mathcal{M}$ then requires the ability to efficiently evaluate $\mathcal{S}$, which naturally favours approximate representations of $\mathcal{S}$ that depend on as few degrees of freedom as possible. In this sense a computationally handy surrogate for $\mathcal{S}$ can be viewed as a *reduced model*.

A guiding theme in this work is that *accuracy quantification* is an indispensable prerequisite for prediction capability of a simulation task. Moreover, doing so with respect to metrics that are intrinsic to the model is particularly important when dealing with inverse problems. However, quantifying errors for the high-dimensional approximation problems with DNNs in operator learning is a widely open problem. One major obstacle is that the only practically relevant way of constructing approximations by DNNs is via nonconvex optimization, which entails a significant uncertainty in optimization success. As a consequence, the theoretically valid expressive power of a given DNN model may not be practically realizable. From the perspective of model reduction it is not an option to resort to generous *over-parameterizations* of DNNs that are common in machine learning, so as to drive the underlying loss function to zero, and thus ensuring that local minima are actually global ones.

To cope with these issues from the viewpoint of accuracy quantification, the central goal in this article is to develop loss functions for the underlying learning problem that we call *variationally correct*. This

means the following:

At any stage of an associated optimization process the current loss is uniformly proportional to the squared error of the corresponding estimator with respect to a metric that is adapted to the problem. (VC)

Thus, although one may not guarantee in practice that the optimization gets close to a minimizer, the terminal loss becomes a rigorous *a posteriori* error bound certifying the result. The concepts proposed in this paper do not target a specific PDE model, but are designed to cover a possibly wide scope of PDE types covering, for example, dissipative elliptic or parabolic problems, transport problems, but also dispersive models.

While so-called purely data-driven regression in the trial space $\mathbb{U}$ is by definition variationally correct, its significant computational cost motivates looking for residual-type loss functions that avoid the computation of many high-fidelity training samples. A key ingredient in rendering such loss functions variationally correct are *stable variational formulations* for (1.1). This readily leads to idealized formulations in terms of residuals. However, these generally need to be measured in dual norms whose practical evaluation is usually problematic. We highlight different basic strategies towards computable variationally correct residual loss functions, focusing mainly on the combination of deep networks with *first-order system least-squares* and with *discontinuous Petrov–Galerkin* (DPG) formulations for different types of parameter-dependent PDEs.

Concerning the connection to operator learning, our focus here is on the *parameter-to-solution operator*, where the parameter dimension d_p can be large, but remains finite. In general, operator learning instead addresses situations where $\mathcal{S}$ acts on (compact) sets of functions of input data. Corresponding meta-architectures for representing $\mathcal{S}$ therefore include a first ‘discretization layer’ that generates finite-dimensional input parameters for a subsequent layer that approximates $\mathcal{S}$. Of course, with increasing accuracy demands, the dimension of input parameters generally has to grow. Mainly to avoid an additional level of technicalities, we skip here such a discretization step and view $\mathcal{S}$ as a mapping on a parameter domain of potentially high, but finite dimension. Our priority here is to learn $\mathcal{S}$ with respect to the correct model-compliant metrics, as explained later in technical terms, which is of particular importance for stability estimates in inverse problems.

1.1 Beyond projection-based model reduction

Finding computationally efficient surrogates for $\mathcal{M}$ or $\mathcal{S}$ can be viewed as a task of *model order reduction*. Well-established strategies for achieving this include the *reduced basis method* (RBM; see, for example, Rozza *et al.* (2008), Binev *et al.* (2011), Dahmen *et al.* (2014) and Rozza (2014)) and *proper orthogonal decomposition* (POD; see, for example, Kahlbacher & Volkwein (2007)). These methods are based on generating a *linear* (or affine) subspace of $\mathbb{U}$ such that $\mathcal{S}$ can be replaced for each parameter p by a suitable projection (in the case of the RBM, a Galerkin projection) to this subspace. Such methods can in principle offer some degree of error control; while in the RBM one usually aims at approximation in $L_\infty(P; \mathbb{U})$, which can be difficult to achieve in practice for high-dimensional parameter domains (see Cohen *et al.* (2020)), POD is particularly suitable for controlling errors in $L_2(P; \mathbb{U})$.

For constructing the corresponding subspace in both cases, one accepts a major computational *offline cost* that has to be spent only once, while subsequent parameter queries at an *online* stage can be carried out very efficiently—in particular, much faster than a single high-fidelity solve of the PDE. To achieve this efficiency, the dimension of the reduced space required for a certain error needs to be small. In

the case of RBM a lower bound on the error achievable by an n -dimensional subspace is given by the Kolmogorov n -widths

$$d_n(\mathcal{M})_{\mathbb{U}} := \inf_{\dim \mathbb{W}=n} \sup_{u \in \mathcal{M}} \inf_{w \in \mathbb{W}} \|u - w\|_{\mathbb{U}}. \quad (1.2)$$

For many elliptic problems where $\mathcal{S}$ depends holomorphically on the parameters these n -widths decay rapidly even for large parametric dimensions. This can be a consequence of anisotropies in the parametric dependence (see, for example, Cohen & DeVore (2015) and the literature cited there) or of additional structural features (Bachmayr & Cohen, 2017).

Why use DNNs? Unfortunately, for PDEs with little or no dissipative effects, such as hyperbolic problems, this is no longer the case, and methods based on a selection of linear subspaces become very inefficient. In particular, for solutions that exhibit jump discontinuities (such as shocks) with parameter-dependent locations, the n -widths generally decay only algebraically, with a rate that deteriorates exponentially with respect to the dimension of P . Many other methods that build explicit representations of parameter-dependent solutions, such as sparse polynomial expansions in p or low-rank tensor representations (see, for example, Bachmayr et al. (2018) and the references given there), are subject to the same restrictions in terms of the decay of n -widths or related quantities, since they eventually also rely on choices of linear subspaces.

A natural question is thus whether in such cases $\mathcal{M}$ still exhibits some sort of structural sparsity that can be exploited by other types of *nonlinear reduced models*. DNNs can provide the required flexibility. In particular, the solution manifold of pure *linear transport equations* with parameter-dependent convection field has been shown to be approximable (under certain structural assumptions on the convection field, such as affine parameter dependence) by certain DNNs without a curse of dimensionality (Dahmen, 2023). Hence, although the use of DNNs by itself does not guarantee good performance in practice, such nonlinear approximations by compositions have the potential to avoid the basic restrictions of the n -widths (1.2).

1.2 Model reduction by empirical risk minimization

A common, purely data-driven approach to approximating $\mathcal{S}$ (pursued, for example, in Dal Santo et al. (2020), Kaushik et al. (2021), Li et al. (2021) and Dahmen et al. (2022)) is based on generating a training set of synthetic data $\tilde{u}(p) \approx u(p)$ for parameter samples $p \in \hat{P}$, where $\hat{P} \subset P$ is a sufficiently large yet finite subset. Thus, acquiring synthetic data means approximately solving (1.1) for each $p \in \hat{P}$. These synthetic data are then used to construct a surrogate for $\mathcal{S}$ by regression, as an instance of what has become known as *operator learning*. A typical way of training a surrogate $u^{\mathcal{N}}(\cdot, p; \theta^*) \approx u(\cdot, p)$, as a function of spatio-temporal variables, as well as parametric variables representing p , is to minimize a mean-square *empirical risk* functional

$$\|\tilde{u} - u^{\mathcal{N}}(\theta)\|_{\ell_2(\hat{P}; \mathbb{U})}^2 := \frac{1}{\#\hat{P}} \sum_{p \in \hat{P}} \|\tilde{u}(p) - u^{\mathcal{N}}(p, \theta)\|_{\mathbb{U}}^2 \quad (1.3)$$

over $\theta \in \Theta \subset \mathbb{R}^D$, where Θ is a set of trainable weights. This terminology reflects the fact that $\|\tilde{u} - u^{\mathcal{N}}(\theta)\|_{\ell_2(\hat{P}; \mathbb{U})}^2$ may be regarded as an approximation of the ‘ideal risk’

$$\|u - u^{\mathcal{N}}(\theta)\|_{L_2(P; \mathbb{U})}^2 = \int_P \|u - u^{\mathcal{N}}(\theta)\|_{\mathbb{U}}^2 d\mu(p) = \mathbb{E}_{\mu} \left[\|u - u^{\mathcal{N}}(\theta)\|_{\mathbb{U}}^2 \right], \quad (1.4)$$

which is indeed the expected error incurred by $u^{\mathcal{N}}(\theta)$ with respect to the probability measure μ on P . If $\hat{P}$ is a set of samples from this measure (1.3) is a standard Monte Carlo approximation to the expectation.

A rationale behind this approach is therefore to treat the dependence on the spatio-temporal variables deterministically, whereas p is viewed as a random variable (where P is endowed with a probability measure). Moreover, a mean-square error metric is amenable to standard gradient descent optimization techniques and leads to more tractable problems in high parametric dimensions (see, for example, Grohs & Voigtlaender (2023)) than error measures in terms of L_∞ -norms on P . The norm $\|\cdot\|_{\mathbb{U}}$ appearing in the loss function typically involves integration and derivatives. Hence, $\|u^{\mathcal{N}}(\theta)\|_{\mathbb{U}}$ can usually not be evaluated exactly. Here, and in what follows, we ignore the additional perturbations incurred by quadrature, which we assume to cause negligible consistency errors.

A resulting approximation $\tilde{\mathcal{J}}$ to $\mathcal{J}$, based on a surrogate $u^{\mathcal{N}}(\cdot, p; \theta^*)$ with

$$\theta^* \in \arg \min_{\theta \in \Theta} \|\tilde{u} - u^{\mathcal{N}}(\theta)\|_{\ell_2(\hat{P}, \mathbb{U})}^2, \quad (1.5)$$

may be viewed as providing a reduced model for $\mathcal{M}$ provided that the allocated budget Θ of degrees of freedom has a feasible size so as to render the evaluation of $u^{\mathcal{N}}(\cdot, p; \theta^*)$ significantly more efficient than resorting to a highly accurate solver for each instance of p .

Certifying such a reduced model requires quantifying the accuracy of the approximation $\tilde{\mathcal{J}}$. In the light of the preceding discussion a suitable metric for measuring errors is given by (1.4) with a choice of the probability measure μ on P that may depend on additional background information.

In summary, (1.3) can be viewed as a classical regression problem and one can tap into results of machine learning to bound the deviation of the empirical loss in (1.3) and the ideal loss in (1.4),

$$\left| \|u - u^{\mathcal{N}}(\theta)\|_{L_2(P, \mathbb{U})}^2 - \|\tilde{u} - u^{\mathcal{N}}(\theta)\|_{\ell_2(\hat{P}, \mathbb{U})}^2 \right| \quad (1.6)$$

in expectation or probability (see, for example, Barron *et al.* (2008)). Corresponding bounds can be based on concentration inequalities (depending on the measure μ) as well as on complexity bounds for the hypothesis class $\mathcal{H}(\Theta) = \{u^{\mathcal{N}}(\cdot; \theta) : \theta \in \Theta\}$ and the solution manifold $\mathcal{M}$. We defer more detailed discussions of this issue to forthcoming work.

Obvious advantages of this data-driven approach are its conceptual simplicity and that approximations naturally respect the model-compliant metric $\|\cdot\|_{\mathbb{U}}$. In addition, once one has an accurate solver at hand the approach works equally well for linear as well as nonlinear models (1.1).

However, the achievable accuracy depends on $\#\hat{P}$ as well as on the accuracy of the synthetic data $\tilde{u}(p)$. Hence, the fact that this requires computing a potentially large number $\#\hat{P}$ of training snapshots approximating $u(p)$, each one requiring a sufficiently accurate discrete approximate solution of (1.1), may be seen as a serious disadvantage. A remedy resulting in a significant complexity reduction is—very much resembling the construction of reduced bases—to train the reduced model on residual-type loss functions, an approach that has been termed *physics-informed neural networks* (PINNs) (Dongkun *et al.*, 2019; Kharazmi *et al.*, 2019; Raissi *et al.*, 2019; Mishra & Molinaro, 2021; De Ryck & Mishra, 2022).

1.3 Limitations of basic PINN formulations

In its basic version PINN is based on minimizing an empirical risk involving *pointwise* samples of a residual in space and parameter domain. The striking and indeed very tempting point is that this avoids

the need for computing high-fidelity solution snapshots and has almost black-box character, enabling the use of generic public domain software packages for linear as well as nonlinear PDE models. In the above terms a typical residual loss function, for $\mathcal{F}$ as in the stationary case of the general parametric problem (1.1), reads

$$\frac{1}{MN} \sum_{i,j=1}^{M,N} |\mathcal{F}(u^{\mathcal{N}}(x^i, p^j; \theta); p^j)|^2. \quad (1.7)$$

Such a loss is amenable to standard machine learning strategies based, for example, on stochastic gradient descent with gradient evaluation by backpropagation.

However, (1.7) amounts to treating residuals as elements of an L_2 -space, which is appropriate only under quite specific assumptions on $\mathcal{F}$. The fundamental issues with such a formulation can be seen in the basic example

$$\mathcal{F}(u; p) = f + \operatorname{div}(a(p)\nabla u) = 0 \quad \text{in } \Omega, \quad u|_{\partial\Omega} = 0.$$

The standard loss function in (1.7) then reads

$$\frac{1}{N} \sum_{j=1}^N \left\{ \frac{|\Omega|}{M_\Omega} \sum_{i=1}^{M_\Omega} \left| f + \operatorname{div} \left(a(p^j) \nabla u^{\mathcal{N}}(x^i, p^j; \theta) \right) \right|^2 + \frac{|\partial\Omega|}{M_{\partial\Omega}} \sum_{i=1}^{M_{\partial\Omega}} \left| u^{\mathcal{N}}(x^i, p^j; \theta) \right|^2 \right\}, \quad (1.8)$$

where p^j are samples in Y and in the two inner summations, x^i are samples in Ω and $\partial\Omega$, respectively. As $N, M_\Omega, M_{\partial\Omega} \rightarrow \infty$, the left part of (1.8) will tend to $\|f + \operatorname{div}(a\nabla u^{\mathcal{N}}(\theta))\|_{L_2(\Omega \times Y)}^2$, and the right part to $\|u^{\mathcal{N}}(\theta)\|_{L_2(\partial\Omega)}^2$. On the one hand, it is well known that the domain term is too strong and may not even define a bounded expression for nonsmooth domains or diffusion coefficients, or when f is a proper functional in $H^{-1}(\Omega)$ (such as a trace integral). On the other hand, the boundary term is too weak: a well-posed variational formulation would require an $H^{1/2}(\partial\Omega)$ -norm. This mismatch implies that the size of the loss (1.8) does not provide any certifiable bound for the error $u - u^{\mathcal{N}}(\theta)$, which, here, one would typically aim to measure in $\mathbb{U} = H_0^1(\Omega)$.

In the recent work, [Bonito et al. \(2025\)](#), the authors propose and analyse a variant of PINN for the (parameter-independent) Poisson problem where the loss does provide also a rigorous computable upper bound for the current error in the $H^1(\Omega)$ -norm. There, the sample-based ℓ_2 -norm of the residual is replaced by an ℓ_p -norm where $p \leq 2$ depends on the regularity of the data. The main differences from the present approach are as follows: (i) the method relies on excess regularity assumptions; (ii) it only provides an upper bound for the $H^1(\Omega)$ -error; (iii) it is to serve as a stand-alone solver for a particular given PDE rather than as a model reduction concept.

1.4 Proportionality of error and residual

As already emphasized earlier, in view of the uncertainty of optimization success and the lack of any realistic *a priori* error bounds for DNN approximations of PDEs, a minimal goal regarding accuracy control is to infer from the current loss, at any stage of the optimization, the size of the error in the model-compliant metric $\|\cdot\|_{\mathbb{U}}$. That is, any residual-based empirical loss along the lines of (1.7) or (1.8) using parametric samples from $\hat{P}$ should remain proportional to $\|u - u^{\mathcal{N}}(\theta)\|_{\ell_2(\hat{P}, \mathbb{U})}^2$ uniformly in

the number of samples $\#\widehat{P}$. Hence, the remarks in Section 1.2 still apply to relate a residual loss, via $\|u - u^{\mathcal{N}}(\theta)\|_{\ell_2(\widehat{P}, \mathbb{U})}^2$, to $\|u - u^{\mathcal{N}}(\theta)\|_{L_2(P, \mathbb{U})}^2$ in expectation or with high probability.

Our strategy for devising residual loss functionals with this property is based on a weak form of (1.1) for each $p \in P$: viewing the residual as a functional, find $u(p) \in \mathbb{U}$ such that

$$\mathcal{F}(u(p); p)(v) = 0 \quad \text{for all } v \in \mathbb{V} \quad (1.9)$$

with a suitable test space $\mathbb{V}$. This test space should be chosen (depending on $\mathbb{U}$) such that $\mathcal{F}(\cdot; p): \mathbb{U} \rightarrow \mathbb{V}'$ satisfies

$$\|\mathcal{F}(w; p)\|_{\mathbb{V}'} \approx \|w - u\|_{\mathbb{U}} \quad (1.10)$$

with constants that are uniform in p . Let us assume for the moment that the spaces $\mathbb{U}, \mathbb{V}$ can be chosen to be *independent* of $p \in P$. Then, (1.10) implies

$$\|\mathcal{F}(w)\|_{L_2(P, \mathbb{V}')}^2 = \int_P \|\mathcal{F}(w; p)\|_{\mathbb{V}'}^2 d\mu(p) \approx \int_P \|u - w\|_{\mathbb{U}}^2 d\mu(p) = \|w - u\|_{L_2(P, \mathbb{U})}^2, \quad (1.11)$$

which suggests interpreting $\mathcal{F}$ as a mapping from $\mathbb{X} = L_2(P, \mathbb{U})$ to $\mathbb{Y}' = (L_2(P; \mathbb{V}))' = L_2(P, \mathbb{V}')$.

As will be seen later, to guarantee the validity of (1.9), it is not always possible to choose a pair of trial and test space $\mathbb{U}, \mathbb{V}$ that is independent of p . In fact, they may even have to differ as *sets* for different p , which we refer to as ‘*essentially different*’. We will show in Section 2.3 that under mild assumptions on the p -dependence integration in (1.11) is well defined, and one still arrives at an error–residual relation of a *single variational formulation* over a pair of suitably extended versions of the spaces $\mathbb{X}$ and $\mathbb{Y}$, defined above.

At any rate, the corresponding loss functional in the spirit of (1.7) thus takes the form

$$\frac{1}{N} \sum_{i=1}^N \|\mathcal{F}(u^{\mathcal{N}}(\cdot, p^i; \theta); p^i)\|_{\mathbb{V}'}^2.$$

Unless (1.10) holds when $\mathbb{V}$ itself is also an L_2 -space the obvious difficulty in evaluating the dual norms

$$\|\mathcal{F}(w)\|_{\mathbb{V}'} = \max_{v \in \mathbb{V}} \frac{\mathcal{F}(w)(v)}{\|v\|_{\mathbb{V}}}$$

is the required maximization over the entire Hilbert space $\mathbb{V}$.

1.5 Relation to existing results

Whenever the solution of a PDE is characterized by an energy minimization principle (the elliptic case) a natural way of producing approximate solutions in the right metric is to minimize the corresponding energy functional over the chosen hypothesis class. In a finite element context this can actually also be used to produce reliable and efficient posteriori error bounds. Unfortunately, these concepts do not carry over to more general hypothesis classes such as DNNs, so that this approach does in general not come with accuracy quantification. This remains crucial because of the issues inherent in the minimization of

such quadratic functionals over highly nonlinear hypothesis classes. Moreover, this approach is limited to elliptic problems.

The idea of minimizing instead a variationally correct residual applies to a wider scope and has been considered in several previous works, sometimes referred to as weak adversarial networks. More specifically, the authors in [Bao et al. \(2020\)](#), [Zang et al. \(2020\)](#) and [Chen et al. \(2023\)](#) discuss specific PDE models, namely so-called Friedrichs systems as well as second-order elliptic problems, primarily not for the purpose of model reduction or operator learning, but to devise DNN-based PDE solvers. This has been the primary focus (albeit not always the original motivation) also in several more recent works such as [Berrone et al. \(2022\)](#), [Ernst & Urban \(2024\)](#), [Taylor et al. \(2023\)](#) and [Bertoluzza et al. \(2024\)](#), where the considered PDEs are posed on domains of small dimension. In all these works the need to minimize a dual norm has been identified as a major theoretical as well as practical hurdle. For instance, the natural idea of solving $\min_{\theta \in \Theta} \|\mathcal{F}(u^{\mathcal{N}}(\theta))\|_{\mathbb{V}}$ by repeatedly performing some gradient ascent steps on

$$\frac{\mathcal{F}(u^{\mathcal{N}}(\theta))(v^{\mathcal{N}}(\psi))}{\|v^{\mathcal{N}}(\psi)\|_{\mathbb{V}}} \quad (1.12)$$

over elements $\psi \in \Psi$ of a suitable, sufficiently rich test system Ψ , followed by some descent steps over the primal system Θ , is problematic and not very reliable. Again, ensuring convergence of the optimization over the test space is an issue. For low-dimensional problems some of these works consider therefore more manageable test systems like wavelet systems ([Ernst & Urban, 2024](#)), finite elements ([Berrone et al., 2022](#)) or trigonometric systems ([Taylor et al., 2023](#)). These remedies are unfortunately not feasible in the present high-dimensional parametric setting. Finally, enforcing essential boundary conditions in the correct norms for DNN approximations poses some challenges ([Berrone et al., 2022](#)).

Note that the obstruction of having to evaluate a (nontrivial) dual norm disappears when the test spaces $\mathbb{V}$ for the fibre problems happen to be self-dual, in other words, when $\mathbb{V}$ (and hence $\mathbb{Y}$) is an L_2 -type space. This is taken up in Section 3.1, where we show under which circumstances directly computable L_2 -least squares losses are variationally correct. We note that in [Muzalevsky & Repin \(2021\)](#) and [Fanaskov et al. \(2024\)](#) such loss functions are used for a parameter-independent problem on the spatial domain in a first-order system formulation of Poisson's equation. There, the flux variable σ is interpreted as a 'certificate' that leads to a strict upper bound of the error in u .

1.6 Hybrid approximation format

Due to the issues mentioned above, rather than approximating parameter-dependent solutions directly by DNNs in all variables, we focus here on a hybrid approximation format that uses piecewise polynomial approximations in the spatio-temporal variables combined with DNN approximations of the parameter-dependent coefficients. In doing so, we exploit finite element methods in small spatial dimensions, where concepts for the reliable evaluation of dual norms are available.

While the specific choice of finite element basis functions depends on the PDE model at hand, the general structure can be described as follows. Given finite element basis functions $\{\phi_i\}_{i \in \mathcal{I}}$ on the spatio-temporal domain under consideration, we use approximations of the form

$$u(t, x, p) \approx \sum_{i \in \mathcal{I}} \mathbf{u}_{\text{NN},i}(p; \theta) \phi_i(t, x), \quad (1.13)$$

where $p \mapsto \mathbf{u}_{\text{NN}}(p; \theta) = (\mathbf{u}_{\text{NN},i}(p; \theta))_{i \in \mathcal{I}}$ is realized by a neural network with trainable parameters θ . Representations of this form have been tested for parametric elliptic problems, for example, in [Geist et al. \(2021\)](#) and [Hei et al. \(2023\)](#); different from our approach these approximations were trained on samples of finite element solutions as in (1.5), where the loss function quantifies the relative error in energy norm to the parameter-dependent discrete solutions.

Let us note that instead in our approach, by the use of appropriate residual-based loss functions, we obtain estimates of the error (up to data oscillation) with respect to the *exact solution* of the underlying continuous problem.

The format (1.13) is instrumental for treating dual norms in variationally correct loss functions without the use of min-max optimization as in (1.12). It also offers a number of further significant advantages. First, rigorously enforcing essential boundary conditions for DNNs is an issue, see [Berrone et al. \(2022\)](#) for various strategies. The hybrid format instead greatly facilitates incorporating essential (spatial) boundary conditions. Second, while highly nonconvex optimization problems will always remain problematic, we view this format as an essential strategic ingredient that helps mitigate the challenges of nonconvex optimization. In fact, since the input parameters of the coefficient neural networks are just the model parameters training requires only efficient backpropagation. When using DNNs with both spatio-temporal and parametric variables as input parameters one needs to compute higher-order derivatives with respect to all variables. This has to be done with great care to avoid a significant loss of efficiency, see, e.g., [Taylor et al. \(2023\)](#). Finally, the hybrid format favours the use of *nested refinement* strategies that help rendering the optimization process more robust.

1.7 Novelty and organization of the material

The central contribution in this work is to propose and analyse concepts for operator learning for a wide and diverse scope of PDE models with rigorous accuracy control in model-compliant metrics that are feasible in high-dimensional parametric regimes. The main conceptual constituents can be summarized as follows:

1. Reformulate the given family of parametric PDEs as a *single variational problem* posed over Hilbert spaces of functions of spatio-temporal and parametric variables, see Section 2. Roughly, this entails treating spatio-temporal variables deterministically while a probabilistic interpretation of the parametric domain allows us to resort to Monte Carlo sampling as a basis for a learning approach.
2. Establish *well-posedness* of the single high-dimensional problem (in the sense of the Babuka-Neas Theorem) from judiciously chosen *stable variational formulations* of the parametric (low-dimensional) fibre problems, see Theorem 3. This gives rise to an *ideal residual loss function*, which is indeed variationally correct. The term ‘ideal’ reflects that this is a preliminary step because such residual loss functions typically involve a *dual norm* whose computational evaluation is problematic. Hence, this serves as the starting point for deriving from such ideal losses that are still variationally correct, but are also practically feasible. We stress that the scope of PDE models to which our approach applies corresponds in essence to the ability of contriving stable variational formulations.
3. Section 3 is then devoted to *practical variationally correct residual loss functions*. Specifically, this is based on first reformulating, if necessary, a given parametric PDE as *first-order systems*. We then distinguish in Section 3.1 two scenarios given in (S) and (UW), respectively: In the first scenario (S) we identify stable variational formulations inducing an operator that maps (an appropriate subspace of) its graph space as trial space bijectively onto an L_2 -type space. In the second scenario (UW),

an L_2 -type space is taken as trial space and the concept of an *optimal test norm* is used to define a test space leading to an induced operator with condition number equal to one. This is referred to as *ultra-weak formulation*.

The relevance of (S) is clear. Since the residual is measured in L_2 the ideal residual loss is also practical. We demonstrate this variant in Sections 3.2.1 and 3.2.2 for a classical second-order elliptic problem with parameter-dependent diffusion coefficient and an analogous parabolic initial-boundary value problem with a space-time variational formulation. In Section 5.2.1 we present some numerical experiments for the elliptic example.

In (UW) we still have to deal with a nontrivial dual norm in the ideal loss function. The point is that this ultra-weak formulation conveniently accommodates stable DPG formulations, see Section 3.3.2. It is the combination of DPG concepts with the hybrid format of our hypothesis classes, as described in Section 1.6, that allows us to derive practical variationally correct loss functions. This is exemplified for linear transport equations with parameter-dependent convection field in Section 3.3.1. Corresponding numerical experiments are presented in Section 5.3. The (VC) loss functions based on the neural DPG approach do entail a higher computational cost than standard PINN losses based on spatial point samples of the residual, as discussed in Appendix C. However, this affects primarily the *offline cost*. We regard this as an inevitable price for rigorous accuracy control, independent of optimization success.

4. The hybrid format mentioned above is essential for scenario (UW). Roughly, the elements in corresponding hypothesis classes are piecewise polynomials as functions of spatio-temporal variables with parameter-dependent coefficients that may, for instance, be represented as DNNs. This is not strictly needed for scenario (S). We have nevertheless tested it in comparison with a standard PINN approach for (S) as well. Moreover, we discuss sparsification formats that do reduce representation complexity significantly. Aside from this, it offers several major advantages: it facilitates naturally enforcing essential boundary conditions, and it does not require expensive forward differentiation with respect to the spatio-temporal variables in the training process. In the absence of parameters one falls back on well-established and accurate PDE solvers.

A few words on the choice of the test cases are in order. A major distinction between the elliptic case and transport equations lies in the fact in the former case trial and test spaces can be chosen as *independent* of the parameters. In this case, graph spaces as trial spaces are appropriate, which is scenario (S). By contrast, for transport equations either trial or test space depends *essentially* on the parameters, that is, the spaces vary *as sets* when the parameters vary. L_2 is then a space that robustly accommodates the solutions for all permissible parameters. For this reason we resort to scenario (UW) in this case. As different as (S) and (UW) appear to be they are nevertheless closely related as argued in Section 3.1. Moreover, when finally resorting to DPG concepts one relies on elements of both scenarios, because the additional auxiliary skeleton variables can be interpreted as approximations in the graph space (see also the proof of well-posedness in Broersen *et al.* (2018)).

2. Variational formulations and residual loss functions

2.1 Parametric fibre problems

As indicated earlier, the validity of (1.10) hinges on a suitable variational formulation for each $p \in P$. For linear problems (1.9) takes for each $p \in P$ the form: find $u(p) \in \mathbb{U}_p$ such that

$$b(u(p), v; p) - f(v) = 0 \quad \text{for all } v \in \mathbb{V}_p, \quad (2.1)$$

viewed as a family of linear operator equations $\mathcal{B}_p u(p) = f$ in $\mathbb{V}'_p$, where $(\mathcal{B}_p w)(v) = b(w, v; p)$, $v \in \mathbb{V}_p$. As mentioned earlier, the spaces $\mathbb{U}_p$ and $\mathbb{V}_p$ are to be chosen so as to ensure well-posedness of (2.1). By the celebrated Babuška–Nečas Theorem (see, e.g., Braess (1997)) this is equivalent to the validity of

$$\sup_{w \in \mathbb{U}_p} \sup_{v \in \mathbb{V}_p} \frac{b(w, v; p)}{\|w\|_{\mathbb{U}_p} \|v\|_{\mathbb{V}_p}} \leq C_b, \quad \inf_{w \in \mathbb{U}_p} \sup_{v \in \mathbb{V}_p} \frac{b(w, v; p)}{\|w\|_{\mathbb{U}_p} \|v\|_{\mathbb{V}_p}} \geq c_b, \quad (2.2)$$

for each $v \in \mathbb{V}_p \exists w_v \in \mathbb{U}_p$ such that $b(w_v, v; p) \neq 0$,

for some constants $0 < c_b = c_b(p)$, $C_b = C_b(p) < \infty$, which generally may depend on p . We remark that respective stable variational formulations are known for a wide scope of problems covering dissipative, indefinite, as well as dispersive models; see, for example, Carstensen *et al.* (2016).

REMARK 2.1. Given a (linear) PDE $\mathcal{B}_p u = f$, underlying (2.1), in strong form, it should be viewed as part of the problem to identify a suitable pair of trial and test space $\mathbb{U}_p, \mathbb{V}_p$ (first on the infinite-dimensional level) for which (2.2) holds. If the problem is *elliptic* (or *coercive*) this choice is easy, namely $\mathbb{U}_p = \mathbb{V}_p$, where $\mathbb{U}_p$ is the energy space. In general, when the problem is indefinite, nonsymmetric, or singularly perturbed, one may have to accept $\mathbb{U}_p \neq \mathbb{V}_p$ in order to warrant stability, see Proposition 3.3 and Remark 3.4 for a ‘roadmap’ of how to proceed.

As explained in Section 1.4, if all the spaces $\mathbb{U}_p, \mathbb{V}_p$ agree as sets $\mathbb{U}, \mathbb{V}$, with respective equivalent norms $\|\cdot\|_{\mathbb{U}}, \|\cdot\|_{\mathbb{V}}$, the Bochner spaces $\mathbb{X} = L_2(P, \mathbb{U})$, $\mathbb{Y} = L_2(P, \mathbb{V})$ provide a natural pair of trial and test spaces for a single variational problem determining the fibre solutions of (2.1) as functions of spatial and parametric variables. We refer to this as the *lifted problem*.

Unfortunately, important examples that are discussed below reveal that to ensure well-posedness in the sense of (2.2) it may be necessary to accept families of trial and test pairs $\mathbb{U}_p, \mathbb{V}_p$ that differ for different p in an *essential way*, by which we mean that they differ as sets with nonequivalent norms.

To be still able to recast the family (2.1) of weak formulations as a *single space-parameter* weak formulation we need to address two issues: first, what are suitable replacements of the classical Bochner spaces as trial and test spaces and second, how does well-posedness of fibre problems (2.1) relate to the well-posedness of the resulting lifted problem?

Regarding the first issue the remarks in Section 1.2 suggest choosing $\mathbb{X}$ as a suitable space of functions $v: p \mapsto v(p) \in \mathbb{U}_p$ such that the quantities

$$\int_P \|v(p)\|_{\mathbb{U}_p}^2 d\mu(p) < \infty,$$

and similarly for $\mathbb{Y}$ with $\mathbb{U}_p$ replaced by $\mathbb{V}_p$, remain well defined. To ensure that resulting notions of spaces $\mathbb{X}$ and $\mathbb{Y}$ are meaningful we need some mild assumptions on the dependence of the fibre spaces $\mathbb{U}_p$ and $\mathbb{V}_p$ on the parameter p that warrant measurability of elements in the Cartesian products $\prod_{p \in P} \mathbb{U}_p, \prod_{p \in P} \mathbb{V}_p$. This leads us to the notion of *direct integrals* discussed next.

2.2 Direct integrals

We now consider a general notion of measurability, going back to von Neumann (Von Neumann, 1949), of mappings $p \mapsto v(p) \in \mathbb{U}_p$ as discussed above. The resulting Hilbert spaces can be regarded as generalizations of the direct sums of Hilbert spaces to the case of uncountable index sets.

DEFINITION 1. Let $(\mathbb{W}_p)_{p \in P}$ be a family of separable Hilbert spaces. We then call a sequence $(\xi_n)_{n \in \mathbb{N}}$ in the Cartesian product $\prod_{p \in P} \mathbb{W}_p$ a *fundamental sequence of μ -measurable sections* if for all n and m the function $P \ni p \mapsto \langle \xi_n(p), \xi_m(p) \rangle_{\mathbb{W}_p}$ is μ -measurable, and for each $p \in P$ the finite linear combinations of $\xi_n(p)$, $n \in \mathbb{N}$, are dense in $\mathbb{W}_p$. Under these conditions we say that $v \in \prod_{p \in P} \mathbb{W}_p$ is μ -measurable if $p \mapsto \langle v(p), \xi_n(p) \rangle_{\mathbb{W}_p}$ is μ -measurable for $n \in \mathbb{N}$.

The following result is shown in (Takesaki, 1979, chapter IV.8) (see also (Dautray & Lions, 1985, chapter VIII.3)).

THEOREM 1. Let $(\mathbb{W}_p)_{p \in P}$ be a family of separable Hilbert spaces such that $\prod_{p \in P} \mathbb{W}_p$ contains a fundamental sequence of μ -measurable sections. Then, the spaces $L_2(P, (\mathbb{W}_p)_{p \in P})$ defined by

$$L_2(P, (\mathbb{W}_p)_{p \in P}) = \left\{ v \in \prod_{p \in P} \mathbb{W}_p \text{ } \mu\text{-measurable} : \int_{p \in P} \|v(p)\|_{\mathbb{W}_p}^2 d\mu(p) < \infty \right\},$$

where elements that agree μ -almost everywhere are identified, endowed with the inner product

$$\langle v, w \rangle_{L_2(P, (\mathbb{W}_p)_{p \in P})} = \int_P \langle v(p), w(p) \rangle_{\mathbb{W}_p} d\mu(p),$$

are Hilbert spaces.

The spaces provided by Theorem 1 are called *direct integrals* (or also *Hilbert integrals*). The simplest scenario where Theorem 1 applies concerns the case where $\mathbb{W}_p = \mathbb{W}$, with equivalent norms, uniformly in $p \in P$.

REMARK 2.2. If $\mathbb{W}$ is a separable Hilbert space, the Bochner space $\mathbb{X} = L_2(P, \mathbb{W})$ fulfils Assumption 2.3, see (Dautray & Lions, 1985, VIII.3, remark 2). This justifies the previous discussion leading to (1.11).

We will also need an analogous characterization of the dual of $L_2(P, (\mathbb{W}_p)_{p \in P})$, which will be applied, in particular, to the space $(L_2(P, (\mathbb{V}_p)_{p \in P}))'$. This is instrumental for arriving at a well-defined ideal residual loss function also in scenarios where the fibre test spaces $\mathbb{V}_p$ differ essentially.

In general, for a Hilbert space $\mathbb{H}$ with inner product $\langle \cdot, \cdot \rangle_{\mathbb{H}}$, we denote by $\mathcal{R} = \mathcal{R}_{\mathbb{H}}: \mathbb{H}' \rightarrow \mathbb{H}$ the inverse of the Riesz isometry identifying $\mathbb{H}$ with its dual $\mathbb{H}'$, which for $\ell \in \mathbb{H}'$ is given by

$$\langle \mathcal{R}\ell, v \rangle_{\mathbb{H}} = \ell(v), \quad v \in \mathbb{H}. \quad (2.3)$$

For direct integrals $\mathbb{H} = L_2(P, (\mathbb{W}_p)_{p \in P})$ this mapping can be characterized in terms of the fibre Riesz lifts $\mathcal{R}_p: \mathbb{W}'_p \rightarrow \mathbb{W}_p$ defined for each $p \in P$ and $\ell_p \in \mathbb{W}'_p$ by $\langle \mathcal{R}_p \ell_p, v \rangle_{\mathbb{W}_p} = \ell_p(v)$ for each $v \in \mathbb{W}_p$.

THEOREM 2. For $\mathbb{H} = L_2(P, (\mathbb{W}_p)_{p \in P})$ we have

$$\mathbb{H}' = \left\{ \ell \in \prod_{p \in P} \mathbb{W}'_p \text{ } \mu\text{-measurable} : \int_{p \in P} \|\ell(p)\|_{\mathbb{W}'_p}^2 d\mu(p) < \infty \right\} \quad (2.4)$$

with inner product

$$\langle \ell, \tilde{\ell} \rangle_{\mathbb{H}'} = \int_P \langle \ell(p), \tilde{\ell}(p) \rangle_{\mathbb{W}_p'} d\mu(p), \quad \ell, \tilde{\ell} \in \mathbb{H}',$$

and for each $\ell \in \mathbb{H}'$,

$$\langle \mathcal{R}\ell, v \rangle_{\mathbb{H}} = \int_P \langle \mathcal{R}_p \ell(p), v(p) \rangle_{\mathbb{W}_p} d\mu(p), \quad (2.5)$$

which says that for each $\ell \in \mathbb{H}'$

$$\ell = \prod_{p \in P} \ell(p) \quad \text{such that} \quad (\mathcal{R}\ell)(p) = \mathcal{R}_p \ell(p), \quad p \in P. \quad (2.6)$$

Proof. Let $\ell \in \mathbb{H}'$. Then, for all $v \in \mathbb{H}$,

$$\begin{aligned} \ell(v) &= \langle \mathcal{R}\ell, v \rangle_{\mathbb{H}} = \int_P \langle (\mathcal{R}\ell)(p), v(p) \rangle_{\mathbb{W}_p} d\mu(p) \\ &= \int_P \mathcal{R}_p^{-1}((\mathcal{R}\ell)(p))(v(p)) d\mu(p) = \int_P \ell_p(v(p)) d\mu(p). \end{aligned}$$

Defining $\ell_p := \mathcal{R}_p^{-1}((\mathcal{R}\ell)(p)) \in \mathbb{W}_p'$ we therefore have

$$\langle \mathcal{R}\ell, v \rangle_{\mathbb{H}} = \int_P \langle \mathcal{R}_p \ell_p, v(p) \rangle_{\mathbb{W}_p} d\mu(p). \quad (2.7)$$

Given now a fundamental sequence of μ -measurable sections $\phi_n(p) \in \mathbb{W}_p$ let $\xi_n(p) := \mathcal{R}_p^{-1}\phi_n(p)$, $n \in \mathbb{N}$. Since $\mathcal{R}_p$ are isometries the $\xi_n(p)$ are dense in $\mathbb{W}_p'$. Noticing that $p \mapsto \langle \xi_n(p), \xi_m(p) \rangle_{\mathbb{W}_p'} = \langle \phi_n(p), \phi_m(p) \rangle_{\mathbb{W}_p}$ it follows that the $\xi_n(p)$ form a fundamental system in $\mathbb{W}_p'$, $p \in P$. Moreover,

$$\langle \ell_p, \xi_n(p) \rangle_{\mathbb{W}_p'} = \langle \mathcal{R}_p^{-1}(\mathcal{R}\ell)(p), \xi_n(p) \rangle_{\mathbb{W}_p'} = \mathcal{R}_p^{-1}(\mathcal{R}\ell)(p)(\phi_n(p)) = \langle (\mathcal{R}\ell)(p), \phi_n(p) \rangle_{\mathbb{W}_p}$$

shows measurability of $p \mapsto \langle \ell_p, \xi_n(p) \rangle_{\mathbb{W}_p'}$. Thus, ℓ can be identified with a μ -measurable function $\ell \in \prod_{p \in P} \mathbb{W}_p'$ with $\ell(p) = \ell_p$. In particular, (2.5) follows from (2.7). Furthermore, we observe that $\|\ell(p)\|_{\mathbb{W}_p'} = \|(\mathcal{R}\ell)(p)\|_{\mathbb{W}_p}$ and

$$\int_P \|\ell(p)\|_{\mathbb{W}_p'}^2 d\mu(p) = \int_P \|(\mathcal{R}\ell)(p)\|_{\mathbb{W}_p}^2 d\mu(p) < \infty,$$

and we thus obtain (2.4). $\square$

An application of these concepts to the families $(\mathbb{U}_p)_{p \in P}$, $(\mathbb{V}_p)_{p \in P}$ in (2.1), where these spaces may differ essentially from each other, requires verifying the validity of the following:

PROPERTY 2.3. For the families $(\mathbb{U}_p)_{p \in P}$, $(\mathbb{V}_p)_{p \in P}$ of separable Hilbert spaces there exist fundamental sequences of μ -measurable sections $(\phi_n)_{n \in \mathbb{N}}$ and $(\psi_n)_{n \in \mathbb{N}}$, respectively.

A convenient criterion that applies in our context later below can be formulated as follows.

PROPOSITION 2.4. If $\bigcap_{p \in P} \mathbb{U}_p$ and $\bigcap_{p \in P} \mathbb{V}_p$ are dense in each $\mathbb{U}_p$ and $\mathbb{V}_p$, respectively, then [Property 2.3](#) is satisfied.

This observation is of interest also in the following respect. The above notion of direct integral is in some sense very flexible and permits counterintuitive facts. For instance, if some family of Hilbert spaces $\mathbb{W}_p$, $p \in P$, gives rise to a direct integral $L_2(P, (\mathbb{W}_p)_{p \in P})$, not every choice of closed subspaces $\mathbb{V}_p \subset \mathbb{W}_p$ induces a direct integral structure $L_2(P, (\mathbb{V}_p)_{p \in P}) \subset L_2(P, (\mathbb{W}_p)_{p \in P})$. We refer the interested reader to [Example B1](#) in [Appendix B](#).

However, under the assumption in [Proposition 2.4](#), the expected subspace relation holds. The precise circumstances are given in [Proposition B.1](#) that can also be found in [Appendix B](#).

2.3 The lifted problem and its well-posedness

We are now ready to reformulate [\(2.1\)](#) as a single linear operator equation

$$\mathcal{B}u = F, \quad (2.8)$$

involving functions of spatial and parametric variables, so as to ensure validity of a tight error-residual relation.

In general, for such a linear operator equation, the question for which data F [\(2.8\)](#) has a unique solution u in a suitable trial space $\mathbb{X}$ depends on the *mapping properties* of $\mathcal{B}$, viewed as a mapping from some (infinite-dimensional) *trial space* $\mathbb{X}$ onto a suitable target space, accommodating the data F . If [Property 2.3](#) holds, which we assume throughout the remainder of this section, natural candidates for a reformulation of [\(2.1\)](#) are the spaces

$$\mathbb{X} = L_2\left(P, (\mathbb{U}_p)_{p \in P}\right) \quad \text{and} \quad \mathbb{Y} = L_2\left(P, (\mathbb{V}_p)_{p \in P}\right) \quad (2.9)$$

with respective fundamental sequences $(\phi_n)_{n \in \mathbb{N}}$, $(\psi_n)_{n \in \mathbb{N}}$. In addition, we use the following property that needs to be verified for the given problem.

PROPERTY 2.5. Assuming [Property 2.3](#), for each $p \in P$, let $b(\cdot, \cdot; p): \mathbb{U}_p \times \mathbb{V}_p \rightarrow \mathbb{R}$ be a bounded bilinear form such that for all $n, m \in \mathbb{N}$, the mapping $P \ni p \mapsto b(\phi_n(p), \psi_m(p); p)$ is measurable.

Under these conditions the quantities

$$b(u, v) := \int_P b(u(p), v(p); p) \, d\mu(p), \quad F(v) = \int_P f(v(p)) \, d\mu(p), \quad (2.10)$$

are well defined. Solving [\(2.8\)](#) for $\mathcal{B}$, defined by $(\mathcal{B}w)(v) = b(w, v)$, $w \in \mathbb{X}$, $v \in \mathbb{Y}$, boils down to finding $u \in \mathbb{X}$ such that for $F \in \mathbb{Y}'$

$$b(u, v) = F(v), \quad \forall v \in \mathbb{Y}. \quad (2.11)$$

As has been used already for the fibre problems well-posedness of (2.11) is *characterized* by the Babuška–Nečas Theorem through the existence of constants $0 < c_b \leq C_b < \infty$ such that

$$\sup_{w \in \mathbb{X}} \sup_{v \in \mathbb{Y}} \frac{b(w, v)}{\|w\|_{\mathbb{X}} \|v\|_{\mathbb{Y}}} \leq C_b, \quad \inf_{w \in \mathbb{X}} \sup_{v \in \mathbb{Y}} \frac{b(w, v)}{\|w\|_{\mathbb{X}} \|v\|_{\mathbb{Y}}} \geq c_b, \quad (2.12)$$

for each $v \in \mathbb{Y} \exists w_v \in \mathbb{X}$ such that $b(w_v, v) \neq 0$.

REMARK 2.6. An equivalent set of condition is obtained when replacing the surjectivity condition in the second line by an inf-sup condition with the roles of $\mathbb{X}$ and $\mathbb{Y}$ being interchanged, since this is equivalent to the dual $\mathcal{B}' : \mathbb{Y} \rightarrow \mathbb{X}'$ being an isomorphism.

Here, we emphasize the equivalence of (2.12) to the fact that $\mathcal{B} : \mathbb{X} \rightarrow \mathbb{Y}'$ is an *isomorphism*, which in turn means that

$$c_b \|u - w\|_{\mathbb{X}} \leq \|\mathcal{B}w - F\|_{\mathbb{Y}'} \leq C_b \|u - w\|_{\mathbb{X}}, \quad \forall w \in \mathbb{X}. \quad (2.13)$$

Hence, $\mathcal{B} : \mathbb{X} \rightarrow \mathbb{Y}'$ has a bounded *condition number*

$$\kappa_{\mathbb{X}, \mathbb{Y}}(\mathcal{B}) := \|\mathcal{B}\|_{\mathbb{X} \rightarrow \mathbb{Y}'} \|\mathcal{B}^{-1}\|_{\mathbb{Y}' \rightarrow \mathbb{X}} \leq \frac{C_b}{c_b}, \quad (2.14)$$

so that the smaller the ratio C_b/c_b of continuity and inf-sup constant the tighter is the relation of an error in $\mathbb{X}$ to the residual in $\mathbb{Y}'$. We refer to Section 3.1 for corresponding strategies for identifying suitable pairs of trial and test spaces, including cases where $\frac{C_b}{c_b} = 1$, that is, the variational problem has an ideal condition.

We are now ready to clarify how well-posedness of the fibre problems (2.1) relates to well-posedness of the lifted problem (2.11).

THEOREM 3. If Properties 2.3 and 2.5 hold the operator $\mathcal{B} : \mathbb{X} \rightarrow \mathbb{Y}'$ defined in (2.10), with $\mathbb{X}, \mathbb{Y}$ as in (2.9), is an isomorphism if and only if there exist uniform constants $c_b, C_b > 0$ such that for μ -almost all $p \in P$ the conditions (2.2) hold. In this case, $\mathcal{B}$ satisfies (2.12) with the same constants c_b, C_b .

Proof. First, let (2.2) hold for μ -almost all $p \in P$. This is equivalent to $\mathcal{B}_p : \mathbb{U}_p \rightarrow \mathbb{V}'_p$ defined by $\mathcal{B}_p w = b(w, \cdot; p) \in \mathbb{V}'_p$ being an isomorphism for μ -almost all $p \in P$ with

$$c_b \|w\|_{\mathbb{U}_p} \leq \|\mathcal{B}_p w\|_{\mathbb{V}'_p} = \sup_{0 \neq v \in \mathbb{V}_p} \frac{b(w, v; p)}{\|v\|_{\mathbb{V}_p}} \leq C_b \|w\|_{\mathbb{U}_p} \quad \text{for all } w \in \mathbb{U}_p.$$

Moreover, for $v \in \mathbb{Y}$,

$$b(w, v) \leq \int_P C_b \|w(p)\|_{\mathbb{U}_p} \|v(p)\|_{\mathbb{V}_p} d\mu(p) \leq C_b \|w\|_{\mathbb{X}} \|v\|_{\mathbb{Y}},$$

confirming continuity of $b(\cdot, \cdot)$. Regarding the inf-sup condition we note that by boundedness and measurability of $b(\cdot, \cdot)$, we have for $w \in \mathbb{X}$ that

$$\mathcal{B}w = b(w, \cdot) \in \mathbb{Y}',$$

and thus we have, by Lemma 2,

$$\langle \mathcal{R}\mathcal{B}w, v \rangle = \int_P b(w(p), v(p), p) \, d\mu(p) = \int_P \langle \mathcal{R}_p \mathcal{B}_p w(p), v(p) \rangle_{\mathbb{V}_p} \, d\mu(p).$$

Therefore, $v = \mathcal{R}\mathcal{B}w$ satisfies

$$\begin{aligned} b(w, v) &= \int_P b(w(p), \mathcal{R}_p \mathcal{B}_p w(p); p) \, d\mu(p) = \int_P \|\mathcal{B}_p w(p)\|_{\mathbb{V}_p}^2 \, d\mu(p) \\ &\geq c_b \|w\|_{\mathbb{X}} \left(\int_P \|v(p)\|_{\mathbb{V}_p}^2 \, d\mu(p) \right)^{1/2}. \end{aligned}$$

The third condition in (2.12) then follows in the same manner by Remark 2.6, using that also $\mathcal{B}'_p$ is an isomorphism for μ -almost all $p \in P$.

To prove the converse assume now that there are constants c_b, C_b such that (2.12) holds, but there are no uniform lower and upper bounds for $c_b(p), C_b(p)$ in (2.2). Observe that as $\hat{w}_b = \mathcal{B}w = b(w, \cdot) \in \mathbb{Y}'$ for all $w \in \mathbb{X}$ we have measurability of

$$p \mapsto \|\hat{w}_b(p)\|_{\mathbb{V}_p} = \sup_{v \in \mathbb{V}_p} \frac{b(w(p), v; p)}{\|v\|_{\mathbb{V}_p}} = \|(\mathcal{B}w)(p)\|_{\mathbb{V}'_p}.$$

Consider the set $P_{c_b, w} = \{p \in P : \|\hat{w}_b(p)\|_{\mathbb{V}'_p} < \frac{c_b}{2} \|w(p)\|_{\mathbb{U}_p}\}$. If $\mu(P_{c_b, w}) = 0$ for all $w \in \mathbb{X}$ there is a uniform lower bound. Else take $w \in \mathbb{X}$ such that $\mu(P_{c_b, w}) > 0$. Then, choose

$$v_w = \mathcal{R} \left(\hat{w}_b \mathbf{1}_{P_{c_b, w}} \right) \in \mathbb{Y},$$

where $\mathbf{1}_{P_{c_b, w}}$ is the indicator function of $P_{c_b, w}$. Then, we have, by definition of $P_{c_b, w}$,

$$b \left(w \mathbf{1}_{P_{c_b, w}}, v_w \right) = \|\hat{w}_b \mathbf{1}_{P_{c_b, w}}\|_{\mathbb{Y}'} < \frac{c_b}{2} \|w \mathbf{1}_{P_{c_b, w}}\|_{\mathbb{X}} \|v_w\|_{\mathbb{Y}}.$$

Since $w \in \mathbb{X}$ implies $w \mathbf{1}_{P_{c_b, w}} \in \mathbb{X}$ and

$$\frac{b \left(w \mathbf{1}_{P_{c_b, w}}, v_w \right)}{\|v_w\|_{\mathbb{Y}}} = \sup_{v \in \mathbb{Y}} \frac{b \left(w \mathbf{1}_{P_{c_b, w}}, v \right)}{\|v\|_{\mathbb{Y}}}$$

we arrive at a contradiction to the inf-sup condition (2.12) for $b(\cdot, \cdot)$.

Next, we use Remark 2.6 and note that interchanging the roles of w and v , the same reasoning yields the contradiction

$$\inf_{v \in \mathbb{Y}} \sup_{w \in \mathbb{X}} \frac{b(w, v)}{\|w\|_{\mathbb{X}} \|v\|_{\mathbb{Y}}} < \frac{c_b}{2}.$$

Finally, defining the set

$$Q_{C_b, w} = \left\{ p \in P : \|\hat{w}_b(p)\|_{\mathbb{V}'_p} > 2C_b \|w(p)\|_{\mathbb{U}_p} \right\}$$

we derive in the same fashion, for any $w \in \mathbb{X}$ such that $\mu(Q_{C_b, w}) > 0$, that

$$b\left(w \mathbf{1}_{Q_{C_b, w}}, v_w\right) = b\left(w \mathbf{1}_{Q_{C_b, w}}, \mathcal{R}\left(\hat{w}_b \mathbf{1}_{Q_{C_b, w}}\right)\right) > 2C_b \|w \mathbf{1}_{Q_{C_b, w}}\|_{\mathbb{X}} \|v_w\|_{\mathbb{Y}},$$

finishing the proof. $\square$

Recall that the integration over P is replaced in computations by a weighted summation over discrete subset $\hat{P} \subset P$. Introducing corresponding analogous objects $\hat{b}(\cdot, \cdot)$, $\hat{\mathcal{B}}$, $\hat{F}$ in (2.10), this suggests defining in analogy to (2.9) the *parameter-discrete* direct sum Hilbert spaces

$$\hat{\mathbb{X}} = \ell_2\left(\hat{P}, (\mathbb{U}_p)_{p \in \hat{P}}\right), \quad \hat{\mathbb{Y}} = \ell_2\left(\hat{P}, (\mathbb{V}_p)_{p \in \hat{P}}\right). \quad (2.15)$$

We record the following trivial fact.

REMARK 2.7. When the fibre problems (2.1) satisfy (2.2) uniformly in p , the problem $\hat{b}(u, v) = \hat{F}(v)$ for all $v \in \hat{\mathbb{Y}}$ is well posed for any subset $\hat{P}$ of P as well.

In summary, once (2.12) has been established the continuous and parameter-discrete error–residual relations

$$\|u - w\|_{\mathbb{X}} \approx \|F - \mathcal{B}w\|_{\mathbb{Y}'} \quad \text{and} \quad \|u - w\|_{\hat{\mathbb{X}}} \approx \|F - \mathcal{B}w\|_{\hat{\mathbb{Y}}'}, \quad w \in \mathbb{X}, \quad (2.16)$$

suggest viewing $\|F - \mathcal{B}w\|_{\mathbb{Y}'}$ and $\|F - \mathcal{B}w\|_{\hat{\mathbb{Y}}'}$ as *ideal* loss functions. We proceed discussing next the numerical evaluation and approximation of such ideal loss functions.

2.4 Obstructions and main goal

In a classical (low-dimensional) finite-element context, once (2.12) has been established, one then proceeds solving a discretized version of the linear problem (2.11), typically in terms of a Galerkin or Petrov–Galerkin schemes. The quantity $\|\mathcal{B}w - F\|_{\mathbb{Y}'}$ can be viewed as an ideal *a posteriori* error bound that is efficient as well as reliable. Such a reduction to large linear algebra problems seems less viable when using nonlinear global trial systems. In such scenarios the reformulation of a PDE as an optimization problem becomes relevant.

There are many ways of reformulating PDEs as a variational problems. For elliptic problems *energy minimization* suggests itself because it naturally respects intrinsic model metrics. This works for a restricted class and corresponding loss functions; however, the value of the energy at termination does not reflect the remaining estimation error.

Alternatively for linear problems, the above considerations show that u solves (2.11) if and only if minimizes the mean square loss

$$u = \arg \min_{w \in \mathbb{X}} \|\mathcal{B}w - F\|_{\mathbb{Y}'}^2, \quad (2.17)$$

which is not at all restricted to elliptic problems. The advantage of this choice is that, at any stage of a minimization over a finitely parameterized hypothesis class $\mathcal{H}(\Theta)$, determined by a budget Θ of trainable weights, by (2.13) the size of the $\|\mathcal{B}u^{\mathcal{N}}(\theta) - F\|_{\mathbb{Y}'}$ is indeed uniformly proportional to the error $\|u - u^{\mathcal{N}}(\theta)\|_{\mathbb{X}}$, and hence complies with our quest for variational correctness, see (VC) in Section 1. Moreover, minimizing $\|\mathcal{B}w - F\|_{\mathbb{Y}'}$ over any given finitely parameterized hypothesis class retains quasi-best approximation properties:

$$u_{\mathcal{H}} \in \arg \min_{w \in \mathcal{H}} \|\mathcal{B}w - F\|_{\mathbb{Y}'} \quad \Rightarrow \quad \|u - u_{\mathcal{H}}\|_{\mathbb{X}} \leq \frac{C_b}{c_b} \min_{w \in \mathcal{H}} \|u - w\|_{\mathbb{X}}. \quad (2.18)$$

Of course, there are obvious practical issues with (2.13), and hence with (2.17) and (2.18), in the finite element context as well as in our present setting: unless $\mathbb{Y} = \mathbb{Y}'$ is self-dual (a product of L_2 -spaces) the quantity $\|\mathcal{B}u - F\|_{\mathbb{Y}'}$ involves a supremum over an infinite-dimensional Hilbert space

$$\|\mathcal{B}w - F\|_{\mathbb{Y}'} = \sup_{v \in \mathbb{Y}} \frac{(\mathcal{B}w - F)(v)}{\|v\|_{\mathbb{Y}}} = \sup_{v \in \mathbb{Y}} \frac{b(w, v) - F(v)}{\|v\|_{\mathbb{Y}}}, \quad (2.19)$$

which cannot be evaluated exactly. Hence, (2.17) in the form

$$\theta^* \in \arg \min_{\theta \in \Theta} \max_{v \in \mathbb{Y}} \frac{b(u^{\mathcal{N}}(\theta), v) - F(v)}{\|v\|_{\mathbb{Y}}} \quad (2.20)$$

is practically infeasible. This leads to our **central aim** of devising *computable* expressions that are *uniformly* proportional to $\|\mathcal{B}w - F\|_{\mathbb{Y}'}$, and hence to the error $\|u - w\|_{\mathbb{X}}$.

We can hope to achieve this up to unavoidable *data oscillation errors*, which may occur whenever the data are subjected to a projection to a finitely parameterized subset of $\mathbb{Y}'$.

A first natural strategy for tightly approximating $\|\mathcal{B}w - F\|_{\mathbb{Y}'}$ is to build directly on (2.20) by restricting the maximization to some finite set Ψ of trainable parameters, so that (2.20) becomes

$$\theta^* \in \arg \min_{\theta \in \Theta} \max_{\psi \in \Psi} \frac{b(u^{\mathcal{N}}(\theta), v^{\mathcal{N}}(\psi)) - F(v^{\mathcal{N}}(\psi))}{\|v^{\mathcal{N}}(\psi)\|_{\mathbb{Y}}}. \quad (2.21)$$

In the language of *generative adversarial networks*, the approximation $u^{\mathcal{N}}$ to the solution $u \in \mathbb{X}$ plays the role of a *generative network*, while $v^{\mathcal{N}}$ is the *discriminator* or *adversarial network* that is to make sure $u^{\mathcal{N}}$ obeys the correct optimization criterion.

From a practical perspective it seems natural to alternately perform gradient descent and ascent steps on the quotient in (2.21). This approach has been followed for special cases, namely for elliptic problems and Friedrichs systems in Bao *et al.* (2020), Zang *et al.* (2020), Berrone *et al.* (2022), Chen *et al.* (2023), Taylor *et al.* (2023) and Bertoluzza *et al.* (2024). One senses from the reported results that the success of this strategy depends strongly on the inner maximization problem being solved with sufficient accuracy.

To understand this and to pave the way for alternatives, recall the mapping $\mathcal{R} = \mathcal{R}_{\mathbb{Y}}: \mathbb{Y}' \rightarrow \mathbb{Y}$ defined in (2.3), which for $\ell \in \mathbb{Y}'$ is given by

$$(\mathcal{R}\ell, v)_{\mathbb{Y}} = \ell(v), \quad v \in \mathbb{Y}.$$

Thus, the *Riesz lift* $\mathcal{R}\ell \in \mathbb{Y}$ of $\ell \in \mathbb{Y}'$ is the solution of an *elliptic* variational problem in the test space $\mathbb{Y}$, regardless of the type of the primal problem (1.1). As an immediate consequence of (2.3)

$$\|\mathcal{R}(\mathcal{B}w - F)\|_{\mathbb{Y}} = \arg \max_{v \in \mathbb{Y}} \frac{(\mathcal{B}w - F)(v)}{\|v\|_{\mathbb{Y}}}.$$

Moreover, one readily checks the relation

$$\|\mathcal{R}\ell\|_{\mathbb{Y}}^2 = \ell(\mathcal{R}\ell) = \|\ell\|_{\mathbb{Y}'}^2, \quad \ell \in \mathbb{Y}'. \quad (2.22)$$

REMARK 2.8. Thus, assuming that the $\mathbb{Y}$ -norm can be evaluated directly finding an expression that is uniformly proportional to $\|\mathcal{B}w - F\|_{\mathbb{Y}'}$ amounts to evaluating the Riesz lift of $\mathcal{B}w - F$ up to a *uniform relative error*.

Specifically, whenever an approximation $v_w \in \mathbb{Y}$ to $\mathcal{R}(\mathcal{B}w - F)$ satisfies

$$\|\mathcal{R}(\mathcal{B}w - F) - v_w\|_{\mathbb{Y}} \leq a \|\mathcal{B}w - F\|_{\mathbb{Y}'} \quad \text{for some } a < 1 \quad (2.23)$$

one readily derives from (2.22) that

$$(1 - a) \|\mathcal{B}w - F\|_{\mathbb{Y}'}^2 \leq (\mathcal{B}w - F)(v_w) \leq \|\mathcal{B}w - F\|_{\mathbb{Y}'}^2, \quad (2.24)$$

so that $(\mathcal{B}w - F)(v_w)$ is (up to quadrature) a computable variationally correct loss. This shows though that, with an improved accuracy of the generative network, the demands on the adversarial network increase as it has to become a better and better approximation to the Riesz lift of the residual. This explains, in particular, why working with fixed budgets for the generative and adversarial networks is problematic especially in view of an uncertain optimization success.

3. Computable variationally correct loss functions

3.1 Conceptual prerequisites

Given an operator equation (2.8) in strong form there usually exist several different pairs of trial and test spaces $\mathbb{U}, \mathbb{V}$ that give rise to a stable variational formulation. In this section we discuss how to exploit this freedom to devise computable variationally correct loss functions. In what follows, we use the notation L_2 to denote an L_2 -type space by which mean in general a product of L_2 -spaces. Specifically, we present two scenarios that can be viewed as extreme points in a spectrum of possibilities. To explain this we suppress first any dependence on parameters p .

We emphasize that both scenarios refer to the (linear) PDE $\mathcal{B}u = f$ in the form of a *first-order system* where $\mathcal{B} : \text{dom } \mathcal{B} \rightarrow L_2$ is a closed operator. If the PDE is initially given as a higher order PDE the first step in both cases is to rewrite it as a first-order system (for simplicity again denoted by $\mathcal{B} = f$), see the examples in subsequent sections. For simplicity of exposition we assume that arising essential boundary conditions are homogeneous.

Recall further that the graph space of $\mathcal{B}$ is defined by $\mathbb{G}(\mathcal{B}) := \{w \in L_2 : \mathcal{B}w \in L_2\}$ and $\|w\|_{\mathbb{G}(\mathcal{B})}^2 := \|w\|_{L_2}^2 + \|\mathcal{B}w\|_{L_2}^2$. Moreover, let $\mathbb{G}(\mathcal{B})_0$ denote the (closed) subspace of $\mathbb{G}(\mathcal{B})$ with built-in essential boundary conditions.

We consider then two choices of the trial space $\mathbb{U}$,

$$\mathbb{U} = \mathbb{G}(\mathcal{B})_0 \quad (\text{strong formulation}) \quad (\text{S})$$

$$\mathbb{U} = L_2 \quad (\text{ultra-weak formulation}) \quad (\text{UW})$$

and then look for respective test spaces $\mathbb{V}$ that give rise to a stable variational formulation.

As explained in more detail later, (S) aims at having an L_2 -least squares functional as loss function, which is therefore easily evaluated (up to quadrature errors). In fact, residuals belong to L_2 because $\mathcal{B}$ maps $\mathbb{U}$, by definition of the graph space, boundedly into L_2 .

Instead, (UW) accepts an idealized loss function in terms of a nontrivial dual norm, but, as explained below, offers advantages in case of an essential parameter dependence of the trial spaces $\mathbb{U}_p$ in (2.9). Moreover, in combination with DPG concepts it still allows one to devise computable explicitly given variationally correct loss functions.

We pause to point out how the underlying seemingly very different scenarios are interrelated.

PROPOSITION 3.1. Let $\mathbb{U} \subseteq \mathbb{G}(\mathcal{B})_0$ be a closed subspace. Then, the mapping $\mathcal{B} : \mathbb{U} \rightarrow L_2$ is bijective if and only if for every $f \in L_2$ there exists a $u \in \mathbb{U}$ such that

$$b(u, v) := (\mathcal{B}u)(v) = f(v) \quad \forall v \in L_2, \quad (3.1)$$

and

$$\|\mathcal{B}w - f\|_{L_2} \approx \|w - u\|_{\mathbb{U}}, \quad w \in \mathbb{U}. \quad (3.2)$$

In this case, (3.1) is a stable variational formulation and $\mathbb{U} = \mathbb{G}(\mathcal{B})_0$.

Proof. First, let $\mathcal{B} : \mathbb{U} \rightarrow L_2$ be bijective. Then, by definition of the graph norm, $\|\mathcal{B}w\|_{L_2} \leq \|w\|_{\mathbb{G}(\mathcal{B})}$, $w \in \mathbb{U}$. Since $\mathbb{U} \subseteq \mathbb{G}(\mathcal{B})_0$ one has $\mathcal{B}(\mathbb{U}) \subset L_2$. Since $\mathcal{B}$ maps $\mathbb{U}$ onto L_2 boundedness of $\mathcal{B}^{-1}$ follows from the open mapping theorem. Moreover, $\mathbb{U} = \mathbb{G}(\mathcal{B})_0$, since otherwise a nontrivial orthogonal complement would contradict injectivity. The Babuška–Nečas Theorem ensures then that the bilinear form $(\mathcal{B}w)(v) : \mathbb{U} \times L_2 \rightarrow \mathbb{R}$ yields stable variational formulation, which implies (3.1) and (3.2). The converse is a trivial consequence of the Babuška–Nečas Theorem. $\square$

REMARK 3.2. When using the strong formulation (S) under the hypotheses in Proposition 3.1, a Monte Carlo approximation of the residual in an L_2 -space, as in standard versions of PINN, does not incur any variational crime.

We proceed with a few comments on scenario (UW) concerning ultra-weak formulations.

PROPOSITION 3.3. Let $\mathcal{B}^*$ denote the formal adjoint of $\mathcal{B}$, i.e., for smooth functions ϕ with compact support in the computational domain Ω one has $\langle \mathcal{B}w, \phi \rangle_{L_2} = \langle w, \mathcal{B}^*\phi \rangle_{L_2}$. Let

$$\mathbb{V} := \text{clos}_{\|\mathcal{B}^*\cdot\|_{L_2}} \left\{ \phi \in \mathbb{G}(\mathcal{B}^*) : \langle \mathcal{B}w, \phi \rangle_{L_2} - \langle w, \mathcal{B}^*\phi \rangle_{L_2} = 0 \quad \forall w \in \mathbb{U} = \mathbb{G}(\mathcal{B})_0 \right\}. \quad (3.3)$$

The following statements are equivalent:

- (i) $\mathcal{B}^* : \mathbb{V} \rightarrow L_2$ is bijective.

(ii) $\|v\|_{\mathbb{V}} := \|\mathcal{B}^*v\|_{L_2}$ is a norm on $\mathbb{V}$ and the *ultra-weak* variational formulation

$$b(u, v) := (\mathcal{B}u)(v) = f(v), \quad \forall v \in \mathbb{V},$$

is stable with constants $c_b = C_b = 1$.

Proof. If (i) holds $\|\cdot\|_{\mathbb{V}}$ is a norm. Obviously, the bilinear form $b(\cdot, \cdot) : L_2 \times \mathbb{V} \rightarrow \mathbb{R}$ satisfies $|b(u, v)| = |(\mathcal{B}u)(v)| = |(\mathcal{B}^*v)(u)| \leq \|\mathcal{B}^*v\|_{L_2} \|u\|_{L_2}$, which means that $C_b \leq 1$. Moreover,

$$\inf_{v \in \mathbb{V}} \sup_{w \in L_2} \frac{b(w, v)}{\|v\|_{\mathbb{V}} \|w\|_{L_2}} = \inf_{v \in \mathbb{V}} \frac{\|\mathcal{B}^*v\|_{L_2}}{\|\mathcal{B}^*v\|_{L_2}} = 1$$

and for each $w \in L_2$ there exists a $v_w \in \mathbb{V}$ such that $b(w, v_w) = (\mathcal{B}^*v_w)(w) \neq 0$ because $\mathcal{B}^*$ is surjective. This confirms the claim. The converse follows directly from the Babuška–Nečas Theorem. $\square$

The key in both scenarios is to establish bijectivity of a linear operator as a mapping to L_2 , keeping in mind that $\mathcal{B}$ and $\mathcal{B}^*$ are of the same ‘type’, so the difficulties in both cases are very similar. Moreover, as a consequence of Proposition 3.3(ii), $\mathcal{B}$ is an isomorphism from L_2 to $\mathbb{V}'$. Hence, its dual $\mathcal{B}'$ (by self-duality of L_2) is an isomorphism from $\mathbb{V}'$ to L_2 , which agrees with $\mathcal{B}^*$ on dense smooth subsets. So, roughly, the dual of $\mathcal{B}^*$ agrees with $\mathcal{B}$ as a mapping from L_2 to $\mathbb{V}'$ (which could be viewed as a continuous extension of the operator $\mathcal{B} : \mathbb{U} \rightarrow L_2$, defined in scenario (S)).

Returning to *parameter dependence*, whenever Proposition 3.3(ii) applies, when employing the optimal test norm, the stability is *uniform* in p , warranting an optimal condition of the parametric problems as well as of the lifted one (2.11), see Theorem 3.

In scenario (S), *uniform* boundedness of $\mathcal{B} = \mathcal{B}_p$ with respect $p \in P$ is ensured by definition. If, in addition, the graph spaces $\mathbb{G}(\mathcal{B}_p)$ agree with equivalent norms for *all* p we conclude that (2.2) is valid. By Theorem 3, this implies well-posedness of the lifted problem (2.11). One then still ends up with a computable L_2 -least squares functional as variationally correct loss function.

If on the other hand, even as sets the graph spaces $\mathbb{G}(\mathcal{B}_p)$ vary with p , we say that $\mathbb{G}(\mathcal{B}_p)$ depends *essentially* on $p \in P$. Thus, the membership of approximate solutions of the fibre problems to the respective parameter-dependent trial spaces is not robust. We see this as a reason to opt for scenario (UW), based on *ultra-weak formulations*. How to proceed in this case can be summarized as follows.

REMARK 3.4. For a wide variety of PDE models one can proceed as follows to obtain suitable test spaces $\mathbb{V}_p$ for the choice $\mathbb{U}_p$: If (2.8) is not already given in this form rewrite it as a *first-order system* of PDEs (introducing auxiliary variables if needed), again denoted for simplicity by $\mathcal{B}u = f$. We assume homogeneous essential boundary conditions, so that $\mathcal{B}$ is indeed linear. With the norm

$$\|v\|_{\mathbb{V}_p} := \|\mathcal{B}_p^*v\|_{L_2}, \quad (3.4)$$

define the spaces $\mathbb{V}_p$ for each parameter p as closures with respect to this norm according to (3.3). As a crucial step one then needs to show that for each $p \in P$ the adjoint $\mathcal{B}_p^*$ is a bijective mapping from $\mathbb{V}_p$ onto L_2 .

Once this has been confirmed one defines $\mathbb{X} := L_2(P \times \Omega)$, $\mathbb{Y} := L_2(P, (\mathbb{V}_p)_{p \in P})$ and $b : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$ as in (2.10), where by integration by parts, all derivatives are applied to the test function $v \in \mathbb{V}_p$. For the

norm of $\mathbb{Y}$, we thus obtain

$$\|v\|_{\mathbb{Y}} = \|\mathcal{B}^*v\|_{L_2(P \times \Omega)}. \quad (3.5)$$

Given $F \in \mathbb{Y}'$, finding $u \in \mathbb{X} := L_2(P \times \Omega)$ satisfying the *ultra-weak* formulation

$$b(u, v) = F(v), \quad v \in \mathbb{Y}, \quad (3.6)$$

is then a well-posed problem in the sense of (2.12).

Once Proposition 3.3(i) has been verified, Proposition 3.3(ii) applies and says that the fibre problems have condition number *equal to one*, uniformly in $p \in P$, that is,

$$\|u - w\|_{\mathbb{X}} = \|\mathcal{F}(w)\|_{\mathbb{Y}'}. \quad (3.7)$$

We therefore call the right-hand side an ‘ideal residual loss’ because the respective norms of errors and residuals agree, and the variational problem has been *optimally preconditioned* already on the continuous level. Note also that in such ultra-weak formulations of the form (3.6), essential boundary conditions can no longer be imposed on elements in L_2 . In fact, essential boundary conditions become natural ones.

On the other hand, one still faces the problem that the resulting ideal residual loss function $\|\mathcal{B}u - F\|_{\mathbb{Y}'}^2$ involves a nontrivial dual norm.

What we have gained, though, is that the ultra-weak formulation (UW) is perhaps the most convenient starting point for deriving stable DPG formulations. For the convenience of the reader we briefly recall in Section 3.3.2 the relevant DPG concepts. In combination with the hybrid representation format as in Section 1.6, which is discussed further in Section 4.1, DPG formulations will be shown below to lead to computable variationally correct loss functions.

Stable DPG formulations have by now been established for a wide range of PDE models including Maxwell’s equations, dispersive models or singularly perturbed problems, primarily in the context of DPG methods based on well-posed variational formulations of the above type; see, for example, Cohen *et al.* (2012), Dahmen *et al.* (2012), Carstensen *et al.* (2016), Demkowicz *et al.* (2017), Broersen *et al.* (2018), Dahmen *et al.* (2019) and Gopalakrishnan & Sepúlveda (2019). We proceed discussing examples for both strong and ultra-weak formulations.

3.2 Strong formulation: first-order system least-squares methods in L_2

We consider two different model problems where $\mathbb{U}$ can be chosen as a parameter-independent graph space as in (S).

3.2.1 An elliptic model problem. Although, in what follows, one could include lower-order terms, we focus for simplicity of exposition on the classical Poisson equation (without reaction or convection term) with homogeneous Dirichlet conditions

$$-\operatorname{div}(a(p)\nabla u) = f \quad \text{on } \Omega \quad \text{and} \quad u|_{\partial\Omega} = 0, \quad (3.8)$$

where we consider parameter-dependent diffusion coefficients satisfying the *uniform ellipticity conditions* (UE):

$$\exists r, R, \quad 0 < r \leq R < \infty \quad \text{such that } r \leq a(x, p) \leq R, \quad (x, p) \in \Omega \times P. \quad (3.9)$$

We suppress in what follows the dependence on the parameter p where it does not matter. We say that $u \in H_0^1(\Omega)$ is a solution of (3.8) if

$$\langle a \nabla u, \nabla v \rangle_{L_2(\Omega)} = f(v), \quad v \in H_0^1(\Omega), \quad (3.10)$$

which is known to be stably solvable for any $f \in H^{-1}(\Omega) := (H_0^1(\Omega))'$. Next, recall that

$$H(\operatorname{div}, \Omega) := \{\mathbf{w} \in L_2(\Omega, \mathbb{R}^d) : \operatorname{div} \mathbf{w} \in L_2(\Omega)\}, \quad \|\mathbf{w}\|_{H(\operatorname{div}, \Omega)}^2 := \|\mathbf{w}\|_{L_2(\Omega; \mathbb{R}^d)}^2 + \|\operatorname{div} \mathbf{w}\|_{L_2(\Omega)}^2.$$

Also recall that every $f \in H^{-1}(\Omega)$ can be written as $f = f_2 + \operatorname{div} \mathbf{f}_1$, for some $f_2 \in L_2(\Omega)$, $\mathbf{f}_1 \in L_2(\Omega; \mathbb{R}^d)$. Hence, if $u \in H_0^1(\Omega)$ is the solution of (3.10) for this f , one has

$$\boldsymbol{\sigma} := a \nabla u + \mathbf{f}_1 \in L_2(\Omega) \quad \text{solves} \quad -\operatorname{div} \boldsymbol{\sigma} = f_2. \quad (3.11)$$

In other words, defining

$$\mathcal{B}_p : \mathbf{u} = [\boldsymbol{\sigma}, u] \mapsto \mathcal{B}_p \mathbf{u} = \begin{pmatrix} \operatorname{id} & -a(p) \nabla \\ -\operatorname{div} & 0 \end{pmatrix} \begin{pmatrix} \boldsymbol{\sigma} \\ u \end{pmatrix} = \begin{pmatrix} \boldsymbol{\sigma} - a(p) \nabla u \\ -\operatorname{div} \boldsymbol{\sigma} \end{pmatrix}, \quad (3.12)$$

$\mathbf{u} := [\nabla u, u]$ solves $\mathcal{B}_p \mathbf{u} = \mathbf{f} = (\mathbf{f}_1, f_2)$. Note that $\nabla u \in H(\operatorname{div}; \Omega)$ and $u \in H_0^1(\Omega)$. Thus,

$$\mathbb{U} := H(\operatorname{div}; \Omega) \times H_0^1(\Omega) \quad (3.13)$$

is a Hilbert space contained in $\mathbb{G}(\mathcal{B}_p)_0$, and the previous observation shows that $\mathcal{B}_p : \mathbb{U} \rightarrow L_2$ is surjective. As shown in Cai *et al.* (1994) the symmetric bilinear form

$$B(\mathbf{u}, \mathbf{w}; p) := \langle \mathcal{B}_p \mathbf{u}, \mathcal{B}_p \mathbf{w} \rangle_{L_2(\Omega)} : \mathbb{U} \times \mathbb{U} \rightarrow \mathbb{R} \quad (3.14)$$

is $\mathbb{U}$ -elliptic, that is,

$$\|\mathbf{w}\|_{\mathbb{G}(\mathcal{B}_p)_0}^2 = B(\mathbf{w}, \mathbf{w}; p) \approx \|\mathbf{w}\|_{\mathbb{U}}^2, \quad \mathbf{w} \in \mathbb{U}, \quad \text{uniformly in } p \text{ obeying (3.9)}. \quad (3.15)$$

This implies, in particular, injectivity, and hence bijectivity of $\mathcal{B}_p : \mathbb{U} \rightarrow L_2$.

Hence, Proposition 3.1 applies and says that $\mathbb{U} = \mathbb{G}(\mathcal{B}_p)_0$ and $\mathcal{B}_p : \mathbb{G}(\mathcal{B}_p)_0 \rightarrow L_2$, defined by (3.12), is bounded and boundedly invertible. This, in turn, means that

$$u(p) = \arg \min_{\mathbf{w} \in \mathbb{U}} \|\mathcal{B}_p \mathbf{w} - \mathbf{f}\|_{L_2}^2 \quad (3.16)$$

is the unique solution of $\mathcal{B}_p u(p) = f$ and, on account of (3.15),

$$\|\mathcal{B}_p w - f\|_{L_2} \approx \|u - w\|_{\mathbb{U}}$$

holds uniformly in p .

The formulation (3.16) can be regarded as a particular instance of the framework of Section 2, with $\mathbb{U}_p = \mathbb{U}$ as in (3.13), $\mathbb{V}_p = \mathbb{V} = L_2$ and the p -dependent bilinear form

$$b(u, v; p) := (\mathcal{B}_p u)(v) = \langle \sigma - a(p) \nabla u, \tau \rangle_{L_2(\Omega)} - \langle \operatorname{div} \sigma, v \rangle_{L_2(\Omega)}, \quad u \in \mathbb{U}, v = [\tau, v] \in L_2. \quad (3.17)$$

In fact, $u(p)$ is characterized by the normal equations $B(u, w; p) = \langle f, \mathcal{B}_p w \rangle_{L_2(\Omega)}$, $w \in \mathbb{U}$. Again, since $\mathcal{B}_p : \mathbb{U} \rightarrow L_2$ is an isomorphism, this is equivalent to

$$b(u, v; p) = B\left(u, \mathcal{B}_p^{-1} v; p\right) = \langle f, v \rangle_{L_2(\Omega)} = f(v),$$

posed over $\mathbb{U} \times L_2$.

REMARK 3.5. By (3.15) and (3.9) the above well-posedness is uniform in p , so that the above trial and test spaces $\mathbb{U} = H(\operatorname{div}; \Omega) \times H_0^1(\Omega)$ and $\mathbb{V} = L_2$ are indeed independent of p . Thus, Theorem 3 applies (see Remark 2.2) and we can readily obtain from (3.17) a weak formulation with bilinear form $b(\cdot, \cdot) : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$ with $\mathbb{X} = L_2(P, \mathbb{U})$, $\mathbb{Y}' = L_2(P, L_2) = L_2(P) \times L_2$, where $L_2 = L_2(\Omega, \mathbb{R}^d) \times L_2(\Omega)$.

REMARK 3.6. We emphasize that the formulations (3.10) and (3.16) are equivalent in the following sense: As noted above, if u solves (3.10) then $u = [a \nabla u + f_1, u]$ solves $\mathcal{B}_p u = f$. Conversely, when $u = [\sigma, u]$ solves $\mathcal{B}_p u = f$ then u solves (3.10) with $\sigma = a \nabla u + f_1$.

3.2.2 A parabolic model problem. We close the discussion of scenario (S) by applying the above principles to a time-dependent problem, leading to space-time variational formulations. We recall the following standard well-posed space-time variational formulation for parabolic initial boundary value problems: for $I = (0, T)$ find

$$u \in L_2\left(I, H_0^1(\Omega)\right) \cap H^1\left(I, H^{-1}(\Omega)\right)$$

such that

$$\int_{I \times \Omega} \partial_t u v + a \nabla_x u \cdot \nabla_x v \, dt \, dx = \int_{I \times \Omega} f v \, dt \, dx, \quad \forall v \in L_2(I, H_0^1(\Omega)), \quad (3.18)$$

where $f \in L_2(I, H^{-1}(\Omega))$ and a may depend also on parameters p such that (3.9) holds. It is well known that this (asymmetric) variational formulation is stable, which remains true when the roles of trial and test space are interchanged, with a terminal condition at T on the test space, see, e.g., Chegini & Stevenson (2011).

Here, both trial and test space involve nontrivial dual norms, which motivates interest in computationally more friendly formulations. A formulation similar to (3.12) for the elliptic case has been

introduced in [Führer & Karkulik \(2021\)](#). Writing $f \in L_2(I, H^{-1}(\Omega))$ again as $f = f_2 + \operatorname{div}_x \mathbf{f}_1$, with $\mathbf{f}_1 \in L_2(I \times \Omega; \mathbb{R}^d)$, consider the first-order system

$$\mathcal{B}u := \begin{pmatrix} -\sigma - a \nabla_x u \\ \partial_t u + \operatorname{div}_x \sigma \\ u(0, \cdot) \end{pmatrix} = \begin{pmatrix} \mathbf{f}_1 \\ f_2 \\ u_0 \end{pmatrix}, \quad u_2|_{I \times \partial\Omega} = 0. \quad (3.19)$$

In this case, as shown in [Gantner & Stevenson \(2021\)](#), the graph space is isomorphic to

$$\mathbb{U} = \left(L_2(I \times \Omega, \mathbb{R}^d) \times L_2(I, H_0^1(\Omega)) \right) \cap H(\operatorname{div}; I \times \Omega), \quad (3.20)$$

where $H(\operatorname{div}; I \times \Omega)$ is defined in terms of the space-time divergence $\operatorname{div}[\sigma, u] = \operatorname{div}_x \sigma + \partial_t u$. In particular, no explicit treatment of dual spaces is required in this formulation. As shown in [Gantner & Stevenson \(2021\)](#) the mapping $\mathcal{B}$ is boundedly invertible from $\mathbb{U}$ to the L_2 -space

$$\mathbb{V} = L_2(I \times \Omega; \mathbb{R}^d) \times L_2(I \times \Omega) \times L_2(\Omega),$$

so that one is in the same situation as in the elliptic case, that is, for any $F = (\mathbf{f}_1, f_2, u_0) \in \mathbb{V}$ there exists a unique $u \in \mathbb{U}$ such that $\mathcal{B}u = F$ and

$$u = \arg \min_{\mathbf{w} \in \mathbb{U}} \|\mathcal{B}\mathbf{w} - F\|_{\mathbb{V}}^2, \quad \|u - \mathbf{w}\|_{\mathbb{U}} \approx \|F - \mathcal{B}\mathbf{w}\|_{\mathbb{V}}, \quad \mathbf{w} \in \mathbb{U}.$$

The remarks on the elliptic case in the previous section apply analogously to the current situation as well.

3.3 Neural DPG methods for ultra-weak formulations (UW)

3.3.1 Linear transport equations. Let $\Omega \subset \mathbb{R}^d$ be a domain and consider

$$\mathcal{B}_p u(x, p) := \mathbf{b}(x, p) \nabla u(x) + c(x, p) u(x) = f(x, p) \quad \text{in } \Omega, \quad u|_{\Gamma_{\text{in}}(p)} = g, \quad (3.21)$$

where $\Gamma_{\text{in}}(p) = \operatorname{clos}\{x \in \partial\Omega : \mathbf{n}(x) \cdot \mathbf{b}(x, p) < 0\}$ denotes the *inflow boundary* of Ω . Here, we assume that $\partial\Omega$ is smooth enough so that the outer normal $\mathbf{n}(x)$ exists a.e. on $\partial\Omega$. Analogously, we define the *outflow boundary*

$$\Gamma_{\text{out}}(p) = \operatorname{clos}\{x \in \partial\Omega : \mathbf{n}(x) \cdot \mathbf{b}(x, p) > 0\},$$

and furthermore the *characteristic boundary* $\Gamma_0(p) = \partial\Omega \setminus (\Gamma_{\text{in}}(p) \cup \Gamma_{\text{out}}(p))$. Note that the formulation (3.21) includes the case of instationary problems, where one simply views the first component of x as the time variable and takes the first component $\mathbf{b}_1(x, p)$ to be equal to one. We postpone specifying conditions on the convection field and first put this example into context.

Although the concepts used in what follows apply in much larger generality, we discuss this specific example for several reasons. First, it is of interest, for instance, as a core constituent of more involved kinetic models. Second, despite its ‘analytical simplicity’, it exhibits several difficulties when dealing with parameter-dependent convection fields. In fact, when Ω is a polyhedral domain the inflow boundary

may flip abruptly when the convection direction passes through specific parameter instances. This entails an intrinsic discontinuity with respect to the parameter dependence.

To that end, since $\mathcal{B}_p$ is already of order one, one could adopt scenario (S), and look for solutions in the graph space of $\mathcal{B}_p$ comprising those functions w in $L_2(\Omega)$ for which $\mathcal{B}_p w$ is also in $L_2(\Omega)$. For a function w to belong to the graph space, this requires directional derivatives of w along characteristics to belong to $L_2(\Omega)$, which requires continuity only along the characteristics. Thus, even under small perturbations of the convection field, the graph space changes, not only with respect to the norm, but already as a set. Therefore, we opt for ultra-weak formulations according to scenario (UW) and seek solutions just in $L_2(\Omega)$, which is indifferent under varying convection.

As shown below the fibre test spaces then depend on p in an essential way, which is more tolerant to perturbations, see Broersen *et al.* (2018). Thus, ‘lifting’ the family of fibre problems (3.21) to a single variational problem over a pair of classical Bochner spaces would not work, and one needs to resort to the type of direct integrals defined in (2.9), in particular, to arrive at well-defined dual norms underlying an ideal loss function.

To exhibit appropriate weak formulations of the fibre problems (3.21), we briefly sketch the relevant arguments from Dahmen *et al.* (2012) and Broersen *et al.* (2018). We consider first a fixed parameter p and recall for convenience some relevant facts from Dahmen *et al.* (2012) and Broersen *et al.* (2018). We follow Proposition 3.3 to determine $\mathcal{B}_p^*$ and the test space $\mathbb{V}_p$. To that end, integration by parts yields

$$\begin{aligned} 0 &= \int_{\Omega} -u(x) [\operatorname{div}(\mathbf{b}(x, p)v)(x) - c(x, p)v(x)] \, dx \\ &\quad + \int_{\partial\Omega} \mathbf{n}(x) \cdot \mathbf{b}(x, p)u(x)v(x) \, d\Gamma(x) - \int_{\Omega} f(x)v(x) \, dx. \end{aligned} \quad (3.22)$$

Note that when $u \in L_2(\Omega)$ traces are not defined. But when inflow boundary conditions g are imposed (3.22) can be rewritten as

$$\begin{aligned} 0 &= - \int_{\Omega} u(x) [\operatorname{div}(\mathbf{b}(x, p)v)(x) - c(x, p)v(x)] \, dx \\ &\quad - \int_{\Gamma_{\text{in}}(p)} |\mathbf{n}(x) \cdot \mathbf{b}(x, p)|g(x, p)v(x) \, d\Gamma(x) - \int_{\Omega} f(x)v(x) \, dx \\ &=: b(u, v; p) - \ell(v; p), \end{aligned} \quad (3.23)$$

where

$$\ell(v, p) := \int_{\Gamma_{\text{in}}(p)} |\mathbf{n}(x) \cdot \mathbf{b}(x, p)|g(x, p)v(x) \, d\Gamma(x) + \int_{\Omega} f(x)v(x) \, dx, \quad (3.24)$$

provided that the test function v vanishes on the outflow boundary $\Gamma_{\text{out}}(p)$. Defining further

$$\mathbb{V} = \mathbb{V}_p := \operatorname{clos}_{\|\cdot\|_{\mathbb{V}_p}} \left\{ v \in C^1(\overline{\Omega}) : v|_{\Gamma_{\text{out}}(p)} = 0 \right\}, \quad (3.25)$$

where

$$\|v\|_{\mathbb{V}_p} := \|\operatorname{div}(\mathbf{b}(\cdot, p)v) - c(\cdot, p)v\|_{L_2(\Omega)}, \quad (3.26)$$

follows exactly the recipe (3.3) in Proposition 3.3 with the norm given in Proposition 3.3(ii).

For the linear functional ℓ to belong to $\mathbb{V}_p$, we need that $g, v|_{\Gamma_{\text{in}}(p)}$ belong the weighted L_2 space $L_2(|\mathbf{b}(\cdot, p) \cdot \mathbf{n}(\cdot)|, \Gamma_{\text{in}}(p))$. It is well known that the elements of $\mathbb{V}_p$ indeed have a trace in this space (see the discussion in Dahmen *et al.* (2012)). Note that essential inflow boundary conditions in the strong formulation have become natural ones appearing as part of the right-hand side.

Moreover, it has also been shown in Broersen *et al.* (2018) that $\mathcal{B}_p^* = \operatorname{div}(\mathbf{b}(\cdot, p)v) - c(\cdot, p)v$ is indeed an isomorphism as a mapping from $\mathbb{V}_p$ onto $L_2(\Omega)$ under mild assumptions on the convection field $\mathbf{b}(\cdot, p)$, which we assume to be valid in what follows.

It is difficult to characterize well-posedness of (3.23) in terms of concise properties of the convection field and one usually has to be content with sufficient conditions, see, e.g., related discussions in Dahmen *et al.* (2012) and Broersen *et al.* (2018). Throughout the remainder of this section we work under the following assumption that guarantees, in particular, well-posedness of (3.23) for each $p \in P$.

ASSUMPTION 3.7. We assume for almost all $p \in P$ that $\mathbf{b}(\cdot; p)$ is Lipschitz on $\bar{\Omega}$ and $|\mathbf{b}(\cdot; p)| \geq \beta > 0$ for some fixed β , that is, the convection is not allowed to degenerate. Moreover, we either require for each $p \in P$ that $\mathbf{b}(\cdot; p)$ is in $C^1(\Omega; \mathbb{R}^d)$ or that there exists a positive constant κ such that $c - \frac{1}{2}\operatorname{div} \mathbf{b}(\cdot; p) \geq \kappa$. Finally, we assume that $\mathbf{b}$ is jointly measurable in x and p .

These conditions ensure, in particular, that for $v \in \mathbb{V}_p$ one has $\|v\|_{L_2} \leq C\|\mathcal{B}_p^*v\|_{L_2}$, which means that $\mathcal{B}_p^*$ is a closed operator. Since the structure of $\mathcal{B}_p$ is the same, $\mathcal{B}_p$ is closed for each $p \in P$ as well. The main consequences can be summarized as follows.

REMARK 3.8. Employing the test-norms (3.26) the family of fibre problems are well posed for the pair $\mathbb{U}_p = L_2(\Omega), \mathbb{V}_p$, given by (3.25), for each $p \in P$, i.e., (2.2) is valid with $c_b = C_b = 1$, and thus uniformly in p .

REMARK 3.9. Under Assumption 3.7 on $\mathbf{b}(\cdot, p)$ Property 2.5 is satisfied.

To be able to formulate the corresponding lifted problem and to invoke Theorem 3 (and take advantage of the tight error-residual relation (3.7)), the following claim is crucial.

THEOREM 4. Under Assumption 3.7 the direct integral $\mathbb{Y} = L_2(P, (\mathbb{V}_p)_{p \in P})$ is a well-defined Hilbert space with dual $\mathbb{Y}' = L_2(P, (\mathbb{V}'_p)_{p \in P})$.

With this result at hand the concepts presented in Section 2 can be applied.

In view of Theorem 1 the claim in Theorem 4 is an immediate consequence of the following Lemma.

LEMMA 3.10. There is a fundamental μ -measurable sequence $\prod_{p \in P} \mathbb{V}_p$ such that the corresponding direct integral $\mathbb{Y} = L_2(P, (\mathbb{V}_p)_{p \in P})$ satisfies

$$L_2\left(P, (\mathbb{V}_p)_{p \in P}\right) = \operatorname{clos}_{\|\cdot\|_{\mathbb{Y}}} \left\{ v \in L^2(P \times \Omega) : v(p, \cdot) \in C^1(\Omega), v(p, \cdot)|_{\Gamma_{\text{out}}(p)} = 0 \text{ for a.e. } p \in P \right\}$$

with

$$\|v\|_{\mathbb{Y}}^2 = \int_P \|v(p, \cdot)\|_{\mathbb{V}_p}^2 d\mu(p).$$

We defer the proof of this lemma to Appendix A. Here, we briefly comment on the main steps. The first auxiliary ingredient is the notion of graph space, which we briefly recall

$$\mathbb{G}_p^* := \mathbb{G}(\mathcal{B}_p^*) = \left\{ v \in L_2(\Omega) : \|v\|_{L_2(\Omega)}^2 + \|\mathcal{B}_p^* v\|_{L_2(\Omega)}^2 := \|v\|_{\mathbb{G}(\mathcal{B}_p^*)}^2 < \infty \right\}.$$

By our previous comments $\mathcal{B}_p^*$ is closed, and hence $\mathbb{G}_p^*$ is a Banach space.

An important observation is that the Sobolev space $H^1(\Omega)$ is dense in every $\mathbb{G}_p^*$ for each $p \in P$. A fundamental system for $\prod_{p \in P} \mathbb{V}_p$ can then be constructed by first taking (any) fundamental system in $H^1(\Omega)$ and multiplying each element in this system by a judiciously chosen p -dependent *clipping function* whose properties allow one to show that the resulting system is fundamental in $\prod_{p \in P} \mathbb{V}_p$.

3.3.2 DPG discretization. For the convenience of the reader we briefly recall features of DPG methodology that matter for the present purposes. DPG methods are relevant when well-posedness necessitates $\mathbb{U} \neq \mathbb{V} \neq L_2$. It is derived from a *well-posed* variational formulation

$$\text{find } u \in \mathbb{U} \text{ such that } \quad b(u, v) = f(v), \quad v \in \mathbb{V}. \quad (3.27)$$

For the moment we suppress for convenience dependence on any parameters $p \in P$. Although not necessary, DPG methods are most conveniently explained when (3.27) is an ultra-weak formulation, which is henceforth assumed, see scenario (UW).

Given a (shape regular) partition or mesh $\mathcal{T}$ of the spatial domain Ω , the first step is to derive from (3.27) an additional *mesh-dependent* well-posed variational formulation—first on the continuous level

$$\text{find } u \in \mathbb{U}_{\mathcal{T}} \text{ such that } \quad b_{\mathcal{T}}(u, v) = f(v), \quad v \in \mathbb{V}_{\mathcal{T}}, \quad (3.28)$$

where its ingredients are explained in a moment. First, elementwise integration by parts introduces trace terms, which for $u \in L_2(\Omega)$ are not defined. This necessitates introducing, as auxiliary unknowns, *skeleton components* living on $\partial\mathcal{T} := \bigcup_{K \in \mathcal{T}} \partial K$. As the notation u indicates the unknown u has therefore several components $u = (u_b, u_s) \in \mathbb{U}_{\mathcal{T}} = \mathbb{U}_{\mathcal{T},b} \times \mathbb{U}_{\mathcal{T},s}$, where

$$\mathbb{U}_{\mathcal{T},b} := \prod_{K \in \mathcal{T}} L_2(K), \quad \mathbb{U}_{\mathcal{T},s} = \left\{ u_s \in L_2(\partial\mathcal{T}) : \|u_s\|_{\mathbb{U}_{\mathcal{T},s}} := \inf_{w \in \mathbb{G}(\mathcal{B}) : w|_{\partial\mathcal{T}} = u_s} \|\mathcal{B}w\|_{L_2(\Omega)} < \infty \right\}. \quad (3.29)$$

The bilinear form $b_{\mathcal{T}}(u, v)$ has then the structure

$$b_{\mathcal{T}}(u, v) = \sum_{K \in \mathcal{T}} b_K(u_b, v) + \langle u_s, v \rangle_{\partial K} =: \sum_{K \in \mathcal{T}} B_K(u, v), \quad (3.30)$$

and finally, perhaps most importantly, the test space $\mathbb{V}_{\mathcal{T}}$ is a *broken space*

$$\mathbb{V}_{\mathcal{T}} = \prod_{K \in \mathcal{T}} \mathbb{V}_K, \quad \|v\|_{\mathbb{V}_{\mathcal{T}}}^2 = \sum_{K \in \mathcal{T}} \|v\|_{\mathbb{V}_K}^2, \quad (3.31)$$

where the $\mathbb{V}_K$, $K \in \mathcal{T}$, are ‘localized’ versions of the test space $\mathbb{V}$ in the underlying ‘conforming’ formulation (3.27). A simple, but crucial consequence of this latter feature, is the product structure of the Riesz lift $\mathcal{R}_{\mathcal{T}} : \mathbb{V}'_{\mathcal{T}} \rightarrow \mathbb{V}_{\mathcal{T}}$

$$\mathbb{V}'_{\mathcal{T}} = \prod_{K \in \mathcal{T}} \mathbb{V}'_K, \quad \mathcal{R}_{\mathcal{T}} = \prod_{K \in \mathcal{T}} \mathcal{R}_K, \quad (3.32)$$

where as before $\langle \mathcal{R}_{\mathcal{T}} \ell, v \rangle_{\mathbb{V}_{\mathcal{T}}} = \ell(v)$, $v \in \mathbb{V}_{\mathcal{T}}$ and $\langle \mathcal{R}_K \ell, v \rangle_{\mathbb{V}_K} = \ell(v)$, $v \in \mathbb{V}_K$.

Once well-posedness of (3.28) has been established (to be assumed in what follows), upon introducing the DPG residual $\mathcal{F}_{\mathcal{T}}(\mathbf{w})(v) := f(v) - b_{\mathcal{T}}(\mathbf{w}, v)$, we have by (3.31)

$$\|\mathcal{F}_{\mathcal{T}}(\mathbf{w})(v)\|_{\mathbb{V}'_{\mathcal{T}}}^2 = \sum_{K \in \mathcal{T}} \|\mathcal{F}_K(\mathbf{w})\|_{\mathbb{V}'_K}^2, \quad \|\mathcal{F}_K(\mathbf{w})\|_{\mathbb{V}'_K} = \sup_{v \in \mathbb{V}_K} \frac{f(v) - b_K(\mathbf{w}, v)}{\|v\|_{\mathbb{V}_K}}.$$

Thus, well-posedness of (3.28) says that the exact solution $u \in \mathbb{U}_{\mathcal{T}}$ of (3.28) satisfies

$$\|u - \mathbf{w}\|_{\mathbb{U}_{\mathcal{T}}}^2 \approx \|\mathcal{F}_{\mathcal{T}}(\mathbf{w})\|_{\mathbb{V}'_{\mathcal{T}}}^2 = \sum_{K \in \mathcal{T}} \|\mathcal{F}_K(\mathbf{w})\|_{\mathbb{V}'_K}^2, \quad \mathbf{w} \in \mathbb{U}_{\mathcal{T}}. \quad (3.33)$$

Although this is the starting point for *a posteriori* error estimation, in the above form it is not practically feasible, since each term $\|\mathcal{F}_K(\mathbf{w})\|_{\mathbb{V}'_K}^2$ still involves a supremization over the infinite-dimensional local test space $\mathbb{V}_K$. The key step towards a practical DPG method is to identify for a given trial spaces $\mathbb{U}_{\mathcal{T}}^{\delta} \subset \mathbb{U}_{\mathcal{T}}$ and a data space $\mathbb{F}^{\delta}$ of piecewise polynomials on $\mathcal{T}$ of a *fixed maximal order*, local *test-search spaces* $\tilde{\mathbb{V}}_K^{\delta}$, $K \in \mathcal{T}$, such that

$$\inf_{\mathbf{w}^{\delta} \in \mathbb{U}_{\mathcal{T}}^{\delta}} \sup_{f^{\delta} \in \mathbb{F}^{\delta}} \frac{b_{\mathcal{T}}(\mathbf{w}^{\delta}, v^{\delta}) - f^{\delta}(v^{\delta})}{\|\mathbf{w} - \mathcal{B}_{\mathcal{T}}^{-1} f^{\delta}\|_{\mathbb{U}_{\mathcal{T}}} \|v^{\delta}\|_{\mathbb{V}_{\mathcal{T}}}} \geq \beta > 0 \quad (3.34)$$

holds uniformly in $\mathcal{T}$. Of course, by our assumption on well-posedness of (3.28), (3.34), holds trivially when $\tilde{\mathbb{V}}_{\mathcal{T}}^{\delta}$ is replaced by $\mathbb{V}_{\mathcal{T}}$. Note that upon taking $f^{\delta} = 0$ (3.34) requires uniform inf-sup stability of the Petrov–Galerkin scheme with trial space $\mathbb{U}_{\mathcal{T}}^{\delta}$ and test space

$$\mathbb{V}_{\mathcal{T}}^{\delta} := \prod_{K \in \mathcal{T}} \mathfrak{T}_K^{\delta} \mathbb{U}_{\mathcal{T}}^{\delta}, \quad \langle \mathfrak{T}_{B_K}^{\delta} \mathbf{w}^{\delta}, v^{\delta} \rangle_{\mathbb{V}_K} = B_K(\mathbf{w}^{\delta}, v^{\delta}), \quad \forall v^{\delta} \in \tilde{\mathbb{V}}_K^{\delta}.$$

Thus, calculating $\mathfrak{T}_{B_K}^{\delta} \mathbf{w}^{\delta}$ amounts to solving a symmetric positive definite system of size $\dim \tilde{\mathbb{V}}_K^{\delta}$. In complete analogy the trial-to-residual map $\mathfrak{T}_{\mathcal{T}}^{\delta}$ is defined by

$$\langle \mathfrak{T}_{\mathcal{T}}^{\delta} \mathbf{w}^{\delta}, v^{\delta} \rangle_{\mathbb{V}_{\mathcal{T}}} = \mathcal{F}_{\mathcal{T}}(\mathbf{w}^{\delta}, v^{\delta}), \quad v^{\delta} \in \tilde{\mathbb{V}}_K^{\delta}, \quad K \in \mathcal{T}. \quad (3.35)$$

Hence, the game is to choose, on the one hand, the $\tilde{\mathbb{V}}_K^{\delta}$ large enough to satisfy (3.34). On the other hand, one wishes to keep $\tilde{n} := \dim \tilde{\mathbb{V}}_K^{\delta}$ as small as possible, ideally *uniformly bounded* with respect to K , $\mathcal{T}$. If this is the case, although quite expensive in quantitative terms, the size of discrete DPG systems

still remains uniformly proportional to $\#\mathcal{T}$, which is a cost-lower bound for all discretizations based on the partition $\mathcal{T}$. This programme has by now been carried out for a wide scope of PDE models, see, e.g., Carstensen *et al.* (2016), Demkowicz *et al.* (2017), Dahmen *et al.* (2019), Dahmen & Stevenson (2019), Dahmen *et al.* (2021) and Gantner & Stevenson (2021).

Given $\tilde{\mathbb{V}}_K^\delta$, the right-hand side f_K , restricted as a functional to $\mathbb{V}_K$, may have a part that cannot be seen when restricted to $\tilde{\mathbb{V}}_K^\delta$, so that in general one misses a *data-oscillation error*. We denote by f_K^δ a suitable projection to $\tilde{\mathbb{V}}_K^\delta$, which differs from f_K by such a data oscillation. Since under mild assumptions on f these errors are negligible (compared with approximation to u), and we ignore, and such terms in what follows for simplicity of exposition. We can summarize these findings as follows.

REMARK 3.11. Validity of (3.34) implies that for any $\mathbf{w}^\delta \in \mathbb{U}_{\mathcal{T}}^\delta$

$$\|\mathcal{F}_K(\mathbf{w}^\delta)\|_{\mathbb{V}_K'}^2 \approx \|\mathcal{F}_K(\mathbf{w}^\delta)\|_{(\mathbb{V}_K^\delta)'}^2 \quad (3.36)$$

(up to data oscillation) so that

$$\|\mathbf{u} - \mathbf{w}^\delta\|_{\mathbb{U}_{\mathcal{T}}}^2 \approx \|\mathcal{F}_{\mathcal{T}}(\mathbf{w}^\delta)(v)\|_{(\mathbb{V}_{\mathcal{T}}^\delta)'}^2 = \sum_{K \in \mathcal{T}} \|\mathcal{F}_K(\mathbf{w}^\delta)\|_{\mathbb{V}_K}^2, \quad \mathbf{w}^\delta \in \mathbb{U}_{\mathcal{T}}^\delta, \quad (3.37)$$

where the quantities $\|\mathcal{F}_K(\mathbf{w}^\delta)\|_{\mathbb{V}_K}$ are now computable.

REMARK 3.12. We emphasize that the equivalence (3.37) holds for *any* $\mathbf{w}^\delta \in \mathbb{U}_{\mathcal{T}}^\delta$, not just for the current Petrov–Galerkin solution, as is the case for classical residual error estimators for Galerkin methods, see, e.g., Verfürth (2013). This is essential in the present context.

REMARK 3.13. Note that well-posedness of (3.28) draws on both the bijectivity of $\mathcal{B}$ and $\mathcal{B}^*$ underlying the scenarios (S), (UW).

Returning to the parameter-dependent case, note first that now, due to the appearance of the graph norm $\|\mathcal{B}_p \cdot\|_{L_2(\Omega)}$ in the definition of the skeleton, the factor $\mathbb{U}_{\mathcal{T},s,p}$ may depend on p and so does $\mathbb{U}_{\mathcal{T},p} := \mathbb{U}_{\mathcal{T},b} \times \mathbb{U}_{\mathcal{T},s,p}$. Likewise, the test spaces $\mathbb{V}_{\mathcal{T},p} = \prod_{K \in \mathcal{T}} \mathbb{V}_{K,p}$ will generally depend on $p \in P$. The roles of $\mathbb{X}, \mathbb{Y}$ from (2.9) are then played by

$$\mathbb{X}_{\mathcal{T}} := L_2\left(P, (\mathbb{U}_{\mathcal{T},p})_{p \in P}\right) \quad \text{and} \quad \mathbb{Y}_{\mathcal{T}} := L_2\left(P, (\mathbb{V}_{\mathcal{T},p})_{p \in P}\right), \quad (3.38)$$

and analogously for their semidiscrete counterparts $\widehat{\mathbb{X}}_{\mathcal{T}}$ and $\widehat{\mathbb{Y}}_{\mathcal{T}}$.

REMARK 3.14. If the spaces $\mathbb{X}$ and $\mathbb{Y}$ associated with the (parameter-dependent) conforming problem (3.27) are well defined as direct integrals, i.e., Property 2.3 holds, it is not hard to show that this carries over to the mesh-dependent variants $\mathbb{X}_{\mathcal{T}}$ and $\mathbb{Y}_{\mathcal{T}}$, which are also well-defined direct integrals. In the same manner Property 2.5 carries over to the broken spaces, so that Theorems 2 and 3 remain applicable.

The error–residual relation asserted by Remark 3.11 holds for any piecewise polynomial in the finite-dimensional DPG trial space $\mathbb{U}_{\mathcal{T}}^\delta$. It immediately carries over to the parameter-dependent case, as long as the trial functions are piecewise polynomials in $\mathbb{U}_{\mathcal{T}}^\delta$ as functions on Ω . This motivates the notion of *hybrid* hypothesis classes $\mathcal{H}(\Theta)$. By this we mean approximation systems that comprise functions $\mathbf{w}^\delta(x, p; \theta)$, $\theta \in \Theta$, that are linear combinations of piecewise polynomial basis functions $\phi_i(x)$, $i \in \mathcal{I}_{\mathcal{T}}$,

with coefficients represented as neural networks with input parameters $p \in P$ and trainable weight vectors θ as outlined in Section 1.6. Hence, under the given assumptions, Remark 3.11 implies that

$$\|u - \mathfrak{w}^\delta(\theta)\|_{\mathbb{X}, \mathcal{T}}^2 \approx \sum_{K \in \mathcal{T}} \int_P \|\mathfrak{T}_{\mathcal{F}_K}^\delta \mathfrak{w}(\cdot, p; \theta)\|_{\mathbb{V}_{K,p}}^2 d\mu(p). \quad (3.39)$$

Accordingly, for any finite training set $\widehat{P} \subset P$ and any $\mathfrak{w}^\delta(\cdot, p; \theta) \in \mathcal{H}(\Theta) \subset \mathbb{U}_{\mathcal{T},p}^\delta$ one has

$$\frac{1}{\#\widehat{P}} \sum_{p \in \widehat{P}} \|u(\cdot, p) - \mathfrak{w}^\delta(\cdot, p; \theta)\|_{\mathbb{U}_{\mathcal{T},p}}^2 \approx \frac{1}{\#\widehat{P}} \sum_{p \in \widehat{P}} \sum_{K \in \mathcal{T}} \|\mathfrak{T}_{\mathcal{F}_K}^\delta(\mathfrak{w}^\delta(\cdot, p; \theta))\|_{\mathbb{V}_{K,p}}^2, \quad (3.40)$$

where the right-hand side is the desired variationally correct computable loss function.

The minimization of the loss in the right-hand side of (3.40) requires evaluating the local terms $\|\mathfrak{T}_{\mathcal{F}_K}^\delta(\mathfrak{w}^\delta(\cdot, p; \theta))\|_{\mathbb{V}_{K,p}}^2$, $K \in \mathcal{T}$. Although this has been well documented in DPG literature, the dependence on the model parameters p , as well as on the trainable weights θ , justifies a few further comments on assembling these quantities, given in Appendix C. As can be seen there the computation of each term $\|\mathfrak{T}_{\mathcal{F}_K}^\delta(\mathfrak{w}^\delta(\cdot, p; \theta))\|_{\mathbb{V}_{K,p}}^2$ is of uniformly bounded cost. However, the underlying constant depends algebraically on the polynomial degrees in trial and test spaces, and can thus be substantial. This affects primarily the potentially very large number of θ -evaluations during optimization. On the other hand, this is part of the offline cost, which tends to be substantial in any model reduction approach.

4. Network architectures and computational strategies

4.1 Low-rank ResNet hybrid architectures

We next discuss in further detail the hybrid representation format introduced in Section 1.6 for approximating solutions $p \mapsto u(p)$ to (2.11). Separating spatio-temporal and parametric variables, we consider piecewise polynomials as functions of spatio-temporal variables with *parameter-dependent* coefficient functions as in (1.13) that can, for instance, be represented as DNNs. In doing so, we exploit the systematic convergence of finite element methods in small spatial dimensions in combination with reliable concepts for evaluating dual norms.

This format is instrumental for strategy (UW) to construct computable variationally correct loss functions without the use of min-max optimization. It is not strictly necessary for scenario (S), but still offers several significant advantages there as well. First, rigorously enforcing essential boundary conditions for DNNs is an issue, see Berrone *et al.* (2022) for various strategies. The hybrid format greatly facilitates incorporating essential (spatial) boundary conditions. Second, since the input parameters of the coefficient neural networks are just the model parameters, training requires only efficient backpropagation, and the neural network activation function need not have higher than first-order weak derivatives. When using DNNs with both spatio-temporal and parametric variables as input parameters, spatial-parametric mixed derivatives are required. This has to be done with great care to avoid a significant loss of efficiency in the evaluation of gradients with respect to neural network parameters; see, for example, Taylor *et al.* (2023).

Recall that, as noted in Section 1.1, we are especially interested in nonlinear approximations that can be more efficient than reduced models based on projection to some fixed linear subspace, such as RBMs. This rules out some typical elements of neural network architectures, such as linear output layers, where

the output of the previous layers (of length n) is multiplied onto some fixed matrix $A^{N \times n}$ to produce the final output, where N is the size of the spatio-temporal discretization. Such a structure is used for parametric problems, for example, in Geist et al. (2021). In this case, however, the columns of A play the role of a reduced basis, and the remaining layers of the neural network simply produce coefficients for these basis elements. Moreover, instead of optimizing heavily overparameterized neural networks as common in machine learning in our setting, we are rather interested in networks that yield data-sparse approximations.

For the spatial or spatio-temporal discretization we assume a hierarchy of meshes $\mathcal{T}_1, \dots, \mathcal{T}_J$ generated by uniform subdivisions, and we approximate the solution on the finest mesh $\mathcal{T}_J$. For $j = 1, \dots, J$ we have bases $\{\phi_i^{(j)}\}_{i \in \mathcal{I}_j}$ of the respective finite element spaces. As outlined in (1.13) we approximate solutions, depending on the parameter p in the form

$$p \mapsto \sum_{i \in \mathcal{I}_J} \mathbf{u}_{\text{NN},i}(p; \theta) \phi_i^{(J)},$$

where the vector $\mathbf{u}_{\text{NN}}(\cdot; \theta)$ is a neural network with trainable parameters θ .

We next discuss the particular network architectures used for the coefficients $\mathbf{u}_{\text{NN}}$, which are based on *residual networks* (ResNets). These are widely used in scientific computing due to their favourable numerical stability properties. We use residual networks with layers $\Phi_{(A,W,b)}: \mathbb{R}^n \rightarrow \mathbb{R}^n$, with n to be determined, of the particular form

$$z \mapsto \Phi_{(A,W,b)}(z) = z + A\sigma(Wz + b), \quad (4.1)$$

where $A \in \mathbb{R}^{n \times r}$, $W \in \mathbb{R}^{r \times n}$ and $b \in \mathbb{R}^n$, with componentwise application of the activation function $\sigma: \mathbb{R} \rightarrow \mathbb{R}$. Here, we aim for $r \ll n$ to achieve a data-sparse representation, and by analogy to low-rank matrix representations we will call r the *rank* of the layer.

With $n_j = \#\mathcal{I}_j$, let the (sparse) prolongation matrices $P_j \in \mathbb{R}^{n_j \times n_{j-1}}$ be given for $j = 1, \dots, J$ by

$$\sum_{i \in \mathcal{I}_{j-1}} \mathbf{v}_i \phi_i^{(j-1)} = \sum_{i \in \mathcal{I}_j} (P_j \mathbf{v})_i \phi_i^{(j)}, \quad \mathbf{v} \in \mathbb{R}^{n_{j-1}}. \quad (4.2)$$

For any matrix $Q \in \mathbb{R}^{n_j \times n_{j-1}}$ with the same sparsity pattern as P_j we accordingly introduce the *prolongation layers* $\Pi_Q: \mathbb{R}^{n_{j-1}} \rightarrow \mathbb{R}^{n_j}$, $z \mapsto Qz$.

Starting from a fully connected layer Λ_0 given by

$$\mathbb{R}^{d_p} \ni p \mapsto \Lambda_0(p) = \sigma(W_0 p + b_0) \in \mathbb{R}^{n_1}$$

the resulting neural networks are composed of layers $\Lambda_j: \mathbb{R}^{n_j} \rightarrow \mathbb{R}^{n_j}$ of the form

$$\Lambda_j = \Phi_{(A_{j,L_j}, W_{j,L_j}, b_{j,L_j})} \circ \dots \circ \Phi_{(A_{j,2}, W_{j,2}, b_{j,2})} \circ \Phi_{(A_{j,1}, W_{j,1}, b_{j,1})} \circ \Pi_{Q_j} \quad (4.3)$$

with some $L_j \in \mathbb{N}$, for $j = 1, \dots, J$. The approximation of $\mathbf{u}_{\text{NN}}$ is then of the form

$$\mathbf{u}_{\text{NN}}(\cdot; \theta) = \Lambda_J \circ \dots \circ \Lambda_1 \circ \Lambda_0,$$

with trainable parameters

$$\theta = (A_{J,L_J}, W_{J,L_J}, b_{J,L_J}, \dots, A_{J,1}, W_{J,1}, b_{J,1}, Q_J, \dots, \dots, A_{1,L_1}, W_{1,L_1}, b_{1,L_1}, \dots, A_{1,1}, W_{1,1}, b_{1,1}, Q_1, W_0, b_0).$$

Note that multilevel structure in discretizations is also used, for example, in [Heiß *et al.* \(2023\)](#) with a more specialized architecture for elliptic parametric problems.

4.2 Computational strategies for training

In the experiments in each test we use a fixed rank parameter r for all ResNet layers as in (4.1) and vary the depth of the neural network. We exploit the multilevel structure of the architecture in performing a gradual refinement by subsequently inserting additional ResNet layers into the layer groups Λ_j as in (4.3), which amounts to increasing L_j . These additional ResNet layers are initialized with zero weights. The layer Λ_0 is initialized with a random initialization with Gaussian entries of vanishing mean and variance 10^{-3} . The matrices $Q_1, \dots, Q_J$ are initialized as $Q_j = P_j$, with P_j as in (4.2), for each j . As common in the context of neural network approximations for PDEs we use L-BFGS for minimizing the loss functions.

5. Numerical experiments

5.1 Validation

In subsequent experiments we will *not* prescribe the ‘exact solution’, typically as an analytically representable smooth function, but rather discuss scenarios that highlight the characteristic features of the problem under consideration. In the elliptic case this permits low regularity of solutions due to the contrast in the diffusion coefficients. In the case of transport equations, the tests include the appearance of shear layers. Hence, we do not have the ‘ground truth’ at hand by which we could validate numerical results. Instead, we exploit that, for each parameter p , we can apply a finite element method, which we know is based on a stable variational formulation so that corresponding residuals in the right norm tightly reflect the achieved accuracy with respect to the exact solution.

Just for the sake of illustrating the performance of the proposed estimation method, we compute as ‘ground truth’ in addition finite element solution coefficients $\mathbf{u}_{\text{FE}}(p^i)$ (with respect to the appropriate problem-dependent framework) for a set $\hat{P}_{\text{test}} \subset P$ of N_{test} test samples p^i . To assess the corresponding ‘generalization error’ of FE-solutions, with respect to the parameters, we evaluate the respective parameter-dependent variationally correct loss function at these test samples. Since we have made sure that the underlying variational formulation is stable, we know that these quantities tightly reflect the error with respect to the respective trial norm $\|\cdot\|_{\mathbb{U}_p}$.

As will be seen in more detail later, for the hybrid format the loss function can, with the notation

$$\|\mathbf{x}\|_{\mathbf{G}}^2 = \langle \mathbf{x}, \mathbf{G}\mathbf{x} \rangle,$$

be represented as $\|\mathbf{B}_p \mathbf{u}_{\text{FE}}(p) - \mathbf{f}\|_{\mathbf{G}}^2$, where $\mathbf{B}_p$ is the representation of the operator $\mathcal{B}_p$ on the FE space and $\mathbf{G}$ represents the Riesz lift.

The first validation quantity is the relative error achieved by the FE solution over $\hat{P}_{\text{test}}$ in the mean square sense,

$$\varepsilon_{\text{ref}} = \sqrt{\frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \frac{\|\mathbf{B}_{p^i} \mathbf{u}_{\text{FE}}(p^i) - \mathbf{f}\|_{\mathbf{G}}^2}{\|\mathbf{f}\|_{\mathbf{G}}^2}}. \quad (5.1)$$

Note that the size of this quantity is determined by the spatial resolution of the FE space (provided N_{test} is large enough) and reflects the achievable spatial discretization error.

We denote by $u_{\text{net}}(\cdot; p)$ the result of our training over the given hybrid hypothesis class, which we can now compare with $u_{\text{FE}}(\cdot; p)$ for each test sample $p^i \in \hat{P}_{\text{test}}$. Rather than computing the norms of the various solution components in $\mathbb{X}$ (which is a product space), we exploit that errors in $\mathbb{X}$ are equivalent to residuals in $\mathbb{Y}$, which in turn can be assessed through the variationally correct loss function.

Finally, we monitor the generalization error of the prediction u_{NN} directly via the variationally correct loss function, tested over $\hat{P}_{\text{test}}$,

$$\varepsilon_{\text{pred}} = \sqrt{\frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \frac{\|\mathbf{B}_{p^i} \mathbf{u}_{\text{NN}}(p^i) - \mathbf{f}\|_{\mathbf{G}}^2}{\|\mathbf{f}\|_{\mathbf{G}}^2}}. \quad (5.2)$$

This is the quantity that would be evaluated in an application as a certification. Throughout this section *level* refers to the number J as in Section 4.1 of uniform refinements used in the spatial discretization starting from two triangles, that is, a triangulation of level j has $2^j + 1$ grid points in each spatial variable. *Rank* refers to the width of the ResNet layers as in (4.1).

With the architectures described in Section 4 we consider numerical experiments for two different types of PDEs. Throughout our tests for the hybrid models we used a LeakyReLU activation function $\sigma(x) = \max\{x, 10^{-3}x\}$. In all experiments (unless indicated otherwise) we used L-BFGS optimization, retaining a maximum of nine previous BFGS updates, with a maximal number of 5000 iteration for the single-parameter case and 20000 iterations for the other experiment.

5.2 An elliptic problem in scenario (S)

We deliberately choose an example with nonsmooth diffusion coefficients with solutions of low spatial regularity that generally do not belong to $H^2(\Omega)$. Specifically, we consider (3.8) with $d_p = 4$ parameters, where $f = 1$ and $a(p) = \sum_{i=1}^4 (p_i \alpha_i(x))$, and α_i is the indicator function of one of the four squares of a uniform subdivision. The indicator function is ordered such that α_1 is in the top left and α_4 is in the bottom right.

As in (3.12), we consider the corresponding first-order operator

$$\mathcal{B}_p : [\sigma, u] \mapsto \begin{pmatrix} \sigma - a(p) \nabla u \\ -\operatorname{div} \sigma \end{pmatrix},$$

which maps $[\sigma, u] \in H(\operatorname{div}, \Omega) \times H_0^1(\Omega) =: \mathbb{U}$ to $L_2 = L_2(\Omega; \mathbb{R}^2) \times L_2(\Omega)$. Thus, the trial spaces do not depend on the parameters in an essential way and we can work in scenario (S). The Riesz lift in this case is the identity, and thus $\mathbf{G}$ is the scaled identity matrix. For the spatial discretization, we

TABLE 1 *The relative norms of the residuals for two different refinement levels for a network with 15 layers and 1000 training samples*

	Level 3, Rank 40	Level 4, Rank 60
ε_{ref}	0.055	0.028
$\varepsilon_{\text{pred}}$	0.057	0.030




FIG. 1. Snapshot of the LSG solution, the prediction by the neural network (with rank 60 and 15 layers trained on 1000 samples) and the pointwise absolute value of the difference for a level four refinement evaluated for the parameter $p = [0.65, 1.45, 1.45, 0.65]$ corresponding to a checkerboard configuration.

use piecewise linear elements in $H_0^1(\Omega)$ for u and lowest-order Raviart–Thomas elements in $H(\text{div}; \Omega)$ for σ .

For N training points $p^i \in \widehat{P} \subset P$ the loss function reads

$$\frac{1}{N} \sum_{i=1}^N \left(\|\sigma(p^i) - a(p_i) \nabla u(p^i)\|_{(L_2(\Omega))^d}^2 + \|\text{div } \sigma(p^i) + f\|_{L_2(\Omega)}^2 \right).$$

As mentioned earlier, in Muzalevsky & Repin (2021) and Fanaskov *et al.* (2024), a similar loss function has been used on the spatial domain for a parameter-independent Poisson problem, where the flux variable σ is interpreted as a ‘certificate’ that leads to a strict upper bound on the error in u .

5.2.1 Numerical results in the hybrid case. Table 1 displays the relative residuals ε_{ref} of the Least-Squares Galerkin (LSG) solution and $\varepsilon_{\text{pred}}$ of the neural network prediction, where the neural network has 15 layers in total. In Figs 1 and 2, two snapshots based on the DPG solution and the network prediction are displayed. As one can see the hybrid method captures the essential features of the FEM solutions and also provides an error reduction under refinement.

5.2.2 Comparison with a vanilla PINN approach. Recall from Remark 3.2 that, in the present setting, a PINN approach applied to the *first-order system* does not incur any variational crime; see, for example, Cai *et al.* (2020) for earlier work in this regard. This is in contrast to PINN applied to the original second-order formulation. However, one still has to be careful in enforcing essential boundary conditions, where a weighted ℓ_2 -term is not strictly correct, and we refer to Berrone *et al.* (2022) and Bertoluzza *et al.* (2024) for respective alternative strategies. We compare below two options.




FIG. 2. Snapshot of the LSG solution, the prediction by the neural network (with rank 60 and 15 layers trained on 1000 samples) and the pointwise absolute value of the difference for a level four refinement and given parameter $p = [0.53, 1.09, 0.84, 0.82]$.

TABLE 2 *The relative dual norm of the residual in the elliptic problem of the prediction by a hybrid Low-rank ResNet (H-Net) for refinements of levels 3 and 4, respectively, as well as for a PINN approach with weakly enforced boundaries (PINN-B) and with trimming function (PINN-T) with three and four dense layers, respectively (PINN-T/B-1, PINN-T/B-2)*

	H-Net Level 3	H-Net Level 4	PINN-B-1	PINN-B-2	PINN-T-1	PINN-T-2
$\varepsilon_{\text{pred}}$	0.057	0.030	0.05	0.059	0.036	0.047
dof	26852	165575	119364	175800	119364	175800

To that end, we take separate neural networks for u , σ_1 and σ_2 , each of them a fully connected neural network from $\mathbb{R}^6$ to $\mathbb{R}$ with a final linear layer. Thus, we have $2 + 4$ variables in the input for the full parametric problem. We use the sigmoid activation function. Furthermore, the networks with maximal width 196 are trained with ADAM with decreasing step sizes from 0.05 to 0.005 and a maximum of 2500 iterations. Here, we used the NeuralPDE.jl package (Zubov *et al.*, 2021) in its default settings; in particular, we use numerical differentiation for the derivative with respect to the spatial variables. The architectures PINN-B uses an L_2 -like loss to enforce the boundary conditions, whereas PINN-T uses the trimming function $x(1-x)y(1-y)$ multiplied to the network output to enforce the boundary conditions.

In Table 2, one can see that with a similar number of weights and biases the hybrid architectures can perform at the same level as the PINNs or even outperform them. It is important to note, however, that the hybrid approximations offer a more systematic way to increase the network size and then actually decrease the loss. In our experiments the hybrid approach exhibits a significantly enhanced optimization efficacy and predictability.

We therefore tested the sensitivity of the training success of the neural networks with respect to the initialization. Specifically, we trained the neural network architectures with 10 randomly initialized collections of weights and biases for a fixed number of epochs. As can be seen in Fig. 3, for the hybrid network architectures in combination with nested refinement, the training success is significantly more robust with respect to the random initialization of the weights. In fact, for the H-Net Level 3, 60% of the random initializations of weights and biases converged to a parameter setting that provides a relative test error of the residual between 0.04 and 0.08, while 20% of the total runs achieved a relative error below 0.06. Note that, by design, the H-Net Level 3 cannot realize a relative error below 0.055, due to the resolution of the finite elements. This can be remedied by the additional level of spatial refinement




FIG. 3. Histogram describing how many initializations of the neural networks achieved a relative test error in the regimes $(0.02, 0.04]$, $(0.04, 0.08]$, $(0.08, 0.16]$ and $(0.16, \infty)$, respectively, after training a fixed number of epochs for the H-Net Level 3, H-Net Level 4, PINN-B-2 and PINN-T-2.

in H-Net Level 4, which achieves 100% of training runs below a relative error of 0.08, with 60% of the runs below an error of 0.04. Note that, again due to the spatial resolution, for this architecture one cannot expect an error below 0.028.

In the case of the fully connected architecture, used for PINN-B2 and PINN-T-2 in connection with the standard setting of the library (Zubov *et al.*, 2021), none of the initializations achieved an accuracy below 0.04. Only 20% and 40%, respectively, achieved an accuracy between 0.04 and 0.08, while the remaining 80% and 60%, respectively, fail to achieve an error below 0.08, where 60% even surpass an error of 0.2. These observations seem to indicate that the combination of the hybrid network architecture with nested refinement significantly increases training success in comparison with the vanilla PINN setting.

5.3 Linear transport with parameter-dependent convection field in scenario (UW)

We consider the transport model from Section 3.3.1. Note that now a rigorous interpretation of the continuous reference loss as an expectation requires resorting to the concept of direct integrals, see Section 3.3. Moreover, such models exhibit several significant obstructions. First, standard Galerkin or discontinuous Galerkin discretizations are not uniformly inf-sup stable. Second, one faces a generally nonsmooth parameter dependence, as well as an essential dependence of the graph spaces and boundary conditions on the parameters. In this case, we find scenario (UW) more adequate. In order to highlight related effects, we consider the simple case of a single parameter in the convection field as well as the case of a convection field, depending on several parameters.

For the spatial discretization we employ a DPG finite element method with piecewise constant elements on the triangles, as well as piecewise linear elements on the skeleton in the trial space, and piecewise quadratic elements for the test space as in Broersen *et al.* (2018).

Denote by $\mathbf{u}_{\text{net}}$ the prediction of the neural network and by $\mathbf{u}_{\text{DPG}}(p)$ the DPG solution for given parameter p . Let $\mathbf{B}_{p_i}$ be the representation of the operator $\mathcal{B}_{p_i} = \mathbf{b}(\cdot, p_i) \cdot \nabla$ in the finite element spaces,

TABLE 3 $\varepsilon_{\text{pred}}$: the relative dual norm of the residual of the prediction by a hybrid low-rank ResNet for the transport problem with one parameter and a refinement of level 4

Rank	Layers	1000 samples	1500 samples	2000 samples
10	11	0.263	0.2618	0.262
	15	0.252	0.255	0.258
15	11	0.251	0.252	0.255
	15	0.25	0.25	0.25
20	11	0.253	0.253	0.253
	15	0.25	0.25	0.25
30	11	0.251	0.249	0.252
	15	0.249	0.249	0.249
relative DPG residual ε_{ref}		0.248		

let $\mathbf{G}_{p^i}$ be the parameter-dependent representation of the Riesz lift and $\mathbf{f}$ be the evaluation of $\langle f, \cdot \rangle$ at the basis functions of the discretized test space. Then, given training samples p^i , the loss function reads

$$\frac{1}{N} \sum_{i=1}^N \|\mathbf{B}_{p^i} \mathbf{u}_{\text{net}}(p^i) - \mathbf{f}\|_{\mathbf{G}_{p^i}}^2.$$

We analyse the quality of the obtained predictions again by monitoring the quantities $\varepsilon_{\text{pred}}$ and ε_{ref} introduced in Section 5.1 for the above loss function.

5.3.1 One parameter linear transport equation. First, we consider a linear transport equation in two spatial variables parameterized by the angle of the velocity field, that is,

$$[\cos(\pi p), \sin(\pi p)] \cdot \nabla u = f$$

with $f = \chi_{[0.25, 0.5]^2}$ and $p \in \left(0, \frac{1}{2}\right)$.

Table 3 presents the accuracy achieved by the predictions. It displays the values of $\varepsilon_{\text{pred}}$ for various training settings. The value in the bottom row shows the accuracy ε_{ref} achieved by the individual DPG solutions, see (5.1). We see that both quantities ε_{ref} and $\varepsilon_{\text{pred}}$ are very close, reflecting a very good estimation quality in this case. Of course, the rather coarse spatial discretization, the very low regularity of the exact solutions and the low order discretization cause an overall poor accuracy. In Fig. 4, the decrease of the error with respect to the refinement level for the FEM solution and the neural network prediction is displayed verifying that one obtains a consistent error reduction as expected. In Figs. 5 and 6 two snapshots of the FEM solution and the neural network prediction for two distinct parameters are shown.

5.4 Multi-parametric linear transport with piecewise linear vector field

Second, we consider a linear transport equation in two physical dimensions, that is, one temporal and one spatial dimension, where the convection field is parameterized by hat functions in space with coefficients




FIG. 4. Visualization of the relative dual norm of the residual of the prediction by the neural network versus the DPG solution for a network with rank 20 trained on 1000 samples.




FIG. 5. Snapshot for $p = 0.2$ of piecewise constant elements for the DPG solution and the prediction by the neural network with rank 20 and 15 layers, trained on 1000 samples for level 6.




FIG. 6. Snapshot for $p = 0.1$ of piecewise constant elements for the DPG solution and the prediction by the neural network with rank 20 and 15 layers, trained on 1000 samples for level 6.

TABLE 4 $\varepsilon_{\text{pred}}$: the relative dual norm of the residual of the prediction by a hybrid low-rank ResNet for the transport equation with seven parameters and a refinement of level 4

Rank	11 Layer	13 Layer	15 Layer
10	0.454	0.447	0.44
15	0.442	0.431	0.428
20	0.481	0.464	0.455
relative DPG residual ε_{ref}		0.337	

constant in time. The problem thus takes the form

$$\left[1, \sum_{i=1}^7 a_i(x) p_i \right] \cdot \nabla_{t,x} u = f$$

with $f = \chi_{[0.25, 0.5]^2}$, where a_i are hat functions in space constant in time on the interior nodes of a grid with nine nodes and $p_i \in [0.3, 1.3]$. We use the same error measures as above in (5.1) and (5.2).

Table 4 shows the results for ε_{ref} and $\varepsilon_{\text{pred}}$ of analogous tests for the case of seven parameters, in addition to the two spatio-temporal dimensions. In principle, the results are very similar, except for a slightly lower agreement between the trained prediction and the DPG solutions evaluated at the test samples. In view of the high dimensionality this is perhaps not surprising. We observe that increasing the rank does not necessarily improve the achieved accuracy. However, this is likely due to the increased complexity of the optimization task, where it is generally difficult to assess whether a given optimization schedule stalls on a mediocre plateau or has converged.

6. Summary and concluding remarks

Our focus in this work has been the approximation of solution manifolds of different classes of parameter-dependent PDEs, with dissipative elliptic or parabolic problems as well as linear transport problems as particular examples. A central issue when using nonlinear model classes, such as neural networks, for such approximations is how one can assess the quality of approximate solutions using only computable information. Such control of errors is of crucial importance for surrogate models produced by operator learning, where known training methods lack reliability.

Our approach to controlling errors centres on what we call *variationally correct loss functions*. That is, at any stage of an optimization process, the current magnitude of the loss is a tight bound for the error in a model-compliant norm, without imposing additional regularity requirements, while at the same time the loss remains computationally realizable. Based on reformulations as first-order systems of PDEs, we have considered formulations where the test space is an L_2 -space, as well as DPG methods based on ultra-weak formulations with L_2 -spaces as trial spaces. Such formulations are available for wide classes of problems beyond the examples considered here.

The common conceptual starting point is an ideal residual loss resulting from the reformulation of the family of parameter-dependent problems as a *single* well-posed variational problem. Analysing such losses, as limits of empirical versions based on parameter samples, involves integrals over the parameter domain, also in cases where trial and test spaces differ essentially under parameter variation. In order to still obtain a generally applicable notion of measurable parameter-dependent functions, leading to

expectations as limit loss, we resort to the concept of *direct integrals* of Hilbert spaces extending the notion of Bochner spaces. The concepts put forward here also provide a basis for treating parametric nonlinear PDEs, which will be a subject of further work.

To make the computation of these residual loss functions possible, in particular for the DPG-type methods, we have considered a hybrid approximation format combining finite element discretizations on the spatio-temporal domain with neural network approximation of the parameter dependence. Compared with a neural network approximation in all variables (as in standard PINN or operator learning approaches), such a neural network approximation of finite element basis coefficients as functions on the parameter domain has several practical advantages. However, it requires neural networks with large output dimension—as we have noted this imposes some strong requirements on possible sparse network architectures. With the particular architectures that we use, we may avoid a curse of dimensionality for large parametric dimensions, a point that requires further investigation.

Aspects that we have not considered here are the approximation of the ideal loss by the corresponding empirical risk as in (1.6), which depends also on the process of choosing the samples, and the related problem of estimating generalization errors. These questions will be addressed in more detail in future work.

Our numerical experiments are only intended to provide a proof of concept on comparably small-scale examples for the combination of variationally correct residual losses with hybrid approximations. In the test cases, due to the proportionality between residuals and error norms, we can assess the total solution error at each sampling point also in the absence of a reference solution, and thus, rather than resorting to manufactured solutions, we can choose examples reflecting the typical difficulties arising in each problem. As we have noted, optimization strategies commonly used in machine learning are not necessarily the most effective for the regression problems that need to be solved here, and we thus use quasi-Newton methods in our numerical tests. The further study of solvers is another direction of future work.

In summary, we can conclude that the proposed mechanisms can enable accuracy control in approximating solution manifolds of parameter-dependent problems using nonlinear model classes such as neural networks. This is of particular importance in the application of such surrogate models for the underlying solution operators in regression problems in state estimation and parameter identification problems.

Funding

Deutsche Forschungsgemeinschaft (DFG, German Research Foundation); National Science Foundation (DMS 2038080, DMS-2012469, DMS-2245097).

REFERENCES

- ADAMS, R. & FOURNIER, J. (2003) *Sobolev Spaces, Volume 140 of Pure and Applied Mathematics (Amsterdam)*, 2nd edn. Amsterdam: Elsevier/Academic Press.
- BACHMAYR, M. & COHEN, A. (2017) Kolmogorov widths and low-rank approximations of parametric elliptic PDEs. *Math. Comp.*, **86**, 701–724.
- BACHMAYR, M., COHEN, A. & DAHMEN, W. (2018) Parametric PDEs: sparse or low-rank approximations? *IMA J. Numer. Anal.*, **38**, 1661–1708.
- BAO, G., YE, X., ZANG, Y. & ZHOU, H. (2020) Numerical solution of inverse problems by weak adversarial networks. *Inverse Probl.*, **36**, 115003.
- BARRON, A., COHEN, A., DAHMEN, W. & DEVORE, R. (2008) Approximation and learning by greedy algorithms. *Ann. Stat.*, **3**, 64–94.

- BERRONE, S., CANUTO, C. & PINTORE, M. (2022) Variational physics informed neural networks: the role of quadratures and test functions. *J. Sci. Comput.*, **92**, 100.
- BERTOLUZZA, S., BURMAN, E. & HE, C. (2024) WAN discretization of PDEs: best approximation, stabilization and essential boundary conditions. *SIAM J. Sci. Comput.*, **46**, C688–C715.
- BINEV, P., COHEN, A., DAHMEN, W., DEVORE, R., PETROVA, G. & WOJASZCZYK, P. (2011) Convergence rates for greedy algorithms in reduced basis methods. *SIAM J. Math. Anal.*, **43**, 1457–1472.
- BONITO, A., DEVORE, R., PETROVA, G. & SIEGEL, J. W. (2025) Convergence and error control of consistent PINNs for elliptic PDEs. *IMA Journal of Numerical Analysis*, draf008.
- BRAESS, D. (1997) *Finite Elements: Theory, Fast Solvers, and Applications in Solid Mechanics*. Berlin, Heidelberg: Springer Berlin Heidelberg.
- BROERSEN, D., DAHMEN, W. & STEVENSON, R. (2018) On the stability of DPG formulations of transport equations. *Math. Comp.*, **87**, 1051–1082.
- CAI, Z., CHEN, J., LIU, M. & LIU, X. (2020) Deep least-squares methods: an unsupervised learning-based numerical method for solving elliptic PDEs. *J. Comput. Phys.*, **420**, 109707.
- CAI, Z., LAZAROV, R., MANTEUFFEL, T. A. & MCCORMICK, S. F. (1994) First-order system least squares for second-order partial differential equations: Part I. *SIAM J. Numer. Anal.*, **31**, 1785–1799.
- CARSTENSEN, C., DEMKOWICZ, L. & GOPALAKRISHNAN, J. (2016) Breaking spaces and forms for the DPG method and applications including Maxwell equations. *Comput. Math. Appl.*, **72**, 494–522.
- CHEGINI, N. & STEVENSON, R. (2011) Adaptive wavelet schemes for parabolic problems: sparse matrices and numerical results. *SIAM J. Numer. Anal.*, **49**, 182–212.
- CHEN, F., HUANG, J., WANG, C. & YANG, H. (2023) Friedrichs learning: weak solutions of partial differential equations via deep learning. *SIAM J. Sci. Comput.*, **45**, A1271–A1299.
- COHEN, A., DAHMEN, W., DEVORE, R. & NICHOLS, J. (2020) Reduced basis greedy selection using random training sets. *ESAIM: M2AN*, **54**, 1509–1524.
- COHEN, A., DAHMEN, W. & WELPER, G. (2012) Adaptivity and variational stabilization for convection-diffusion equations. *ESAIM Math. Model. Numer. Anal.*, **46**, 1247–1273.
- COHEN, A. & DEVORE, R. (2015) Approximation of high-dimensional parametric PDEs. *Acta Numer.*, **24**, 1–159.
- DAHMEN, W. (2023) Compositional sparsity, approximation classes, and parametric transport equations. *Constr. Approx.*, **61**, 219–283.
- DAHMEN, W., GRUBER, F. & MULA, O. (2019) An adaptive nested source term iteration for radiative transfer equations. *Math. Comp.*, **89**, 1605–1646.
- DAHMEN, W., HUANG, C., SCHWAB, C. & WELPER, G. (2012) Adaptive Petrov–Galerkin methods for first order transport equations. *SIAM J. Numer. Anal.*, **50**, 2420–2445.
- DAHMEN, W., MONSUUR, H. & STEVENSON, R. (2023) Least squares solvers for ill-posed PDEs that are conditionally stable. *ESAIM: M2AN*, **57**, 2227–2255.
- DAHMEN, W., PLESKEN, C. & WELPER, G. (2014) Double greedy algorithms: reduced basis methods for transport dominated problems. *ESAIM: M2AN*, **48**, 623–663.
- DAHMEN, W. & STEVENSON, R. (2019) Adaptive strategies for transport equations. *Comput. Methods Appl. Math.*, **19**, 431–464.
- DAHMEN, W., STEVENSON, R. & WESTERDIEP, J. (2021) Accuracy controlled data assimilation for parabolic problems. *Math. Comp.*, **91**, 557–595.
- DAHMEN, W., WANG, M. & WANG, Z. (2022) Nonlinear reduced DNN models for state estimation. *Commun. Comput. Phys.*, **32**, 1–40.
- DAL SANTO, N., DEPARIS, S. & PEGOLOTTI, L. (2020) Data driven approximation of parametrized PDEs by reduced basis and neural networks. *J. Comput. Phys.*, **416**, 109550.
- DAUTRAY, R. & LIONS, J.-L. (1985) *Analyse mathématique et Calcul numérique Pour les Sciences et les Techniques. Tome 3. Collection du Commissariat à L'énergie Atomique: Série Scientifique. [Collection of the Atomic Energy Commission: Science Series]*. Paris: Masson With the collaboration of Michel Artola, Claude Bardos, Michel Cessenat, Alain Kavenoky, Hélène Lanchon, Patrick Lascaux, Bertrand Mercier, Olivier Pironneau, Bruno Scheurer and Rémi Sentis.

- DEMOKOWICZ, L., GOPALAKRISHNAN, J., NAGARAJ, S. & SEPÚLVEDA, P. (2017) A spacetime DPG method for the Schrödinger equation. *SIAM J. Numer. Anal.*, **55**, 1740–1759.
- DiPERNA, R. & LIONS, P. L. (1989) Ordinary differential equations, transport theory and Sobolev spaces. *Invent. Math.*, **98**, 511–547.
- DONGKUN, Z., LU, L., GUO, L. & KARNIADAKIS, G. (2019) Quantifying total uncertainty in physics-informed neural networks for solving forward and inverse stochastic problems. *J. Comput. Phys.*, **397**, 108850.
- ERNST, L. & URBAN, K. (2024) A certified wavelet-based physics-informed neural network for the solution of parameterized partial differential equations. *IMA J. Numer. Anal.*, **45**, 494–515.
- FANASKOV, V., RUDIKOV, A. & OSELEDETS, I. (2024) Neural functional a posteriori error estimates. arXiv preprint arXiv:2402.05585.
- FÜHRER, T. & KARKULIK, M. (2021) Space-time least-squares finite elements for parabolic equations. *Comput. Math. Appl.*, **92**, 27–36.
- GANTNER, G. & STEVENSON, R. (2021) Further results on a space-time FOSLS formulation of parabolic PDEs. *ESAIM: M2AN*, **55**, 283–299.
- GEIST, M., PETERSEN, P., RASLAN, M., SCHNEIDER, R. & KUTYNIOK, G. (2021) Numerical solution of the parametric diffusion equation by deep neural networks. *J. Sci. Comput.*, **88**, Paper No. 22, 37.
- GOPALAKRISHNAN, J. & SEPÚLVEDA, P. (2019) *4. A Space-Time DPG Method for the Wave Equation in Multiple Dimensions*. Berlin, Boston: De Gruyter, pp. 117–140.
- GROHS, P. & VOIGTLAENDER, F. (2023) Proof of the theory-to-practice gap in deep learning via sampling complexity bounds for neural network approximation spaces. *Found. Comput. Math.*, **24**, 1085–1143.
- HEISS, C., GÜHRING, I. & EIGEL, M. (2023) Multilevel CNNs for parametric PDEs. *J. Mach. Learn. Res.*, **24**, 1–42.
- KAHLBACHER, M. & VOLKWEIN, S. (2007) Galerkin proper orthogonal decomposition methods for parameter dependent elliptic systems. *Discussiones Mathematicae: Differential Inclusions, Control and Optimization*, **27**, 95–117.
- KAUSHIK, B., HOSSEINI, B., KOVACHKI, N. & STUART, A. (2021) Model reduction and neural networks for parametric PDEs. *SMAI J. Comput. Math.*, **7**, 121–157.
- KHARAZMI, E., ZHANG, Z. & KARNIADAKIS, G. (2019) *Variational physics-informed neural networks for solving partial differential equations*.
- LI, Z., KOVACHKI, N., AZIZZADENESHELI, K., LIU, B., BHATTACHARYA, K., STUART, A. & ANANDKUMAR, A. (2021) Fourier neural operator for parametric partial differential equations. *International Conference on Learning Representations*.
- MISHRA, S. & MOLINARO, R. (2021) Physics informed neural networks for simulating radiative transfer. *J. Quant. Spectrosc. Radiat. Transfer*, **270**, 107705.
- MUZALEVSKY, A. V. & REPIN, S. I. (2021) A posteriori error control of approximate solutions to boundary value problems constructed by neural networks. *Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI)*, **499**, 77–104.
- RAISSI, M., PERDIKARIS, P. & KARNIADAKIS, G. E. (2019) Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.*, **378**, 686–707.
- ROZZA, G. (2014) *Fundamentals of Reduced Basis Method for Problems Governed by Parametrized PDEs and Applications*. Vienna: Springer Vienna, pp. 153–227.
- ROZZA, G., HUYNH, D. & PATERA, A. (2008) Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations. *Arch. Comput. Methods Eng.*, **15**, 229–275.
- DE RYCK, T. & MISHRA, S. (2022) Error analysis for physics-informed neural networks (PINNs) approximating Kolmogorov PDEs. *Adv. Comput. Math.*, **48**, 79.
- TAKESAKI, M. (1979) *Theory of Operator Algebras*. New York, NY: Springer New York.
- TAYLOR, J., PARDO, D. & MUGA, I. (2023) A deep Fourier residual method for solving PDEs using neural networks. *Comput. Methods Appl. Mech. Eng.*, **405**, 115850.
- VERFÜRTH, R. (2013) *A Posteriori Error Estimation Techniques for Finite Element Methods*. Oxford: Oxford University Press.
- VON NEUMANN, J. (1949) On rings of operators. Reduction theory. *Ann. Math.*, **50**, 401–485.

- ZANG, Y., BAO, G., YE, X. & ZHOU, H. (2020) Weak adversarial networks for high-dimensional partial differential equations. *J. Comput. Phys.*, **411**, 109409.
- ZUBOV, K., MCCARTHY, Z., MA, Y., CALISTO, F., PAGLIARINO, V., AZEGLIO, S., BOTTERO, L., LUJÁN, E., SULZER, V., BHARAMBE, A., VINCHHI, N., BALAKRISHNAN, K., UPADHYAY, D. & RACKAUCKAS, C. (2021) NeuralPDE: automating physics-informed neural networks (PINNs) with error approximations.

Appendix A. Proof of Lemma 3.10

The proof of Lemma 3.10 will be done in several steps whose proofs, in turn, use standard techniques concerning denseness results. For the convenience of the reader we present here a coherent version. Throughout this section we assume that Assumption 3.7 is valid.

The first step ensures denseness of smooth functions in $\mathbb{G}_p^*$ for elements that have support compactly contained in Ω . From (DiPerna & Lions, 1989, lemma II.1(i)) we obtain the following result.

LEMMA A.1. Let $\psi_\varepsilon(x) = \varepsilon^{-d} \psi(\varepsilon^{-1}x)$ with $\psi \in C^\infty(\mathbb{R}^d)$, $\text{supp } \psi \subseteq B_1$ (where in general B_τ denotes the ball of radius τ in $\mathbb{R}^d$) be a mollifier, and let $v \in \mathbb{G}_p^*$. Let Ω' be any subdomain with compact closure in Ω . Then, for $\varepsilon \leq \varepsilon(\Omega') := \text{dist}(\Omega', \partial\Omega)$ the convolution $\psi_\varepsilon * v$ is well defined on Ω' and belongs to $\mathbb{G}_p^*(\Omega') := \mathbb{G}_p^*|_{\Omega'}$, with the canonical understanding of $\|w\|_{\mathbb{G}_p^*(\Omega')}^2 := \|w\|_{L_2(\Omega')}^2 + \|\mathcal{B}_p^* w\|_{L_2(\Omega')}^2$. Moreover,

$$\|\mathcal{B}_p^*(\psi_\varepsilon * v - v)\|_{L_2(\Omega')} \rightarrow 0 \quad \text{as } \varepsilon \rightarrow 0. \quad (\text{A.1})$$

As a second step, we use that $\mathbb{V}_p \subset \mathbb{G}_p^*$, $p \in P$, and that we can construct a p -independent fundamental system of sections for $\prod_{p \in P} \mathbb{G}_p^*$ due to the fact that $H^1(\Omega)$ is dense in $\mathbb{G}_p^*$ for each p , as shown next.

LEMMA A.2. Assume that $\Omega \subset \mathbb{R}^d$ is a simply connected Lipschitz domain.

- (i) The embedding $H^1(\Omega) \subset \mathbb{G}_p^*$ is continuous, that is, there exists a $C < \infty$ such that for all $v \in H^1(\Omega)$ and every $p \in P$

$$\|v\|_{\mathbb{G}_p^*} \leq C \|v\|_{H^1(\Omega)}. \quad (\text{A.2})$$

- (ii) $H^1(\Omega)$ is dense in $\mathbb{G}_p^*$ for each $p \in P$.

Proof. When reference to Ω is clear we sometimes write briefly $\|\cdot\| := \|\cdot\|_{L_2(\Omega)}$. Regarding (i), observe that $\|\mathcal{B}_p^* v\|^2 \leq (\|c\|_\infty^2 + \|\text{div}(\mathbf{b})\|_\infty^2) \|v\|^2 + \|\mathbf{b}\|_\infty^2 \|\nabla v\|^2$. Choose $C = \max\{1, \|c\|_\infty^2 + \|\text{div}(\mathbf{b})\|_\infty^2, \|\mathbf{b}\|_\infty^2\}$ (note that since $\mathbf{b} \in W_\infty^1$, we have $\text{div}(\mathbf{b}(\cdot; p)) \in L_\infty(\Omega)$) and the claim follows. For (ii), we use an immediate adaptation of the standard argument given in (Adams & Fournier, 2003, theorem 3.17), based on the partition of unity from (Adams & Fournier, 2003, theorem 3.15). $\square$

We are now prepared to show Lemma 3.10 by constructing a fundamental system of sections for $\prod_{p \in P} \mathbb{V}_p$.

Proof. (Proof of Lemma 3.10) The first ingredient is a p -dependent cut-off function that reduces smooth functions to elements in $\mathbb{V}_p$. To that end, it is convenient to abbreviate in what follows the directional derivatives $D_p v := \mathbf{b}(\cdot; p) \cdot \nabla v$. Consider the boundary value problem

$$-D_p \phi(\cdot; p) = 1, \quad \text{in } \Omega, \quad \phi|_{\Gamma_{\text{out}}(p)} = 0. \quad (\text{A.3})$$

The solution has an explicit representation that we describe next. To this end consider the p -dependent family of characteristics $z(t) = z(t; x_0; p)$, $t \in \mathbb{R}$, defined by the initial value problems

$$\dot{z}(t) = -\mathbf{b}(z(t; x_0; p); p), \quad z(0, x_0; p) = x_0. \quad (\text{A.4})$$

The Lipschitz properties on the convection field in Assumption 3.7 ensure the unique existence of trajectories for all times.

As a next step we infer then measurability of the solution $\phi(\cdot; p)$ to (A.3) from these properties of the characteristics. These will be based on an explicit representation of $\phi(\cdot; p)$.

To describe the solution $\phi(\cdot; p)$ to (A.3) in terms of the characteristics, let the time of escape from a point $x \in \Omega$ along a characteristic path be

$$t(x; p) := \inf\{s > 0 : z(s; x; p) \notin \Omega\}, \quad x_+(x; p) = z(t(x; p); x; p) \in \Gamma_{\text{out}}(p). \quad (\text{A.5})$$

For a given $p \in P$ we can express any $x \in \Omega$ in terms of characteristic coordinates as

$$x = z(-t(x; p); x_+(x; p); p), \quad (\text{A.6})$$

that is, $x_+(x; p) \in \Gamma_{\text{out}}(p)$ is the intersection of the curve $\{z(t; x; p) : t \geq 0\}$, with $\Gamma_{\text{out}}(p)$. In these terms

$$\phi(x, p) := t(x; p) \quad (\text{A.7})$$

is the time spent on a characteristic path from x in Ω when travelling with constant speed one, i.e., the length of the corresponding characteristic segment. Hence,

$$D_p \phi(x; p) = \lim_{h \rightarrow 0} \frac{1}{h} (\phi(z(h; x; p); p) - \phi(x; p)) = \lim_{h \rightarrow 0} \frac{1}{h} (t(z(h; x; p); p) - t(x; p)) = 1,$$

which means that $\phi(x; p)$ solves (A.3) μ -a.e. in P . Thus, measurability of $\mathbf{b}(\cdot; p)$ implies measurability of $z(\cdot; \cdot; p)$, which in turn implies measurability of $t(\cdot; p)$, and hence that of $\phi(\cdot; p)$. Moreover, by the nondegeneracy assumption on the convection field, for any $x_0 \in \Gamma_{\text{out}}(p)$ we have

$$|\phi(z(-s; x_0; p); p)| \approx s, \quad s \leq t_{x_0}, \quad (\text{P.1})$$

where t_{x_0} is the maximal travel time when starting at $x_0 \in \Gamma_{\text{out}}(p)$. Hence, we get

$$\phi(\cdot, p)|_{\Gamma_{\text{out}}(p)} = 0.$$

Moreover, there is a constant M (depending on the convection field and Ω) such that

$$\|\phi(\cdot; p)\|_{L^\infty(\Omega)} \leq M, \quad \|D_p \phi(\cdot; p)\|_{L^\infty(\Omega)} \leq M. \quad (\text{P.2})$$

As a second step assume that $\{\zeta_n\}_{n \in \mathbb{N}}$ is total in $H^1(\Omega)$, and hence in $\mathbb{G}(\mathcal{B}_p^*)$ for every $p \in P$. To proceed let $C_{0,p}^\infty(\Omega) = \{v \in C^\infty(\Omega) : \text{supp } v \cap \Gamma_{\text{out}}(p) = \emptyset\}$. By the same reasoning as in the proof of Lemma A.2(ii) one sees that $C_{0,p}^\infty(\Omega)$ is dense in $\mathbb{V}_p$. Thus, it suffices to show that the system $\phi(\cdot; p)\zeta_n$, $n \in \mathbb{N}$, is dense in $C_{0,p}^\infty(\Omega)$. For any given $v \in C_{0,p}^\infty(\Omega)$ we have

$$v - \sum_{n \in \mathbb{N}} \alpha_n \phi(p) \zeta_n = \phi(\cdot; p) \left(\phi(\cdot; p)^{-1} v - \sum_{n \in \mathbb{N}} \alpha_n \zeta_n \right).$$

Thus, by (P2),

$$\begin{aligned}
\left\| D_p \left(v - \sum_{n \in \mathbb{N}} \alpha_n \phi(\cdot; p) \zeta_n \right) \right\|_{L_2(\Omega)} &\leq \left\| (D_p \phi(\cdot; p)) \left(\phi(\cdot; p)^{-1} v - \sum_{n \in \mathbb{N}} \alpha_n \zeta_n \right) \right\|_{L_2(\Omega)} \\
&\quad + \left\| \phi(\cdot; p) D_p \left(\phi(\cdot; p)^{-1} v - \sum_{n \in \mathbb{N}} \alpha_n \zeta_n \right) \right\|_{L_2(\Omega)} \\
&\leq M \left(\left\| \left(\phi(\cdot; p)^{-1} v - \sum_{n \in \mathbb{N}} \alpha_n \zeta_n \right) \right\|_{L_2(\Omega)} \right. \\
&\quad \left. + \left\| D_p \left(\phi(\cdot; p)^{-1} v - \sum_{n \in \mathbb{N}} \alpha_n \zeta_n \right) \right\|_{L_2(\Omega)} \right).
\end{aligned}$$

Now we will show that $\phi(\cdot, p)^{-1} v \in \mathbb{G}_p^*$. For $v \in C_{0,p}^\infty(\Omega)$ there is a $\delta > 0$ and neighbourhood $\mathcal{N}_\delta$ of width δ such that $\phi(\cdot; p)^{-1} v$ remains bounded in $\Omega \setminus \mathcal{N}_\delta$ because by (P1) $\|\phi(\cdot; p)^{-1}\|_{L_\infty(\Omega \setminus \mathcal{N}_\delta)} \lesssim \delta^{-1}$. Since $\phi(\cdot; p)$ solves (A.3) and $|\phi(\cdot; p)|_{\Omega \setminus \mathcal{N}_\delta}$ remains lower bounded by a constant times δ^{-1} we get by product rule that $\|D_p \phi(\cdot; p)^{-1}\|_{L_\infty(\Omega \setminus \mathcal{N}_\delta)} < \infty$. Therefore, $\phi^{-1}(\cdot; p)v \in \mathbb{G}_p^*$ for each $p \in P$. As ξ_n are chosen to be total in $\mathbb{G}_p^*$ for all $p \in \mathbb{G}_p^*$ we can choose the α_n to make the term $\left\| D_p \left(\phi(\cdot; p)^{-1} v - \sum_{n \in \mathbb{N}} \alpha_n \zeta_n \right) \right\|_{L_2(\Omega)}$ as small as we wish, and thus $\phi(\cdot, p)\zeta_n$ is total in $\mathbb{V}_p$.

Since the ζ_n are independent of p and in $H^1(\Omega)$ and ϕ is jointly measurable in x and p , the functions $\xi_n = \phi \zeta_n$ are measurable on $\Omega \times P$. Now consider the mapping

$$p \mapsto \langle \xi_n, \xi_m \rangle_{\mathbb{V}_p} = \langle \mathcal{B}_p^* \xi_n, \mathcal{B}_p^* \xi_m \rangle_{L_2(\Omega)},$$

which is measurable since by product rule the map

$$(x, p) \mapsto \mathcal{B}_p^* \xi_n(p, x) = \operatorname{div}(\mathbf{b}(x; p))\phi(x, p)\zeta_n(x) + \mathbf{b}(x, p) \cdot (\zeta_n(x)\nabla\phi(x, p) + \phi(x, p)\nabla\zeta_n(x))$$

is measurable by assumption on $\mathbf{b}$ and the measurability of $\phi, \zeta_n, D_p \phi$ and $D_p \zeta_n$. Hence, we have found a fundamental sequence of μ -measurable sections in $\prod_{p \in P} \mathbb{V}_p$ so that Theorem 2.2 applies and confirms that $\mathbb{Y} = L_2\left(P, (\mathbb{V}_p)_{p \in P}\right)$ is a well-defined direct integral with norm $\|v\|_{\mathbb{Y}}^2 = \int_P \|v\|_{\mathbb{V}_p}^2 d\mu(p)$.

Moreover, $\|\cdot\|_{L_2(\Omega)} \leq C\|\cdot\|_{\mathbb{V}_p}$ (see the comment following Assumption 3.7) implies that $\mathbb{Y}$ is continuously embedded in $L_2(P \times \Omega) = L_2(P, L_2(\Omega))$, and, as a Hilbert space, is closed. Here, it is understood that $P \times \Omega$ is viewed as measure space with the product measure whose marginals are μ and the Lebesgue measure, respectively. This completes the proof of Lemma 3.10. $\square$

Appendix B. Additional remarks on direct integrals

EXAMPLE B1. Let Γ_1, Γ_2 be two distinct parts of $\Gamma = \partial\Omega$ for some bounded, simply connected domain $\Omega \subset \mathbb{R}^d$. Define the space

$$H_{\Gamma_i} = \{v \in H^1(\Omega) : v|_{\Gamma_i} = 0\}.$$

Let $P \subset \mathbb{R}^{d_p}$ be measurable and consider $P^1 \subset P$ nonmeasurable such that all measurable subsets have vanishing measure, such as a Vitali set. Note that $P^2 = P \setminus P^1$ is then also nonmeasurable. Define

$$\mathbb{W}_p = \begin{cases} H_{\Gamma_1} & \text{if } p \in P^1, \\ H_{\Gamma_2} & \text{else.} \end{cases}$$

Let $(\xi_n(p))_{n \in \mathbb{N}}$ be an orthonormal basis in $\mathbb{W}_p$. Observe that we can choose $\xi_n(p) = \xi_n^{1,2}$ to be constant when restricted to P^1 and P^2 , respectively.

Since by construction

$$p \mapsto \langle \xi_n(p), \xi_m(p) \rangle_{H^1(\Omega)} = \delta_{m,n},$$

$(\xi_n)_{n \in \mathbb{N}}$ is a μ -measurable sequence so that the direct integral is well defined. Observe, $v \in \prod_{p \in P} \mathbb{W}_p$, with $v(p) = \sum_{n \in \mathbb{N}} v_n(p) \xi_n(p)$ is μ -measurable if and only if for all $n \in \mathbb{N}$ the coefficient map

$$p \mapsto v_n(p)$$

is measurable.

Choose $v \in L_2(P, (\mathbb{W}_p)_{p \in P})$ as $v(p) = \xi_1(p)$, i.e., $v_1(p) = 1$ and $v_n(p) = 0$ for all $n \geq 2$ and $p \in P$. Then, for all $n \in \mathbb{N}$ the coefficient map is constant, and hence measurable and square integrable. Since $\bigcap_{p \in P} \mathbb{W}_p$ is not dense in any $\mathbb{W}_p$, we can assume (after reordering of the basis $\xi_n^{1,2}$) that $\xi_1^{1,2} \notin \bigcap_{p \in P} \mathbb{W}_p$, and thus for all $p \in P$ we have $v(p) \notin \bigcap_{p \in P} \mathbb{W}_p$. Furthermore, let $\tilde{\xi}_n$ be an orthonormal basis of $H^1(\Omega)$ with $\tilde{\xi}_1 = \xi_1(p)$ for some $p \in P^1$, and thus a fundamental system of $\prod_{p \in P} H^1(\Omega)$.

Then, we have for all $p \in P^1$ that

$$\langle v(p), \tilde{\xi}_1 \rangle_{H^1(\Omega)} = \langle \xi_1(p), \tilde{\xi}_1 \rangle_{H^1(\Omega)} = \langle \tilde{\xi}_1, \tilde{\xi}_1 \rangle_{H^1(\Omega)} = 1$$

as well as for all $\tilde{p} \in P^2$ we get

$$\langle v(\tilde{p}), \tilde{\xi}_1 \rangle = \langle \xi_1(\tilde{p}), \tilde{\xi}_1 \rangle_{H^1(\Omega)} < 1$$

as $\xi_1(p) \notin H_{\Gamma_2}$ for $p \in P^1$, and thus $\langle \xi_1(p_1), \xi_1(p_2) \rangle < 1$ for $p_1 \in P^1$ and $p_2 \in P^2$ (recall that $\xi(p)$ is an orthonormal basis for all $p \in P$). Then,

$$\Psi : p \mapsto \langle v(p), \tilde{\xi}_1 \rangle$$

is not measurable since $\Psi^{-1}(\{1\}) = P^1$, and thus $v \notin L_2(P, (H^1(\Omega))_{p \in P})$.

However, if there is a fundamental sequence in $\prod_{p \in P} \mathbb{V}_p$ that is measurable in $\prod_{p \in P} \mathbb{W}_p$, we have the subspace relations we expect.

PROPOSITION B.1. Let $\mathbb{W}_p$ be a separable Hilbert space such that $\bigcap_{p \in P} \mathbb{W}_p$ is dense in each $\mathbb{W}_p$, and equip it for each p with the same fundamental sequence $(\tilde{\xi}_n)_{n \in \mathbb{N}}$ with $\tilde{\xi}_n \in \bigcap_{p \in P} \mathbb{W}_p$ for all $n \in \mathbb{N}$. Consider closed subspaces $\mathbb{V}_p$ endowed with the scalar product of $\mathbb{W}_p$. Assume that there is a fundamental sequence $(\xi_n)_{n \in \mathbb{N}}$ of $\prod_{p \in P} \mathbb{V}_p$ that is μ -measurable with respect to $(\tilde{\xi}_n)_{n \in \mathbb{N}}$. Then, we have

$$L_2(P, (\mathbb{V}_p)_{p \in P}) \subset L_2(P, (\mathbb{W}_p)_{p \in P}) \quad \text{as a closed subspace.} \quad (\text{B.1})$$

Conversely, if (B.1) holds then the fundamental sequence of $\prod_{p \in P} \mathbb{V}_p$ is μ -measurable in $\prod_{p \in P} \mathbb{W}_p$.

Proof. By definition we have $\prod_{p \in P} \mathbb{V}_p \subset \prod_{p \in P} \mathbb{W}_p$ as sets. By (Takesaki, 1979, lemma 8.12) we can assume that $(\xi_n(p))_{n \in \mathbb{N}}$ forms an orthonormal basis in $\mathbb{V}_p$ for all $p \in P$, and thus μ -measurability in $\prod_{p \in P} \mathbb{V}_p$ reduces to measurability of the coefficient map. As the map

$$p \mapsto \langle \xi_n(p), \tilde{\xi}_m \rangle_{\mathbb{W}_p}$$

is measurable for all $n, m \in \mathbb{N}$, we have that any $v \in \prod_{p \in P}$ that is μ -measurable with respect to ξ_n is also μ -measurable with respect to $\tilde{\xi}_n$ as the coefficient maps are measurable. As the fibre spaces share the same bilinear form and as both Hilbert integrals are complete, we get by definition

$$L_2(P, (\mathbb{V}_p)_{p \in P}) \subset L_2(P, (\mathbb{W}_p)_{p \in P})$$

as a closed subspace.

On the other hand, if $L_2(P, (\mathbb{V}_p)_{p \in P}) \subset L_2(P, (\mathbb{W}_p)_{p \in P})$ we get by definition of the Hilbert integral that

$$p \mapsto \langle \xi_n(p), \tilde{\xi}_m \rangle_{\mathbb{W}_p}$$

has to be measurable. □

Appendix C. Computing the DPG loss terms

A generally applicable way to compute the quantities $\mathfrak{T}_K^\delta(\mathfrak{w}^\delta(\cdot, p; \theta))$ in (3.40) is, in view of (3.35), to solve for each $K \in \mathcal{T}$ a Galerkin problem of uniformly bounded (small) size. In fact, recall that this realizes the Riesz lift to $\tilde{\mathbb{V}}^\delta$ of the local residual $\mathcal{F}_K(\mathfrak{w}^\delta)$, defined by (see (3.30))

$$\mathcal{F}_K(\mathfrak{w}^\delta)(v^\delta) := B_K(\mathfrak{w}^\delta, v^\delta) - f(v^\delta), \quad v^\delta \in \tilde{\mathbb{V}}_K^\delta.$$

We say ‘generally applicable’ because in specific cases there could be alternatives based on explicitly determining ‘near optimal test functions’ as is the case for linear transport, see Broersen *et al.* (2018) and Dahmen & Stevenson (2019). Nevertheless, quantitatively best stability properties are ensured by Galerkin solves because they provide the best test functions in the given test search space $\tilde{\mathbb{V}}_K^\delta$. Recall also that, once the homework of establishing uniform stability of the DPG scheme is done, the test search spaces $\tilde{\mathbb{V}}_K^\delta$ are indeed independent of p, θ . The precise form of $\tilde{\mathbb{V}}^\delta$ depends on the underlying PDE model. In some instances, it has been shown that it suffices to take $\tilde{\mathbb{V}}^\delta$ as a space of polynomials of some fixed degree m_v over K , to warrant uniform inf-sup stability. In other cases, e.g., when dealing with transport equations, $\tilde{\mathbb{V}}_K^\delta$ can be chosen as a space of *piecewise polynomials* over K with some fixed polynomial degree and some fixed *subgrid depth*. In either case we can fix a basis $\Psi_K := \{\psi_{K,i} : i = 1, \dots, n_K\}$ for $\tilde{\mathbb{V}}_K^\delta$, supported in K , where n_K typically exceeds the number of degrees of freedom in $\mathfrak{w}^\delta|_K$. Finally, since the ultra-weak formulation is based on the optimal test norm $\|\mathcal{B}_p^* v\|_{L_2(\Omega)}$, the local test-inner products are given by

$$\langle v, z \rangle_{\mathbb{V}_{K,p}} := \langle \mathcal{B}_{K,p}^* v, \mathcal{B}_{K,p}^* z \rangle_K,$$

where

$$(\mathcal{B}_{K,p}^* v^\delta)(\mathfrak{w}^\delta) = B_K(\mathfrak{w}^\delta, v^\delta; p), \quad \mathfrak{w}^\delta \in \mathbb{U}_{\mathcal{T}}^\delta, v^\delta \in \tilde{\mathbb{V}}_K^\delta.$$

For instance, in the special case of the transport equation (3.21), we recall from (Dahmen *et al.*, 2021, (4.2)) that for $\mathbf{w} = (w_b, w_s) \in \mathbb{U}_{\mathcal{T}}$ (see (3.29))

$$\begin{aligned} B_K(\mathbf{w}, v; p) &= \int_K (cv - \operatorname{div}(\mathbf{b}(p)v))w_b \, dx + \int_{\partial K} \mathbf{b}(p) \cdot \mathbf{n}_K v w_s \, d\Gamma \\ &= \int_K (c - \operatorname{div} \mathbf{b}(p))v w_b + (w_s - w_b)\mathbf{b}(p) \cdot \nabla v + v\mathbf{b}(p) \cdot \nabla w_s + v w_s \operatorname{div} \mathbf{b}(p) \, dx, \end{aligned} \quad (\text{C.1})$$

where the last equality shows how to replace the skeleton components by functions on the domain Ω .

In these terms our goal is to determine $\mathbf{r} = (r_1, \dots, r_{n_K})^\top = \mathbf{r}(p, \theta)$ such that

$$\mathfrak{T}_{\mathcal{F}_K}^\delta(\mathbf{w}^\delta(\cdot, p; \theta)) = \mathbf{r}(p, \theta)^\top \Psi_K = \sum_{i=1}^{n_K} r_i(p, \theta) \psi_{K,i} \in \widetilde{\mathbb{V}}_K^\delta, \quad (\text{C.2})$$

and

$$\|\mathfrak{T}_{\mathcal{F}_K}^\delta(\mathbf{w}^\delta(\cdot, p; \theta))\|_{\widetilde{\mathbb{V}}_K^\delta}^2 = \mathbf{r}(p, \theta)^\top \mathbf{M}_{K,p} \mathbf{r}(p, \theta), \quad (\text{C.3})$$

where

$$\mathbf{M}_{K,p} := \langle \Psi_K, \Psi_K \rangle_{\mathbb{V}_{K,p}} = \left(\langle \psi_{K,i}, \psi_{K,j} \rangle_{K,p} \right)_{i,j=1}^{n_K} = \left(\langle \mathcal{B}_{K,p}^* \psi_{K,i}, \mathcal{B}_{K,p}^* \psi_{K,j} \rangle_K \right)_{i,j=1}^{n_K}.$$

The vector $\mathbf{r}(p, \theta)$ in turn solves the local Galerkin system

$$\mathbf{M}_{K,p} \mathbf{r} = \mathbf{F}_K(p, \theta), \quad (\text{C.4})$$

where the entries of the load vector $\mathbf{F}_K(p, \theta) \in \mathbb{R}^{n_K}$ are the coefficients

$$\begin{aligned} F_{K,i}(p, \theta) &:= \mathfrak{T}_{K,p}(\mathbf{w}^\delta(\cdot, p; \theta))(\psi_{K,i}) \\ &= B_K(\mathbf{w}^\delta(\cdot, p; \theta), \psi_{K,i}; p) - f(\psi_{K,i}; p), \quad i = 1, \dots, n_K. \end{aligned} \quad (\text{C.5})$$

Hence, (C.3) can be rewritten as

$$\|\mathfrak{T}_{\mathcal{F}_K}^\delta(\mathbf{w}^\delta(\cdot, p; \theta))\|_{\widetilde{\mathbb{V}}_K^\delta}^2 = \mathbf{F}_K(p, \theta)^\top \mathbf{M}_{K,p}^{-1} \mathbf{F}_K(p, \theta), \quad (\text{C.6})$$

showing how the numerator in (5.2) is composed in this case.

Note that the dependence on the weights θ is solely through the load vector $\mathbf{F}_K(p, \theta)$. Since optimization over θ requires frequent computations of $\mathbf{r}(p, \theta)$, it is worth computing for each training point $p \in \widehat{P}$ the inverse $\mathbf{M}_{K,p}^{-1}$ (or its Cholesky factorization). An efficient implementation would attempt to precompute possibly many terms in the inner products $\langle \mathcal{B}_{K,p}^* \psi_{K,i}, \mathcal{B}_{K,p}^* \psi_{K,j} \rangle_K$, which depends on the particular PDE at hand. For instance, in the case of (C.1), this involves, in particular, evaluating quantities like

$$\int_K (c\psi_{K,i} - \operatorname{div}(\mathbf{b}(p)\psi_{K,i}))(c\psi_{K,j} - \operatorname{div}(\mathbf{b}(p)\psi_{K,j})) \, dx.$$

For instance, when $\mathbf{b}(x, p) = \sum_{k=1}^{d_p} p_k b_k(x)$ or when $\mathbf{b}$ does not depend on x , $\mathbf{M}_{K,p}$ can be efficiently assembled for different p . Analogous remarks apply to assembling $\mathbf{F}_K(p, \theta)$, where now an additional dependence on θ comes into play. If we assume for simplicity that, as in the numerical example, f in (C.5) does not depend on p , the vectors $\mathbf{f}_K = (f(\psi_{K,1}), \dots, f(\psi_{K,n_K}))^\top \in \mathbb{R}^{n_K}$ can be precomputed for

a given $\mathcal{T}$. The remaining part $\mathbf{B}_K(p, \theta) := \mathbf{F}_K(p, \theta)$ has, in view of (C.5), the form

$$\mathbf{B}_K(p, \theta) = \langle \mathcal{B}_{K,p}^* \Psi_K, \Phi_K \rangle_K \mathbf{c}_K(p; \theta). \quad (\text{C.7})$$

Here, $\mathbf{c}_K(p; \theta) = \begin{pmatrix} \mathbf{c}_{K,b} \\ \mathbf{c}_{K,s} \end{pmatrix}$ is composed of the coefficient vector for the (polynomial) bulk-basis functions $\phi_{K,i}$, and of the coefficient vector for skeleton basis functions $\varphi_{K,s,i}$ appearing in the representation (see (3.29))

$$\begin{aligned} w_{\mathcal{T},b}^\delta(x, p; \theta^b) &= \sum_{K \in \mathcal{T}} \sum_{i \in \mathcal{I}_K} c_{K,b,i}(p; \theta^b) \phi_{K,i}(x), \\ w_{\mathcal{T},s}^\delta(x, p; \theta^s) &= \sum_{v \in \mathcal{V}(\mathcal{T})} \sum_{i \in \mathcal{I}_v} c_{v,s,i}(p; \theta^s) \varphi_{v,i}(x), \end{aligned} \quad (\text{C.8})$$

where $\mathcal{V}(\partial \mathcal{T})$ is the set of nodes in the skeleton $\partial \mathcal{T}$. The sub-array $\mathbf{c}_{K,s}$ corresponds to those basis functions $\varphi_{K,s,i}$ whose supports overlap ∂K . Finally, the assembly of the matrix $\langle \mathcal{B}_{K,p}^* \Psi_K, \Phi_K \rangle_K = (\langle \mathcal{B}_{K,p}^* \psi_{K,i}, \phi_{K,\ell} \rangle_K)$ with $\phi_{K,\ell} \in \{\phi_{K,b,i} : i \in \mathcal{I}_K\} \cup \{\varphi_{K,s,i} : i \in \mathcal{V}(\partial \mathcal{T}), |\text{supp } \varphi_{K,s,i} \cap \partial K| > 0\}$ is analogous to the assembly of $\mathbf{M}_{K,p}$.

In summary, the cost of the computation of each loss function summand in (3.40) for a given p and θ remains *uniformly bounded*. However, the bounding constant depends algebraically on the degrees in the trial and test functions.