

Neural Q-learning for solving PDEs

Samuel N. Cohen

*Mathematical Institute, University of Oxford
Oxford, OX2 6GG, UK*

COHENS@MATHS.OX.AC.UK

Deqing Jiang

*Mathematical Institute, University of Oxford
Oxford, OX2 6GG, UK*

DEQING.JIANG@MATHS.OX.AC.UK

Justin Sirignano

*Mathematical Institute, University of Oxford
Oxford, OX2 6GG, UK*

SIRIGNANO@MATHS.OX.AC.UK

Editor: Michael Mahoney

Abstract

Solving high-dimensional partial differential equations (PDEs) is a major challenge in scientific computing. We develop a new numerical method for solving elliptic-type PDEs by adapting the Q-learning algorithm in reinforcement learning. To solve PDEs with Dirichlet boundary condition, our “Q-PDE” algorithm is mesh-free and therefore has the potential to overcome the curse of dimensionality. Using a neural tangent kernel (NTK) approach, we prove that the neural network approximator for the PDE solution, trained with the Q-PDE algorithm, converges to the trajectory of an infinite-dimensional ordinary differential equation (ODE) as the number of hidden units $\rightarrow \infty$. For monotone PDEs (i.e. those given by monotone operators, which may be nonlinear), despite the lack of a spectral gap in the NTK, we then prove that the limit neural network, which satisfies the infinite-dimensional ODE, strongly converges in L^2 to the PDE solution as the training time $\rightarrow \infty$. More generally, we can prove that any fixed point of the wide-network limit for the Q-PDE algorithm is a solution of the PDE (not necessarily under the monotone condition). The numerical performance of the Q-PDE algorithm is studied for several elliptic PDEs.

Keywords: Deep learning, neural networks, high-dimensional PDEs, high-dimensional learning, Q-learning.

1. Introduction

High dimensional partial differential equations (PDEs) are widely used in many applications in physics, engineering, and finance. It is challenging to numerically solve high-dimensional PDEs, as traditional finite difference methods become computationally intractable due to the curse of dimensionality. In the past decade, deep learning has become a revolutionary tool for a number of different areas including image recognition, natural language processing, and scientific computing. The idea of solving PDEs with deep neural networks has been rapidly developed in recent years and has achieved promising performance in solving real-world problems (e.g. Hu et al. (2022), Li et al. (2022), Cai et al. (2021), and Misyris et al. (2020)).

A number of deep-learning-based PDE solving algorithms, for instance the deep Galerkin method (DGM; Sirignano and Spiliopoulos (2018)) and physics-informed neural networks (PINNs) (Raissi et al. (2019)), have been proposed. Inspired by the Q-learning method in reinforcement learning, we propose a new “Q-PDE” algorithm which approximates the solution of PDEs with an artificial neural network. In this paper, we prove that for monotone PDEs the Q-PDE training process gives an approximation which converges to the solution of a certain limit equation as the number of hidden units in a single-layer neural network goes to infinity. Furthermore, we prove that the limit neural network converges strongly to the solution of the PDE that the algorithm is trying to solve.

Q-learning (Watkins and Dayan (1992)) is a well-known algorithm for computing the value function of the optimal policy in reinforcement learning. Deep Q-learning (DQN) (Mnih et al. (2013)), a neural-network-based Q-learning algorithm, has successfully learned to play Atari games (Mnih et al. (2015)) and subsequently has become widely-used for many applications (Zhao et al. (2019)). The Q-learning algorithm updates its value function approximator following a biased gradient estimate computed from input data samples. We propose an algorithm which is similar to Q-learning in the sense that we update the parameters of a neural network via a biased gradient flow in continuous time. As far as the authors are aware, this is the first time that a Q-learning algorithm has been developed to solve PDEs directly.

The universal approximation theorem (Cybenko (1989)) indicates that the family of single-layer neural networks are powerful function approximators. However, the universal approximation theorem only states that there exists a neural network which can approximate continuous functions arbitrarily well; it does not suggest how to identify the parameters for such a neural network. When the number of units in our neural network becomes large, it is, however, possible to obtain asymptotic limits, for example the “neural tangent kernel” limit (Jacot et al. (2018)), which gives a variant of the law of large numbers for infinitely wide neural networks. We will use this approach to study the performance of the biased gradient flow as a training algorithm, and show that the wide-network limit satisfies an infinite-dimensional ordinary differential equation (ODE).

We apply our Q-PDE approach to second order nonlinear PDEs with Dirichlet boundary conditions. In particular, we are able to give strong convergence results when the differential operator satisfies a strong monotonicity condition. Monotone PDEs arise in various applications, particularly in PDEs arising from stochastic modeling – the generators of ergodic stochastic processes are monotone (when evaluated with their stationary distributions), which suggests a variety of possible applications of our approach. Further, the subdifferentials of convex functionals (i.e. of maps from functions to the real line) are monotone; this suggests that monotone PDEs may be a particularly well-suited class of equations for gradient methods, as they correspond to (a generalization of) minimizations of convex functionals. We will see that, given this monotonicity assumption, we can prove that the limit Q-PDE algorithm will converge (strongly in L^2) to the solution of the monotone PDE. More generally, we can prove that any fixed point of the wide-network limit for the Q-PDE algorithm is a solution of the PDE (not necessarily under the monotone condition).

In the remainder of the introduction, we provide a brief survey of relevant literature and some related approaches. The necessary properties of the PDEs which we study, the

architecture of the neural networks we use, and the Q-PDE training algorithm we propose are presented in Section 2. Section 3 derives useful properties of the limiting system, while Section 4 contains a rigorous proof that the training process converges to this limit. Analysis of the limit neural network, in particular the proof of convergence to the solution of the PDE, is presented in Section 5. Numerical results are presented in Section 6.

1.1 Summary

In summary, this paper will provide the following:

- A new neural network training algorithm for second order PDEs (of the form given by Assumption 13), based on Q-learning (Algorithm 20), and designed to ensure Dirichlet boundary conditions are enforced.
- Proofs of various functional-analytic properties of the neural tangent kernel limit of this algorithm (Section 3).
- A rigorous convergence proof of the neural-tangent kernel limit for a wide single-layer network (Section 4). This limit gives a simple deterministic dynamical system characterizing the neural network approximation during training (Definition 43).
- Proof that any fixed point of the limit neural network training dynamics is a solution of the PDE (Theorem 46).
- A strong convergence proof for the limit neural network as the training time $t \rightarrow \infty$, when considering a PDE given by a monotone operator (Section 5).
- Numerical examples to demonstrate effectiveness of the training algorithm (Section 6).

A more precise summary of our mathematical results is given in Section 2.3.1.

1.2 Relevant literature

The idea of solving PDEs with artificial neural networks has been rapidly developed in recent years. Various approaches have been proposed: Lagaris et al. (1998), Lagaris et al. (2000), Lee and Kang (1990), and Malek and Beidokhti (2006) applied neural networks to solve differential equations on an a priori fixed mesh. However, the curse of dimensionality shows that these approaches cannot be extended to high dimensional cases. In contrast, as a meshfree method, DGM (Sirignano and Spiliopoulos (2018)) randomly samples data points in each training step. Based on this random grid, it derives an unbiased estimator of the error of the approximation using the differential operator and the boundary condition of the PDE and then iteratively updates the neural network parameters with stochastic gradient descent. This approach has been successful at solving certain high-dimensional PDEs including free-boundary PDEs, Hamilton–Jacobi–Bellman equations, and Burger’s equation. Unlike DGM, our algorithm is inspired by Q-learning, in the sense that we compute a biased gradient estimator as the search direction for parameter training. The biased gradient estimator does not require taking a derivative of the differential operators of the neural network; only the gradient of the neural network model itself is required.

The simplicity of this gradient estimator, together with the monotonicity of the PDEs we consider, allows a proof of the convergence of the Q-PDE training algorithm to the solution of the PDE.

Another related area of research is PINNs (Raissi et al. (2019)), which train neural networks to merge observed data with PDEs. This enables its users to make use of a priori knowledge of the governing equations in physics. Analysis of solving second order elliptic/parabolic PDEs with PINNs can be found in Shin et al. (2020). DeepONet (Lu et al. (2021)) proposes solving a family of PDEs in parallel. Its architecture is split into two parallel networks: a branch net, which approximates functions related to input PDEs, and a trunk net that maps spatial coordinates to possible functions. Combining these two networks, it is possible to learn the solution of a PDE given enough training data from traditional solvers. Mathematical analysis of DeepONet in detail can be found in Lanthaler et al. (2021). DeepSets (Germain et al. (2021a)) are designed to solve a class of PDEs which are invariant to permutations. It computes simultaneously an approximation of the solution and its gradient. For solving fully nonlinear PDE, Pham et al. (2021) estimates the solution and its gradient simultaneously by backward time induction. For linear PDEs, MOD-Net (Zhang et al. (2021)) uses a DNN to parameterize the Green’s function and approximate the solution.

Further approaches, which go beyond looking for an approximation of the strong solution of the PDE directly, have also been studied. Zang et al. (2020) proposes using generative adversarial networks (GANs) to solve a min-max optimization problem to approximate the weak solution of the PDE. The Fourier neural operator (FNO) (Li et al. (2020)) approach makes use of a Fourier transform and learns to solve the PDE via convolution. The solution of certain PDEs can be expressed in terms of stochastic processes using the Feynman–Kac formula. Grohs and Herrmann (2020) apply neural networks to learn the solution of a PDE via its probabilistic representation and numerically demonstrates the approach for Poisson equations. Based on an analogy between the BSDE and reinforcement learning, E et al. (2017) proposes an algorithm to solve parabolic semilinear PDEs and the corresponding backward stochastic differential equations (BSDEs), where the loss function gives the error between the prescribed terminal condition and the solution of the BSDE. PDE-net and its variants (Long et al. (2018), Long et al. (2019)) attempt to solve inverse PDE problems: a neural network is trained to infer the governing PDE given physical data from sensors. In a recent paper Sirignano et al. (2021), neural network terms are introduced to optimize PDE-constrained models. The parameters in the PDE are trained using gradient descent, where the gradient is evaluated by an adjoint PDE. In Ito et al. (2021), a class of numerical schemes for solving semilinear Hamilton–Jacobi–Bellman–Isaacs (HJBI) boundary value problems is proposed. By policy iteration, the semilinear problem is reduced into a sequence of linear Dirichlet problems. They are subsequently approximated by a feedforward neural network ansatz. For a detailed overview of deep learning for solving PDEs, we refer the reader to E et al. (2022), Beck et al. (2020) and Germain et al. (2021b), and references therein.

Many of these works do not directly study the process of training the neural network. Gradient descent-based training of neural networks has been mathematically analyzed using the neural tangent kernel (NTK) approach (Jacot et al. (2018)). The NTK analysis characterizes the evolution of the neural network during training in the regime when the number of hidden units is large. For instance, Jacot et al. (2018), Lee et al. (2017), and Lee

et al. (2019) study the NTK limit for classical regression problems where a neural network is trained to predict target data given an input. Using an NTK approach, Wang et al. (2020) gives a possible explanation as to why sometimes PINNs fail to solve certain PDEs. Their analysis highlights the difficulty in approximating highly oscillatory functions using neural networks, due to the lack of a spectral gap in the NTK. The NTK approach is not without criticism, particularly when applied to deeper neural networks, where it has been seen to be unable to accurately describe observed performance (see, for example, Ghorbani et al. (2020) or Chizat et al. (2019)). Alternative approaches include mean-field analysis (e.g. Mei et al. (2019)).

In this article, we train a neural network to solve a PDE, using a different training algorithm, for which we give a rigorous proof of the convergence to a NTK limiting regime (as the width of the neural network increases). This NTK approach yields a particularly simple dynamic under our training algorithm, allowing us to prove convergence to the solution of the PDE (as training time increases), despite the poorly behaved spectral properties of the NTK, for those PDEs given by monotone operators. Our Q-PDE algorithm is substantially different to the gradient descent algorithm studied in Jacot et al. (2018), Lee et al. (2017), Lee et al. (2019), since a PDE operator will be directly applied to the neural network to evaluate the parameter updates.

2. The Q-PDE Algorithm

In this section, we first of all state our assumptions on the PDE and its domain. Then we describe the neural network approximator and our Q-PDE algorithm.

2.1 Assumptions on the PDE

We consider a class of second-order nonlinear PDEs on a bounded open domain $\Omega \subset \mathbb{R}^n$ with Dirichlet boundary conditions:

$$\begin{cases} \mathcal{L}u = 0 & \text{in } \Omega \\ u = f & \text{on } \partial\Omega. \end{cases} \quad (1)$$

We assume a measure ¹ μ is given on Ω , with continuous Radon–Nikodym density $d\mu/d\text{Leb}$ bounded away from 0 and ∞ (where Leb denotes Lebesgue measure). It follows that $\mu(\Omega) < \infty$ and $\mu(A) > 0$ for all open sets $A \subset \Omega$; all integrals on Ω will be taken with respect to this measure (unless otherwise indicated).

We will study strong Sobolev solutions to the PDE (1); that is, we are interested in solutions $u \in \mathcal{H}^2$, where

$$\mathcal{H}^p = \left\{ f \in L^2(\Omega, \mu) : \|f\|_{\mathcal{H}^p} := \left(\sum_{|\alpha| \leq p} \|D^\alpha f\|_{L^2} \right) < \infty \right\}, \quad (2)$$

1. For example, μ can be taken to be the Lebesgue measure on $\Omega \subset \mathbb{R}^n$, and the analytic theory we consider is essentially the same as that when using the Lebesgue measure (i.e. the Sobolev norms using μ are equivalent to those using Leb). We allow for more general μ , as this will roughly correspond to a choice of sampling scheme in our numerical algorithm, and gives us a degree of flexibility when considering the monotonicity of \mathcal{L} .

where Du is the weak derivative of u and where the boundary condition $u|_{\partial\Omega} = f$ is understood in the usual trace-operator sense in \mathcal{H}^1 .

For notational simplicity, we write $\mathcal{H}_{(0)}^2 = \mathcal{H}^2 \cap \mathcal{H}_0^1$ for the subspace of functions $u \in \mathcal{H}^2$ such that $u|_{\partial\Omega} = 0$, that is, with boundary trace zero. This remains a Hilbert space under the \mathcal{H}^2 inner product. We write $\|\cdot\|$ (without a subscript) for the L^2 norm and similarly for the inner product. We refer to Evans (2010) or Adams and Fournier (2003) for further details and general theory of these spaces and concepts.

The class of PDEs we will consider are those for which \mathcal{L} , f and Ω satisfy some (strong) regularity conditions, which we now detail. We first of all present the assumptions and lemmas for the domain and the boundary condition.

Assumption 1 *The boundary $\partial\Omega$ is $C^{3,\alpha}$, for some $\alpha \in (0,1)$, that is, three times continuously differentiable, with α -Hölder continuous third derivative.*

Assumption 2 (Auxiliary function η) *There exists a (known) function $\eta \in C_b^3(\mathbb{R}^n)$, which satisfies $0 < \eta < 1$ in Ω , and $\eta = 0$ on $\partial\Omega$. Furthermore, its first order derivative does not vanish at the boundary (that is, for $x \in \partial\Omega$ and \mathbf{n}_x an outward unit normal vector at x , we have $\langle \nabla\eta(x), \mathbf{n}_x \rangle \neq 0$).*

Using Assumptions 1 and 2, we have the following useful result, which allows us to approximate functions in $\mathcal{H}_{(0)}^2$. Its proof is given in Appendix A.1.

Lemma 3 *1. The set of functions $C^3(\overline{\Omega}) \cap C_0(\overline{\Omega})$ is dense in $\mathcal{H}_{(0)}^2 = \mathcal{H}^2 \cap \mathcal{H}_0^1$ (under the \mathcal{H}^2 topology).*

2. For any function $u \in C^3(\overline{\Omega}) \cap C_0(\overline{\Omega})$, the function $\tilde{u} = u/\eta$ is in $C_b^2(\Omega) \subset \mathcal{H}^2$.

In addition to these assumptions on \mathcal{L} , we make the following assumptions on the boundary value.

Assumption 4 (Interpolation of the boundary condition function) *There exists a (known) function $\bar{f} \in \mathcal{H}^2$ such that $\bar{f}|_{\partial\Omega} = f$ (or, more precisely, if $T : \mathcal{H}^1(\Omega) \rightarrow L^2(\partial\Omega)$ is the boundary trace operator, we have $T(\bar{f}) = f$). In the rest of this paper, for notational simplicity, we identify f with its extension \bar{f} defined on $\overline{\Omega}$.*

Remark 5 *The existence of some \bar{f} satisfying Assumption 4 is guaranteed, given a solution $u \in \mathcal{H}^2$ exists (as we could take $\bar{f} = u$). Practically, we assume not only that \bar{f} exists, but that it is possible for us to use it as part of our numerical method to find u . In many cases, this is a mild and natural assumption, for instance when we have $f = 0$, where we can simply take $\bar{f} = 0$.*

Now we present the assumptions of the PDE itself.

Assumption 6 (Lipschitz continuity of \mathcal{L}) *There exists a constant C such that for any $f_1, f_2 \in \mathcal{H}^2 \cap C^2$ and any $x \in \Omega$, the (nonlinear differential) operator \mathcal{L} satisfies*

$$\begin{aligned} |\mathcal{L}f_1(x) - \mathcal{L}f_2(x)| \leq C & \left[|f_1(x) - f_2(x)| + \sum_{i=1}^n |\partial_{x_i} f_1(x) - \partial_{x_i} f_2(x)| \right. \\ & \left. + \frac{1}{2} \sum_{i,j=1}^n |\partial_{x_i x_j}^2 f_1(x) - \partial_{x_i x_j}^2 f_2(x)| \right]. \end{aligned} \tag{3}$$

Assumption 7 (Integrability of \mathcal{L} at zero) *Taking $f_0 \equiv 0$, we have $\mathcal{L}f_0 \in L^2$.*

Remark 8 *Under Assumptions 6 and 7, \mathcal{L} is a Lipschitz operator $\mathcal{L} : \mathcal{H}^2 \rightarrow L^2$.*

Our algorithm will also make use of the following auxiliary function, which allows us to control behaviour of functions near the boundary of Ω .

In order to prove convergence of our scheme for large training time, we will particularly focus on PDEs given by strongly monotone operators (cf. Browder (1967)), which for convenience we take with the following sign convention.

Assumption 9 (Strong L^2 -monotonicity of \mathcal{L}) *There exists a constant $\gamma > 0$ such that for any $f_1, f_2 \in \mathcal{H}_{(0)}^2$ the operator \mathcal{L} satisfies*

$$\langle f_1 - f_2, \mathcal{L}f_1 - \mathcal{L}f_2 \rangle \leq -\gamma \|f_1 - f_2\|_{L^2}^2, \quad (4)$$

where $\langle \cdot, \cdot \rangle$ is the $L^2(\mu)$ inner product.

While these assumptions are somewhat restrictive, they are general enough to allow for the case $\mathcal{L}u = \mathcal{A}u - \gamma u + r$, for r any L^2 function, where \mathcal{A} is the generator of a sufficiently nice Feller process X and $\gamma > 0$, for an appropriate choice of sampling distribution μ ; see Appendix A.7 for further discussion. In this case, the solution u can be expressed (using the Feynman–Kac theorem) in terms of the stochastic process,

$$u(x) = \mathbb{E} \left[e^{-\gamma\tau} f(X_\tau) + \int_0^\tau e^{-\gamma t} r(X_t) dt \mid X_0 = x \right] \quad (5)$$

where $\tau = \inf\{t \geq 0 : X_t \notin \Omega\}$.

Further, these assumptions are also sufficiently general to allow some nonlinear PDEs of interest, for example Hamilton–Jacobi–Bellman equations under a Cordes condition, as discussed by Smears and Süli (2014). The assumption of monotonicity is also connected to Lyapunov stability analysis, as it is easy to check that strong monotonicity is equivalent to stating that $v \mapsto \frac{1}{2}\|v\|_{L^2}^2$ is a Lyapunov function for the infinite-dimensional dynamical system $dv/dt = \mathcal{L}v + \gamma v$. For more general discussion of monotone operators, and their connection to analysis of the traditional Galerkin method, we refer to Zeidler (2013).

Our final assumption on the PDE is that a solution exists.

Assumption 10 (Existence of solutions) *There exists a (unique) solution $u \in \mathcal{H}^2$ to (1).*

Remark 11 *We note that, assuming strong monotonicity (Assumption 9), if a solution exists then it is guaranteed to be unique: if u, u' are two solutions, then $\mathcal{L}u = \mathcal{L}u'$, hence $0 = \langle u - u', \mathcal{L}u - \mathcal{L}u' \rangle \leq -\gamma \|u - u'\|^2$, and so $u = u'$ almost everywhere.*

While existence of solutions in \mathcal{H}^2 is a strong assumption, it is often satisfied for weak solutions to elliptic equations, given the well-known elliptic regularity results which ensure solutions are sufficiently smooth.

Remark 12 *The algorithm we present below can easily be extended to higher-order PDEs. To do this, further smoothness assumptions (on the activation function σ , continuation value f and auxiliary function η) and higher moments of the initial weights w^i in the neural network are needed, and one argument (Lemma 3) needs to be extended using alternative approaches. In this paper, for the sake of concreteness (and to avoid unduly long derivations), we present our results for PDEs up to second-order; the details of the extension to the general case are left as a tedious exercise for the reader.*

Summary of assumptions 13 *For ease of reference, we now summarize the assumptions we have made on our PDE:*

1. \mathcal{L} is a nonlinear second-order differential operator, Lipschitz (as a map $\mathcal{H}^2(\Omega) \rightarrow L^2(\Omega)$), square integrable at zero (Assumptions 6, 7).
2. The boundary value f has a known continuation to a function in $\mathcal{H}^2(\Omega)$ (Assumption 4).
3. $\Omega \subset \mathbb{R}^n$ is bounded and open, and has a $C^{3,\alpha}$ boundary $\partial\Omega$, which is the zero level set of a known auxiliary function $\eta \in C_b^3(\mathbb{R}^n)$, and η has nonzero derivative on the boundary (Assumptions 1, 2)
4. \mathcal{L} is strongly monotone, and there exists an $\mathcal{H}^2(\Omega)$ solution to the PDE (Assumptions 9, 10).

We emphasise that only parts 1–3 of this assumption will be used, except in Section 5, where the strong monotonicity and existence of solutions will also be needed.

2.2 Neural Network Configuration

We will approximate the solution to the PDE (1) with a neural network, which will be trained with a deep Q-learning inspired algorithm. In particular, we will study a single-layer neural network with N hidden units:

$$S^N(x; \theta) = \frac{1}{N^\beta} \sum_{i=1}^N c^i \sigma(w^i \cdot x + b^i), \quad (6)$$

where the scaling factor $\beta \in (\frac{1}{2}, 1)$ and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a non-linear scalar function. The parameters θ are defined as $\theta := \{c^i, w^i, b^i\}_{i=1, \dots, N}$ where $c^i, b^i \in \mathbb{R}$, and $w^i \in \mathbb{R}^n$. The function η can then be used to design a neural network model which automatically satisfies the boundary condition, introduced by McFall and Mahan (2009):

$$Q^N(x; \theta) := S^N(x; \theta)\eta(x) + f(x). \quad (7)$$

Assumption 14 (Activation function) *The activation function $\sigma \in C_b^4(\mathbb{R})$ is non-constant, where C_b^k is the space of functions with k -th order continuous derivatives.*

Remark 15 *This assumption coincides with the condition of Theorem 4 of Hornik (1991), and guarantees that the functions σ generate neural nets which are dense in the Sobolev space \mathcal{H}^2 . (Hornik shows, under this condition, the stronger result of density in \mathcal{H}^4 ; we focus on \mathcal{H}^2 , but require additional bounded derivatives as part of our proof.) We will particularly make use of this to establish Lemma 30.*

Before training begins (i.e. at $t = 0$), the neural network parameters are randomly initialized. The random initialization satisfies the assumptions described below.

Assumption 16 (Neural network initialization) *The initialization of the parameters θ_0 , for all $i \in \{1, 2, \dots, N\}$, satisfies:*

- *The parameters c_0^i, w_0^i, b_0^i are independent random variables.*
- *The random variables c_0^i are bounded, $|c_0^i| < K_0$, and $\mathbb{E}[c_0^i] = 0$.*
- *The distribution of the random variables w_0^i, b_0^i has full support. That is, for any open set $D \subseteq \mathbb{R}^{n+1}$, we have $\mathbb{P}((w_0^i, b_0^i) \in D) > 0$.*
- *For any indices $i, k \in \{1, \dots, N\}$, we have $\mathbb{E}[|(w_0^i)_k|^3] < \infty, \mathbb{E}[|b_0^i|] < \infty$.*

2.3 Training Algorithm

We now present our algorithm for training the neural network $Q^N(x; \theta)$ to solve the PDE (1). At training time t , the parameters of the neural network are denoted $\theta_t = \{c_t^i, w_t^i, b_t^i\}_{i=1, \dots, N}$. For simplicity, we denote $S^N(\cdot; \theta_t)$ as S_t^N , and $Q^N(\cdot; \theta_t)$ as Q_t^N . Then,

$$Q_t^N = S_t^N \eta + f. \quad (8)$$

Our goal is to design an algorithm to train the approximator Q^N to find the solution of the PDE. One approach, see for instance Sirignano and Spiliopoulos (2018), is to train the model Q^N to minimize the average PDE residual:

$$\int_{\Omega} [\mathcal{L}Q^N(x)]^2 d\mu(x). \quad (9)$$

To improve integrability, we will smoothly truncate this objective function, and use the result to motivate our training algorithm.

Definition 17 (Smooth truncation function) *The functions $\{\psi^N\}_{N \in \mathbb{N}}$ are a family of smooth truncation functions if, for some $\delta \in (0, \frac{1-\beta}{2})$,*

- $\psi^N \in \mathbb{C}_b^2(\mathbb{R})$ is increasing on \mathbb{R} .
- $|\psi^N|$ is bounded by $2N^\delta$.
- $\psi^N(x) = x$ for $x \in [-N^\delta, N^\delta]$.
- $|(\psi^N)'| \leq 1$ on \mathbb{R} .
- $F^N := (\psi^N) \cdot (\psi^N)'$ is uniformly Lipschitz continuous on \mathbb{R} for all $N \in \mathbb{N}$.

For simplicity, we will usually describe the family ψ^N as a (smooth) truncation function. A simple example of a truncation function is as follows: take $\delta = (1 - \beta)/4 > 0$ and

$$\psi^N(x) = \int_0^x g^N(y) dy, \quad \text{where} \quad g^N(x) = \begin{cases} e^{-(|x| - N^\delta)^2} & \text{if } |x| \geq N^\delta, \\ 1 & \text{if } |x| < N^\delta. \end{cases} \quad (10)$$

Remark 18 *It is easy to see that, for any smooth truncation function, the related function $F^N(x) = (\psi^N) \cdot (\psi^N)'$ satisfies*

$$|F^N(x) - x| \leq 2|x| \mathbb{1}_{\{|x| \geq N^\delta\}} \text{ for all } x \in \mathbb{R}. \quad (11)$$

We will use a truncation function ψ^N to modify (9), and will consider minimizing the truncated objective:

$$\int_{\Omega} [\psi^N(\mathcal{L}Q^N(x))]^2 d\mu(x). \quad (12)$$

To minimise a loss function like (12), one can apply gradient-descent based methods.

However, in practice, computing the gradient of $\mathcal{L}(Q_t^N(\cdot; \theta))$ with respect to θ does not lead to an algorithm permitting simple analysis, as natural properties of the differential operator \mathcal{L} are not directly preserved. Instead, we introduce a biased gradient estimator as follows:

Definition 19 *The sequence of functions*

$$G^N(\theta_t) = \int_{\Omega} \left[\psi^N(\mathcal{L}Q_t^N(x)) \cdot (\psi^{N'})'(\mathcal{L}Q_t^N(x)) \right] \nabla_{\theta}(-Q_t^N(x)) d\mu(x) \quad (13)$$

are called the biased gradient estimators for the truncated loss. These can be approximated using Monte Carlo sampling, as

$$\hat{G}_M^N(\theta_t) = \frac{1}{M} \sum_{i=1}^M \left[\psi^N(\mathcal{L}Q_t^N(x_i)) \cdot (\psi^{N'})'(\mathcal{L}Q_t^N(x_i)) \right] \nabla_{\theta}(-Q_t^N(x_i)) \quad (14)$$

where x_i are independent samples from the distribution μ .

Our analysis will focus on the biased gradient estimator G^N , however in our numerical implementation (in order to avoid the curse of dimensionality) we will use the Monte Carlo estimates \hat{G}_M^N for large M .

We shall see that this biased gradient is a continuous-space, PDE analog of the classic Q-learning algorithm. In summary, our ‘‘Q-PDE’’ algorithm for solving PDEs is:

Algorithm 20 (Q-PDE Algorithm) *We fix a family of smooth truncation functions $\{\psi^N\}_{N \in \mathbb{N}}$ and, for each value of $N \in \mathbb{N}$, we proceed as follows:*

1. *Randomly initialize the parameters θ_0 , as specified in Assumption 16.*
2. *Train the neural network via the biased gradient flow*

$$\frac{d\theta_t}{dt} = -\alpha_t^N G^N(\theta_t). \quad (15)$$

with G^N as in (13) and learning rate

$$\alpha_t^N = \alpha^N = \alpha N^{2\beta-1}, \quad (16)$$

where α is a positive constant.

The approximate solution to the PDE at training time t is Q_t^N , as defined by (7).

In the remainder of this paper, we will study this Q-PDE algorithm for solving PDEs both mathematically and numerically. First, we prove that the neural network model Q_t^N converges to the solution of an infinite-dimensional ODE as the number of hidden units $N \rightarrow \infty$. Then, we study the limit ODE, and prove that it converges to the true solution of the PDE as the training time $t \rightarrow \infty$. Finally, in Section 6, we demonstrate numerically that the algorithm performs well in practice with a finite number of hidden units. For high dimensional PDEs, it is impossible to form mesh-grids to apply conventional finite difference methods. However, our method is mesh-free. The crucial feature for implementing our algorithm in practice is to evaluate the integral term G^N . As discussed, we can do this using the Monte Carlo estimate \hat{G}_M^N ; therefore our algorithm has the potential to solve high-dimensional PDEs.

Remark 21 *It is worth observing that, given this choice of biased gradient flow, there is generally no guarantee that our algorithm corresponds to stochastic gradient descent applied to any potential function. The key advantage of the use of this biased gradient is that it separates the derivatives in θ from the differential operator \mathcal{L} , which allows the underlying properties of the PDE to be preserved more simply. We will see that this results in particularly simple dynamics in the wide-network limit.*

2.3.1 MAIN RESULTS

Our main mathematical results are the following:

Theorem 22 (cf. Theorem 45) *Define the kernel*

$$B(x, y) = \eta(x)\eta(y)\mathbb{E}_{c,w,b}\left[\sigma(w \cdot x + b)\sigma(w \cdot y + b) + c^2\sigma'(w \cdot x + b)\sigma'(w \cdot y + b)(x \cdot y + 1)\right], \quad (17)$$

where the expectation is taken with respect to the distribution of the random initialization of c , w , and b satisfying Assumption 16, and define the corresponding integral operator $(\mathcal{B}v)(y) = \int_{\Omega} B(x, y)v(x)d\mu(x)$. For a single-hidden layer neural network with N units, as $N \rightarrow \infty$ the PDE approximator Q^N trained using the Q-PDE algorithm converges (in \mathcal{H}^2 , for each time t) to the deterministic \mathcal{H}^2 -valued dynamical system $dQ_t/dt = \alpha\mathcal{B}\mathcal{L}(Q_t)$, with $Q_0 = f$, and $Q_t|_{\partial\Omega} \equiv f$ for all t .

Theorem 23 *The limiting dynamical system Q_t has the following convergence properties:*

- *If Q_t converges in \mathcal{H}^1 to a fixed point Q , then Q is the solution to the PDE (1). (Theorem 46)*
- *If \mathcal{L} is a monotone operator, and u is the solution to the PDE (1), then for almost every sequence $t_k \rightarrow \infty$, we have $Q_{t_k} \rightarrow u$ in L^2 . (Theorem 47)*
- *If \mathcal{L} is a Gateaux differentiable monotone operator then, for almost every sequence $t_k \rightarrow \infty$, we have $\mathcal{B}\mathcal{L}Q_{t_k} \rightarrow 0$. (Theorem 48)*

2.3.2 SIMILARITY TO Q-LEARNING

Our training algorithm’s biased gradient flow is inspired by the classic Q-learning algorithm in a reinforcement learning setting. The Q-learning algorithm approximates the value function of the optimal policy for Markov decision problems (MDPs). It seeks to minimize the error of a parametric approximator, such as a neural network, $R(\cdot, \cdot; \theta)$:

$$L(\theta) = \sum_{(s,a) \in \mathcal{S} \times \mathcal{X}} [Y(s, a; \theta) - R(s, a; \theta)]^2 \pi(s, a). \quad (18)$$

Here \mathcal{S} and \mathcal{X} are the finite state and action spaces of the MDP and π is a probability mass function which is strictly positive for every $(s, a) \in \mathcal{S} \times \mathcal{X}$. Y is the target Bellman function

$$Y(s, a; \theta) = r(s, a) + \gamma \sum_{s' \in \mathcal{S}} \max_{a' \in \mathcal{X}} R(s', a'; \theta) p(s'|s, a). \quad (19)$$

Q-learning updates the parametric approximator R via $\theta_{k+1} = \theta_k + \alpha_k^N U_k$ where U is a biased estimator of the gradient of $(Y - R)^2$. Since the transition density $p(s'|s, a)$ is unknown, $Y - R$ is estimated using random samples from the Markov chain (s_k, a_k) :

$$\tilde{Y}(s_k, a_k) - R(s_k, a_k; \theta_k) := r(s_k, a_k) + \gamma \max_{a' \in \mathcal{X}} R(s_{k+1}, a'; \theta_k) - R(s_k, a_k; \theta_k). \quad (20)$$

The Q-learning algorithm treats Y (and \tilde{Y}) as a constant and takes a derivative only with respect to the last term in (20). The Q-learning update direction U_k is:

$$-U_k =: \left(\tilde{Y}(s_k, a_k) - R(s_k, a_k; \theta_k) \right) \nabla_{\theta_k} [-R(s_k, a_k; \theta_k)] \quad (21)$$

We emphasize that the term $\nabla_{\theta} \tilde{Y}$ is not included, which means U_k is a biased estimator for the direction of steepest descent.

Returning to the objective function for the Q-PDE algorithm, the gradient of the loss function (12) is

$$\int_{\Omega} \left[\psi^N(\mathcal{L}Q_t^N(x)) \psi^{N'}(\mathcal{L}Q_t^N(x)) \right] \nabla_{\theta} \mathcal{L}Q_t^N(x) d\mu(x). \quad (22)$$

The update direction G^N in (13) differs from (22), as $\nabla_{\theta} \mathcal{L}Q_t^N(x)$ has been replaced with $\nabla_{\theta} (-Q_t^N(x))$. This is, conceptually, the same update direction as used in the Q-learning algorithm.

3. Preliminary analysis of the training algorithm

We begin our analysis of the training algorithm (15) by proving several useful bounds, in particular, that the neural network parameters c_t and w_t are bounded in expectation. The proof is provided in Appendix A.2. We recall that n is the dimension of our underlying space, that is, $\Omega \subset \mathbb{R}^n$.

Lemma 24 (Boundedness of parameters) *For all $T > 0$, there exists a deterministic constant $C > 0$ such that, for all $N \in \mathbb{N}$, all $0 \leq t \leq T$, all $i \in \{1, 2, \dots, N\}$, and all $k \in \{1, 2, \dots, n\}$,*

$$|c_t^i| < C, \quad \mathbb{E}|(w_t^i)_k| < C, \quad \text{and} \quad \mathbb{E}|b_t^i| < C. \quad (23)$$

3.0.1 DYNAMICS OF Q_t^N

We next consider the evolution of the output of the neural network Q_t^N as a function of the training time t . Define B_t^N as the symmetric, positive semi-definite kernel function

$$B_t^N(x, y) = \eta(x)\eta(y)A_t^N(x, y), \quad (24)$$

where A_t^N is the *neural tangent kernel* (Jacot et al. (2018))

$$\begin{aligned} A_t^N(x, y) &= \left(\nabla_{\theta} S_t^N(x) \right)^{\top} \left(\nabla_{\theta} S_t^N(y) \right) \\ &= \frac{1}{N} \sum_{i=1}^N \left[\sigma(w_t^i \cdot x + b_t^i) \sigma(w_t^i \cdot y + b_t^i) \right. \\ &\quad \left. + c_t^{i2} \sigma'(w_t^i \cdot x + b_t^i) \sigma'(w_t^i \cdot y + b_t^i) (x \cdot y + 1) \right]. \end{aligned} \quad (25)$$

For any $y \in \bar{\Omega}$, we can derive the evolution of value $Q_t^N(y)$ using the chain rule:

$$\frac{dQ_t^N(y)}{dt} = \nabla_{\theta} Q_t^N(y) \cdot \frac{d\theta_t}{dt}. \quad (26)$$

The RHS of (26) can be evaluated using (7), (15), and Fubini's theorem, which yields

$$\frac{dQ_t^N(y)}{dt} = \alpha \int_{\Omega} F^N \left(\mathcal{L}Q_t^N(x) \right) B_t^N(x, y) d\mu(x), \quad (27)$$

where F^N is as in Definition 17.

3.0.2 LIMIT KERNEL

Equation (27) shows that the kernel B^N has a key role in the dynamics of Q^N . We now characterize the limit of B^N as $N \rightarrow \infty$. In Section 4.1, we will study this convergence in detail.

At time $t = 0$, the parameters c_0^i, w_0^i, b_0^i are independently sampled. Therefore, for each $(x, y) \in \Omega$, by the strong law of large numbers $A_0^N(x, y)$ converges almost surely to $A(x, y)$, where

$$A(x, y) = \mathbb{E}_{c,w,b} \left[\sigma(w \cdot x + b) \sigma(w \cdot y + b) + c^2 \sigma'(w \cdot x + b) \sigma'(w \cdot y + b) (x \cdot y + 1) \right]. \quad (28)$$

It follows that $B_0^N(x, y)$ converges almost surely to $B(x, y)$:

$$\lim_{N \rightarrow \infty} B_0^N(x, y) = B(x, y) := \eta(x)\eta(y)A(x, y). \quad (29)$$

We can write

$$B(x, y) = \eta(x)\eta(y) \mathbb{E}_{c,w,b} \left[\sigma(w \cdot x + b) \sigma(w \cdot y + b) + c^2 U(x) \cdot U(y) \right], \quad (30)$$

where the (w, b) -dependent vector function U is given by $U(x) = \sigma'(w \cdot x + b)(x + 1/\sqrt{n})$. As $\sigma \in C_b^4(\mathbb{R})$, we know that A, A^N, B and B^N are all uniformly continuous in (x, y) , so the above a.s. convergences hold for all $(x, y) \in \Omega$ simultaneously.

Lemma 25 *There exists a constant $M > 0$ such that, for any $(x, y) \in \bar{\Omega} \times \bar{\Omega}$,*

$$|B(x, y)| + \sum_k |\partial_{y_k} B(x, y)| + \sum_{k,l} |\partial_{y_k y_l}^2 B(x, y)| \leq M. \quad (31)$$

Proof Taking B as defined in (29), as $\eta \in C_b^3$, $\sigma \in C_b^4$ and each term of $\{c_t^i\}_{i \in \mathbb{N}}$ is uniformly bounded, we have $|B(x, y)| < C$. The partial derivative of B is

$$\begin{aligned} & \partial_{y_k} B(x, y) \\ &= \eta(x) (\partial_{y_k} \eta(y)) A(x, y) + \eta(x) \eta(y) (\partial_{y_k} A(x, y)) \\ &= \eta(x) \partial_{y_k} \eta(y) A(x, y) + \eta(x) \eta(y) \mathbb{E}_{c, w, b} \left[\sigma(w \cdot x + b) \sigma'(w \cdot y + b) w_k \right. \\ & \quad \left. + c^2 \sigma'(w \cdot x + b) [w_k \sigma''(w \cdot y + b) x \cdot y + \sigma'(w \cdot y + b) x_k] \right]. \end{aligned} \quad (32)$$

As before, our assumptions on η , σ and $\{c_t^i\}_{i \in \mathbb{N}}$ guarantee that the terms in $\partial_{y_k} B(x, y)$ are uniformly bounded for all $x, y \in \bar{\Omega}$, so $|\partial_{y_k} B(x, y)| < C$; similarly for $\partial_{y_k y_l}^2 B(x, y)$. The result follows. \blacksquare

Given we have a kernel B , it is natural to define the corresponding integral operator, $v \mapsto \int_{\Omega} B(x, y) v(x) d\mu(x)$. It will be convenient to define this on a slightly larger space than L^2 , to account for the effect of the function η .

Definition 26 *We define the function space $L_{\eta}^2 := \{f : \Omega \rightarrow \mathbb{R} \mid \eta f \in L^2\}$, and observe (as η is bounded) that $L^2 \subset L_{\eta}^2$.*

Definition 27 *The operator $\mathcal{B} : L_{\eta}^2 \rightarrow L^2$ is defined by $(\mathcal{B}v)(y) = \int_{\Omega} B(x, y) v(x) d\mu(x)$ with B as in (29).*

Lemma 28 *For any $v \in L_{\eta}^2$, we have $\mathcal{B}v \in \mathcal{H}_{(0)}^2$, or more specifically, $\mathcal{B}v \in C_b^2(\bar{\Omega})$ and $(\mathcal{B}v)(y) = 0$ for all $y \in \partial\Omega$.*

Proof From the definition of $\mathcal{B}v$, we have $(\mathcal{B}v)(y) = \eta(y) \int_{\Omega} A(x, y) \eta(x) v(x) d\mu(x)$. Lemma 25 and the smoothness of σ and η clearly imply that $\mathcal{B}v \in C_b^2(\bar{\Omega}) \subset \mathcal{H}^2$. We know that $\eta(x) = 0$ for $x \in \partial\Omega$, so it is clear that $(\mathcal{B}v)(y) = 0$ for all $y \in \partial\Omega$. As \mathcal{H}_0^1 is the kernel (in \mathcal{H}^1) of the boundary trace operator, and for smooth functions the trace operator is simply evaluation of the function on the boundary, it follows that $\mathcal{B}v \in \mathcal{H}_0^1$. We conclude that $\mathcal{B}v \in \mathcal{H}_{(0)}^2 = \mathcal{H}^2 \cap \mathcal{H}_0^1$. \blacksquare

Lemma 29 *The linear map $\mathcal{B} : L_{\eta}^2 \rightarrow \mathcal{H}^2$ is Lipschitz, in particular, there exists $C > 0$ such that, for any $v \in L_{\eta}^2$, we have $\|\mathcal{B}v\|_{\mathcal{H}^2} \leq C \|\eta \cdot v\|_{L^2}$.*

Proof By definition, with D^{α} denoting the differential operator with respect to the y argument, $\|\mathcal{B}v\|_{\mathcal{H}^2}^2 = \sum_{|\alpha| \leq 2} \|D^{\alpha} \mathcal{B}v\|_{L^2}^2$. By the boundedness of η , A and their derivatives, there exists a constant K such that for any multi-index α , $|D^{\alpha}[\eta(y)A(x, y)]| < K$ for any

$x, y \in \Omega$. Hence, exchanging the order of integration (with respect to x) and differentiation (with respect to y) by the dominated convergence theorem,

$$\begin{aligned} |(D^\alpha \mathcal{B}v)(y)| &= \left| D^\alpha \left[\int_{\Omega} \eta(y) A(x, y) \eta(x) v(x) d\mu(x) \right] \right| = \left| \int_{\Omega} D^\alpha [\eta(y) A(x, y)] \eta(x) v(x) d\mu(x) \right| \\ &\leq \int_{\Omega} \left| D^\alpha [\eta(y) A(x, y)] \eta(x) v(x) \right| d\mu(x) \leq K \int_{\Omega} |\eta(x) v(x)| d\mu(x). \end{aligned} \quad (33)$$

By Jensen's inequality, there is a constant $C > 0$ such that

$$|(D^\alpha \mathcal{B}v)(y)|^2 \leq K^2 \left(\int_{\Omega} |\eta(x) v(x)| d\mu(x) \right)^2 \leq C \int_{\Omega} |\eta(x) v(x)|^2 d\mu(x) = C \|\eta \cdot v\|_{L^2}^2. \quad (34)$$

Integrating over Ω , we see that $\|D^\alpha \mathcal{B}v\|_{L^2}^2 \leq (C\mu(\Omega)) \|\eta \cdot v\|_{L^2}^2$. Summing over indices α yields the result. \blacksquare

Our next challenge is to prove that the image of \mathcal{B} is dense in $\mathcal{H}_{(0)}^2$. This will allow us to represent the solutions to our PDE in terms of \mathcal{B} , and is closely related to the universal approximation theorems of Cybenko (1989), Hornik (1991), and others.

Lemma 30 *For a function $g \in \mathcal{H}^2(\Omega)$, if the equality*

$$\int_{\Omega} \sum_{|\alpha| \leq 2} D^\alpha \sigma(w \cdot x + b) D^\alpha g(x) d\mu(x) = 0 \quad (35)$$

holds for all $w \in \mathbb{R}^n$, $b \in \mathbb{R}$, then $g \equiv 0$.

Proof Since μ is a finite measure and $\sigma \in C_b^2$ (cf. Assumption 14) and non-constant, from Theorem 4 in Hornik (1991) we have that the linear span of $\{\sigma(w \cdot x + b)\}_{w, b \in \mathbb{R}}$ is dense in \mathcal{H}^2 . Hence, there is no nontrivial element of \mathcal{H}^2 orthogonal to $\sigma(w \cdot x + b)$ for all w, b . In other words, for $g \in \mathcal{H}^2$, if the inner product $\langle \sigma_{w, b}, g \rangle_{\mathcal{H}^2} = 0$ for all w, b , then $g = 0$, which is the stated result. \blacksquare

Lemma 31 *Define $S : \mathcal{H}^2 \rightarrow \mathcal{H}^2$ by $(Sh)(y) := \int_{\Omega} \sum_{|\alpha| \leq 2} D_x^\alpha h(x) D_x^\alpha A(x, y) d\mu(x)$. Then $Sh = 0$ if and only if $h = 0$.*

Proof We first observe that, as $\sigma \in C_b^4(\mathbb{R})$, $D_x^\alpha A(x, y)$ is clearly a C_b^2 function of y for all x . Consider the inner product between Sh and h on \mathcal{H}^2 :

$$\begin{aligned} \langle Sh, h \rangle_{\mathcal{H}^2} &= \sum_{|\alpha| \leq 2} \langle D^\alpha Sh, D^\alpha h \rangle_{L^2} \\ &= \int_{\Omega^2} \sum_{|\alpha_1| \leq 2, |\alpha_2| \leq 2} D_x^{\alpha_1} h(x) D_y^{\alpha_2} h(y) \left(D_x^{\alpha_1} D_y^{\alpha_2} A(x, y) \right) d\mu(x) d\mu(y) \\ &= \mathbb{E}_{c, w, b} \left[\int_{\Omega^2} \sum_{|\alpha_1| \leq 2, |\alpha_2| \leq 2} \left\{ D_x^{\alpha_1} h(x) D_y^{\alpha_2} h(y) \right. \right. \\ &\quad \left. \left. \times \left(D_x^{\alpha_1} D_y^{\alpha_2} \left[\sigma(w \cdot x + b) \sigma(w \cdot y + b) + c^2 U(x) \cdot U(y) \right] \right) d\mu(x) d\mu(y) \right\} \right]. \end{aligned} \quad (36)$$

Notice that

$$\begin{aligned}
 0 &\leq \mathbb{E}_{c,w,b} \left[\left\| c \int_{\Omega} \sum_{|\alpha| \leq 2} D_x^\alpha h(x) D_x^\alpha U(x) d\mu(x) \right\|^2 \right] \\
 \mathbb{E}_{c,w,b} &\left[\int_{\Omega^2} \sum_{|\alpha_1| \leq 2, |\alpha_2| \leq 2} \left\{ D_x^{\alpha_1} h(x) D_y^{\alpha_2} h(y) \left(D_x^{\alpha_1} D_y^{\alpha_2} (c^2 U(x) \cdot U(y)) \right) d\mu(x) d\mu(y) \right\} \right]
 \end{aligned} \tag{37}$$

Combining (36) with (37) we derive

$$\langle Sh, h \rangle_{\mathcal{H}^2} \geq \mathbb{E}_{c,w,b} \left[\left(\int_{\Omega} \sum_{|\alpha| \leq 2} D_x^\alpha h(x) D_x^\alpha \sigma(w \cdot x + b) d\mu(x) \right)^2 \right] \geq 0. \tag{38}$$

By Lemma 30, for $h \neq 0$, there exists $\epsilon > 0$, $w^* \in \mathbb{R}^n$, $b^* \in \mathbb{R}$ such that $0 < \epsilon = \left| \int_{\Omega} \sum_{|\alpha| \leq 2} D_x^\alpha h(x) D_x^\alpha \sigma(w^* \cdot x + b^*) d\mu(x) \right|$. As this integral is continuous with respect to the parameters w, b , there exists $\delta > 0$ such that for any $(w, b) \in R_\delta := \{(w, b) : \|w - w^*\| + |b - b^*| \leq \delta\}$, we know $\left| \int_{\Omega} \sum_{|\alpha| \leq 2} D_x^\alpha h(x) D_x^\alpha \sigma(w \cdot x + b) d\mu(x) \right| \geq \frac{\epsilon}{2}$. Therefore, from (38) we have

$$\langle Sh, h \rangle_{\mathcal{H}^2} \geq \mathbb{E}_{c,w,b} \left[\left(\int_{\Omega} \sum_{|\alpha| \leq 2} D_x^\alpha h(x) D_x^\alpha \sigma(w \cdot x + b) d\mu(x) \right)^2 \right] \geq \mathbb{E}_{c,w,b} \left[\frac{\epsilon^2}{4} \mathbb{1}_{R_\delta} \right] > 0. \tag{39}$$

■

Lemma 32 *Define the operator $\mathcal{A} : L^2 \rightarrow \mathcal{H}^2$ by $(\mathcal{A}f)(y) = \int_{\Omega} A(x, y) f(x) d\mu(x)$. The image of \mathcal{A} is dense in \mathcal{H}^2 , that is, for any $u \in \mathcal{H}^2$, there exists a sequence $\{v_k\}_{k \in \mathbb{N}}$ in L^2 such that $\lim_{k \rightarrow \infty} \|\mathcal{A}v_k - u\|_{\mathcal{H}^2} = 0$.*

Proof By the smoothness of the kernel A , we have $\text{im}(\mathcal{A}) := \{\mathcal{A}v : v \in L^2\} \subset \mathcal{H}^2$. By definition, for any $h, g \in \mathcal{H}^2$, the inner product between $\mathcal{A}h$ and g in \mathcal{H}^2 space is

$$\begin{aligned}
 \langle \mathcal{A}g, h \rangle_{\mathcal{H}^2} &= \sum_{|\alpha| \leq 2} \langle D^\alpha \mathcal{A}g, D^\alpha h \rangle_{L^2} = \int_{\Omega} \sum_{|\alpha| \leq 2} D_x^\alpha h(x) \int_{\Omega} D_x^\alpha A(x, y) g(y) d\mu(y) d\mu(x) \\
 &= \int_{\Omega} g(y) \int_{\Omega} \sum_{|\alpha| \leq 2} D_x^\alpha h(x) D_x^\alpha A(x, y) d\mu(x) d\mu(y) = \langle g, Sh \rangle_{L^2}.
 \end{aligned} \tag{40}$$

We introduce the adjoint operator of \mathcal{A} on \mathcal{H}^2 , denoted \mathcal{A}^* , and write

$$\langle \mathcal{A}g, h \rangle_{\mathcal{H}^2} = \langle g, \mathcal{A}^* h \rangle_{\mathcal{H}^2} = \langle g, Sh \rangle_{L^2}. \tag{41}$$

By Lemma 31, $Sh = 0$ if and only if $h = 0$. Therefore, by setting $g = Sh$ in (41), we have $\langle g, \mathcal{A}^* h \rangle_{\mathcal{H}^2} = \langle Sh, Sh \rangle_{L^2} > 0$ for any non-zero h . Thus, $\ker(\mathcal{A}^*) := \{h \in \mathcal{H}^2 : \mathcal{A}^* h = 0\} = \{0\}$. Write $\overline{\text{im}(\mathcal{A})}$ for the closure in \mathcal{H}^2 of $\text{im}(\mathcal{A})$. We recall that the ortho-complement in \mathcal{H}^2 of $\overline{\text{im}(\mathcal{A})}$ is $\ker(\mathcal{A}^*)$, so we can decompose the space \mathcal{H}^2 as $\mathcal{H}^2 = \overline{\text{im}(\mathcal{A})} \oplus \ker(\mathcal{A}^*)$. We know that $\ker(\mathcal{A}^*) = \{0\}$ in \mathcal{H}^2 , and hence conclude that $\overline{\text{im}(\mathcal{A})}$ is dense in \mathcal{H}^2 . ■

Theorem 33 (Density in $\mathcal{H}_{(0)}^2$) *The image of \mathcal{B} is dense in $\mathcal{H}_{(0)}^2$, that is, for any $u \in \mathcal{H}_{(0)}^2$, there exists a sequence $\{v_k\}_{k \in \mathbb{N}}$ in L_η^2 such that $\lim_{k \rightarrow \infty} \|\mathcal{B}v_k - u\|_{\mathcal{H}^2} = 0$.*

Proof For a function $u \in C^3(\overline{\Omega}) \cap C_0(\overline{\Omega})$, we define the function $\tilde{u} = u/\eta$. By Lemma 3(ii), $\tilde{u} \in C_b^2(\Omega) \subset \mathcal{H}^2$. Since the image of \mathcal{A} is dense in \mathcal{H}^2 , there exists a sequence $w_k \in \mathcal{H}^2$ such that $\|\mathcal{A}w_k - \tilde{u}\|_{\mathcal{H}^2} \rightarrow 0$. Consequently, the boundedness of η and its derivatives shows $\|\eta(\mathcal{A}w_k - \tilde{u})\|_{\mathcal{H}^2} = \|\mathcal{B}(w_k/\eta) - u\|_{\mathcal{H}^2} \rightarrow 0$. Since $C^3(\overline{\Omega}) \cap C_0(\overline{\Omega})$ is dense in $\mathcal{H}_{(0)}^2$ (under the \mathcal{H}^2 topology, Lemma 3(i)), we conclude that the image of \mathcal{B} is dense in $\mathcal{H}_{(0)}^2$. \blacksquare

Remark 34 *The above result shows that it is, in principle, possible to approximate the PDE solution u using an infinite neural network, with the boundary condition enforced by η . It also justifies the boundary-matching method proposed by McFall and Mahan (2009). In particular, as $u - (1-\eta)f \in \mathcal{H}_{(0)}^2$, there exists a sequence $v_k \in L_\eta^2$ such that $\mathcal{B}v_k + (1-\eta)f \rightarrow u$ in \mathcal{H}^2 . Of course, whether a particular training algorithm yields such a sequence is a separate question.*

We collect the remaining key properties of \mathcal{B} into the following lemma.

Lemma 35 *The linear operator $\mathcal{B} : L_\eta^2 \rightarrow \mathcal{H}^2$ has the following properties:*

1. \mathcal{B} is strictly positive definite, and induces a norm $\|\cdot\|_{\mathcal{B}}$ on L_η^2 , given by $\|v\|_{\mathcal{B}} := \sqrt{\langle \mathcal{B}v, v \rangle_{L^2}}$. In particular, $\|v\|_{\mathcal{B}} < \infty$ for all $v \in L_\eta^2$.
2. There exists a constant $\lambda > 0$ such that $\|\mathcal{B}v\|_{L^2}^2 \leq \lambda \|v\|_{\mathcal{B}}^2$.

Proof For $v \in L_\eta^2$, let $g := \eta v \in L^2$. By definition,

$$\begin{aligned}
 \langle v, \mathcal{B}v \rangle &= \langle v, \eta \mathcal{A}g \rangle = \langle g, \mathcal{A}g \rangle = \int_{\Omega^2} g(x)g(y)A(x, y)d\mu(x)d\mu(y) \\
 &= \int_{\Omega^2} g(x)g(y)\mathbb{E}_{c, w, b}[\sigma(w \cdot x + b)\sigma(w \cdot y + b)]d\mu(x)d\mu(y) \\
 &\quad + \int_{\Omega^2} g(x)g(y)\mathbb{E}_{c, w, b}[c^2U(x) \cdot U(y)]d\mu(x)d\mu(y) \\
 &= \mathbb{E}_{c, w, b} \left[\left(\int_{\Omega} \sigma(w \cdot x + b)g(x)d\mu(x) \right)^2 + \left\| c \int_{\Omega} U(x)g(x)d\mu(x) \right\|^2 \right] \\
 &\geq \mathbb{E}_{c, w, b} \left[\left(\int_{\Omega} \sigma(w \cdot x + b)g(x)d\mu(x) \right)^2 \right] \geq 0.
 \end{aligned} \tag{42}$$

For $v \neq 0$, we know $g \neq 0$. With Assumption 14, by Theorem 5 in Hornik (1991), there exists $w^* \in \mathbb{R}^2$, $b^* \in \mathbb{R}$ such that $|\int_{\Omega} \sigma(w^* \cdot x + b^*)g(x)d\mu(x)| = \epsilon > 0$. As the integral is continuous with respect to parameters w, b there exists $\delta > 0$ such that for any (w, b) in the set $R_\delta := \{(w, b) : \|w - w^*\| + |b - b^*| < \delta\}$, we have $|\int_{\Omega} \sigma(w \cdot x + b)g(x)d\mu(x)| > \epsilon/2$. Therefore $\langle v, \mathcal{B}v \rangle \geq \mathbb{E}_{c, w, b} [(\int_{\Omega} \sigma(w \cdot x + b)g(x)d\mu(x))^2] \geq \mathbb{E}_{c, w, b} [\mathbb{1}_{R_\delta} \frac{\epsilon^2}{4}] > 0$; so \mathcal{B} defines a norm as stated.

From the above calculations, we also see that \mathcal{A} is a positive-definite Hilbert–Schmidt integral operator on L^2 . Therefore, its eigenfunctions span the entire L^2 space and the spectral theorem applies, in particular \mathcal{A} has nonnegative eigenvalues bounded above. Suppose λ is the supremum of the eigenvalues of \mathcal{A} , we have $\|\mathcal{A}g\|_{L^2} \leq \lambda\langle g, \mathcal{A}g \rangle$. Thus, as $|\eta(x)| < 1$, we know $\|\mathcal{B}v\|_{L^2}^2 = \|\eta\mathcal{A}(\eta v)\|_{L^2}^2 \leq \|\mathcal{A}(\eta v)\|_{L^2}^2 \leq \lambda\langle \eta v, \mathcal{A}(\eta v) \rangle = \lambda\langle v, \mathcal{B}v \rangle = \lambda\|v\|_{\mathcal{B}}^2$. \blacksquare

4. Convergence to the limit ODE

In this section, we seek to understand the behavior of our approximator Q_t^N when $N \rightarrow \infty$. We prove that the pre-limit process, $\{Q_t^N(\cdot)\}_{0 \leq t \leq T}$, converges to a limiting process $\{Q_t(\cdot)\}_{0 \leq t \leq T}$, in an appropriate space of functions. In this section, we will only use Assumptions 13(i–iii), in particular we will not need the assumption that \mathcal{L} is monotone or that a solution to the PDE exists in \mathcal{H}^2 . The challenge here is that our operator \mathcal{L} is not generally the gradient of any potential function, and is an unbounded operator (in the L^2 or supremum norms), and so some care is needed.

In the following subsections, we first bound the difference between the kernels $B_t^N(x, y)$ and $B(x, y)$ in a convenient sense. We then show that, as N becomes large, the neural network approximator converges to the solution of an infinite dimensional ODE. Our main convergence result is Theorem 45.

4.1 Characterizing the difference between kernels

We characterize the (second order) difference between two kernels at (x, y) by

$$\begin{aligned} H(A_1, A_2)(x, y) &:= |A_1(x, y) - A_2(x, y)| + \sum_k |\partial_{y_k} A_1(x, y) - \partial_{y_k} A_2(x, y)| \\ &+ \sum_{k,l} |\partial_{y_k y_l}^2 A_1(x, y) - \partial_{y_k y_l}^2 A_2(x, y)|. \end{aligned} \tag{43}$$

Note that partial derivatives are only taken with respect to y components. From Lemma 25, there exists a constant M such that, for all $x, y \in \Omega$,

$$H(B, 0)(x, y) \leq M. \tag{44}$$

Similarly, the (second order) difference between two smooth functions is characterized by

$$\begin{aligned} G(f_1, f_2)(x) &:= |f_1(x) - f_2(x)| + \sum_k |\partial_{x_k} f_1(x) - \partial_{x_k} f_2(x)| \\ &+ \sum_{k,l} |\partial_{x_k x_l}^2 f_1(x) - \partial_{x_k x_l}^2 f_2(x)|. \end{aligned} \tag{45}$$

4.1.1 DIFFERENCE BETWEEN KERNEL $B_t^N(x, y)$ AND $B_0^N(x, y)$

In this subsection we characterize the difference between kernel B_t^N and B_0^N . We denote the expectation taken with respect to randomized initialization by $\mathbb{E}_{c_0, w_0, b_0}[\cdot]$. The proofs of the following lemmas are included in the appendix.

Lemma 36 *There exists $C > 0$ such that, for all $i, j \in \{1, 2, \dots, n\}$, all $0 \leq t \leq T$, and all $N > 0$,*

$$|A_t^N(x, y) - A_0^N(x, y)| \leq CN^{\delta+\beta-1}, \quad (46)$$

$$|\partial_{y_i} A_t^N(x, y) - \partial_{y_i} A_0^N(x, y)| \leq CN^{\delta+\beta-1} \left[1 + \frac{\sum_{k=1}^N |(w_0^k)_i|}{N} \right], \quad (47)$$

$$|\partial_{y_i y_j}^2 A_t^N(x, y) - \partial_{y_i y_j}^2 A_0^N(x, y)| \leq CN^{\delta+\beta-2} \left(\sum_{k=1}^N (1 + |(w_0^k)_i|)(1 + |(w_0^k)_j|) \right). \quad (48)$$

Lemma 37 *There exists $C > 0$ such that, for all $0 \leq t \leq T$ and all $N > 0$, the expected difference between A_t^N and A_0^N satisfies*

$$\mathbb{E}_{c_0, w_0, b_0} \left[|H(A_t^N, A_0^N)(x, y)|^2 \right] \leq CN^{2(\delta+\beta-1)}. \quad (49)$$

After characterizing the difference between A_t^N and A_0^N , we estimate the difference between the kernels B_t^N and B_0^N .

Lemma 38 *There exists $C > 0$ such that, for all $0 \leq t \leq T$ and all $N > 0$*

$$\mathbb{E}_{c_0, w_0, b_0} \left[N^{2\delta} \int_{\Omega^2} |H(B_t^N, B_0^N)(x, y)|^2 d\mu(x) d\mu(y) \right] \leq CN^{2(2\delta+\beta-1)}. \quad (50)$$

Proof Since $B_t^N(x, y) = \eta(x)\eta(y)A_t^N(x, y)$, by the Cauchy–Schwarz inequality and the assumption $\eta \in C_b^3$, for some constant $C > 0$,

$$\begin{aligned} |H(B_t^N, B_0^N)(x, y)|^2 &\leq C \left[|B_t^N(x, y) - B_0^N(x, y)|^2 + \sum_k |\partial_{y_k} B_t^N(x, y) - \partial_{y_k} B_0^N(x, y)|^2 \right. \\ &\quad \left. + \sum_{k,l} |\partial_{y_k y_l} B_t^N(x, y) - \partial_{y_k y_l} B_0^N(x, y)|^2 \right]. \end{aligned} \quad (51)$$

By the product rule, we have

$$\partial_{y_i} B_t^N(x, y) = \eta(x)A_t^N(x, y)\partial_{y_i}\eta(y) + \eta(x)\eta(y)\partial_{y_i}A_t^N(x, y), \quad (52)$$

$$\begin{aligned} \partial_{y_i y_j}^2 B_t^N(x, y) &= \eta(x)A_t^N(x, y)\partial_{y_i y_j}^2 \eta(y) + \eta(x)\partial_{y_j}A_t^N(x, y)\partial_{y_i}\eta(y) \\ &\quad + \partial_{y_i}A_t^N(x, y)\partial_{y_j}\eta(y) + \eta(x)\eta(y)\partial_{y_i y_j}^2 A_t^N(x, y). \end{aligned} \quad (53)$$

As $\eta \in C_b^3$, there exists a constant $C > 0$ independent of training time t , such that

$$|B_t^N(x, y) - B_0^N(x, y)| \leq C|A_t^N(x, y) - A_0^N(x, y)|, \quad (54)$$

$$|\partial_{y_i} B_t^N(x, y) - \partial_{y_i} B_0^N(x, y)| \leq C \left[|A_t^N(x, y) - A_0^N(x, y)| + |\partial_{y_i} A_t^N(x, y) - \partial_{y_i} A_0^N(x, y)| \right], \quad (55)$$

and

$$\begin{aligned}
 & |\partial_{y_i y_j}^2 B_t^N(x, y) - \partial_{y_i y_j}^2 B_0^N(x, y)| \\
 & \leq C \left[|A_t^N(x, y) - A_0^N(x, y)| + |\partial_{y_i} A_t^N(x, y) - \partial_{y_i} A_0^N(x, y)| \right. \\
 & \quad \left. + |\partial_{y_j} A_t^N(x, y) - \partial_{y_j} A_0^N(x, y)| + |\partial_{y_i y_j}^2 A_t^N(x, y) - \partial_{y_i y_j}^2 A_0^N(x, y)| \right].
 \end{aligned} \tag{56}$$

Substituting (54), (55) and (56) in (51),

$$\begin{aligned}
 |H(B_t^N, B_0^N)(x, y)|^2 & \leq C \left[|A_t^N(x, y) - A_0^N(x, y)|^2 + \sum_k |\partial_{y_k} A_t^N(x, y) - \partial_{y_k} A_0^N(x, y)|^2 \right. \\
 & \quad \left. + \sum_{k,l} |\partial_{y_k y_l}^2 A_t^N(x, y) - \partial_{y_k y_l}^2 A_0^N(x, y)|^2 \right].
 \end{aligned} \tag{57}$$

By Lemma 37, there exists a constant $C > 0$ such that for all $N \in \mathbb{N}$ and any time $0 \leq t \leq T$ $\mathbb{E}_{c_0, w_0, b_0} [|H(B_t^N, B_0^N)|^2] \leq CN^{2(\delta+\beta-1)}$. Thus, by Tonelli's theorem we have

$$\mathbb{E}_{c_0, w_0, b_0} \left[\int_{\Omega^2} |H(B_t^N, B_0^N)(x, y)|^2 d\mu(x) d\mu(y) \right] \leq CN^{2(\delta+\beta-1)}. \tag{58}$$

Multiplying by $N^{2\delta}$ on both sides concludes the result. \blacksquare

4.1.2 DIFFERENCE BETWEEN $B_0^N(x, y)$ AND $B(x, y)$

We now characterize the difference between the kernels B_0^N and B . The law of large numbers implies B_0^N is converging to B almost surely; here we obtain a bound on its speed of convergence. Again, the proof of the following lemma is given in the appendix.

Lemma 39 *There exists $C > 0$ such that, for any $j, k, l \in \{1, 2, \dots, n\}$, $\forall N > 0$,*

$$\begin{aligned}
 \mathbb{E}_{c_0, w_0, b_0} \left[|A_0^N(x, y) - A(x, y)|^2 \right] & \leq \frac{C}{N}, \\
 \mathbb{E}_{c_0, w_0, b_0} \left[|\partial_{y_i} A_0^N(x, y) - \partial_{y_i} A(x, y)|^2 \right] & \leq \frac{C}{N}, \\
 \mathbb{E}_{c_0, w_0, b_0} \left[|\partial_{y_i y_j}^2 A_0^N(x, y) - \partial_{y_i y_j}^2 A(x, y)|^2 \right] & \leq \frac{C}{N}.
 \end{aligned} \tag{59}$$

Lemma 40 *There exists $C > 0$ such that, for all $N > 0$,*

$$\mathbb{E}_{c_0, w_0, b_0} \left[N^{2\delta} \int_{\Omega^2} |H(B_0^N, B)(x, y)|^2 d\mu(x) d\mu(y) \right] \leq \frac{C}{N^{1-2\delta}}. \tag{60}$$

Proof Similar to (51), we derive

$$\begin{aligned}
 |H(B_0^N, B)(x, y)|^2 & \leq C \left[|A_0^N(x, y) - A(x, y)|^2 + \sum_k |\partial_{y_k} A_0^N(x, y) - \partial_{y_k} A(x, y)|^2 \right. \\
 & \quad \left. + \sum_{k,l} |\partial_{y_k y_l}^2 A_0^N(x, y) - \partial_{y_k y_l}^2 A(x, y)|^2 \right].
 \end{aligned} \tag{61}$$

Therefore, Lemma 39 guarantees $\mathbb{E}_{c_0, w_0, b_0} [|H(B_0^N, B)(x, y)|^2] \leq C/N$. By Tonelli's theorem, we have

$$\begin{aligned} & \mathbb{E}_{c_0, w_0, b_0} \left[N^{2\delta} \int_{\Omega^2} |H(B_0^N, B)(x, y)|^2 d\mu(x) d\mu(y) \right] \\ &= \int_{\Omega^2} N^{2\delta} \mathbb{E}_{c_0, w_0, b_0} [|H(B_0^N, B)(x, y)|^2] d\mu(x) d\mu(y) \leq \frac{C}{N^{1-2\delta}}. \end{aligned} \quad (62)$$

■

4.1.3 DIFFERENCE BETWEEN KERNEL $B_t^N(x, y)$ AND $B(x, y)$

Combining the results from the above two subsections, we provide one of our key lemmas.

Lemma 41 *The kernels B_t^N and B satisfy*

$$\lim_{N \rightarrow \infty} \mathbb{E}_{c_0, w_0, b_0} \left[N^\delta \int_0^T \int_{\Omega^2} |H(B_u^N, B)(x, y)|^2 d\mu(x) d\mu(y) du \right] = 0. \quad (63)$$

Proof By the triangle inequality,

$$H(B_u^N, B)(x, y) \leq H(B_u^N, B_0^N)(x, y) + H(B_0^N, B)(x, y). \quad (64)$$

Therefore,

$$|H(B_u^N, B)(x, y)|^2 \leq 2|H(B_u^N, B_0^N)(x, y)|^2 + 2|H(B_0^N, B)(x, y)|^2. \quad (65)$$

Combining Lemmas 38 and 40, we have, for any $0 \leq t \leq T$,

$$N^{2\delta} \mathbb{E}_{c_0, w_0, b_0} \left[\int_{\Omega^2} |H(B_u^N, B)(x, y)|^2 d\mu(x) d\mu(y) \right] \leq CN^{2(2\delta+\beta-1)} + CN^{2\delta-1}. \quad (66)$$

Integrating with respect to time,

$$\begin{aligned} & \mathbb{E}_{c_0, w_0, b_0} \left[N^\delta \int_0^T \int_{\Omega^2} |H(B_u^N, B)(x, y)|^2 d\mu(x) d\mu(y) du \right] \\ & \leq CT(N^{2(2\delta+\beta-1)} + N^{2\delta-1}), \end{aligned} \quad (67)$$

which converges to zero as $N \rightarrow \infty$. ■

4.2 Convergence of initial approximator Q_0^N to Q_0

In this subsection, we show that the randomly initialized approximator Q_0^N converges to its limit, given by $Q_0 := f$ as the number of hidden units N goes to infinity.

Lemma 42 *The initial approximator Q_0^N satisfies $\lim_{N \rightarrow \infty} \mathbb{E}_{c_0, w_0, b_0} [\|Q_0^N - Q_0\|_{\mathcal{H}^2}^2] = 0$.*

Proof By definition, for any indices k, l ,

$$\begin{aligned} Q_0^N(x) - Q_0(x) &= \frac{\eta(x)}{N^\beta} \sum_{i=1}^N c_0^i \sigma(w_0^i \cdot x + b_0^i), \\ \partial_{x_k}(Q_0^N - Q_0)(x) &= \frac{\eta(x)}{N^\beta} \sum_{i=1}^N c_0^i \sigma'(w_0^i \cdot x + b_0^i)(w_0^i)_k + \frac{\partial_{x_k} \eta(x)}{N^\beta} \sum_{i=1}^N c_0^i \sigma(w_0^i \cdot x + b_0^i), \end{aligned} \quad (68)$$

and

$$\begin{aligned} &\partial_{x_k x_l}^2(Q_0^N - Q_0)(x) \\ &= \frac{\eta(x)}{N^\beta} \sum_{i=1}^N c_0^i \sigma''(w_0^i \cdot x + b_0^i)(w_0^i)_k (w_0^i)_l + \frac{\partial_{x_l} \eta(x)}{N^\beta} \sum_{i=1}^N c_0^i \sigma'(w_0^i \cdot x + b_0^i)(w_0^i)_k \\ &\quad + \frac{\partial_{x_k} \eta(x)}{N^\beta} \sum_{i=1}^N c_0^i \sigma'(w_0^i \cdot x + b_0^i)(w_0^i)_l + \frac{\partial_{x_k x_l}^2 \eta(x)}{N^\beta} \sum_{i=1}^N c_0^i \sigma(w_0^i \cdot x + b_0^i). \end{aligned} \quad (69)$$

As $\{c_0^i, w_0^i, b_0^i\}$ are independent for different $i \in \{1, 2, \dots, N\}$, we have, for a constant $C > 0$ which may vary from line to line,

$$\mathbb{E}_{c_0, w_0, b_0} \left[|Q_0^N(x) - Q_0(x)|^2 \right] = \frac{\eta(x)^2}{N^{2\beta-1}} \mathbb{E}_{c, w, b} [c^2 \sigma(w \cdot x + b)^2] \leq \frac{C}{N^{2\beta-1}}, \quad (70)$$

and

$$\begin{aligned} &\mathbb{E}_{c_0, w_0, b_0} \left[|\partial_{x_k} Q_0^N(x) - \partial_{x_k} Q_0(x)|^2 \right] \\ &\leq \frac{2}{N^{2\beta-1}} \mathbb{E}_{c, w, b} \left[\eta(x)^2 c^2 \sigma'(w \cdot x + b)^2 w_k^2 + (\partial_{x_k} \eta(x))^2 c^2 \sigma(w \cdot x + b)^2 \right] \leq \frac{C}{N^{2\beta-1}}. \end{aligned} \quad (71)$$

Similarly,

$$\mathbb{E}_{c_0, w_0, b_0} \left[|\partial_{x_k x_l}^2 Q_0^N(x) - \partial_{x_k x_l}^2 Q_0(x)|^2 \right] \leq \frac{C}{N^{2\beta-1}}. \quad (72)$$

Summing over all indices, then integrating over Ω , gives $\mathbb{E}_{c_0, w_0, b_0} [\|Q_0^N - Q_0\|_{\mathcal{H}^2}^2] \leq CN^{1-2\beta}$. As $\beta > 1/2$, we have the desired convergence. \blacksquare

4.3 Convergence of the approximator Q_t^N to Q_t

Definition 43 (Wide network limit) For $t \geq 0$, we define the wide network limit Q_t by the infinite-dimensional ODE

$$\frac{dQ_t}{dt}(y) = \alpha \int_{\Omega} \mathcal{L}(Q_t)(x) B(x, y) d\mu(x), \quad \text{for all } y \in \bar{\Omega}, \quad (73)$$

with initial value $Q_0(y) = f(y)$, for all $y \in \bar{\Omega}$. Equivalently, we can write

$$\frac{dQ_t}{dt} = \alpha \mathcal{B} \mathcal{L}(Q_t), \quad Q_0 = f. \quad (74)$$

Theorem 44 *The limit ODE (74) admits a unique solution in \mathcal{H}^2 .*

Proof From Lemma 29, \mathcal{B} is uniformly Lipschitz as a map $L^2 \rightarrow \mathcal{H}^2$. By Assumption 6, \mathcal{L} is uniformly Lipschitz as a map $\mathcal{H}^2 \rightarrow L^2$. Hence, the operator $\mathcal{B}\mathcal{L} : \mathcal{H}^2 \mapsto \mathcal{H}^2$ is uniformly Lipschitz. From the Picard–Lindelöf theorem (see, for example, Theorem 2.2.1 in Kolokoltsov (2019)), we know that (74) admits a unique solution with $Q_t \in \mathcal{H}^2$ for all $t \geq 0$. \blacksquare

We next prove that Q_t^N converges to Q_t , justifying the name ‘wide network limit’.

Theorem 45 (Convergence to the limit process Q_t) *For any fixed time $0 \leq t \leq T$,*

$$\lim_{N \rightarrow \infty} \mathbb{E}_{c_0, w_0, b_0} \left[\|Q_t^N - Q_t\|_{\mathcal{H}^2} \right] = 0. \quad (75)$$

In particular, the boundary value of Q_t is given by $Q_t|_{\partial\Omega} = f$.

Proof From the dynamics of Q_t^N and Q_t ((27) and (73)) we have

$$Q_t^N(y) - Q_0^N(y) = \alpha \int_0^t \int_{\Omega} F^N(\mathcal{L}Q_u^N(x)) B_u^N(x, y) d\mu(x) du, \quad (76)$$

$$Q_t(y) - Q_0(y) = \alpha \int_0^t \int_{\Omega} \mathcal{L}Q_u(x) B(x, y) d\mu(x) du. \quad (77)$$

Differentiate $Q_t^N(y)$ and $Q_t(y)$ with respect to y_k to obtain

$$\partial_{y_k} Q_t^N(y) - \partial_{y_k} Q_0^N(y) = \alpha \int_0^t \int_{\Omega} F^N(\mathcal{L}Q_u^N(x)) \partial_{y_k} B_u^N(x, y) d\mu(x) du, \quad (78)$$

$$\partial_{y_k} Q_t(y) - \partial_{y_k} Q_0(y) = \alpha \int_0^t \int_{\Omega} \mathcal{L}Q_u(x) \partial_{y_k} B(x, y) d\mu(x) du. \quad (79)$$

Twice-differentiate Q_t^N and Q_t with respect to y_k and y_l to obtain:

$$\partial_{y_k y_l}^2 Q_t^N(y) - \partial_{y_k y_l}^2 Q_0^N(y) = \alpha \int_0^t \int_{\Omega} F^N(\mathcal{L}Q_u^N(x)) \partial_{y_k y_l}^2 B_u^N(x, y) d\mu(x) du, \quad (80)$$

$$\partial_{y_k y_l}^2 Q_t(y) - \partial_{y_k y_l}^2 Q_0(y) = \alpha \int_0^t \int_{\Omega} \mathcal{L}Q_u(x) \partial_{y_k y_l}^2 B(x, y) d\mu(x) du. \quad (81)$$

Subtracting (76) from (77), we have

$$\begin{aligned} & |Q_t^N(y) - Q_t(y)| - |Q_0^N(y) - Q_0(y)| \\ & \leq \alpha \int_0^t \left| \int_{\Omega} F(Q_u^N(x)) B_u^N(x, y) - \mathcal{L}Q_u(x) B(x, y) d\mu(x) \right| du \\ & \leq \alpha \int_0^t \left| \int_{\Omega} F^N(\mathcal{L}Q_u^N(x)) (B_u^N(x, y) - B(x, y)) d\mu(x) \right| du \\ & \quad + \alpha \int_0^t \left| \int_{\Omega} (F^N(\mathcal{L}Q_u^N(x)) - F^N(\mathcal{L}Q_u(x))) B(x, y) d\mu(x) \right| du \\ & \quad + \alpha \int_0^t \left| \int_{\Omega} (F^N(\mathcal{L}Q_u(x)) - \mathcal{L}Q_u(x)) B(x, y) d\mu(x) \right| du. \end{aligned} \quad (82)$$

Subtracting (78) from (79), we have

$$\begin{aligned}
 & |\partial_{y_k} Q_t^N(y) - \partial_{y_k} Q_t(y)| - |\partial_{y_k} Q_0^N(y) - \partial_{y_k} Q_0(y)| \\
 & \leq \alpha \int_0^t \left| \int_{\Omega} F^N(\mathcal{L}Q_u^N(x)) \left(\partial_{y_k} B_u^N(x, y) - \partial_{y_k} B(x, y)(x, y) \right) d\mu(x) \right| du \\
 & \quad + \alpha \int_0^t \left| \int_{\Omega} \left(F^N(\mathcal{L}Q_u^N(x)) - F^N(\mathcal{L}Q_u(x)) \right) \partial_{y_k} B(x, y) d\mu(x) \right| du \\
 & \quad + \alpha \int_0^t \left| \int_{\Omega} \left(F^N(\mathcal{L}Q_u(x)) - \mathcal{L}Q_u(x) \right) \partial_{y_k} B(x, y) d\mu(x) \right| du.
 \end{aligned} \tag{83}$$

Similarly, subtracting (80) from (81), we have

$$\begin{aligned}
 & |\partial_{y_k y_l}^2 Q_t^N(y) - \partial_{y_k y_l}^2 Q_t(y)| - |\partial_{y_k y_l}^2 Q_0^N(y) - \partial_{y_k y_l}^2 Q_0(y)| \\
 & \leq \alpha \int_0^t \left| \int_{\Omega} F^N(\mathcal{L}Q_u^N(x)) \left(\partial_{y_k y_l}^2 B_u^N(x, y) - \partial_{y_k y_l}^2 B(x, y)(x, y) \right) d\mu(x) \right| du \\
 & \quad + \alpha \int_0^t \left| \int_{\Omega} \left(F^N(\mathcal{L}Q_u^N(x)) - F^N(\mathcal{L}Q_u(x)) \right) \partial_{y_k y_l}^2 B(x, y) d\mu(x) \right| du \\
 & \quad + \alpha \int_0^t \left| \int_{\Omega} \left(F^N(\mathcal{L}Q_u(x)) - \mathcal{L}Q_u(x) \right) \partial_{y_k y_l}^2 B(x, y) d\mu(x) \right| du.
 \end{aligned} \tag{84}$$

As stated in Definition 17, the function $F^N = (\psi^N) \cdot (\psi^N)'$ satisfies a global Lipschitz condition. Therefore, for some $C > 0$,

$$\begin{aligned}
 & \left| F^N(\mathcal{L}Q_t^N(x)) - F^N(\mathcal{L}Q_t(x)) \right| < C \left| \mathcal{L}Q_t^N(x) - \mathcal{L}Q_t(x) \right| \\
 & \leq C \left(|Q_t^N(x) - Q_t(x)| + \sum_k |\partial_{x_k} Q_t^N(x) - \partial_{x_k} Q_t(x)| + \sum_{k,l} |\partial_{x_k x_l}^2 Q_t^N(x) - \partial_{x_k x_l}^2 Q_t(x)| \right) \\
 & = CG(Q_t^N, Q_t)(x).
 \end{aligned} \tag{85}$$

Summing (82), (83) and (84) over all indices k and l , we have

$$\begin{aligned}
 & G(Q_t^N, Q_t)(y) \\
 & = |Q_t^N(y) - Q_t(y)| + \sum_k |\partial_{x_k} Q_t^N(y) - \partial_{x_k} Q_t(y)| + \sum_{k,l} |\partial_{x_k x_l}^2 Q_t^N(y) - \partial_{x_k x_l}^2 Q_t(y)| \\
 & \leq G(Q_t^N, Q_0)(y) + CN^\delta \int_0^t \int_{\Omega} H(B_u^N, B)(x, y) d\mu(x) du \\
 & \quad + C \int_0^t \int_{\Omega} G(Q_u^N, Q_u)(x) H(B, 0)(x, y) d\mu(x) du \\
 & \quad + C \int_0^t \int_{\Omega} |F^N(\mathcal{L}Q_u(x)) - \mathcal{L}Q_u(x)| H(B, 0)(x, y) d\mu(x) du.
 \end{aligned} \tag{86}$$

Here, the coefficient N^δ comes from the term F^N , as by Assumption 17, $|F^N| = |(\psi^N) \cdot (\psi^N)'| \leq 2N^\delta$. As mentioned in Remark 18, we also know $|F^N(x) - x| \leq 2|x| \mathbb{1}_{\{|x| \geq N^\delta\}}$. As

$H(B, 0)$ is uniformly bounded (Lemma 25), from inequality (86), for some $C > 0$ we have

$$\begin{aligned}
 G(Q_t^N, Q_t)(y) &\leq G(Q_t^N, Q_0)(y) + CN^\delta \int_0^t \int_\Omega H(B_u^N, B)(x, y) d\mu(x) du \\
 &\quad + C \int_0^t \int_\Omega G(Q_u^N, Q_u)(x) d\mu(x) du \\
 &\quad + C \int_0^t \int_\Omega |\mathcal{L}Q_u(x)| \mathbb{1}_{\{|\mathcal{L}Q_u(x)| \geq N^\delta\}} d\mu(x) du.
 \end{aligned} \tag{87}$$

By multiplying $G(Q_t^N, Q_t)(y)$ in inequality (87), integrating on both sides, and applying the Cauchy–Schwarz inequality, from (87) we derive

$$\begin{aligned}
 \int_\Omega |G(Q_t^N, Q_t)(y)|^2 d\mu(y) &\leq C \left(\int_\Omega |G(Q_t^N, Q_t)(y)|^2 d\mu(y) \right)^{\frac{1}{2}} \left[\left(\int_\Omega |G(Q_0^N, Q_0)(y)|^2 d\mu(y) \right)^{\frac{1}{2}} \right. \\
 &\quad + N^\delta \int_0^t \left(\int_{\Omega^2} |H(B_u^N, B)(x, y)|^2 d\mu(x) d\mu(y) \right)^{\frac{1}{2}} du \\
 &\quad + \int_0^t \int_\Omega |G(Q_u^N, Q_u)(x)|^2 d\mu(x) du \\
 &\quad \left. + \int_0^t \int_\Omega |\mathcal{L}Q_u(x)| \mathbb{1}_{\{|\mathcal{L}Q_u(x)| \geq N^\delta\}} d\mu(x) du \right].
 \end{aligned} \tag{88}$$

We write

$$\begin{aligned}
 J^N &:= \left(\int_\Omega |(Q_0^N, Q_0)(y)|^2 d\mu(y) \right)^{\frac{1}{2}} + N^\delta \int_0^T \int_{\Omega^2} |H(B_u^N, B)(x, y)|^2 d\mu(x) d\mu(y) du \\
 &\quad + \int_0^T \int_\Omega |\mathcal{L}Q_u(x)| \mathbb{1}_{\{|\mathcal{L}Q_u(x)| \geq N^\delta\}} d\mu(x) du.
 \end{aligned} \tag{89}$$

Then the inequality (87) can be formulated as

$$\left(\int_\Omega |G(Q_t^N, Q_t)(y)|^2 d\mu(y) \right)^{\frac{1}{2}} \leq J^N + C \int_0^t \left(\int_\Omega |G(Q_u^N, Q_u)(x)|^2 d\mu(x) \right)^{\frac{1}{2}} du. \tag{90}$$

Since

$$G(Q_t^N, Q_t)(y) = |Q_t^N(y) - Q_t(y)| + \sum_k |\partial_{y_k}(Q_t^N - Q_t)(y)| + \sum_{k,l} |\partial_{y_k, y_l}^2(Q_t^N - Q_t)(y)|, \tag{91}$$

by a simple quadratic-mean inequality we have

$$\|Q_t^N - Q_t\|_{\mathcal{H}^2}^2 \leq \int_\Omega |G(Q_t^N, Q_t)(y)|^2 d\mu(y) \leq (n^2 + n + 1) \|Q_t^N - Q_t\|_{\mathcal{H}^2}^2. \tag{92}$$

Consequently, from (90) we derive

$$\|Q_t^N - Q_t\|_{\mathcal{H}^2} \leq J^N + C \int_0^t \|Q_u^N - Q_u\|_{\mathcal{H}^2} du. \quad (93)$$

By applying Grönwall's inequality, (93) yields the exponential bound $\int_0^t \|Q_u^N - Q_u\|_{\mathcal{H}^2} du \leq \frac{J^N}{C}(e^{Ct} - 1)$.

It remains to bound J^N in (89), as N goes to infinity. By Lemma 42 and (92), $\mathbb{E}_{c_0, w_0, b_0} \left[\int_{\Omega} |G(Q_0^N, Q_0)(y)|^2 d\mu(y) \right] \rightarrow 0$. By the dominated convergence theorem, we obtain $\int_0^T \int_{\Omega} |\mathcal{L}Q_u(x)| \mathbb{1}_{\{|\mathcal{L}Q_u(x)| \geq N^\delta\}} d\mu(x) du \rightarrow 0$. Finally, by Lemma 41, we show that $\mathbb{E}_{c_0, w_0, b_0} \left[N^\delta \int_0^T \int_{\Omega^2} |H(B_u^N, B)(x, y)|^2 d\mu(x) d\mu(y) du \right] \rightarrow 0$. Therefore, $\mathbb{E}_{c_0, w_0, b_0} [J^N] \rightarrow 0$ as $N \rightarrow \infty$ and consequently, from the above exponential bound,

$$\lim_{N \rightarrow \infty} \mathbb{E}_{c_0, w_0, b_0} \left[\int_0^T \|Q_u^N - Q_u\|_{\mathcal{H}^2} du \right] = 0. \quad (94)$$

Substituting this result back into inequality (93) and taking the expectation with respect to the random initialization leads to $\lim_{N \rightarrow \infty} \mathbb{E}_{c_0, w_0, b_0} [\|Q_t^N - Q_t\|_{\mathcal{H}^2}] = 0$ for any time $t \leq T$.

Finally, observe that $Q_t^N|_{\partial\Omega} = f$ by construction, so almost sure convergence in \mathcal{H}^2 (for a subsequence of N) and continuity of the boundary trace operator imply $Q_t|_{\partial\Omega} = f$. ■

5. Convergence for large training time

In the previous section, we showed that, as the neural network is made wider, its training process converges to a process that satisfies an infinite-dimensional ODE (74). In particular, the wide-limit of the approximate solution has the dynamics $\frac{dQ_t}{dt} = \alpha \mathcal{B}\mathcal{L}(Q_t)$. This limit process can therefore be regarded as an approximation of the setting where we use a large (single-layer) neural network, and our problem (in the limiting setting) is transformed into the study of this infinite-dimensional dynamical system. The following theorem is an easy consequence.

Theorem 46 *If the wide-network limit converges in \mathcal{H}^1 to a fixed point Q , then Q is a solution to the PDE (1).*

Proof For a fixed point Q , we know $\mathcal{B}\mathcal{L}(Q) = 0$. From Lemma 35, this implies that $\mathcal{L}(Q) = 0$, so the PDE dynamics are satisfied. As Q^N satisfies the boundary conditions $Q_t^N|_{\partial\Omega} = f$, and $Q_t^N \rightarrow Q_t$ in \mathcal{H}^2 for each t , and $Q_t \rightarrow Q$ in \mathcal{H}^1 we see that Q must satisfy the boundary condition also (by \mathcal{H}^1 -continuity of the boundary trace operator). ■

In this section, we will consider the simple case where \mathcal{L} is a monotone operator. This implies that the dynamical system $dv/dt = \mathcal{L}v$ converges in L^2 exponentially quickly, which suggests the dynamics of Q will be well behaved. However, the presence of the operator \mathcal{B} leads to some difficulties in our analysis. Nevertheless, we will show that, in this setting, our approximation Q_t converges to the true solution u of the PDE, at least for a generic subsequence of times.

We assume $\alpha = 1$ in (16) for notational simplicity in this section, without loss of generality (as this is a simple rescaling of time).

Theorem 47 *Let u be the solution to the PDE (1) under our assumptions (Assumption 13, now including part iv), and assume the neural network configuration satisfies Assumptions 14 and 16. Then the wide network limit $\{Q_t\}_{t \geq 0}$ of the Q-PDE algorithm (Definitions 20, 43) satisfies*

$$\frac{1}{t} \int_0^t \|Q_s - u\|_{L^2}^2 ds \rightarrow 0 \text{ as } t \rightarrow \infty. \quad (95)$$

In particular, there exists a set $A \subset [0, \infty)$ with $\lim_{t \rightarrow \infty} (t^{-1} \text{Leb}(A \cap [0, t])) = 0$ such that, for all sequences $t_k \rightarrow \infty$ which do not take values in A , we have the (strong) L^2 convergence

$$\|Q_{t_k} - u\|_{L^2} \rightarrow 0.$$

Proof We begin by supposing we have a decomposition

$$Q_0 = f = u + \mathcal{B}v_0 + g, \quad (96)$$

for some $v_0 \in L^2_\eta$ and $g \in \mathcal{H}^2$. As $u \in \mathcal{H}^2$, we can choose v_0 arbitrarily and then find g using (96). Consider the process v solving the ODE

$$\frac{dv_t}{dt} = \mathcal{L}(\mathcal{B}v_t + u + g) \quad (97)$$

with initial value v_0 as in (96). As detailed in Appendix A.6, as \mathcal{B} and \mathcal{L} are both Lipschitz continuous in appropriate spaces, a standard Picard–Lindelöf argument shows that v is uniquely defined in L^2_η , for all $t \geq 0$.

By differentiating $\mathcal{B}v_t$ (and using dominated convergence to move the derivative through \mathcal{B}), we obtain

$$\frac{d(\mathcal{B}v_t + u + g)}{dt} = \frac{d(\mathcal{B}v_t)}{dt} = \mathcal{B}\mathcal{L}(\mathcal{B}v_t + u + g); \quad \mathcal{B}v_0 + u + g = Q_0. \quad (98)$$

In particular, as the ODE defining Q has a unique solution in \mathcal{H}^2 (Theorem 44), we have the identity $Q_t = \mathcal{B}v_t + u + g$.

Using the Young inequality $2\langle x, y \rangle \leq \gamma\|x\|^2 + (4/\gamma)\|y\|^2$ for $\gamma > 0$, together with the fact $\mathcal{L}(u) = 0$ and our assumption that \mathcal{L} is strongly monotone and Lipschitz, for some constant $k > 0$ we have

$$\begin{aligned} \frac{d}{dt} \|v_t\|_{\mathcal{B}}^2 &= \frac{d}{dt} \langle \mathcal{B}v_t, v_t \rangle = 2\langle \mathcal{B}v_t, \mathcal{L}(\mathcal{B}v_t + u + g) \rangle \\ &= 2\langle \mathcal{B}v_t, \mathcal{L}(\mathcal{B}v_t + u + g) - \mathcal{L}(u + g) \rangle + 2\langle \mathcal{B}v_t, \mathcal{L}(u + g) - \mathcal{L}(u) \rangle \\ &\leq -2\gamma\|\mathcal{B}v_t\|_{L^2}^2 + \gamma\|\mathcal{B}v_t\|_{L^2}^2 + \frac{4}{\gamma}\|\mathcal{L}(u + g) - \mathcal{L}(u)\|_{L^2}^2 \\ &\leq -\gamma\|\mathcal{B}v_t\|_{L^2}^2 + \gamma k\|g\|_{\mathcal{H}^2}^2. \end{aligned} \quad (99)$$

From Lemma 35 there exists $\lambda > 0$ such that, for all $v \in L_\eta^2$, $\|\mathcal{B}v_t\|_{L^2}^2 \leq \lambda \|v_t\|_{\mathcal{B}}$. Using this value of λ , as $\|g\|_{\mathcal{H}^2}^2$ is a constant, $\frac{d}{dt}(\lambda \|v_t\|_{\mathcal{B}}^2 - k \|g\|_{\mathcal{H}^2}^2) \leq -\gamma \lambda (\|\mathcal{B}v_t\|_{L^2}^2 - k \|g\|_{\mathcal{H}^2}^2)$. Using the definition of λ and integrating,

$$\begin{aligned} \left(\|\mathcal{B}v_t\|_{L^2}^2 - k \|g\|_{\mathcal{H}^2}^2 \right) &\leq \left(\lambda \|v_t\|_{\mathcal{B}}^2 - k \|g\|_{\mathcal{H}^2}^2 \right) \\ &\leq \left(\lambda \|v_0\|_{\mathcal{B}}^2 - k \|g\|_{\mathcal{H}^2}^2 \right) - \gamma \lambda \int_0^t \left(\|\mathcal{B}v_s\|_{L^2}^2 - k \|g\|_{\mathcal{H}^2}^2 \right) ds. \end{aligned} \quad (100)$$

Writing $h_t = \int_0^t (\|\mathcal{B}v_s\|_{L^2}^2 - k \|g\|_{\mathcal{H}^2}^2) ds$, we see $dh_t/dt \leq (\lambda \|v_0\|_{\mathcal{B}}^2 - k \|g\|_{\mathcal{H}^2}^2) - \gamma \lambda h_t$, and hence Grönwall's inequality yields

$$h_t = \int_0^t \left(\|\mathcal{B}v_s\|_{L^2}^2 - k \|g\|_{\mathcal{H}^2}^2 \right) ds \leq (\lambda \|v_0\|_{\mathcal{B}}^2 - k \|g\|_{\mathcal{H}^2}^2) \int_0^t e^{-\gamma \lambda (t-s)} ds. \quad (101)$$

Recalling that $Q_t - u = \mathcal{B}v_t + g$ and

$$\|\mathcal{B}v + g\|_{L^2}^2 \leq 2\|\mathcal{B}v\|_{L^2}^2 + 2\|g\|_{\mathcal{H}^2}^2 = 2(\|\mathcal{B}v\|_{L^2}^2 - k \|g\|_{\mathcal{H}^2}^2) + (2k + 2)\|g\|_{\mathcal{H}^2}^2, \quad (102)$$

we conclude that, for some constant $C > 0$,

$$\int_0^t \|Q_s - u\|_{L^2}^2 ds = \int_0^t \|\mathcal{B}v_s + g\|_{L^2}^2 ds \leq Ct \|g\|_{\mathcal{H}^2}^2 + C \|v_0\|_{\mathcal{B}}^2. \quad (103)$$

This inequality must hold for all choices of g and v_0 satisfying (96), and we can choose them to optimize (103). We observe that $Q_0 - u \in \mathcal{H}_{(0)}^2$, as $Q_0, u \in \mathcal{H}^2$ and they have the same boundary value. For every $\epsilon > 0$, by Theorem 33, there exists a choice of $v_0^{(\epsilon)} \in L_\eta^2$ such that $\|Q_0 - u - \mathcal{B}v_0^{(\epsilon)}\|_{\mathcal{H}^2} \leq \epsilon$. We further recall that $\|v_0^{(\epsilon)}\|_{\mathcal{B}}^2 < \infty$ for $v_0^{(\epsilon)} \in L_\eta^2$.

Defining $g^{(\epsilon)} = Q_0 - u - \mathcal{B}v_0^{(\epsilon)}$, (103) yields the limit as $t \rightarrow \infty$

$$\frac{1}{t} \int_0^t \|Q_s - u\|_{L^2}^2 ds \leq C \|g^{(\epsilon)}\|_{\mathcal{H}^2}^2 + \frac{C}{t} \|v_0^{(\epsilon)}\|_{\mathcal{B}}^2 \rightarrow C\epsilon. \quad (104)$$

As $\epsilon > 0$ was arbitrary, we conclude that $\lim_{t \rightarrow \infty} \left\{ \frac{1}{t} \int_0^t \|Q_s - u\|_{L^2}^2 ds \right\} = 0$.

To obtain the final statement, set $\beta(t) = \sup_{T \geq t} \left\{ \frac{1}{T} \int_0^T \|Q_s - u\|_{L^2}^2 ds \right\}$. Observe that β is a nonincreasing positive function with $\lim_{t \rightarrow \infty} \beta(t) = 0$. Defining the set $A = \{t : \|Q_t - u\|^2 > \sqrt{\beta(t)}\}$, Markov's inequality yields

$$\begin{aligned} \frac{1}{t} \text{Leb}(A \cap [0, t]) &= \frac{1}{t} \text{Leb} \left(\left\{ s \leq t : \frac{\|Q_s - u\|_{L^2}^2}{\sqrt{\beta(s)}} > 1 \right\} \right) \leq \frac{1}{t} \int_0^t \frac{\|Q_s - u\|_{L^2}^2}{\sqrt{\beta(s)}} ds \\ &\leq \frac{1}{\sqrt{\beta(t)}} \left(\frac{1}{t} \int_0^t \|Q_s - u\|_{L^2}^2 ds \right) \leq \sqrt{\beta(t)} \rightarrow 0. \end{aligned} \quad (105)$$

■

The above result only considers the convergence of Q . In the case where \mathcal{L} is Gateaux differentiable (which is certainly the case when \mathcal{L} is linear), we can give a similar result for the convergence of $\mathcal{L}Q_t$, involving the positive-definite kernel \mathcal{B} .

Theorem 48 *Suppose $\mathcal{L} : \mathcal{H}^2 \rightarrow L^2$ is Gateaux differentiable and the conditions of Theorem 47 hold. Then there exists $K > 0$ such that $\int_0^t \|\mathcal{B}\mathcal{L}Q_s\|_{L^2}^2 ds \leq K$. In particular, there exists a set $A \subset [0, \infty)$ with $\text{Leb}(A) < \infty$ such that for all sequences $t_k \rightarrow \infty$ which do not take values in A , we have $\|\mathcal{B}\mathcal{L}Q_{t_k}\|_{L^2} \rightarrow 0$.*

Proof For any $w \in L^2$, $h \in \mathcal{H}^2$ we write $\partial\mathcal{L}(w; h) = \lim_{\epsilon \rightarrow 0} \{(\mathcal{L}(w+h) - \mathcal{L}(w))/\epsilon\}$ for the Gateaux derivative in direction h evaluated at w , the limit being taken in L^2 .

We first observe that, as \mathcal{L} is strongly monotone, for any $h \in \mathcal{H}^2$,

$$\langle h, \partial\mathcal{L}(w; h) \rangle = \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \langle h, \mathcal{L}(w + \epsilon h) - \mathcal{L}(w) \rangle \leq -\gamma \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon^2} \|\epsilon h\|_{L^2}^2 = -\gamma \|h\|_{L^2}^2. \quad (106)$$

Now write $v_t = \mathcal{L}Q_t$. As $Q_t \in \mathcal{H}^2$, we know that $\mathcal{B}v_t = \mathcal{B}\mathcal{L}Q_t \in \mathcal{H}^2$ (as $\mathcal{L} : \mathcal{H}^2 \rightarrow L^2$ and $\mathcal{B} : L^2 \rightarrow \mathcal{H}^2$). By the chain rule, we have $dv_t/dt = \partial\mathcal{L}(Q_t; dQ_t/dt) = \partial\mathcal{L}(Q_t; \mathcal{B}\mathcal{L}Q_t) = \partial\mathcal{L}(Q_t; \mathcal{B}v_t)$. Therefore,

$$\frac{d}{dt} \|v_t\|_{\mathcal{B}}^2 = 2 \left\langle \mathcal{B}v_t, \frac{dv_t}{dt} \right\rangle = 2 \left\langle \mathcal{B}v_t, \partial\mathcal{L}(Q_t; \mathcal{B}v_t) \right\rangle \leq -2\gamma \|\mathcal{B}v_t\|_{L^2}^2. \quad (107)$$

Lemma 35 gives the bound $\|\mathcal{B}v_t\|_{L^2}^2 \leq \lambda \|v_t\|_{\mathcal{B}}$, from which we obtain

$$\|\mathcal{B}v_t\|_{L^2}^2 \leq \lambda \|v_t\|_{\mathcal{B}}^2 \leq \lambda \|v_0\|_{\mathcal{B}}^2 - 2\lambda\gamma \int_0^t \|\mathcal{B}v_s\|_{L^2}^2 ds. \quad (108)$$

By Grönwall's inequality, we conclude

$$\int_0^t \|\mathcal{B}v_s\|_{L^2}^2 ds \leq \lambda \|v_0\|_{\mathcal{B}}^2 \int_0^t e^{-2\lambda\gamma(t-s)} ds \leq \frac{\|v_0\|_{\mathcal{B}}^2}{2\gamma} =: K < \infty. \quad (109)$$

Taking $A = \{t : \|\mathcal{B}v_t\|_{L^2}^2 > 1/t\}$, the final stated property is obtained from Markov's inequality, similarly to in Theorem 47. \blacksquare

Corollary 49 *Under the conditions and notation of Theorem 48, the sequence $\mathcal{L}Q_{t_k}$ converges weakly to zero in L^2 if and only if it remains bounded in L^2 .*

Proof Recall that a sequence x_n converges L^2 -weakly to zero if and only if it is L^2 -bounded and there exists an L^2 -dense set \mathcal{Y} such that $\langle y, x_n \rangle \rightarrow 0$ for all $y \in \mathcal{Y}$.

We know from Theorem 48 that $\mathcal{B}\mathcal{L}Q_{t_k} \rightarrow 0$ strongly in L^2 so, for any $v \in L^2$, $\langle \mathcal{B}v, \mathcal{L}Q_{t_k} \rangle = \langle v, \mathcal{B}\mathcal{L}Q_{t_k} \rangle \rightarrow 0$. In particular, we see that $\langle y, Q_{t_k} \rangle \rightarrow 0$ for any $y \in \text{im}(\mathcal{B})$. We know from Theorem 33 that $\text{im}(\mathcal{B})$ is dense in \mathcal{H}^2 , and hence in L^2 . The result follows. \blacksquare

6. Numerical experiments

In this section, we present numerical results where we apply our algorithm to solve a family of partial differential equations. The approximator matches the solution of the differential equation closely in a relatively short period of training time in these test cases.

6.1 Discretization of the continuous-time algorithm

As a reminder, in a continuous time setting, the algorithm follows a biased gradient flow: $d\theta_t/dt = -\alpha_t^N G_t^N(\theta_t)$, based on the biased gradient estimator

$$G_t^N(\theta_t) = \int_{\Omega} 2\psi(\mathcal{L}Q_t^N(x))\psi'(\mathcal{L}Q_t^N(x))\nabla_{\theta}(-\gamma Q_t^N(x))d\mu(x). \quad (110)$$

A natural discretization of our algorithm with which we train our approximator is

$$\theta_{t+1} - \theta_t = -\alpha_t^N G_{M,t}^{\hat{N}}, \quad (111)$$

where

$$G_{M,t}^{\hat{N}} = \frac{1}{M} \sum_{i=1}^M 2\psi(\mathcal{L}Q_t^N(x_t^i))\psi'(\mathcal{L}Q_t^N(x_t^i))\nabla_{\theta}(-\gamma Q_t^N(x_t^i)) \quad (112)$$

is an unbiased estimator of G_t^N , given $\{x_t^i\}_{i=1,2,\dots,M}$ the random grid independently sampled from distribution μ at time step t , and α_t^N is the learning rate. More advanced discretization methods, for example the ADAM adaptive gradient descent rule (Kingma and Ba (2014)) can also be used, and (as is common in gradient based optimization problems) give a noticeable improvement in performance.

The Q-learning algorithm for solving PDEs as follows:

Algorithm 1: Q-PDE Algorithm

Parameters: Hyper-parameters of the single-layer neural network; Domain Ω ; PDE operator \mathcal{L} ; Boundary condition at $\partial\Omega$; Sampling measure μ ; Number of Monte Carlo points M ;

Upper bound of training time T ;

Initialise: Neural net S^N ; Auxiliary function η ; Approximator Q^N based on S^N and η ; Smoothing function ψ^N ; Learning rate scheduler $\{\alpha_t^N\}_{t \geq 0}$; Stopping criteria ϵ ; Current time $t = 0$.

while $err \geq \epsilon$ and $t \leq T$ **do**

Sample M points in Ω using μ , $\{x_i\}$;

Compute biased gradient estimator $G_{M,t}^N$ using (112);

Update neural network parameters via (111);

Compute $err = \frac{1}{M} \sum_{i=1}^M \psi^N(\mathcal{L}Q_t^N(x_i))^2$;

Update time t ;

end

Return approximator Q_t^N .

We implement this algorithm using PyTorch (with backpropagation to compute gradients in both parameters and the state variable x) and the standard ADAM optimization scheduler for α .²

2. The implementation is available at <https://github.com/DeqingJ/QPDE>.

6.2 Test equation: Survival time of a Brownian motion

Suppose Ω is the n -dimensional unit ball. For n -dimensional Brownian motion X_t starting from point $x \in \mathbb{R}^n$, we wish to compute

$$u(x) := \mathbb{E} \left[\int_0^{\tau_\Omega} e^{-\gamma t} dt \mid X_0 = x \right] \tag{113}$$

where τ_Ω is the first exit time of X for domain Ω . By the Feynman–Kac theorem (applied to the process X stopped at τ_Ω), this expectation is given by the solution of the following PDE:

$$\begin{cases} 1 - \gamma u + \frac{1}{2} \Delta u &= 0 & \text{in } \Omega \\ u &= 0 & \text{on } \partial\Omega, \end{cases} \tag{114}$$

where Δu is the Laplacian of u . We consider this problem as a good test case as its explicit solution is known to us, allowing us to check if our algorithm works for a high-dimensional version of this PDE. We initialize a single neural network S^N equipped with sigmoid activation function, and introduce auxiliary function $\eta(x) = 1 - \|x\|^2$. The approximator $Q = S \cdot \eta$ is trained to fit the solution. As shown in the following subsections, through our algorithm, the approximator Q learns to have the solution of the PDE with different discounting factor in different dimensions. For detailed configuration of our approximator, see Appendix A.8.

6.2.1 1-DIMENSIONAL CASE

For dimension $n = 1$, domain $\Omega = (-1, 1)$. The exact solution of this differential equation is $u(x) = \frac{1}{\gamma} + c_1(e^{\sqrt{2\gamma}x} + e^{-\sqrt{2\gamma}x})$ where $c_1 = -1/(\gamma(e^{-\sqrt{2\gamma}} + e^{\sqrt{2\gamma}}))$. For this subsection, set $\gamma = 0.1$. To keep track of the training progress, we monitor the average loss level at time t , that is $e_t \approx \int_\Omega [\mathcal{L}Q_t]^2 d\mu$ where μ is the Lebesgue measure on Ω , which can be estimated using our sample of evaluation points at each time.




Figure 1: Loss level during and after training

The left side of Figure 1 shows how the average operator loss smoothly decays during training. The right side of it plots the operator error at terminal time $\mathcal{L}Q_T^N(x)$. For a

perfect fit, the operator loss $\mathcal{L}Q_T(x)$ should be zero across the entire domain. We observe that for our approximator, the norm of this loss remains small and fluctuates around 0. In Figure 2, we compare our approximator Q_τ^N with the exact solution. Here τ indicates the training time at which we observed the minimal loss e_τ , which occurs for τ close to the terminal time T . (Given the fluctuations observed in the loss, choosing this minimal point can give a noticeable qualitative improvement in performance.) The relative error of our approximation at x , given by $|(Q_\tau^N(x) - u(x))/u(x)|$, remains lower than 0.02%.

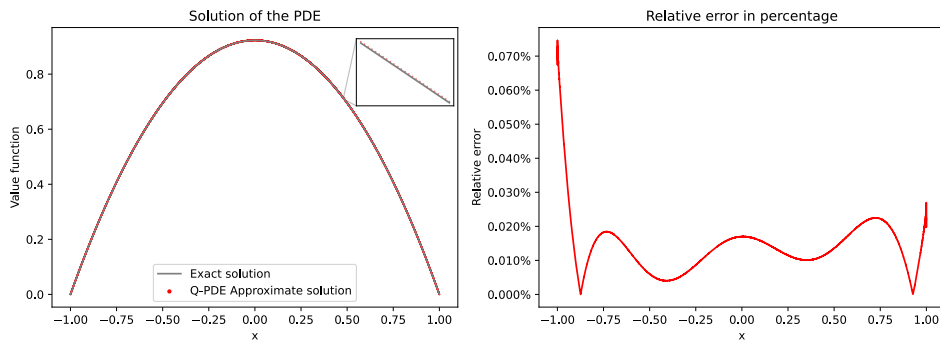


Figure 2: Approximate and exact solutions in 1d case

6.2.2 20-DIMENSIONAL CASE

Now we consider the 20-dimensional case. High-dimensional PDEs are typically hard to solve, as mesh grids are computationally unaffordable in these cases. However, as a mesh-free method, our algorithm can be applied. We numerically solve the PDE (114) for discounting factor $\gamma = 0.2$ via our algorithm and compare the result with the exact solution, given by

$$u(x) = \frac{1}{\gamma} \left(1 - \frac{J_9(i\sqrt{2\gamma}\|x\|)}{\|x\|^9 J_9(i\sqrt{2\gamma})} \right), \quad (115)$$

where J_9 is the ninth-order Bessel function of the first kind. We also compare the training results of our method to the results from the Deep Galerkin Method (which does stochastic gradient descent with an unbiased gradient estimate) as a benchmark.

Our approximator trained with the Q-PDE algorithm is able to accurately approximate the solution of the 20-dimensional PDE. Left of Figure 3 shows that the operator loss is less than 10^{-5} after training for 40,000 epochs, for both Q-PDE and DGM. For both methods, we observe cyclic behavior in the loss during training, which appears similar to the slingshot effect discussed in Thilak et al. (2022), which is potentially a consequence of momentum accumulation of the ADAM optimizer. A full understanding of this effect (which we have not observed when using a simple gradient descent optimizer) is left for future work.




Figure 3: Average loss level during training

The solution u of this PDE is radially symmetric, that is, it can be represented by a function which depends on $\|x\|$ rather than x . For the right side of Figure 3, we verify that our approximate solution is close to this function for different $\|x\| \in [0, 1]$. For $0 \leq r \leq 1$, we randomly generate 100 sample points $\{x_{j,r}\}_{j=1,\dots,100}$ on the sphere $\|x\| = r$ and compute the mean-squared error of the set of values $\{Q_\tau^N(x_{i,r})\}$ as compared to the exact solution. It is shown that for different r , the mean-squared error of the neural network approximator is consistently small. Q-PDE solution obtains smaller error than our DGM solution, but they are at the same scale.

For $0 \leq r \leq 1$, we also consider the maximum relative error, defined as

$$R_r = \max_{x_{i,r}} \left\{ \left| \frac{Q_\tau^N(x_{i,r}) - u(x_{i,r})}{u(x_{i,r})} \right| \right\}. \quad (116)$$

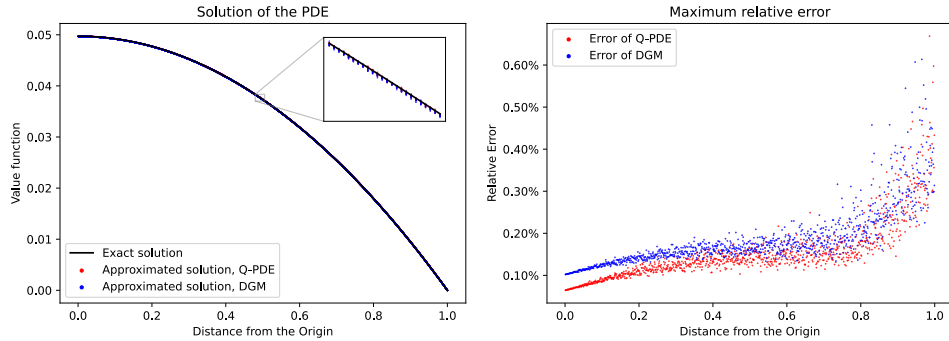


Figure 4: Relative and absolute errors of the approximate solution, in terms of distance from the origin.

The left plot of Figure 4 implies that this error remains small for all $0 \leq r \leq 1$. Again, DGM solution admits slightly larger relative error. Meanwhile, as DGM solution is trained

with exact gradients which are computationally more expensive, consequently, it takes 5 times longer to train. Both DGM and Q-PDE solutions reach high accuracy. We plot a scatter plot for $\{Q_\tau^N(w_i)\}$ in the right side in Figure 4 and compare with the exact solution u . The approximate solutions closely matches the exact solution.

7. Conclusion

In this paper, we propose a novel algorithm which numerically solves elliptic PDEs. The neural network approximator is designed to match the Dirichlet boundary condition of the PDE and is trained to satisfy the PDE operator. We prove that for single-layer neural network approximators, as the number of hidden units $\rightarrow \infty$, the evolution of the approximator during training is characterized by a limit ODE. We further prove that as the training time $\rightarrow \infty$, the solution of the limit ODE converges to the classical solution of the PDE. In addition, we provide a numerical test case to show that our algorithm can numerically solve (high-dimensional) PDEs.

Acknowledgments

Thanks to Endre Süli, Lukasz Szpruch, Christoph Reisinger and Rama Cont, and to anonymous referees, for useful conversations and suggestions.

This publication is based on work supported by the EPSRC Centre for Doctoral Training in Industrially Focused Mathematical Modelling (InFoMM) (EP/L015803/1) in collaboration with the Alan Turing Institute. Samuel Cohen also acknowledges the support of the Oxford–Man Institute for Quantitative Finance, and the UKRI Prosperity Partnership Scheme (FAIR) under the EPSRC Grant EP/V056883/1.

The order of authorship in this paper was determined alphabetically.

Appendix A. Appendix

In our proofs, the constants K, C may vary from line to line.

A.1 Proof of Lemma 3

Proof We first prove³ statement *i*. Suppose that $u \in \mathcal{H}^2(\Omega) \cap \mathcal{H}_0^1(\Omega)$. Consider the Poisson problem solved by u :

$$\begin{cases} -\Delta u = f & \text{in } \Omega, \\ u = 0 & \text{on } \partial\Omega, \end{cases} \quad (117)$$

where $f := -\Delta u \in L^2(\Omega)$. Let $f_n \in C_0^\infty(\Omega)$ be a sequence with $f_n \rightarrow f \in L^2(\Omega)$. Then consider the Poisson problem

$$\begin{cases} -\Delta u_n = f_n & \text{in } \Omega, \\ u_n = 0 & \text{on } \partial\Omega. \end{cases} \quad (118)$$

Since (by Assumption 1) $\partial\Omega \in C^{3,\alpha}$ for a Hölder exponent $\alpha \in (0, 1)$ and $f_n \in C_0^\infty(\Omega) \subset C^{1,\alpha}(\bar{\Omega})$, by the Schauder theory for elliptic boundary-value problems (see, for example, Theorems 6.14 and 6.19 in Gilbarg and Trudinger (1998)), there exists a unique solution $u_n \in C^{3,\alpha}(\bar{\Omega}) \cap C_0(\bar{\Omega})$ to this boundary-value problem. As

$$\begin{cases} -\Delta(u - u_n) = (f - f_n) & \text{in } \Omega, \\ u - u_n = 0 & \text{on } \partial\Omega, \end{cases} \quad (119)$$

and $\partial\Omega \in C^{3,\alpha}$ (and therefore also C^2), and $u - u_n \in \mathcal{H}^2(\Omega)$, it follows from Theorem 3.1.2.1 and Remark 3.1.2.2 in Grisvard (2011), that for some $C > 0$

$$\|u - u_n\|_{\mathcal{H}^2} \leq C \|\Delta(u - u_n)\|_{L^2} = C \|f - f_n\|_{L^2} \rightarrow 0 \quad (120)$$

as $n \rightarrow \infty$. Thus, we have constructed a sequence of functions $u_n \in C^{3,\alpha}(\bar{\Omega}) \cap \mathcal{H}_0^1$ which converges to $u \in \mathcal{H}^2(\Omega)$ in the \mathcal{H}^2 norm. Therefore, $C^{3,\alpha}(\bar{\Omega}) \cap C_0(\bar{\Omega})$ is dense in $\mathcal{H}^2(\Omega) \cap \mathcal{H}_0^1(\Omega)$. As

$$\left(C^{3,\alpha}(\bar{\Omega}) \cap C_0(\bar{\Omega}) \right) \subset \left(C^3(\bar{\Omega}) \cap C_0(\bar{\Omega}) \right), \quad (121)$$

we conclude that $C^3(\bar{\Omega}) \cap C_0(\bar{\Omega})$ is also dense in $\mathcal{H}^2(\Omega) \cap \mathcal{H}_0^1(\Omega)$.

We now prove statement *ii*. Take an arbitrary $u \in C^3(\bar{\Omega}) \cap C_0(\bar{\Omega})$, and define $\tilde{u} = u/\eta$. It is clear that $\tilde{u}|_\Omega$ is in $C^3(\Omega)$, from classical rules of calculus, as $\eta|_\Omega > 0$. In particular, \tilde{u} and its derivatives are bounded on any closed subset of Ω , so we only need to consider the behavior of \tilde{u} in a neighbourhood of $\partial\Omega$.

We know that η has nonvanishing derivative on $\partial\Omega$, and equals zero on $\partial\Omega$, which implies that $|\langle \nabla\eta(x), \mathbf{n}_x \rangle| > 0$ for each $x \in \partial\Omega$, where \mathbf{n}_x is the outward pointing normal at x . As $\partial\Omega$ is compact, this implies that

$$\inf_{x \in \partial\Omega} |\langle \nabla\eta(x), \mathbf{n}_x \rangle| > \epsilon, \quad \text{for some } \epsilon > 0. \quad (122)$$

3. Thanks to Endre Süli for suggesting the proof of statement *i* given here.

As $u \in C^3(\overline{\Omega}) \cap C_0(\overline{\Omega})$, for any $x \in \partial\Omega$, we also have $\limsup_{x' \rightarrow x} |u(x')| = 0$.

Consider a smooth path $\rho : [0, 1) \rightarrow \Omega$ with $\lim_{t \rightarrow 1} \rho(t) = x \in \partial\Omega$. For a function g with domain Ω (e.g. $g = \eta$, u or \tilde{u}), we write g_ρ for the composition $g \circ \rho$ with domain $[0, 1)$. Suppose ρ does not approach $\partial\Omega$ tangentially, in particular,

$$\lim_{t \rightarrow 1} \langle d\rho/dt, \mathbf{n}_x \rangle \geq \epsilon \quad \text{and} \quad \|d\rho/dt\| \equiv 1. \quad (123)$$

As $\langle \nabla\eta(x), y \rangle = 0$ for all $y \perp \mathbf{n}_x$, we have

$$\lim_{t \rightarrow 1} |\eta'_\rho(t)| = |\langle \nabla\eta(x), \mathbf{n}_x \rangle| \langle d\rho/dt, \mathbf{n}_x \rangle > \epsilon^2. \quad (124)$$

Using L'Hôpital's rule, we can determine the behaviour of \tilde{u}_ρ (and its derivatives) as $t \rightarrow 1$. In particular, by repeated application of the rule, as $t \rightarrow 1$

$$\tilde{u}_\rho = \frac{u_\rho}{\eta_\rho} \rightarrow \frac{u'_\rho}{\eta'_\rho} \quad (125)$$

$$\tilde{u}'_\rho = \frac{\eta_\rho u'_\rho - \eta'_\rho u_\rho}{\eta_\rho^2} \rightarrow \frac{u''_\rho}{2\eta'_\rho} - \frac{\eta''_\rho}{2\eta'_\rho} \tilde{u}_\rho \quad (126)$$

$$\tilde{u}''_\rho = \frac{\eta_\rho^2 u''_\rho - \eta_\rho \eta''_\rho u_\rho - 2\eta_\rho \eta'_\rho u'_\rho + 2(\eta'_\rho)^2 u_\rho}{\eta_\rho^3} \rightarrow \frac{u'''_\rho}{3\eta'_\rho} - \frac{\eta'''_\rho}{3\eta'_\rho} \tilde{u}_\rho + \frac{\eta''_\rho}{\eta'_\rho} \tilde{u}'_\rho \quad (127)$$

Considering $t \approx 1$, as $|\eta'_\rho| > \epsilon^2$, the right hand side of each of these terms is bounded, with a uniform bound for all paths under consideration. Writing $D_h(\tilde{u})$ for the directional derivative of \tilde{u} in a direction h , we know $\|\tilde{u}'_\rho\| = \|D_{d\rho/dt}(\tilde{u})\|$ and $\|\tilde{u}''_\rho\| = \|D_{d\rho/dt}^2(\tilde{u})\|$. We can therefore find a single value $K > 0$, and take a ball B_x around each $x \in \partial\Omega$, such that, within each B_x , for all unit vectors h satisfying $\langle h, \mathbf{n}_x \rangle > \epsilon$,

$$|\tilde{u}| + |D_h(\tilde{u})| + |D_h^2(\tilde{u})| \leq K. \quad (128)$$

As $\Omega \setminus (\cup_{x \in \partial\Omega} B_x)$ is compact, we know that, for K sufficiently large, (128) also holds on $\Omega \setminus (\cup_{x \in \partial\Omega} B_x)$, for all unit vectors h . We conclude that (128) holds on all of Ω , that is, \tilde{u} , and its directional first and second derivatives (except possibly in tangent directions at the boundary), are uniformly bounded.

However, as $\partial\Omega$ is a $C^{3,\alpha}$ boundary (in particular, for all ϵ sufficiently small and any $x \in \partial\Omega$ the set $\{y \in \mathbb{R}^n : \langle y, \mathbf{n}_x \rangle \geq \epsilon\}$ is a convex cone containing a linear basis for \mathbb{R}^n), this implies that the first and second derivatives of \tilde{u} are also uniformly bounded on Ω . We conclude that $\tilde{u} \in C_b^2(\Omega)$. \blacksquare

A.2 Proof of Lemma 24

Proof (1) From the training algorithm, calculating the entries in vector $\nabla_\theta Q_t^N(x)$, we have

$$\frac{\partial Q_t^N}{\partial c^i}(x) = \frac{\eta(x) \sigma(w_t^i \cdot x + b_t^i)}{N^\beta}. \quad (129)$$

Therefore, by (15), applying the chain rule we derive

$$\frac{dc_t^i}{dt} = \alpha N^{\beta-1} \int_{\Omega} F^N(\mathcal{L}Q_t^N(x)) \eta(x) \sigma(w_t^i \cdot x + b_t^i) d\mu(x). \quad (130)$$

By the boundedness of the functions F^N , η and σ , the RHS of (130) is bounded:

$$\left| \int_{\Omega} F^N(\mathcal{L}Q_t^N(x)) \eta(x) \sigma(w_t^i \cdot x + b_t^i) d\mu(x) \right| < KN^{\delta}, \quad (131)$$

where K is a constant. Therefore, we conclude that

$$|c_t^i| < |c_0^i| + |c_t^i - c_0^i| < K_0 + \alpha N^{\delta+\beta-1} tK < C. \quad (132)$$

(2) Similar to above, considering the w^i entries in the vector $\nabla_{\theta} Q_t^N(x)$, we have

$$\frac{\partial Q_t^N}{\partial (w_t^i)_k}(x) = \frac{\eta(x) \sigma'(w_t^i \cdot x + b_t^i) c_t^i x_k}{N^{\beta}}. \quad (133)$$

Therefore w_t^i satisfies:

$$\frac{d(w_t^i)_k}{dt} = \alpha N^{\beta-1} \int_{\Omega} F^N(\mathcal{L}Q_t^N(x)) \eta(x) c_t^i \sigma'(w_t^i \cdot x + b_t^i) x_k d\mu(x). \quad (134)$$

By the boundedness of F^N , η , σ' and c_t^i , the RHS of (134) is bounded by

$$\left| \int_{\Omega} F^N(\mathcal{L}Q_t^N(x)) \eta(x) c_t^i \sigma'(w_t^i \cdot x + b_t^i) x_k d\mu(x) \right| < KN^{\delta}, \quad (135)$$

where K is a constant. Therefore,

$$|(w_t^i)_k| < |(w_0^i)_k| + |(w_t^i)_k - (w_0^i)_k| < |(w_0^i)_k| + \alpha N^{\delta+\beta-1} tK < |(w_0^i)_k| + C. \quad (136)$$

Taking an expectation on both sides, we obtain

$$\mathbb{E}|(w_t^i)_k| < C. \quad (137)$$

(3) Considering the b^i entries in the vector $\nabla_{\theta} Q_t^N(x)$, we have

$$\frac{\partial Q_t^N}{\partial b_t^i}(x) = \frac{\eta(x) \sigma'(w_t^i \cdot x + b_t^i) c_t^i}{N^{\beta}}. \quad (138)$$

Therefore, b_t^i satisfies

$$\frac{db_t^i}{dt} = \alpha N^{\beta-1} \int_{\Omega} F^N(\mathcal{L}Q_t^N(x)) \eta(x) c_t^i \sigma'(w_t^i \cdot x + b_t^i) d\mu(x). \quad (139)$$

By the boundedness of F^N , η , σ' and c_t^i , the RHS of (134) is bounded by

$$\left| \int_{\Omega} F^N(\mathcal{L}Q_t^N(x)) \eta(x) c_t^i \sigma'(w_t^i \cdot x + b_t^i) d\mu(x) \right| < KN^{\delta}, \quad (140)$$

where K is a constant. Therefore,

$$|b_t^i| < |b_0^i| + |b_t^i - b_0^i| < |b_0^i| + \alpha N^{\delta+\beta-1} t K < |b_0^i| + C. \quad (141)$$

Taking an expectation on both sides, we obtain

$$\mathbb{E}|b_t^i| < C. \quad (142)$$

■

A.3 Proof of Lemma 36

Proof (1) Notice that

$$\begin{aligned} |A_t^N(x, y) - A_0^N(x, y)| &= \left| \frac{1}{N} \sum_{i=1}^N \left[\sigma(w_t^i \cdot x + b_t^i) \sigma(w_t^i \cdot y + b_t^i) - \sigma(w_0^i \cdot x + b_0^i) \sigma(w_0^i \cdot y + b_0^i) \right. \right. \\ &\quad \left. \left. + x \cdot y (C_t^i)^2 \sigma'(w_t^i \cdot x + b_t^i) \sigma'(w_t^i \cdot y + b_t^i) - x \cdot y (C_0^i)^2 \sigma'(w_0^i \cdot x + b_0^i) \sigma'(w_0^i \cdot y + b_0^i) \right] \right| \\ &\leq \frac{1}{N} \sum_{i=1}^N \left[|\sigma(w_t^i \cdot x + b_t^i) \sigma(w_t^i \cdot y + b_t^i) - \sigma(w_0^i \cdot x + b_0^i) \sigma(w_0^i \cdot y + b_0^i)| \right. \\ &\quad \left. + |x \cdot y| |c_t^i|^2 - c_0^i|^2 \|\sigma'(w_t^i \cdot x + b_t^i) \sigma'(w_t^i \cdot y + b_t^i)\| \right. \\ &\quad \left. + c_0^i|^2 |x \cdot y| \|\sigma'(w_t^i \cdot x + b_t^i) \sigma'(w_t^i \cdot y + b_t^i) - \sigma'(w_0^i \cdot x + b_0^i) \sigma'(w_0^i \cdot y + b_0^i)\| \right]. \end{aligned} \quad (143)$$

For a smooth function $f : \mathbb{R} \rightarrow \mathbb{R}$, applying the mean value theorem with respect to w_t^i , there exists $w^* \in \mathbb{R}^n$ such that

$$\begin{aligned} &f(w_t^i \cdot x + b) f(w_t^i \cdot y + b) - f(w_0^i \cdot x + b) f(w_0^i \cdot y + b) \\ &= (w_t^i - w_0^i) \cdot \left[f'(w^* \cdot x + b) f(w^* \cdot y + b) x + f(w^* \cdot x + b) f'(w^* \cdot y + b) y \right]. \end{aligned} \quad (144)$$

By the fact that $\sigma \in C_b^4(\mathbb{R})$ and $|c_u^i| < C$ for any $u \in [0, T]$, we have

$$\begin{aligned} |A_t^N(x, y) - A_0^N(x, y)| &\leq \frac{C}{N} \sum_{i=1}^N \left[(\|x\|_1 + \|y\|_1) \|w_t^i - w_0^i\|_1 + |x \cdot y| |c_t^i - c_0^i| \right. \\ &\quad \left. + \|w_t^i - w_0^i\|_1 (\|x\|_1 + \|y\|_1) |x \cdot y| + (1 + |x \cdot y|) |b_t^i - b_0^i| \right]. \end{aligned} \quad (145)$$

Therefore, by (131), (140) and (135), the increments of c_t , b_t and w_t are bounded by $CN^{\delta+\beta-1}$. Therefore, we have

$$|A_t^N(x, y) - A_0^N(x, y)| \leq CN^{\delta+\beta-1} \left[(\|x\|_1 + \|y\|_1)(1 + |x \cdot y|) + |x \cdot y| \right] \leq CN^{\delta+\beta-1}. \quad (146)$$

(2) By definition,

$$\begin{aligned} \partial_{y_i} A_t^N(x, y) &= \frac{1}{N} \sum_{k=1}^N \left[\sigma(w_t^k \cdot x + b_t^k) \sigma'(w_t^k \cdot x + b_t^k) (w_t^k)_i + c_t^{k2} \sigma'(w_t^k \cdot x + b_t^k) \sigma'(w_t^k \cdot y + b_t^k) x_i \right. \\ &\quad \left. + c_t^{k2} \sigma'(w_t^k \cdot x + b_t^k) \sigma''(w_t^k \cdot y + b_t^k) x \cdot y (w_t^k)_i \right]. \end{aligned} \quad (147)$$

Similarly to in (143), we split the difference between $\partial_{y_i} A_t^N(x, y)$ and $\partial_{y_i} A_0^N(x, y)$ into terms

$$\begin{aligned} |\partial_{y_i} A_t^N(x, y) - \partial_{y_i} A_0^N(x, y)| &\leq \frac{1}{N} \sum_{k=1}^N \left[\left| \sigma(w_t^k \cdot x + b_t^k) \sigma'(w_t^k \cdot x + b_t^k) (w_t^k)_i \right. \right. \\ &\quad \left. \left. - \sigma(w_0^k \cdot x + b_0^k) \sigma'(w_0^k \cdot x + b_0^k) (w_0^k)_i \right| \right. \\ &\quad \left. + \left| c_t^{k2} \sigma'(w_t^k \cdot x + b_t^k) \sigma'(w_t^k \cdot y + b_t^k) x_i - c_0^{k2} \sigma'(w_0^k \cdot x + b_0^k) \sigma'(w_0^k \cdot y + b_0^k) x_i \right| \right. \\ &\quad \left. + \left| c_t^{k2} \sigma'(w_t^k \cdot x + b_t^k) \sigma''(w_t^k \cdot y + b_t^k) x \cdot y (w_t^k)_i - c_0^{k2} \sigma'(w_0^k \cdot x + b_0^k) \sigma''(w_0^k \cdot y + b_0^k) x \cdot y (w_0^k)_i \right| \right]. \end{aligned} \quad (148)$$

Applying the mean value theorem with respect to w_t^i to each term on the RHS of (148), using the fact that $\sigma \in C_b^4(\mathbb{R})$ and that $|c_t^i| < C$ for any $t \in [0, T]$, we have

$$\begin{aligned} |\partial_{y_i} A_t^N(x, y) - \partial_{y_i} A_0^N(x, y)| &\leq \frac{C}{N} \sum_{k=1}^N \left[|(w_t^k)_i - (w_0^k)_i| + |(w_0^k)_i| (\|x\|_1 + \|y\|_1) \|w_t^k - w_0^k\|_1 \right. \\ &\quad \left. + |c_t^i - c_0^i| |x_i| + |x_i| (\|x\|_1 + \|y\|_1) \|w_t^k - w_0^k\|_1 + |(w_t^k)_i - (w_0^k)_i| |x \cdot y| \right. \\ &\quad \left. + |c_t^i - c_0^i| |(w_0^k)_i| |x \cdot y| + |(w_0^k)_i| |x \cdot y| (\|x\|_1 + \|y\|_1) \|w_t^k - w_0^k\|_1 + (1 + |x \cdot y|) |b_t^k - b_0^k| \right]. \end{aligned}$$

Therefore, by the boundedness of $|c_t^k - c_0^k|$, $|b_t^k - b_0^k|$, $|(w_t^k)_i - (w_0^k)_i|$ and $\|w_t^k - w_0^k\|_1$ we derive

$$\begin{aligned} |\partial_{y_i} A_t^N(x, y) - \partial_{y_i} A_0^N(x, y)| &\leq CN^{\beta-1} \left[1 + \frac{\sum_{k=1}^N |(w_0^k)_i|}{N} (|x \cdot y| + 1) (1 + \|x\|_1 + \|y\|_1) \right. \\ &\quad \left. + |x_i| (1 + \|x\|_1 + \|y\|_1) + |x \cdot y| \right] \leq CN^{\delta+\beta-1}. \end{aligned}$$

By definition,

$$\begin{aligned} \partial_{y_i y_j}^2 A_t^N(x, y) &= \frac{1}{N} \sum_{k=1}^N \left[\sigma(w_t^k \cdot x + b_t^k) \sigma''(w_t^k \cdot x + b_t^k) (w_t^k)_i (w_t^k)_j \right. \\ &\quad \left. + c_t^{k2} \sigma'(w_t^k \cdot x + b_t^k) \sigma''(w_t^k \cdot y + b_t^k) x_i (w_t^k)_j \right. \\ &\quad \left. + c_t^{k2} \sigma'(w_t^k \cdot x + b_t^k) \sigma''(w_t^k \cdot y + b_t^k) x_j (w_t^k)_i \right. \\ &\quad \left. + c_t^{k2} \sigma'(w_t^k \cdot x + b_t^k) \sigma'''(w_t^k \cdot y + b_t^k) x \cdot y (w_t^k)_i (w_t^k)_j \right]. \end{aligned}$$

The difference between $\partial_{y_i y_j}^2 A_t^N$ and $\partial_{y_i y_j}^2 A_0^N$ can be split into terms

$$\begin{aligned}
 |\partial_{y_i y_j}^2 A_t^N(x, y) - \partial_{y_i y_j}^2 A_0^N(x, y)| &\leq \frac{1}{N} \sum_{k=1}^N \left[\left| \sigma(w_t^k \cdot x + b_t^k) \sigma''(w_t^k \cdot x + b_t^k) (w_t^k)_i (w_t^k)_k \right. \right. \\
 &\quad - \left. \left. \sigma(w_0^k \cdot x + b_0^k) \sigma''(w_0^k \cdot x + b_0^k) (w_0^k)_i (w_0^k)_j \right| + \left| c_t^{k2} \sigma'(w_t^k \cdot x + b_t^k) \sigma''(w_t^k \cdot y + b_t^k) x_i (w_t^k)_j \right. \right. \\
 &\quad - \left. \left. c_0^{k2} \sigma'(w_0^k \cdot x + b_0^k) \sigma''(w_0^k \cdot y + b_0^k) x_i (w_0^k)_j \right| + \left| c_t^{k2} \sigma'(w_t^k \cdot x + b_t^k) \sigma''(w_t^k \cdot y + b_t^k) x_j (w_t^k)_i \right. \right. \\
 &\quad - \left. \left. c_0^{k2} \sigma'(w_0^k \cdot x + b_0^k) \sigma''(w_0^k \cdot y + b_0^k) x_j (w_0^k)_i \right| + \left| c_t^{k2} \sigma'(w_t^k \cdot x + b_t^k) \sigma'''(w_t^k \cdot y + b_t^k) x \cdot y (w_t^k)_i (w_t^k)_j \right. \right. \\
 &\quad \left. \left. - c_0^{k2} \sigma'(w_0^k \cdot x + b_0^k) \sigma'''(w_0^k \cdot y + b_0^k) x \cdot y (w_0^k)_i (w_0^k)_j \right| \right].
 \end{aligned} \tag{149}$$

We apply the mean value theorem to each term above, leading to

$$\begin{aligned}
 |\partial_{y_i y_j}^2 A_t^N(x, y) - \partial_{y_i y_j}^2 A_0^N(x, y)| &\leq CN^{\beta-1} \left[\frac{\sum_{k=1}^N |(w_0^k)_i| |(w_0^k)_j|}{N} (|x \cdot y| + 1) (\|x\|_1 + \|y\|_1 + 1) \right. \\
 &\quad + \frac{\sum_{k=1}^N |(w_0^k)_i|}{N} [(1 + |x_j|)(1 + \|x\|_1 + \|y\|_1) + |x \cdot y|] \\
 &\quad \left. + \frac{\sum_{k=1}^N |(w_0^k)_j|}{N} [(1 + |x_i|)(1 + \|x\|_1 + \|y\|_1) + |x \cdot y| + 1 + |x_i| + |x_j| + \cdot y] \right] \leq CN^{\delta+\beta-1}.
 \end{aligned}$$

■

A.4 Proof of Lemma 37

Proof Taking squares and then expectations on both sides of inequalities (46), (47), and (48) gives

$$\begin{aligned}
 \mathbb{E}_{c_0, w_0, b_0} [|A_t^N(x, y) - A_0^N(x, y)|^2] &\leq CN^{2(\delta+\beta-1)}, \\
 \mathbb{E}_{c_0, w_0, b_0} [|\partial_{y_k} A_t^N(x, y) - \partial_{y_k} A_0^N(x, y)|^2] &\leq CN^{2(\delta+\beta-1)}, \\
 \mathbb{E}_{c_0, w_0, b_0} [|\partial_{y_k y_l}^2 A_t^N(x, y) - \partial_{y_k y_l}^2 A_0^N(x, y)|^2] &\leq CN^{2(\delta+\beta-1)}.
 \end{aligned} \tag{150}$$

Summing these inequalities for all $k, l \in \{1, 2, \dots, n\}$ and applying the Cauchy–Schwarz inequality gives the result. ■

A.5 Proof of Lemma 39

Proof First note that

$$\begin{aligned}
 \mathbb{E}_{c_0, w_0, b_0} [A_0^N(x, y)] &= \mathbb{E}_{c_0, w_0, b_0} \left[\frac{1}{N} \sum_{k=1}^N \sigma(w_0^k \cdot x + b_0^k) \sigma(w_0^k \cdot y + b_0^k) \right. \\
 &\quad \left. + c_0^{k2} x \cdot y \sigma'(w_0^k \cdot x + b_0^k) \sigma'(w_0^k \cdot y + b_0^k) \right] = A(x, y).
 \end{aligned} \tag{151}$$

We can also calculate

$$\begin{aligned} \mathbb{E}_{c_0, w_0, b_0} [|A_0^N(x, y) - A(x, y)|^2] &= \text{Var}[A_0^N(x, y)] = \frac{1}{N} \text{Var}[U(x; c, w, b) \cdot U(y; c, w, b)] \\ &\leq \frac{1}{N} \mathbb{E}_{c, w, b} |U(x; c, w, b) \cdot U(y; c, w, b)|^2 \leq \frac{C}{N}. \end{aligned} \quad (152)$$

For the partial derivatives, for any fixed $x, y \in \Omega$,

$$\begin{aligned} &|\partial_{y_i}(\sigma(w \cdot x + b)\sigma(w \cdot y + b) + c^2\sigma'(w \cdot x + b)\sigma'(w \cdot y + b)x \cdot y)| \\ &= |\sigma(w \cdot x + b)\sigma'(w \cdot y + b)w_i \\ &\quad + c^2\sigma'(w \cdot x + b)\sigma'(w \cdot y + b)x_i + c^2x \cdot y\sigma'(w \cdot x + b)\sigma''(w \cdot y + b)w_i| \end{aligned} \quad (153)$$

and

$$\begin{aligned} &|\partial_{y_i y_j}^2(\sigma(w \cdot x + b)\sigma(w \cdot y + b) + c^2\sigma'(w \cdot x + b)\sigma'(w \cdot y + b)x \cdot y)| \\ &= |\sigma(w \cdot x + b)\sigma''(w \cdot y + b)w_i w_j + c^2\sigma'(w \cdot x + b)\sigma''(w \cdot y + b)(x_i w_j + x_j w_i) \\ &\quad + c^2x \cdot y\sigma'(w \cdot x + b)\sigma'''(w \cdot y + b)w_i w_j|. \end{aligned} \quad (154)$$

These are locally bounded around $(x, y) \in \Omega^2$, so by Leibniz's integral rule, we swap the expectation and the differential operator

$$\begin{aligned} \partial_{y_i} A(x, y) &= \partial_{y_i} \mathbb{E}_{c, w, b} [\sigma(w \cdot x + b)\sigma(w \cdot y + b) + c^2\sigma'(w \cdot x + b)\sigma'(w \cdot y + b)x \cdot y] \\ &= \mathbb{E}_{c, w, b} [\partial_{y_i}(\sigma(w \cdot x + b)\sigma(w \cdot y + b) + c^2\sigma'(w \cdot x + b)\sigma'(w \cdot y + b)x \cdot y)] \\ &= \mathbb{E}_{c_0, w_0, b_0} [\partial_{y_i} A_0^N(x, y)], \end{aligned} \quad (155)$$

and

$$\begin{aligned} \partial_{y_i y_j}^2 A(x, y) &= \partial_{y_i y_j}^2 \mathbb{E}_{c, w, b} [\sigma(w \cdot x + b)\sigma(w \cdot y + b) + c^2\sigma'(w \cdot x + b)\sigma'(w \cdot y + b)x \cdot y] \\ &= \mathbb{E}_{c, w, b} [\partial_{y_i y_j}^2(\sigma(w \cdot x + b)\sigma(w \cdot y + b) + c^2\sigma'(w \cdot x + b)\sigma'(w \cdot y + b)x \cdot y)] \\ &= \mathbb{E}_{c_0, w_0, b_0} [\partial_{y_i y_j}^2 A_0^N(x, y)]. \end{aligned} \quad (156)$$

Therefore, we can bound the first-derivative variance $\text{Var}[\partial_{y_i} A_0^N(x, y)]$:

$$\begin{aligned} \mathbb{E}_{c_0, w_0, b_0} |\partial_{y_i} A_0^N(x, y) - \partial_{y_i} A(x, y)|^2 &= \frac{1}{N} \text{Var}[\partial_{y_i}(U(x; c, w, b) \cdot U(y; c, w, b))] \\ &\leq \frac{1}{N} \mathbb{E}_{c, w, b} |\partial_{y_i}(U(x; c, w, b) \cdot U(y; c, w, b))|^2 \leq \frac{1}{N} \mathbb{E}_{c, w, b} [(C|w_i| + C|x_i| + C|x \cdot y||w_i|)^2] \leq \frac{C}{N}. \end{aligned}$$

and similarly the second derivative variance $\text{Var}[\partial_{y_i y_j}^2 A_0^N(x, y)]$:

$$\begin{aligned} \mathbb{E}_{c_0, w_0, b_0} |\partial_{y_i y_j}^2 A_0^N(x, y) - \partial_{y_i y_j}^2 A(x, y)|^2 &= \frac{1}{N} \text{Var}[\partial_{y_i y_j}^2(U(x; c, w, b) \cdot U(y; c, w, b))] \\ &\leq \frac{1}{N} \mathbb{E}_{c, w, b} |\partial_{y_i y_j}^2(U(x; c, w, b) \cdot U(y; c, w, b))|^2 \\ &\leq \frac{1}{N} \mathbb{E}_{c, w, b} [(C|w_i w_j| + C|x_i w_j| + C|x_j w_i| + C|x \cdot y||w_i w_j|)^2] \leq \frac{C}{N}. \end{aligned}$$

■

A.6 Existence of ODE solutions

Lemma 50 *Under the assumptions of Theorem 47, for any $q \in \mathcal{H}^2$, and any initial value $v_0 \in L^2_\eta$, the ODE*

$$\frac{dv_t}{dt} = \mathcal{L}(\mathcal{B}v_t + q) \quad (157)$$

admits a unique solution v taking values in L^2_η .

Proof Consider the process $w_t = \eta v_t$. This has initial value $w_0 = \eta v_0 \in L^2$, and satisfies the dynamics

$$\frac{dw_t}{dt} = \eta \mathcal{L}(\mathcal{B}(w_t/\eta) + q) =: G(w_t). \quad (158)$$

From Lemma 29, the map $w \mapsto \mathcal{B}(w/\eta)$ is uniformly Lipschitz as a map $L^2 \rightarrow \mathcal{H}^2$. By Assumption 6, \mathcal{L} is uniformly Lipschitz as a map $\mathcal{H}^2 \rightarrow L^2$. By Assumption 2, η is bounded on Ω . Therefore, $G : L^2 \rightarrow L^2$ is uniformly Lipschitz.

From the Picard–Lindelöf theorem (see, for example, (Kolokoltsov, 2019, Theorem 2.2.1)), we know that (158) admits a unique solution with $w_t \in L^2$ for all $t \geq 0$. Using this, together with the fact $\eta > 0$ on Ω , we see that there is a unique $v_t = w_t/\eta$ satisfying the original equation, with values in L^2_η . \blacksquare

A.7 Monotonicity for linear elliptic PDE

In this subsection, we give a brief proof of monotonicity of a class of elliptic linear differential operators, with particular choices of sampling measure μ . Our results in this appendix are not exhaustive, but demonstrate some of the flexibility of the equations we have considered.

Lemma 51 *Consider an autonomous Markov process X satisfying the SDE*

$$dX_t = \nu(X_t)dt + \sigma(X_t)dW_t \quad (159)$$

where $\nu : \bar{\Omega} \rightarrow \mathbb{R}^n$ and $\sigma : \bar{\Omega} \rightarrow \mathbb{R}^{n \times n}$ and W is an \mathbb{R}^n -valued standard Brownian motion. Suppose further that ν and σ both vanish outside Ω . For an absolutely continuous measure μ on Ω , with density f_μ , let ν , σ and f_μ be sufficiently smooth that

$$\gamma^* = \sup_{x \in \Omega} \left\{ \frac{-\sum_i \frac{\partial}{\partial x_i} [\nu_i(x) f_\mu(x)] + \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} [(\sigma \sigma^\top)_{ij}(x) f_\mu(x)]}{f_\mu(x)} \right\} < \infty. \quad (160)$$

Then the generator of X satisfies, for all $u \in \mathcal{H}^2$,

$$\left\langle u, \sum_i \nu_i \frac{\partial u}{\partial x_i} + \frac{1}{2} \sum_{i,j} (\sigma \sigma^\top)_{ij} \frac{\partial^2 u}{\partial x_i \partial x_j} \right\rangle \leq \frac{\gamma^*}{2} \|u\|, \quad (161)$$

where the inner product and norm are both taken in $L^2(\Omega, \mu)$.

Proof Consider a copy of X initialized with $X_0 \sim \mu$. The Fokker–Plank equation states that the density of X_0 satisfies $f(x, 0) = f_\mu$ and

$$\frac{\partial}{\partial t} f(x, t) = - \sum_i \frac{\partial}{\partial x_i} [\nu_i(x) f(x, t)] + \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} [(\sigma \sigma^\top)_{ij}(x) f(x, t)] \quad (162)$$

and hence, using the definition of γ^* ,

$$\frac{\partial}{\partial t} (e^{-\gamma^* t} f(x, t)) \leq 0. \quad (163)$$

As $f(x, t)$ is positive, we can define an operator

$$T_t u(x) := \mathbb{E} \left[u(X_t^x) e^{-\frac{\gamma^*}{2} t} \mid X_0 = x \right] \quad (164)$$

which satisfies, by Jensen’s inequality,

$$\|T_t u\|_{L^2(\mu)} \leq \int_\Omega u^2(t) e^{-\gamma^* t} f(t, x) dx \leq \int_\Omega u^2(t) f_\mu(x) dx = \|u\|_{L^2(\mu)}^2 \quad (165)$$

so the operator T_t is a contraction. It is easy to verify that T_t is also a strongly continuous semigroup, and its generator is given by $u \mapsto -\frac{\gamma^*}{2} u + \mathcal{A}u$, where \mathcal{A} is the generator of X . By the Lumer–Phillips theorem (Lumer and Phillips (1961)), we know that the generator of T_t is dissipative on $L^2(\Omega, \mu)$, that is

$$\left\langle u, -\frac{\gamma^*}{2} u + \mathcal{A}u \right\rangle = \left\langle u, -\frac{\gamma^*}{2} u + \sum_i \nu_i \frac{\partial u}{\partial x_i} + \frac{1}{2} \sum_{ij} (\sigma \sigma^\top)_{ij} \frac{\partial^2 u}{\partial x_i \partial x_j} \right\rangle \leq 0. \quad (166)$$

Rearrangement yields the result. ■

Remark 52 *This result immediately shows that if μ is a stationary distribution for the Markov process, then $\gamma^* = 0$ and the generator of X is automatically L^2 -dissipative. (This result is known, see, for example, (Kallenberg, 2002, Chapter 20).) This suggests that this will often be a wise choice for μ in linear problems, if it is known explicitly.*

By manipulating the choice of measure μ , or equivalently setting μ equal to Lebesgue measure and considering the equivalent PDE

$$\tilde{\mathcal{L}}u := f(x) \mathcal{L}u = 0, \quad (167)$$

this result can be leveraged usefully, provided the drift of X does not grow quickly. Similar results for other bounds on b can also be obtained.

Lemma 53 *Consider a process X as in Lemma 51. Suppose that σ is constant and there exist bounded functions $g_i : \mathbb{R} \rightarrow \mathbb{R}$ for $i \leq n$ such that, for all $x \in \Omega$,*

$$\gamma \geq - \sum_i \left(\frac{\partial}{\partial x_i} \nu_i(x) + g(x_i) \nu_i(x) \right) + \sum_{ij} g_i(x_i) g_j(x_j) (\sigma \sigma^\top)_{ij}. \quad (168)$$

Then the generator of X satisfies $\langle u, \mathcal{A}u \rangle \leq \frac{\gamma}{2}\|u\|$, where μ is the measure on Ω with Lebesgue density

$$f_\mu = C \exp\left(\sum_i \int_0^{x_i} g_i(\zeta) d\zeta\right),$$

for C a normalizing constant. In particular, this implies that the operator

$$\mathcal{L}u = r - \gamma u + \mathcal{A}u \tag{169}$$

is strongly monotone, in the sense of Assumption 9.

Proof We calculate

$$\begin{aligned} & - \sum_i \frac{\partial}{\partial x_i} [\nu_i(x) f_\mu(x)] + \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} [(\sigma \sigma^\top)_{ij} f_\mu(x)] \\ &= - \sum_i \left(\frac{\partial}{\partial x_i} \nu_i(x) f_\mu(x) + g_i(x_i) \nu_i(x) f_\mu(x) \right) + \sum_{i,j} g(x_i) g(x_j) (\sigma \sigma^\top)_{ij} f_\mu(x) \\ &\leq \gamma f_\mu(x). \end{aligned} \tag{170}$$

Therefore, in (160) we have $\gamma^* \leq \gamma$, and the result follows from (161). ■

Remark 54 In one dimension, (168) is clearly satisfied whenever we know that $4\sigma^2(\gamma + \nu') \geq -\nu^2$, by taking $g(x) = -\nu(x)/(2\sigma^2)$.

A.8 Configuration of the numerical test cases

A.8.1 TABLE OF HYPER PARAMETERS

Dimension	Method	Layer	Units	Activation	Optimizer	Numer of MC samples
1	Q-PDE	1	64	Sigmoid	ADAM	$l_{MC}=1k, u_{MC}=2k$
20	Q-PDE	1	256	Sigmoid	ADAM	$l_{MC}=2k, u_{MC}=10k$
20	DGM	1	256	Sigmoid	ADAM	$l_{MC}=2k, u_{MC}=10k$

A.8.2 INITIALIZATION OF NEURAL NETWORKS

Parameters of the single-layer net S_0 are randomly sampled: c_0^i are i.i.d sampled from uniform distribution $\mathcal{U}[-1, 1]$; w_0^i and b_0^i are i.i.d sampled from Gaussian distribution $\mathcal{N}(0, I_d)$ and $\mathcal{N}(0, 1)$ where I_d is the identity matrix of dimension d .

A.8.3 LEARNING PROCESS

We use $Q_t := S_t \cdot \eta$ as the approximator, where $\eta(x) := 1 - \|x\|^2$. We apply the built-in ADAM optimizer in our test with initial learning rate $l_0 = 0.5$, and the learning rate decays as $l_t = l_0/(1+t/200)$. In each step, we sample MC samples for gradient estimate. As a larger number of MC sample points reduces the random error, we linearly increase the number M_t of MC points to be sampled at each step as $M_t = \text{round}(l_{MC} + (u_{MC} - l_{MC})t/T)$, where T is the terminal number of training steps.

References

- Robert A. Adams and John J.F. Fournier. *Sobolev spaces*. Elsevier, 2003.
- Christian Beck, Martin Hutzenthaler, Arnulf Jentzen, and Benno Kuckuck. An overview on deep learning-based approximation methods for partial differential equations. *arXiv preprint arXiv:2012.12348*, 2020.
- Felix E. Browder. Existence and perturbation theorems for nonlinear maximal monotone operators in banach spaces. *Bulletin of the American Mathematical Society*, 73(3):322–327, 1967.
- Shengze Cai, Zhiping Mao, Zhicheng Wang, Minglang Yin, and George Em Karniadakis. Physics-informed neural networks (pinns) for fluid mechanics: A review. *Acta Mechanica Sinica*, 37(12):1727–1738, 2021.
- Lénaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché Buc, Edward A. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 2933–2943, 2019. URL <http://papers.nips.cc/paper/8559-on-lazy-training-in-differentiable-programming>.
- George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- Weinan E, Jiequn Han, and Arnulf Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in Mathematics and Statistics*, 5(4):349–380, 2017.
- Weinan E, Jiequn Han, and Arnulf Jentzen. Algorithms for solving high dimensional PDEs: from nonlinear Monte Carlo to machine learning. *Nonlinearity*, 35(1):278–310, 2022.
- Lawrence C. Evans. *Partial differential equations*. American Mathematical Soc., 2010.
- Maximilien Germain, Mathieu Laurière, Huyên Pham, and Xavier Warin. Deepsets and their derivative networks for solving symmetric PDEs. *arXiv preprint arXiv:2103.00838*, 2021a.
- Maximilien Germain, Huyên Pham, and Xavier Warin. Neural networks-based algorithms for stochastic control and PDEs in finance. *arXiv preprint arXiv:2101.08068*, 2021b.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. When do neural networks outperform kernel methods? In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 14820–14830. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/a9df2255ad642b923d95503b9a7958d8-Paper.pdf>.

- David Gilbarg and Neil S. Trudinger. *Elliptic partial differential equations of second order*. Springer, 1998.
- Pierre Grisvard. *Elliptic problems in nonsmooth domains*. SIAM, 2011.
- Philipp Grohs and Lukas Herrmann. Deep neural network approximation for high-dimensional elliptic PDEs with boundary conditions. *arXiv preprint arXiv:2007.05384*, 2020.
- Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.
- Ruimeng Hu, Quyuan Lin, Alan Raydan, and Sui Tang. Higher-order error estimates for physics-informed neural networks approximating the primitive equations. *arXiv preprint arXiv:2209.11929*, 2022.
- Kazufumi Ito, Christoph Reisinger, and Yufei Zhang. A neural network-based policy iteration algorithm with global H^2 -superlinear convergence for stochastic games on domains. *Foundations of Computational Mathematics*, 21(2):331–374, 2021.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *arXiv preprint arXiv:1806.07572*, 2018.
- Olav Kallenberg. *Foundations of Modern Probability*. Springer, 2nd edition, 2002.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Vassili Kolokoltsov. *Differential equations on measures and functional spaces*. Springer, 2019.
- Isaac E. Lagaris, Aristidis Likas, and Dimitrios I. Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE transactions on neural networks*, 9(5):987–1000, 1998.
- Isaac E. Lagaris, Aristidis C. Likas, and Dimitris G. Papageorgiou. Neural-network methods for boundary value problems with irregular boundaries. *IEEE Transactions on Neural Networks*, 11(5):1041–1049, 2000.
- Samuel Lanthaler, Siddhartha Mishra, and George Em Karniadakis. Error estimates for deeponets: A deep learning framework in infinite dimensions. *arXiv preprint arXiv:2102.09618*, 2021.
- Hyuk Lee and In Seok Kang. Neural algorithm for solving differential equations. *Journal of Computational Physics*, 91(1):110–131, 1990.
- Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S. Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as Gaussian processes. *arXiv preprint arXiv:1711.00165*, 2017.

- Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32: 8572–8583, 2019.
- Jian Li, Jing Yue, Wen Zhang, and Wansuo Duan. The deep learning galerkin method for the general stokes equations. *Journal of Scientific Computing*, 93(1):5, 2022.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.
- Zichao Long, Yiping Lu, Xianzhong Ma, and Bin Dong. PDE-net: Learning PDEs from data. In *International Conference on Machine Learning*, pages 3208–3216. PMLR, 2018.
- Zichao Long, Yiping Lu, and Bin Dong. PDE-net 2.0: Learning PDEs from data with a numeric-symbolic hybrid deep network. *Journal of Computational Physics*, 399:108925, 2019.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- Günter Lumer and R.S. Phillips. Dissipative operators in a Banach space. *Pacific Journal of Mathematics*, 11:679–698, 1961.
- Alaeddin Malek and R. Shekari Beidokhti. Numerical solution for high order differential equations using a hybrid neural network–optimization method. *Applied Mathematics and Computation*, 183(1):260–271, 2006.
- Kevin Stanley McFall and James Robert Mahan. Artificial neural network method for solution of boundary value problems with exact satisfaction of arbitrary boundary conditions. *IEEE Transactions on Neural Networks*, 20(8):1221–1233, 2009.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 2388–2464. PMLR, 25–28 Jun 2019.
- George S Misyris, Andreas Venzke, and Spyros Chatzivasileiadis. Physics-informed neural networks for power systems. In *2020 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2020.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Huyên Pham, Xavier Warin, and Maximilien Germain. Neural networks-based backward scheme for fully nonlinear PDEs. *SN Partial Differential Equations and Applications*, 2(1):1–24, 2021.
- Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- Yeonjong Shin, Jerome Darbon, and George Em Karniadakis. On the convergence and generalization of physics informed neural networks. *arXiv preprint arXiv:2004.01806*, 2020.
- Justin Sirignano and Konstantinos Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *Journal of computational physics*, 375:1339–1364, 2018.
- Justin Sirignano, Jonathan MacArt, and Konstantinos Spiliopoulos. Pde-constrained models with neural network terms: Optimization and global convergence. *arXiv preprint arXiv:2105.08633*, 2021.
- Iain Smears and Endre Süli. Discontinuous Galerkin finite element approximation of Hamilton–Jacobi–Bellman equations with Cordes coefficients. *SIAM Journal on Numerical Analysis*, 52(2):993–1016, 2014.
- Vimal Thilak, Etai Littwin, Shuangfei Zhai, Omid Saremi, Roni Paiss, and Joshua Susskind. The slingshot mechanism: An empirical study of adaptive optimizers and the grokking phenomenon. *arXiv preprint arXiv:2206.04817*, 2022.
- Sifan Wang, Xinling Yu, and Paris Perdikaris. When and why PINNs fail to train: A neural tangent kernel perspective. *arXiv preprint arXiv:2007.14527*, 2020.
- Christopher J.C.H. Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- Yaohua Zang, Gang Bao, Xiaojing Ye, and Haomin Zhou. Weak adversarial networks for high-dimensional partial differential equations. *Journal of Computational Physics*, 411:109409, 2020.
- Eberhard Zeidler. *Nonlinear Functional Analysis and Its Applications: II/B: Nonlinear Monotone Operators*. Springer, 2013.
- Lulu Zhang, Tao Luo, Yaoyu Zhang, Zhi-Qin John Xu, and Zheng Ma. MOD-net: A machine learning approach via model-operator-data network for solving PDEs. *arXiv preprint arXiv:2107.03673*, 2021.

Xiangyu Zhao, Long Xia, Jiliang Tang, and Dawei Yin. Deep reinforcement learning for search, recommendation, and online advertising: a survey. *ACM SIGWEB Newsletter*, (Spring):1–15, 2019.