

Global Convergence of Adjoint-Optimized Neural PDEs

Konstantin Riedl

*University of Oxford, Mathematical Institute
Andrew Wiles Building, Oxford, OX2 6GG, UK*

KONSTANTIN.RIEDL@MATHS.OX.AC.UK

Justin Sirignano

*University of Oxford, Mathematical Institute
Andrew Wiles Building, Oxford, OX2 6GG, UK*

JUSTIN.SIRIGNANO@MATHS.OX.AC.UK

Konstantinos Spiliopoulos

*Boston University, Department of Mathematics & Statistics
665 Commonwealth Ave, Boston, MA 02215, USA*

KSPILIOP@BU.EDU

Editor: Jianfeng Lu

Abstract

Many engineering and scientific fields have recently become interested in modeling terms in partial differential equations (PDEs) with neural networks, which requires solving the inverse problem of learning neural network terms from observed data in order to approximate missing or unresolved physics in the PDE model. The resulting neural-network PDE model, being a function of the neural network parameters, can be calibrated to the available ground truth data by optimizing over the PDE using gradient descent, where the gradient is evaluated in a computationally efficient manner by solving an adjoint PDE. These neural PDE models have emerged as an important research area in scientific machine learning. In this paper, we study the convergence of the adjoint gradient descent optimization method for training neural PDE models in the limit where both the number of hidden units and the training time tend to infinity. Specifically, for a general class of nonlinear parabolic PDEs with a neural network embedded in the source term, we prove convergence of the trained neural-network PDE solution to the target data (i.e., a global minimizer). The global convergence proof poses a unique mathematical challenge that is not encountered in finite-dimensional neural network convergence analyses due to (i) the neural network training dynamics involving a non-local neural network kernel operator in the infinite-width hidden layer limit where the kernel lacks a spectral gap for its eigenvalues and (ii) the nonlinearity of the limit PDE system, which leads to a non-convex optimization problem in the neural network function even in the infinite-width hidden layer limit (unlike in typical neural network training cases where the optimization problem becomes convex in the large neuron limit). The theoretical results are illustrated and empirically validated by numerical studies.

Keywords: neural PDEs, neural-network PDEs, nonlinear PDEs, neural network terms, inverse problem, adjoint gradient descent method, infinite-width hidden layer limit

1. Introduction

Motivated by the remarkable successes of machine learning and deep learning (LeCun et al., 2015) in speech and image recognition (Hinton et al., 2012; Krizhevsky et al., 2012), computer vision (Krizhevsky et al., 2012), natural language processing (Vaswani et al.,

2017), and biology (Jumper et al., 2021), researchers and practitioners have sought to leverage and extend machine learning techniques to scientific disciplines, leading to the emergence of the field of scientific machine learning (SciML) (Karniadakis et al., 2021), where data-driven methods are integrated with physics-based modeling. SciML seeks to develop machine learning methods with physics-based modeling, integrating the modeling flexibility of neural networks (NNs) and large real-world datasets with well-established partial differential equation (PDE) models derived from physics. The field has developed a variety of different approaches. Physics-informed neural networks (PINNs) (Raissi et al., 2019; Karniadakis et al., 2021; Lu et al., 2021b; Cuomo et al., 2022; Wang et al., 2023; Jiang et al., 2023; Zhao and Luo, 2025), the deep Galerkin method (Sirignano and Spiliopoulos, 2018), the deep Ritz method (E and Yu, 2018), and neural Q-learning (Cohen et al., 2023a) exploit the property of an NN as a universal function approximator to approximate the PDE solution of a known PDE with an NN by training the NN parameters to satisfy the differential operator as well as initial and boundary conditions. In some scenarios, however, the PDE governing the physical phenomenon of interest is either entirely unknown or only partially accessible, leading to an inverse problem. When the PDE is unknown, operator learning (Boullé and Townsend, 2024; Li et al., 2021) attempts to learn, leveraging the universality of neural operators (Lu et al., 2021a; Kovachki et al., 2023) as approximators of nonlinear mappings, the PDE solution operator from observed data by minimizing a suitable loss. In the case of hidden, incomplete or unclosed physics, as in the setting of this paper, the PDE structure is known while certain terms are unknown (such as coefficients or source terms of the PDE). Leveraging again the property of NNs as universal function approximators, NNs can be trained to model the unknown PDE terms using observed data (Brunton et al., 2016; Schaeffer, 2017; Brunton and Kutz, 2019; Champion et al., 2019; Duraisamy et al., 2019; Brunton et al., 2020; Sun et al., 2020; Srivastava and Duraisamy, 2021; Duraisamy, 2021; Dong et al., 2022; Sirignano et al., 2023; Aarset et al., 2023; Dong et al., 2024; Holler and Morina, 2024).

Integration of machine learning with PDEs in science and engineering, combined with the growing availability of large datasets from field measurements, experiments, and high-fidelity numerical simulations, can yield more accurate engineering models and inform physical models with data-driven insights (Brunton et al., 2016; Schaeffer, 2017; Champion et al., 2019; Brunton et al., 2020; Brunton and Kutz, 2019; Sirignano et al., 2023) across a diverse range of application domains. NN terms in the PDE can be trained to learn missing, unknown, or unrepresented physics and correct numerical discretization errors. In computational fluid dynamics, for instance, NNs are introduced into the governing equations to represent the unclosed terms in PDE models of turbulent flows such as Reynolds-averaged Navier-Stokes (RANS) and Large-eddy Simulation (LES) equations (Tracey et al., 2015; Duraisamy et al., 2019; Srivastava and Duraisamy, 2021; Kochkov et al., 2021; Duraisamy, 2021; Sirignano and MacArt, 2023; Hickling et al., 2024; Kakka and MacArt, 2025; Nair et al., 2023, 2024). Neural PDE or neural-network PDE (NN-PDE) models also have applications in finance, economics, and biology (Gierjatowicz et al., 2020; Goswami et al., 2021; Cohen et al., 2023b; Fan and Sirignano, 2024).

Since the solution of the NN-PDE is a function of parameters of the NN which models certain terms/coefficients in the PDE, the NN parameters must be calibrated such that the NN-PDE solution matches the available ground truth data as closely as possible. Such

target data may come from real-world experiments or high-fidelity numerical simulations. In order to solve this inverse problem, i.e., calibrate the NN parameters, one must optimize over the NN-PDE via gradient descent on a suitable loss that quantifies the discrepancy between the NN-PDE solution and the target data. The gradient descent algorithm requires evaluating the gradient of the objective function, which depends upon the solution of the NN-PDE. Therefore, the gradient of the NN-PDE solution with respect to the NN parameters needs to be evaluated, which is computationally challenging. However, this gradient can be efficiently evaluated using the adjoint PDE of the NN-PDE. Adjoint optimization methods have been developed and applied to NN-PDEs in (Sirignano et al., 2023; Sirignano and Spiliopoulos, 2022; MacArt et al., 2021; Sirignano et al., 2020; Holland et al., 2019; Srivastava and Duraisamy, 2021). More generally, adjoint optimization has been widely used for PDE optimization (Bosse et al., 2014; Brandenburg et al., 2009; Bueno-Orovio et al., 2012; Cagnetti et al., 2013; Duta et al., 2002; Gauger et al., 2012; Giles and Pierce, 2000; Giles and Ulbrich, 2010a,b; Hazra and Schulz, 2004; Hazra, 2007; Hinze et al., 2008; Jameson et al., 1998; Jameson, 2003; Jameson and Kim, 2003; Kaland et al., 2014; Knopoff et al., 2013; Pierce and Giles, 2000; Nadarajah and Jameson, 2000, 2001; Reuther et al., 1996). The adjoint gradient descent algorithm solves a linear adjoint PDE at each optimization iteration, evaluates the gradient of the objective function with respect to the NN parameters using the adjoint PDE solution, and then takes a gradient descent step to update the NN parameters.

While adjoint methods have demonstrated effectiveness and efficiency across a wide range of applications, including design and shape optimization, aerodynamics, combustion, and tumor growth modeling in medicine, a rigorous mathematical analysis of adjoint gradient descent optimization methods in the setting of nonlinear NN-PDEs remains absent. The analysis in this paper is focused on this topic and provides first-of-its-kind convergence results in the nonlinear NN-PDE regime, where training the NN model leads to a non-convex optimization problem in the NN function, even in the large neuron limit. This is because the NN-PDE solution, which the loss is a function of, depends nonlinearly on the NN function, which is the design variable of our problem formulation and the quantity being trained. Our global convergence proof must address unique mathematical challenges that are not encountered in finite-dimensional NN convergence analyses due to (i) the NN training dynamics involving a non-local NN kernel operator in the infinite-width hidden layer limit where the kernel lacks a spectral gap for its eigenvalues and (ii) the nonlinearity of the limit PDE system, which leads to a non-convex optimization problem in the NN function even in the infinite-width hidden layer limit (unlike in typical NN training cases where the optimization problem becomes convex in the large neuron limit (Jacot et al., 2018; Chizat et al., 2019)).

We prove that the NN-PDE solution converges weakly to the target data (i.e., a global minimizer) during training. The first step is to prove that the adjoint PDE solution vanishes in the weak sense as the training time tends to infinity, which in turn requires establishing that a quadratic functional of the adjoint involving the positive definite non-local NN kernel operator converges to zero. The latter is proven by applying a cycle of stopping times analysis. The cycle of stopping times analysis requires the development of a novel approach for obtaining a regularity bound for this quadratic functional in terms of the learning rate, which is based on a careful PDE analysis of an adjoint PDE system associated with the quadratic functional, thus the analysis of the *adjoint of the original adjoint PDE*. The

derivation of bounds on several norms of the different adjoint PDEs is also required; see Sections 1.1 and 2.6 for more details. Developing these mathematical methods allows us to prove a much more general and stronger convergence result than done by prior analyses that only proved a substantially weaker notion of convergence for linear PDEs and a very restrictive class of objective functions (Sirignano et al., 2023); see Remark 36 for the details.

In our mathematical analysis, we will consider the second-order semi-linear parabolic neural-network partial differential equation (NN-PDE)

$$\begin{aligned} \partial_t u_\theta^N + \mathcal{L}u_\theta^N - q(u_\theta^N) &= g_\theta^N && \text{in } D_T, \\ u_\theta^N &= 0 && \text{on } [0, T] \times \partial D, \\ u_\theta^N &= f && \text{on } \{0\} \times D, \end{aligned} \tag{1}$$

where \mathcal{L} denotes a second-order linear PDE operator given in divergence form as

$$\mathcal{L}u = - \sum_{i,j=1}^d \partial_{x_j} (a^{ij}(t,x) \partial_{x_i} u) + \sum_{i=1}^d b^i(t,x) \partial_{x_i} u + c(t,x)u \tag{2}$$

with PDE coefficients $a^{ij}, b^i, c : [0, T] \times D \rightarrow \mathbb{R}$, and where $q : [0, T] \times D \times \mathbb{R} \rightarrow \mathbb{R}$ denotes the nonlinearity of the PDE. For notational convenience, we will sometimes (as done in (1)) omit writing the physical time and space dependency for the PDE solution as well as the PDE coefficients and terms, i.e., we write u instead of $u(t, x)$ or $q(u)$ instead of $q(t, x, u(t, x))$. The PDE operator \mathcal{L} together with its coefficients a^{ij} , b^i and c as well as the nonlinearity q are assumed to satisfy Assumptions A3, A4 and W1 as well as Assumptions A5, A6, W2 and W3, respectively. On the time-space domain $[0, T] \times D \subset \mathbb{R} \times \mathbb{R}^d$ we moreover impose Assumptions A1 and A2, and on the initial condition f Assumption W4 (see Section 2.1).

The PDE (1) is driven by an NN with parameters θ in the source term. We design it to be a fully-connected NN $g_\theta^N = g_\theta^N(t, x)$ with a single hidden layer consisting of N neurons, i.e., it takes the form

$$g_\theta^N(t, x) = \frac{1}{N^\beta} \sum_{i=1}^N c^i \sigma(w^{t,i}t + (w^i)^T x + \eta^i), \tag{3}$$

where the NN parameters of the i th neuron/unit are collected in the weight vector $\theta^i = (c^i, w^{t,i}, w^i, \eta^i) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}$ and where $\theta = (\theta^i)_{i=1, \dots, N}$ denotes the collection of all NN parameters, which are initialized independently according to $\theta_0^i = (c_0^i, w_0^{t,i}, w_0^i, \eta_0^i) \sim \mu_0$ with a measure $\mu_0 \in \mathcal{P}(\mathbb{R} \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R})$ obeying Assumption B3. The factor $1/N^\beta$ in (3) with $\beta \in (1/2, 1)$ is a normalization/scaling, and the NN nonlinearity σ satisfies Assumptions B1 and B2.

The aforementioned assumptions on the PDE are collected in Assumptions 6 and 7, and the ones on the NN in Assumption 9. They are assumed to hold throughout the manuscript.

Given a continuous target function $h \in L_2(D_T)$ corresponding to or representing measured or observed data (the ground truth), we wish to solve the inverse problem of calibrating the NN parameters θ so that the solution $u_\theta^N = u_\theta^N(t, x)$ to the NN-PDE (1) closely approximates the prescribed target data h . For this purpose, we seek to minimize the loss (least squares

loss)

$$\mathcal{J}_\theta^N = \frac{1}{2} \int_0^T \int_D (u_\theta^N(t, x) - h(t, x))^2 dx dt \quad (4)$$

by training¹ the NN parameters θ via continuous-time gradient descent

$$\frac{d}{d\tau} \theta_\tau = -\alpha_\tau^N \nabla_\theta \mathcal{J}_\theta^N, \quad (5)$$

where $\alpha_\tau^N = \frac{\alpha_\tau}{N^{1-2\beta}}$ denotes the learning rate which is assumed to be decreasing in the training time τ and satisfies the well-known Robbins-Monro conditions (Robbins and Monro, 1951)

$$\int_0^\infty \alpha_\tau d\tau = \infty \quad \text{and} \quad \int_0^\infty \alpha_\tau^2 d\tau < \infty. \quad (6)$$

Monotonicity as well as the conditions (6) are standard requirements for learning rates in machine learning (Bertsekas and Tsitsiklis, 2000). A suitable and classical example fulfilling those conditions is given by $\alpha_\tau = \frac{1}{1+\tau}$. Computing the gradient $\nabla_\theta \mathcal{J}_\theta^N$ w.r.t. the NN parameters θ in (5) is computationally challenging due to its dependency on the solution u_θ^N of the PDE (1). Deriving naively a PDE for $\nabla_\theta u_\theta^N$ by applying the gradient to (1) yields a PDE whose dimension is equal to the number of the NN parameters θ . Typically, the number of NN parameters is very large (thousands or even hundreds of thousands), leading to a high-dimensional PDE which is computationally costly to solve numerically. Analogously, trying to estimate the gradient by numerical differentiation with finite differences amounts to an equivalently expensive and thus infeasible task. A computationally efficient way to evaluate the gradient $\nabla_\theta \mathcal{J}_\theta^N$, however, is given by the adjoint method (Giles and Pierce, 2000), which can be regarded as a continuous PDE version of the usual backpropagation algorithm. As we verify in Lemma 10, the gradient $\nabla_\theta \mathcal{J}_\theta^N$ w.r.t. the NN parameters θ can be computed according to

$$\nabla_\theta \mathcal{J}_\theta^N = \int_0^T \int_D \nabla_\theta g_\theta^N(t, x) \widehat{u}_\theta^N(t, x) dx dt, \quad (7)$$

which requires solving the adjoint PDE of (1), which is a second-order linear parabolic backward PDE of the form

$$\begin{aligned} -\partial_t \widehat{u}_\theta^N + \mathcal{L}^\dagger \widehat{u}_\theta^N - q_u(u_\theta^N) \widehat{u}_\theta^N &= (u_\theta^N - h) && \text{in } D_T, \\ \widehat{u}_\theta^N &= 0 && \text{on } [0, T] \times \partial D, \\ \widehat{u}_\theta^N &= 0 && \text{on } \{T\} \times D, \end{aligned} \quad (8)$$

where \mathcal{L}^\dagger denotes the adjoint of \mathcal{L} which is given by

$$\mathcal{L}^\dagger \widehat{u} = - \sum_{i,j=1}^d \partial_{x_i} (a^{ij}(t, x) \partial_{x_j} \widehat{u}) - \sum_{i=1}^d b^i(t, x) \partial_{x_i} \widehat{u} + \left(c(t, x) - \sum_{i=1}^d \partial_{x_i} b^i(t, x) \right) \widehat{u}, \quad (9)$$

and where $q_u = \partial_u q : [0, T] \times D \times \mathbb{R} \rightarrow \mathbb{R}$ denotes the partial derivative of q w.r.t. u , i.e., the function $(t, x, u) \mapsto q_u(t, x, u) = \partial_u q(t, x, u)$.

1. Note that the training time, denoted by τ , is distinct from and unrelated to the physical PDE time t .

Adjoint optimization is widely used in practice due to its computational efficiency for evaluating the gradient of a PDE objective function. Computing (7) requires solving only one linear PDE of the same dimension as the original PDE (1) that is being optimized, which is computationally tractable. Evaluating the gradient $\nabla_{\theta} g_{\theta}^N$ of the NN g_{θ}^N itself w.r.t. the NN parameters θ is typically done efficiently using automatic differentiation frameworks such as TensorFlow or PyTorch.

Before discussing our analytical setup, let us address the well-posedness (existence and uniqueness) of the PDE system (1) & (8) coupled with the gradient descent update (5) during training. As we rigorously prove in Lemma 11, there exists, for any finite training time horizon \mathcal{T} , a unique weak solution $((u_{\theta_{\tau}}^N, \widehat{u}_{\theta_{\tau}}^N))_{\tau \in [0, \mathcal{T}]}$ in the sense of Definitions 4 and 5 in the space $\mathcal{C}([0, \mathcal{T}], \mathcal{S} \times \mathcal{S})$, where $\mathcal{S} := L_2([0, T], H^1(D)) \cap L_{\infty}([0, T], L_2(D))$.

For training nonlinear parabolic NN-PDE models of the form (1), we study in this paper the global convergence of the adjoint gradient descent optimization method (5), where the gradient of the loss \mathcal{J}_{θ}^N w.r.t. the NN parameters θ is computed by solving the adjoint PDE (8) and evaluating formula (7). We consider the theoretical limit where both the number of neurons N in the NN g_{θ}^N in (3) and the training time τ in (5) tend to infinity.

As a first step, let us derive and theoretically justify the limiting training dynamics in the infinite-width hidden layer limit, i.e., as the number of hidden units $N \rightarrow \infty$. Therefore, denote by $\mu_{\tau}^N = \frac{1}{N} \sum_{i=1}^N \delta_{c_{\tau}^i, w_{\tau}^{t,i}, w_{\tau}^i, \eta_{\tau}^i}$ the empirical measure at training time τ of the NN parameters of our fully-connected NN (3) with a single hidden layer with N neurons and their parameters $\theta_{\tau} = (c_{\tau}^i, w_{\tau}^{t,i}, w_{\tau}^i, \eta_{\tau}^i)_{i=1, \dots, N}$. By computing with chain rule

$$\begin{aligned} \frac{d}{d\tau} g_{\theta_{\tau}}^N(t, x) &= \nabla_{\theta} g_{\theta_{\tau}}^N(t, x) \cdot \frac{d}{d\tau} \theta_{\tau} \\ &= -\alpha_{\tau}^N \int_0^T \int_D \nabla_{\theta} g_{\theta_{\tau}}^N(t, x) \cdot \nabla_{\theta} g_{\theta_{\tau}}^N(t', x') \widehat{u}_{\theta_{\tau}}^N(t', x') dx' dt', \end{aligned} \tag{10}$$

as done in detail in (41)–(43), we obtain by the fundamental theorem of calculus for the training time evolution of the NN function $g_{\theta_{\tau}}^N$ that

$$g_{\theta_{\tau}}^N(t, x) = g_{\theta_0}^N(t, x) - \int_0^{\tau} \alpha_s [T_{B(\mu_s^N)} \widehat{u}_{\theta_s}^N](t, x) ds \tag{11}$$

with the NN integral operator $T_{B(\mu)}$ defined as

$$[T_{B(\mu)} \widehat{u}](t, x) = \int_0^T \int_D B(t, x, t', x'; \mu) \widehat{u}(t', x') dx' dt' \tag{12}$$

and where the symmetric non-local NN kernel (also known as the neural tangent kernel (NTK) (Jacot et al., 2018)) is given by

$$B(t, x, t', x'; \mu) = \langle k(t, x, t', x'; c, w^t, w, \eta), \mu(dc, dw^t, dw, d\eta) \rangle \tag{13}$$

with

$$\begin{aligned} k(t, x, t', x'; c, w^t, w, \eta) &= \sigma(w^t t + w^T x + \eta) \sigma(w^t t' + w^T x' + \eta) \\ &\quad + c^2 \sigma'(w^t t + w^T x + \eta) \sigma'(w^t t' + w^T x' + \eta) (t t' + x^T x' + 1). \end{aligned} \tag{14}$$

That means, the NN function $g_{\theta_\tau}^N$ follows during training the kernel gradient of the least squares loss (4) using the pre-limit NTK $B(\mu_\tau^N)$, which is random at initialization and varies during training, as can be seen from (11). In contrast, in the infinite-width hidden layer limit, i.e., as the number of hidden units $N \rightarrow \infty$ in (11), the kernel becomes deterministic and converges to a limit NTK, which remains constant during training, as can be seen from (64). This is similar to and in line with the overparameterized training phenomenon (Jacot et al., 2018; Chizat et al., 2019) observed for certain scalings in (3), yet requires, due to the nonlinear PDE setting considered in this manuscript, detailed computations, which we provide in the proof of Theorem 1 below. Theorem 1 allows us to represent the limit NN function g_τ^* during training by the integro-differential equation²

$$g_\tau^*(t, x) = - \int_0^\tau \alpha_s [T_{B_0} \hat{u}_s^*](t, x) ds \quad (15)$$

with the constant limit NTK $B_0 = B(\mu_0) = B(\bullet, \bullet, \bullet, \bullet; \mu_0)$, where μ_0 is the probability distribution for the parameter initialization of the NN. The representation (15) of the NN function g_τ^* during training reveals a linearization of the NN training dynamics around their initialization. In particular, while the learning rate for individual NN parameters is $\frac{\alpha_\tau}{N^{1-\beta}}$, as can be seen from (24), and thus converges to zero as the number of parameters $N \rightarrow \infty$, the NN function g_τ^* itself has the non-zero learning rate α_τ , as apparent from (15). Thus, due to the large number of degrees of freedom in the overparameterized regime, the individual parameters are required to move smaller and smaller distances from their initial locations to achieve a given magnitude change in the neural network output as $N \rightarrow \infty$.

As the number of hidden units $N \rightarrow \infty$, the NN source term $g_{\theta_\tau}^N$ of the PDE (1) converges to g_τ^* while the PDE solution $u_{\theta_\tau}^N$ and the solution to the adjoint PDE $\hat{u}_{\theta_\tau}^N$ converge in $L_2([0, T], H^1(D))$ - and $L_\infty([0, T], L_2(D))$ -norm to functions u_τ^* and \hat{u}_τ^* solving the PDE system

$$\begin{aligned} \partial_t u_\tau^* + \mathcal{L}u_\tau^* - q(u_\tau^*) &= g_\tau^* && \text{in } D_T, \\ u_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\ u_\tau^* &= f && \text{on } \{0\} \times D, \end{aligned} \quad (16)$$

and

$$\begin{aligned} -\partial_t \hat{u}_\tau^* + \mathcal{L}^\dagger \hat{u}_\tau^* - q_u(u_\tau^*) \hat{u}_\tau^* &= (u_\tau^* - h) && \text{in } D_T, \\ \hat{u}_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\ \hat{u}_\tau^* &= 0 && \text{on } \{T\} \times D, \end{aligned} \quad (17)$$

which is coupled with the integro-differential equation (15) for g_τ^* . Before making this joint convergence as the number of neurons N tends to infinity mathematically precise in Theorem 1, let us address the well-posedness of the PDE system (16)–(17) coupled with (15). As we rigorously prove in Lemma 13, there exists, for any finite training time horizon \mathcal{T} , a unique weak solution $((u_\tau^*, \hat{u}_\tau^*))_{\tau \in [0, \mathcal{T}]}$ in the sense of Definitions 4 and 5 in the space $\mathcal{C}([0, \mathcal{T}], \mathcal{S} \times \mathcal{S})$.

2. Note that equation (15) can be written after taking the training time derivative equivalently as the infinite-dimensional ODE $\frac{d}{dt} g_\tau^*(t, x) = -\alpha_\tau T_{B_0} \hat{u}_\tau^* = -\alpha_\tau \int_0^t \int_D B(t, x, t', x'; \mu_0) \hat{u}_\tau^*(t', x') dx' dt'$ with \hat{u}_τ^* depending nonlinearly on g_τ^* according to the nonlinear PDE system (16)–(17).

Theorem 1 (Overparameterized training regime) *Assume that the learning rate satisfies additionally $\int_0^\infty \alpha_\tau^{4/3} d\tau < \infty$. Let $\mathcal{T} < \infty$ be a given training time horizon. For each N , let us denote by $((u_{\theta_\tau}^N, \widehat{u}_{\theta_\tau}^N))_{\tau \in [0, \mathcal{T}]} \in \mathcal{C}([0, \mathcal{T}], \mathcal{S} \times \mathcal{S})$ the unique weak solution to the PDE system (1)–(8) coupled with the gradient descent update (5) in the sense of Lemma 11, and let us denote by $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in [0, \mathcal{T}]} \in \mathcal{C}([0, \mathcal{T}], \mathcal{S} \times \mathcal{S})$ the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) in the sense of Lemma 13. Then, as the number of hidden units $N \rightarrow \infty$,*

$$\sup_{\tau \in [0, \mathcal{T}]} \mathbb{E} \left[\|u_{\theta_\tau}^N - u_\tau^*\|_{L_2([0, \mathcal{T}], H^1(D))} + \|u_{\theta_\tau}^N - u_\tau^*\|_{L_\infty([0, \mathcal{T}], L_2(D))} \right] \rightarrow 0, \quad (18a)$$

$$\sup_{\tau \in [0, \mathcal{T}]} \mathbb{E} \left[\|\widehat{u}_{\theta_\tau}^N - \widehat{u}_\tau^*\|_{L_2([0, \mathcal{T}], H^1(D))} + \|\widehat{u}_{\theta_\tau}^N - \widehat{u}_\tau^*\|_{L_\infty([0, \mathcal{T}], L_2(D))} \right] \rightarrow 0, \quad (18b)$$

$$\sup_{\tau \in [0, \mathcal{T}]} \mathbb{E} \|g_{\theta_\tau}^N - g_\tau^*\|_{L_2(D_{\mathcal{T}})} \rightarrow 0. \quad (18c)$$

Here, the expectation \mathbb{E} is taken w.r.t. the random initialization of the NN parameters. (The only source of randomness is the random initialization of the NN parameters before training begins.)

Theorem 1 proves the convergence of the solution $((u_{\theta_\tau}^N, \widehat{u}_{\theta_\tau}^N))_{\tau \in [0, \mathcal{T}]}$ to $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in [0, \mathcal{T}]}$ as the number of hidden units N tends to infinity for any finite training time horizon $\mathcal{T} < \infty$. Therefore, the NN-PDE trained with adjoint gradient descent optimization converges to the solution of the limit PDE system (16)–(17) coupled with the integro-differential equation (15) as the number of hidden units $N \rightarrow \infty$. A detailed proof of Theorem 1 is presented in Section 2.5.

Remark 2 *Theorem 1 requires the well-posedness (existence and uniqueness) of both the pre-limit PDE system (1)–(8) coupled with the gradient descent update (5) and the limit PDE system (16)–(17) coupled with the integro-differential equation (15). We rigorously state those results in Lemmas 11 and 13, respectively, with their detailed proofs given in Appendix A.*

The additional assumption $\int_0^\infty \alpha_\tau^{4/3} d\tau < \infty$ on the learning rate, which is satisfied by typical learning rates, is exclusively required for the well-posedness of the pre-limit PDE system (1)–(8) as stated in Lemma 11 below and commented on in more details thereafter in Remark 12. This additional integrability assumption is not used elsewhere in the paper.

In this large neuron limit, we show as a second step that the dynamics (16)–(17) coupled with (15) converges to a global minimizer of the loss

$$\mathcal{J}_\tau^* = \frac{1}{2} \int_0^{\mathcal{T}} \int_D (u_\tau^*(t, x) - h(t, x))^2 dx dt \quad (19)$$

as the training time τ tends to infinity. We thus prove the convergence of the NN-PDE solution u_τ^* to the target data h (i.e., a global minimizer) as $\tau \rightarrow \infty$. To be more precise, the following main convergence result about the adjoint gradient descent optimization method is proven in this paper.

Theorem 3 (Global convergence of NN-PDE) *Let $((u_\tau^*, \hat{u}_\tau^*))_{\tau \in [0, \infty)} \in \mathcal{C}([0, \infty), \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) in the sense of Lemma 13 and Remark 14 on the training time interval $[0, \infty)$. Then, the loss \mathcal{J}_τ^* defined in (19) is monotonically decreasing with $\frac{d}{d\tau} \mathcal{J}_\tau^* = -\alpha_\tau \mathcal{Q}_\tau^* = -\alpha_\tau (\hat{u}_\tau^*, T_{B_0} \hat{u}_\tau^*)_{L_2(D_T)} \leq 0$, and the solution u_τ^* to (16) converges weakly to the target h in L_2 as $\tau \rightarrow \infty$, i.e.,*

$$u_\tau^* \rightharpoonup h \text{ in } L_2 \quad \text{as } \tau \rightarrow \infty. \quad (20)$$

The statement follows from Propositions 26 and 35. A detailed proof sketch of Theorem 3 is presented in Section 2.6.

1.1 Contributions

Motivated by the popularity and effectiveness of the adjoint gradient descent optimization method (5) for training NN-PDE models and thereby solving the inverse problem of learning the neural network modeled terms from observed data as demonstrated in the literature, our paper develops a rigorous global convergence analysis of this machine learning algorithm for a general class of nonlinear parabolic NN-PDEs of the form (1). To calibrate the NN-PDE to available data, the method trains the NN parameters θ embedded within the PDE by running gradient descent on the least squares loss (L_2 -loss) \mathcal{J}_θ^N with the gradient being evaluated in a computationally efficient manner by solving an associated adjoint PDE. This is a highly non-convex optimization problem and therefore, for a finite number of hidden units N , the trained NN-PDE may only converge to a local minimizer of the objective function. We study the algorithm’s asymptotic convergence behavior in the limit where both the number of hidden units N of the NN g_θ^N in (3) and the training time τ in continuous-time gradient descent (5) tend to infinity. First-of-its-kind convergence results to a global minimizer are proven in the nonlinear setting, which go significantly beyond previous analyses that considered much more restrictive classes of linear PDEs, a very restrictive class of objective functions, and a substantially weaker notion of convergence.

Our first result is about the convergence to the infinite-width hidden layer limit as the number N of neurons tends to infinity. We prove that, as $N \rightarrow \infty$, the NN function g_θ^N converges to its infinite-width hidden layer limit g^* in (15), which can be represented during training by an integro-differential equation involving a positive definite non-local NN kernel operator (12) that remains constant during training but lacks a spectral gap; that is, its eigenvalues do not have a uniform positive lower bound. The NN-PDE solution u_θ^N and the adjoint \hat{u}_θ^N , which solve the PDE system (1) & (8) coupled with the gradient descent update (5), converge to u^* and \hat{u}^* solving the PDE system (16)–(17) coupled with the integro-differential equation (15).

Our second result proves global convergence of the trained NN-PDE solution u_τ^* to the target data h (i.e., a global minimizer) as the training time τ goes to infinity. Due to the PDE system (16)–(17) coupled with the integro-differential equation (15) being both nonlinear and non-local, several mathematical challenges need to be addressed. In particular, due to the nonlinearity of the PDE, training the NN model leads to a non-convex optimization problem in the NN function g_τ^* even in the large neuron limit. This is very different from typical NN limits (e.g., gradient descent training of a standard feedforward fully-connected

network) where the training of the infinite-width NN is shown to satisfy the gradient flow of a convex function.

Furthermore, the aforementioned lack of a spectral gap in the NN kernel of the non-local NN kernel operator in the infinite-width hidden layer limit poses a unique technical complication that is not encountered in finite-dimensional NN convergence analyses (where the eigenvalues of the NN kernel matrix have a positive lower bound). By showing that the quadratic functional $\mathcal{Q}_\tau^* = (\hat{u}_\tau^*, T_{B_0} \hat{u}_\tau^*)_{L_2(D_T)}$ of the PDE adjoint \hat{u}_τ^* , which appears in the training time derivative $\frac{d}{d\tau} \mathcal{J}_\tau^* = -\alpha_\tau \mathcal{Q}_\tau^*$ of the loss \mathcal{J}^* and involves the positive definite non-local NN kernel operator T_{B_0} , converges to zero as the training time $\tau \rightarrow \infty$, we establish the weak convergence of the solution \hat{u}_τ^* to the adjoint PDE (17) to zero as $\tau \rightarrow \infty$. The weak convergence of the adjoint PDE solution can then be used to prove that the original NN-PDE solution u_τ^* to the nonlinear PDE (16) converges weakly to the target data h as $\tau \rightarrow \infty$.

To prove that the functional \mathcal{Q}_τ^* of the adjoint vanishes as the training time τ tends to infinity, we apply a cycle of stopping times analysis. This technique crucially requires the development of a novel approach for obtaining a regularity bound for the functional \mathcal{Q}_τ^* in terms of the learning rate α_τ , which involves the analysis of an adjoint associated with the functional \mathcal{Q}_τ^* , thus the analysis of an adjoint PDE system of the adjoint PDE (17). This is a second-level adjoint system of the original adjoint PDE. In addition, the proof requires carefully establishing uniform (in the training time τ) bounds on several norms of the different adjoint PDEs.

We expect that the developed mathematical methods can be applied to other PDEs and NN architectures in scientific machine learning. For example, we prove a result of independent interest that (strong) limit points of the trained NN-PDE solution are global minimizers of the loss \mathcal{J}^* for an even more general class of second-order parabolic NN-PDEs.

Numerical studies that illustrate and support our theoretical findings are also presented in the paper.

1.2 Organization

In Section 2, we discuss in detail the main contributions of this paper. Therefore, after collecting all assumptions made throughout this paper in Section 2.1, we derive in Section 2.2 formula (7) for $\nabla_\theta \mathcal{J}_\theta^N$, before providing in Sections 2.3 and 2.4, respectively, well-posedness results for the NN-PDE training dynamics in both the finite-width hidden layer regime and the infinite-width hidden layer limit, i.e., for the PDE system (1) & (8) coupled with the gradient descent update (5) and the PDE system (16)–(17) coupled with the integro-differential equation (15). Their proofs are provided in Appendix A. Afterwards, we elaborate on and prove in Section 2.5 our first main theoretical result, Theorem 1, which is about the convergence of $(g_{\theta_\tau}^N, u_{\theta_\tau}^N, \hat{u}_{\theta_\tau}^N)$ to their infinite-width hidden layer limit counterparts $(g_\tau^*, u_\tau^*, \hat{u}_\tau^*)$ as the number N of neurons tends to infinity. We conclude with Section 2.6, where we discuss and provide an insightful proof sketch of our second main theoretical result, Theorem 3, which is concerned with the convergence of the limit NN-PDE solution u_τ^* to the target data h during training, i.e., as the training time $\tau \rightarrow \infty$. Its proof is based on several auxiliary results which we discuss in detail in the thematic Sections 4 to 9.

Section 3 contains numerical examples demonstrating the theoretical results of the paper. We provide the code implementing the adjoint gradient descent optimization method in the GitHub repository <https://github.com/KonstantinRiedl/NNPDEs>.

As discussed, the proofs of the main results are contained in Sections 4 to 9. Section 4 is dedicated to presenting the mathematical tools related to the NN. In Section 5, we compute the training time derivative of the loss \mathcal{J}_τ^* and show that $\frac{d}{d\tau}\mathcal{J}_\tau^* = -\alpha_\tau\mathcal{Q}_\tau^*$ with the quadratic functional $\mathcal{Q}_\tau^* = (\widehat{u}_\tau^*, T_{B_0}\widehat{u}_\tau^*)_{L_2(D_T)}$ of the adjoint. This implies in particular that \mathcal{J}_τ^* is monotonically non-increasing. Leveraging this property, we provide in Section 6 uniform (in the training time τ) bounds on several norms of the PDE solution u_τ^* and the adjoint \widehat{u}_τ^* , which eventually, by analyzing a second-level adjoint system of the original adjoint PDE in Section 7, permit to establish a regularity bound for the functional \mathcal{Q}_τ^* in terms of the learning rate α_τ . Adapting a cycle of stopping times analysis (Bertsekas and Tsitsiklis, 2000; Sirignano and Spiliopoulos, 2022) while leveraging the aforementioned regularity bound, we eventually prove in Section 8 that $\frac{d}{d\tau}\mathcal{J}_\tau^* = -\alpha_\tau\mathcal{Q}_\tau^*$ implies $\mathcal{Q}_\tau^* \rightarrow 0$ as the training time $\tau \rightarrow \infty$. With the positive definiteness of the NN kernel operator T_{B_0} we therefrom infer in Section 9 the weak convergence $\widehat{u}_\tau^* \rightharpoonup 0$ and thus $u_\tau^* \rightharpoonup h$ by definition of the adjoint PDE (17). We conclude Section 9 by proving as a result of independent interest, that (strong) limit points u_∞^* of the trained NN-PDE solution u_τ^* satisfy $u_\infty^* \equiv h$ a.e., thus being global minimizers of the loss \mathcal{J}^* .

1.3 Notation

We denote by $D \subset \mathbb{R}^d$ the spatial domain of the considered parabolic PDE. Its boundary is ∂D . T denotes the physical time horizon of the PDE. $D_T := (0, T) \times D$ denotes the time-space domain. Its lateral surface is $\partial D_T := [0, T] \times \partial D$, and $\Gamma_T := \partial D_T \cup \{(t, x) : t = 0, x \in D\}$. Moreover, for $\Delta T' > 0$, we introduce the notation $D_{T', T'+\Delta T'} := (T', T'+\Delta T') \times D$.

For a spatial domain D , the spaces $L_p(D)$ and $W_p^k(D)$ denote the classical Lebesgue and Sobolev spaces. They contain all measurable functions $u : D \rightarrow \mathbb{R}$ with finite corresponding norm. For the norms on those spaces it holds $\|u\|_{L_p(D)}^p = \int_D |u(x)|^p dx$ and $\|u\|_{W_p^k(D)}^p = \sum_{|\alpha| \leq k} \|D^\alpha u\|_{L_p(D)}^p$ or $\|u\|_{L_\infty(D)} = \text{ess sup}_{x \in D} |u(x)|$ and $\|u\|_{W_\infty^k(D)} = \max_{|\alpha| \leq k} \|D^\alpha u\|_{L_\infty(D)}$ in the case $p = \infty$. We abbreviate $H^k(D) = W_2^k(D)$ and denote by $H_0^1(D)$ the space of all functions in $H^1(D)$ with zero trace. $H^{-1}(D)$ denotes the dual space of $H_0^1(D)$.

For a time-space domain D_T , the spaces $L_p(D_T)$ denote the classical Lebesgue spaces. They contain all measurable functions $u : D_T \rightarrow \mathbb{R}$ with finite corresponding norm. For the norms on those spaces it hold $\|u\|_{L_p(D_T)}^p = \int_0^T \int_D |u(t, x)|^p dx dt$ and $\|u\|_{L_\infty(D_T)} = \text{ess sup}_{(t, x) \in D_T} |u(t, x)|$ in the case $p = \infty$.

For a function space \mathcal{X} on the space D , the spaces $L_p([0, T], \mathcal{X})$ denote the Bochner spaces. Let us associate with a function $u : D_T \rightarrow \mathbb{R}$ the mapping $\mathbf{u} : [0, T] \rightarrow \mathcal{X}$ defined by $\mathbf{u}(t) := u(t, \bullet)$. In what follows we may abuse notation and write u in place of \mathbf{u} . The Bochner spaces contain all strongly (Bochner) measurable functions with finite Bochner norm. For those norms it hold $\|u\|_{L_p([0, T], \mathcal{X})}^p = \int_0^T \|u(t, \bullet)\|_{\mathcal{X}}^p dt$ or $\|u\|_{L_\infty([0, T], \mathcal{X})} = \text{ess sup}_{t \in [0, T]} \|u(t, \bullet)\|_{\mathcal{X}}$ in the case $p = \infty$, see (Evans, 2010, Section 5.9.2).

A weak solution to the nonlinear parabolic PDE (1) in the sense of (Evans, 2010, Chapter 7) is defined as follows.

Definition 4 (Weak solution of (1)) A function $u_\theta^N \in L_2([0, T], H_0^1(D))$ with weak derivative $\partial_t u_\theta^N \in L_2([0, T], H^{-1}(D))$ is a weak solution of the PDE (1) provided

$$(i) \quad \langle \partial_t u_\theta^N(t, \bullet), v \rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B}[u_\theta^N(t, \bullet), v; t] - (q(u_\theta^N(t, \bullet)), v)_{L_2(D)} = (g_\theta(t, \bullet), v)_{L_2(D)}$$

for each $v \in H_0^1(D)$ and a.e. time $t \in [0, T]$, where the bilinear form \mathcal{B} is given by

$$\mathcal{B}[u, v; t] := \int_D \sum_{i,j=1}^d a^{ij}(t, x) \partial_{x_i} u \partial_{x_j} v + \sum_{i=1}^d b^i(t, x) \partial_{x_i} u v + c(t, x) u v \, dx, \quad (21)$$

and

$$(ii) \quad u_\theta^N(0, \bullet) = f.$$

Analogously, we define a weak solution to the linear PDE (8) as follows.

Definition 5 (Weak solution of (8)) A function $\hat{u}_\theta^N \in L_2([0, T], H_0^1(D))$ with weak derivative $\partial_t \hat{u}_\theta^N \in L_2([0, T], H^{-1}(D))$ is a weak solution of the adjoint PDE (8) (parabolic backward PDE) provided

$$(i) \quad \langle -\partial_t \hat{u}_\theta^N(t, \bullet), v \rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B}^\dagger[\hat{u}_\theta^N(t, \bullet), v; t] - (q_u(u_\theta^N(t, \bullet)) \hat{u}_\theta^N(t, \bullet), v)_{L_2(D)} \\ = (u_\theta^N(t, \bullet) - h, v)_{L_2(D)}$$

for each $v \in H_0^1(D)$ and a.e. time $t \in [0, T]$, where \mathcal{B}^\dagger denotes the adjoint bilinear form satisfying $\mathcal{B}^\dagger[\hat{u}, u; t] = \mathcal{B}[u, \hat{u}; t]$, and

$$(ii) \quad \hat{u}_\theta^N(T, \bullet) = 0.$$

Since we investigate the evolution of the PDE solutions to (1) & (8) during training (see (5)), we are interested in their training time trajectories which we denote by $((u_{\theta_\tau}^N, \hat{u}_{\theta_\tau}^N))_{\tau \in [0, T]}$. The function space $\mathcal{C}([0, T], \mathcal{S} \times \mathcal{S})$ denotes the space of all such continuous trajectories, i.e., the space of all continuous functions mapping from $[0, T]$ to $\mathcal{S} \times \mathcal{S}$.

By C we typically denote generic constants, which may vary throughout the proof. To keep the notation concise, we indicate by α their dependency on α_0 or $\int_0^\infty \alpha_\tau^2 d\tau$ (see (6)), by D their dependency on $\text{vol}(D)$ or $|D|$ (see Assumption A2), by \mathcal{L} their dependency on ν , some norms of a^{ij} , b^i , c , as well as their partial space derivatives, or some norms of f and h (see Assumptions A3, A4, W1 and W4), by q their dependency on properties of q (see Assumptions A5, A6, W2 and W3), by σ their dependency on properties of the NN nonlinearity σ (see Assumptions B1 and B2), and by μ_0 their dependency on properties of μ_0 (see Assumption B3).

2. Discussion of the Main Results

This section is dedicated to the discussion of the main theoretical contributions of this paper.

2.1 Assumptions

Let us start by stating all assumptions used throughout this manuscript. We cluster them into assumptions related to the PDE (1), which we summarize in Assumptions 6 and 7, and assumptions on the NN listed thereafter in Assumption 9.

Assumption 6 (Second-order semi-linear parabolic PDE (1)) *Throughout we assume that the time horizon T of the PDE (1) is finite and that the spatial domain $D \subset \mathbb{R}^d$ of the PDE (1)*

A1 is an open connected set with a C^2 smooth boundary ∂D ,

A2 has finite volume $\text{vol}(D)$ and is bounded by $|D|$.

Moreover, we assume that

A3 the parabolic PDE operator $\partial_t + \mathcal{L}$ is uniformly parabolic, i.e., there exists $\nu > 0$ such that $\sum_{i,j=1}^d a^{ij}(t,x)\xi_i\xi_j \geq \nu \|\xi\|^2$ for all $(t,x) \in \overline{D_T}$ and $\xi \in \mathbb{R}^d$,

A4 the coefficients $a^{ij}, b^i, c \in L_\infty(D_T)$ and $\partial_{x_k} a^{ij}, \partial_{x_k} b^i \in L_\infty(D_T)$,

A5 the nonlinearity q is such that $|q_u| \leq c_q$ for a constant $c_q > 0$,

A6 the nonlinearity q is such that $|q_{uu}| \leq c'_q$ for a constant $c'_q > 0$.

Assumption 7 (Well-posedness of second-order semi-linear parabolic PDE (1))

Moreover, we assume that

W1 the coefficients $a^{ij}, b^i, c \in L_\infty(D_T)$ and $\partial_{x_k} a^{ij}, \partial_{x_k} b^i \in L_\infty(D_T)$ are $(\gamma_1/2, \gamma_1)$ -Hölder continuous in (t,x) with $\gamma_1 > 0$,

W2 the nonlinearity q is $(\gamma_1/2, \gamma_1)$ -Hölder continuous in (t,x) with $\gamma_1 > 0$ and such that $|q(u)| \leq C_q(1 + |u|)$ for any $u \in \mathbb{R}$ and for a constant $C_q > 0$,

W3 the nonlinearity q is such that q_u is continuous,

W4 the initial condition $f \in C^2(D)$ with $f|_{\partial D} = 0$ is γ_2 -Hölder continuous.

Remark 8 *The conditions of Assumption 7 are required only for the well-posedness proof in Lemmas 11 and 13. If well-posedness as below can be shown under a different set of assumptions, these new assumptions would replace Assumption 7.*

Assumption 9 (Neural network in (3)) *Throughout we assume that the NN is such that*

B1 the nonlinearity σ of the NN is non-constant, bounded (i.e., $|\sigma| \leq C_\sigma$), and L_σ -Lipschitz continuous,

B2 the derivative σ' of the nonlinearity σ of the NN is bounded (i.e., $|\sigma'| \leq C_{\sigma'}$) and $L_{\sigma'}$ -Lipschitz continuous,

B3 the randomly initialized NN parameters $\theta_0^i = (c_0^i, w_0^{t,i}, w_0^i, \eta_0^i)$ are i.i.d. and drawn from a distribution $\mu_0 \in \mathcal{P}(\mathbb{R} \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R})$ which is such that

- (i) c_0^i is independent from $(w_0^{t,i}, w_0^i, \eta_0^i)$,
- (ii) the marginal distribution $\mu_{0,c}$ of c_0^i is mean-zero and compactly supported,
- (iii) the marginal distribution $\mu_{0,(w^t,w,\eta)}$ of $(w_0^{t,i}, w_0^i, \eta_0^i)$ has bounded k th-order moments $\mathcal{M}_k(\mu_{0,(w^t,w,\eta)})$ for $k = \max\{4, d + 2\}$,
- (iv) the marginal distribution $\mu_{0,(w^t,w,\eta)}$ of $(w_0^{t,i}, w_0^i, \eta_0^i)$ assigns positive probability to every set with positive Lebesgue measure.

2.2 A Computationally Efficient Formula for $\nabla_\theta \mathcal{J}_\theta^N$

As pointed out in the introduction, the practicability of the adjoint gradient descent method (5) is thanks to an efficient computation of the gradient $\nabla_\theta \mathcal{J}_\theta^N$. The following result proves (7). Its proof is given at the end of Section 5.

Lemma 10 *Let $u_\theta^N, \nabla_\theta g_\theta^N \in L_2(D_T)$ and let \widehat{u}_θ^N denote a weak solution to (8) in the sense of Definition 5, which satisfies $\partial_t \widehat{u}_\theta^N(t, \bullet) \in L_2(D)$ for a.e. $t \in [0, T]$. Define the loss \mathcal{J}_θ^N as in (4). Then, the gradient $\nabla_\theta \mathcal{J}_\theta^N$ w.r.t. the NN parameters θ can be written as (7).*

2.3 Well-Posedness of the NN-PDE Training Dynamics in the Finite-Width Hidden Layer Regime

We first provide a result about the well-posedness of the NN-PDE training dynamics in the finite-width hidden layer regime, i.e., for the PDE system (1) & (8) coupled with the gradient descent update (5) for the NN parameters of the NN function $g_{\theta_\tau}^N$.

Lemma 11 (Well-posedness of NN-PDE training dynamics (1) & (8)) *Let $N \in \mathbb{N}$ be fixed and let $\theta_0 = (\theta_0^i)_{i=1,\dots,N}$ be initialized such that $\theta_0^i \sim \mu_0$ for each $i = 1, \dots, N$. Assume that the learning rate satisfies additionally $\int_0^\infty \alpha_\tau^{4/3} d\tau < \infty$. Let $\mathcal{T} < \infty$ be a given training time horizon. Then there exists a unique weak solution*

$$((u_{\theta_\tau}^N, \widehat{u}_{\theta_\tau}^N)_{\tau \in [0, \mathcal{T}]}) \in \mathcal{C}([0, \mathcal{T}], \mathcal{S} \times \mathcal{S}) \quad (22)$$

to the PDE system (1) & (8) coupled with the gradient descent update (5) in the sense of Definitions 4 and 5 which satisfies $(\partial_t u_{\theta_\tau}^N(t, \bullet), \partial_t \widehat{u}_{\theta_\tau}^N(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$ and for every $\tau \in [0, \mathcal{T}]$.

The proof is based on a fixed point argument which allows one to decouple the PDE system (1) & (8) from the gradient descent update (5). After invoking classical existence results for the nonlinear PDE system (1) & (8) from (Ladyženskaja et al., 1968) for given NN parameter updates $\widetilde{c}_\tau^i, \widetilde{w}_\tau^{t,i}, \widetilde{w}_\tau^i$, and $\widetilde{\eta}_\tau^i$, we eventually employ the Banach fixed point theorem to prove well-posedness of the PDE system (1) & (8) coupled with the gradient descent update (5) on a local training time domain, which is eventually extended by a bootstrapping argument.

Remark 12 *The additional assumption $\int_0^\infty \alpha_\tau^{4/3} d\tau < \infty$ on the learning rate required in Lemma 11 is slightly stronger than (6) but satisfied by typical learning rates such as $\alpha_\tau = \frac{1}{1+\tau}$. We leverage this assumption in the proof of Lemma 11 to deal with the technical challenges arising from the NN integral operator varying during training.*

More precisely, when establishing contractivity estimates for the application of the Banach fixed point theorem, this nonlinearity leads to the appearance of higher-order (up to fourth-order) product terms between the NN parameters and the adjoint PDE solution, see, e.g., (263), where the bound on the right-hand side scales as $C(1 + M)^4 \mathcal{T}$ with M denoting a bound on the NN parameters and the adjoint PDE solution up to time \mathcal{T} . When extending well-posedness by a bootstrapping argument from a training time interval $[0, \mathcal{T}_{k-1}]$ to $[0, \mathcal{T}_k = \mathcal{T}_{k-1} + \Delta \mathcal{T}_k]$, this requires us (see particularly (271)) to choose time intervals of the form $\Delta \mathcal{T}_k \propto 1/(C(1 + M_k^4))$, more precisely (269). One can see from (270) that assuming only (6) on the learning rate would inevitably lead to the choice $M_k \propto Ck^{1/2} + C(h, f, \theta_0)$. However, in that case, we notice that the series $\sum_{k=0}^{\infty} \Delta \mathcal{T}_k$ is essentially a geometric series, which does not diverge. Therefore, this would not allow to extend well-posedness to arbitrary time horizons. In contrast, the additional assumption on the learning rate allows us to control more stringently the worst-case growth of the NN parameters, the NN function, and the NN-PDE solution during training, see, e.g., (252) and (268), where we estimate that they grow as $M_k \propto Ck^{1/4} + C(h, f, \theta_0)$. This can be exploited to balance the appearance of the higher-order product terms with the slower growth, enabling us to eventually get that the series $\sum_{k=0}^{\infty} \Delta \mathcal{T}_k$ is essentially a harmonic series, which diverges, therefore allowing for a global well-posedness result, see, e.g., (272).

Due to its technical nature, the proof is deferred to Appendix A.2.

2.4 Well-Posedness of the NN-PDE Training Dynamics in the Infinite-Width Hidden Layer Limit

Let us now provide a result about the well-posedness of the NN-PDE training dynamics in the infinite-width hidden layer limit, i.e., for the PDE system (16)–(17) coupled with the integro-differential equation (15) for g_τ^* .

Lemma 13 (Well-posedness of NN-PDE training dynamics (16)–(17)) *Let $\mathcal{T} < \infty$ be a given training time horizon. Then there exists a unique weak solution*

$$((u_\tau^*, \hat{u}_\tau^*)_{\tau \in [0, \mathcal{T}]} \in \mathcal{C}([0, \mathcal{T}], \mathcal{S} \times \mathcal{S}) \quad (23)$$

to the PDE system (16)–(17) coupled with (15) in the sense of Definitions 4 and 5 which satisfies $(\partial_t u_\tau^*(t, \bullet), \partial_t \hat{u}_\tau^*(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$ and for every $\tau \in [0, \mathcal{T}]$.

The proof resembles the one of Lemma 11 and is based again on a fixed point argument which allows decoupling the PDE system (16)–(17) from the integro-differential equation (15) as before. After invoking classical existence results for the nonlinear PDE system (16)–(17) from (Ladyženskaja et al., 1968) for a given right-hand side \tilde{g}_τ , we eventually employ the Banach fixed point theorem to prove well-posedness of the PDE system (16)–(17) coupled with the integro-differential equation (15) on a local training time domain, which is eventually extended by a bootstrapping argument.

Due to its technical nature, the proof is deferred to Appendix A.1.

Remark 14 *With the statements of Lemma 11 and Lemma 13 being valid for arbitrary training time horizons \mathcal{T} , we can infer well-posedness of the NN-PDE training dynamics*

(1) & (8) as well as well-posedness of the NN-PDE training dynamics (16)–(17) on the infinite training time interval $[0, \infty)$. In particular, it is proven in Step 2e in the proofs of Lemma 11 and Lemma 13 that the Banach fixed point theorem gives existence of the corresponding solution globally in the training time.

2.5 Infinite-Width Neural Network Perspective

Our first main theoretical result, Theorem 1, rigorously proves that the PDE system (16)–(17) coupled with the integro-differential equation (15) is indeed the correct limit of the PDE system (1) & (8) coupled with the gradient descent update (5) as the number of neurons $N \rightarrow \infty$.

Proof [Proof of Theorem 1] Step 1: Boundedness of gradient descent updates. Let us first prove that the gradient descent updates (5) are uniformly bounded in N and in the training time τ for $\tau \in [0, \mathcal{T}]$. According to formula (5) and using the definition of the learning rate α_τ^N it holds for continuous-time gradient descent $\frac{d}{d\tau}\theta_\tau = -\frac{\alpha_\tau}{N^{1-2\beta}}\nabla_\theta\mathcal{J}_{\theta_\tau}^N$, which allows to explicitly derive expressions for $\frac{d}{d\tau}c_\tau^i$, $\frac{d}{d\tau}w_\tau^{t,i}$, $\frac{d}{d\tau}w_\tau^i$, and $\frac{d}{d\tau}\eta_\tau^i$. With the fundamental theorem of calculus we therefrom infer

$$c_\tau^i = c_0^i - \frac{1}{N^{1-\beta}} \int_0^\tau \alpha_s \int_0^T \int_D \sigma(w_s^{t,i}t + (w_s^i)^T x + \eta_s^i) \widehat{u}_{\theta_s^N}^N(t, x) dx dt ds, \quad (24a)$$

$$w_\tau^{t,i} = w_0^{t,i} - \frac{1}{N^{1-\beta}} \int_0^\tau \alpha_s \int_0^T \int_D c_s^i \sigma'(w_s^{t,i}t + (w_s^i)^T x + \eta_s^i) t \widehat{u}_{\theta_s^N}^N(t, x) dx dt ds, \quad (24b)$$

$$w_\tau^i = w_0^i - \frac{1}{N^{1-\beta}} \int_0^\tau \alpha_s \int_0^T \int_D c_s^i \sigma'(w_s^{t,i}t + (w_s^i)^T x + \eta_s^i) x \widehat{u}_{\theta_s^N}^N(t, x) dx dt ds, \quad (24c)$$

$$\eta_\tau^i = \eta_0^i - \frac{1}{N^{1-\beta}} \int_0^\tau \alpha_s \int_0^T \int_D c_s^i \sigma'(w_s^{t,i}t + (w_s^i)^T x + \eta_s^i) \widehat{u}_{\theta_s^N}^N(t, x) dx dt ds. \quad (24d)$$

Exploiting that σ is bounded as of Assumption B1 and that the domain D has bounded volume as of Assumption A2, we can use (24a) to bound with Cauchy-Schwarz inequality

$$\begin{aligned} |c_\tau^i - c_0^i| &\leq \frac{1}{N^{1-\beta}} \int_0^\tau \alpha_s \sqrt{\int_0^T \int_D (\sigma(w_s^{t,i}t + (w_s^i)^T x + \eta_s^i))^2 dx dt} \|\widehat{u}_{\theta_s^N}^N\|_{L_2(D_T)} ds \\ &\leq \frac{C}{N^{1-\beta}} \int_0^\tau \alpha_s \|\widehat{u}_{\theta_s^N}^N\|_{L_2(D_T)} ds \end{aligned} \quad (25)$$

for a constant $C = C(T, D, \sigma)$.

By following the computations of Step 1c in the proof of Lemma 13 in Appendix A that lead to (179) and (186) for the solutions to the PDE system (1) & (8) coupled with the gradient descent update (5), we obtain the bounds

$$\|u_{\theta_\tau}^N\|_{L_2([0, T], H^1(D))} + \|u_{\theta_\tau}^N\|_{L_\infty([0, T], L_2(D))} \leq C \left(\|f\|_{L_2(D)} + \|g_{\theta_\tau}^N\|_{L_2(D_T)} + 1 \right) \quad (26)$$

and

$$\|\widehat{u}_{\theta_\tau}^N\|_{L_2([0, T], H^1(D))} + \|\widehat{u}_{\theta_\tau}^N\|_{L_\infty([0, T], L_2(D))} \leq C \left(\|u_{\theta_\tau}^N\|_{L_2(D_T)} + \|h\|_{L_2(D_T)} \right) \quad (27)$$

for a constant $C = C(T, \mathcal{L}, q)$ which is in particular independent of N . Using the definition of the NN (3) we can estimate with Jensen's inequality

$$\begin{aligned} \|g_{\theta_\tau}^N\|_{L_2(D_T)}^2 &= \int_0^T \int_D \left(\frac{1}{N^\beta} \sum_{i=1}^N c_\tau^i \sigma(w_\tau^{t,i} t + (w_\tau^i)^T x + \eta_\tau^i) \right)^2 dx dt \\ &\leq C \frac{1}{N^{2\beta-2}} \frac{1}{N} \sum_{i=1}^N (c_\tau^i)^2 = C \frac{1}{N^{2\beta-2}} \gamma_\tau^N, \end{aligned} \quad (28)$$

for $C = C(T, D, \sigma) < \infty$, where we used the boundedness of σ as of Assumptions B1 and that the domain D has bounded volume as of Assumption A2. In the last step, we introduced the notation $\gamma_\tau^N := \frac{1}{N} \sum_{i=1}^N (c_\tau^i)^2$. Combining (26)–(28), we end up with the bound

$$\|\widehat{u}_{\theta_\tau}^N\|_{L_2([0,T], H^1(D))}^2 + \|\widehat{u}_{\theta_\tau}^N\|_{L_\infty([0,T], L_2(D))}^2 \leq C \left(\frac{1}{N^{2\beta-2}} \gamma_\tau^N + \|f\|_{L_2(D)}^2 + \|h\|_{L_2(D_T)}^2 + 1 \right) \quad (29)$$

for a constant $C = C(T, D, \mathcal{L}, q, \sigma)$.

After squaring both sides of (25) and using Cauchy-Schwarz inequality we obtain

$$\begin{aligned} |c_\tau^i - c_0^i|^2 &\leq \frac{C}{N^{2(1-\beta)}} \int_0^\tau \alpha_s^2 ds \int_0^\tau \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^2 ds \\ &\leq C \int_0^\tau \gamma_s^N ds + \frac{C\tau}{N^{2(1-\beta)}} \left(\|f\|_{L_2(D)}^2 + \|h\|_{L_2(D_T)}^2 + 1 \right), \end{aligned} \quad (30)$$

where we inserted (29) and used the second part of (6) in the second step. Summing over $i = 1, \dots, N$ and normalizing by N we can bound

$$\begin{aligned} \gamma_\tau^N &\leq 2\gamma_0^N + \frac{2}{N} \sum_{i=1}^N |c_\tau^i - c_0^i|^2 \\ &\leq 2\gamma_0^N + C \int_0^\tau \gamma_s^N ds + \frac{C\tau}{N^{2(1-\beta)}} \left(\|f\|_{L_2(D)}^2 + \|h\|_{L_2(D_T)}^2 + 1 \right). \end{aligned} \quad (31)$$

Since γ_0^N is compactly supported due to Assumption B3, an application of Grönwall's inequality gives the estimate

$$\sup_{\tau \in [0, \mathcal{T}]} \gamma_\tau^N \leq C \quad (32)$$

for a constant $C = C(\alpha, \mathcal{T}, T, D, \mathcal{L}, q, \sigma, \mu_0)$ which is in particular independent of N . Employing (32) in (30) shows after using Young's inequality that

$$|c_\tau^i|^2 \leq 2|c_0^i|^2 + C + \frac{C}{N^{2(1-\beta)}} \left(\|f\|_{L_2(D)}^2 + \|h\|_{L_2(D_T)}^2 + 1 \right), \quad (33)$$

for some other, potentially larger, constant C . Recalling that the parameters c_0^i are initialized with compact support as of Assumption B3, (33) proves that

$$\sup_{N \in \mathbb{N}} \sup_{i=1, \dots, N, \tau \in [0, \mathcal{T}]} |c_\tau^i| \leq C_c, \quad (34)$$

for a constant $C_c = C_c(\alpha, \mathcal{T}, T, D, \mathcal{L}, q, \sigma, \mu_0)$ which is in particular independent of N .

Leveraging that the NN parameters c_τ^i are uniformly bounded and exploiting that σ' is bounded as of Assumption B2 and that the domain D is bounded as of Assumption A2, we can use (24b)–(24d) to bound (analogously to (25)) with Cauchy-Schwarz inequality

$$|w_\tau^{t,i} - w_0^{t,i}| + \|w_\tau^i - w_0^i\| + |\eta_\tau^i - \eta_0^i| \leq \frac{C}{N^{1-\beta}} \int_0^\tau \alpha_s \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)} ds \quad (35)$$

for a constant $C = C(T, D, \sigma, C_c)$. After squaring both sides of (35) and using Cauchy-Schwarz inequality we obtain

$$\begin{aligned} |w_\tau^{t,i} - w_0^{t,i}|^2 + \|w_\tau^i - w_0^i\|^2 + |\eta_\tau^i - \eta_0^i|^2 &\leq \frac{C}{N^{2(1-\beta)}} \int_0^\tau \alpha_s^2 ds \int_0^\tau \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^2 ds \\ &\leq C \int_0^\tau \gamma_s^N ds + \frac{C\tau}{N^{2(1-\beta)}} \left(\|f\|_{L_2(D)}^2 + \|h\|_{L_2(D_T)}^2 + 1 \right), \end{aligned} \quad (36)$$

where we inserted (29) and used the second part of (6) in the second step. Employing (32) in (36) shows after using Young's inequality that for $\tau \in [0, \mathcal{T}]$ with $\mathcal{T} < \infty$

$$|w_\tau^{t,i}|^2 \leq 2|w_0^{t,i}|^2 + C + \frac{C}{N^{2(1-\beta)}} \left(\|f\|_{L_2(D)}^2 + \|h\|_{L_2(D_T)}^2 + 1 \right), \quad (37)$$

$$\|w_\tau^i\|^2 \leq 2\|w_0^i\|^2 + C + \frac{C}{N^{2(1-\beta)}} \left(\|f\|_{L_2(D)}^2 + \|h\|_{L_2(D_T)}^2 + 1 \right), \quad (38)$$

$$|\eta_\tau^i|^2 \leq 2|\eta_0^i|^2 + C + \frac{C}{N^{2(1-\beta)}} \left(\|f\|_{L_2(D)}^2 + \|h\|_{L_2(D_T)}^2 + 1 \right), \quad (39)$$

for a constant $C = C(\alpha, \mathcal{T}, T, D, \mathcal{L}, q, \mu_0, C_c)$ which is in particular independent of N .

Recalling that the parameters $w_0^{t,i}$, w_0^i , and η_0^i are initialized according to the measure μ_0 , whose marginal distribution $\mu_{0,(w^t,w,\eta)}$ of $(w_0^{t,i}, w_0^i, \eta_0^i)$ has bounded moments as of Assumption B3(iii), (37)–(39) prove that

$$\sup_{N \in \mathbb{N}} \sup_{i=1, \dots, N, \tau \in [0, \mathcal{T}]} \mathbb{E} \left[|w_\tau^{t,i}| + \|w_\tau^i\| + |\eta_\tau^i| \right] \leq C_{w,\eta} \quad (40)$$

for a constant $C_{w,\eta} = C_{w,\eta}(\alpha, \mathcal{T}, T, D, \mathcal{L}, q, \sigma, \mu_0, C_c)$ which is in particular independent of N .

Step 2: Boundedness of the NN (3). Combining the explicit expressions for $\frac{d}{d\tau} c_\tau^i$, $\frac{d}{d\tau} w_\tau^{t,i}$, $\frac{d}{d\tau} w_\tau^i$, and $\frac{d}{d\tau} \eta_\tau^i$, we obtain for $\frac{d}{d\tau} g_{\theta_\tau}^N$ by taking the training time derivative in (3) that

$$\begin{aligned} \frac{d}{d\tau} g_{\theta_\tau}^N(t, x) &= \frac{1}{N^\beta} \sum_{i=1}^N \left(\frac{d}{d\tau} c_\tau^i \right) \sigma(\star) + c_\tau^i \sigma'(\star) \left(\left(\frac{d}{d\tau} w_\tau^{t,i} \right) t + \left(\frac{d}{d\tau} w_\tau^i \right)^T x + \frac{d}{d\tau} \eta_\tau^i \right) \\ &= -\frac{\alpha_\tau}{N} \sum_{i=1}^N \int_0^T \int_D (\sigma(\star) \sigma(\star') + (c_\tau^i)^2 \sigma'(\star) \sigma'(\star') (tt' + x^T x' + 1)) \widehat{u}_{\theta_\tau}^N(t', x') dx' dt', \end{aligned} \quad (41)$$

where we abbreviated $\star = w_\tau^{t,i} t + (w_\tau^i)^T x + \eta_\tau^i$ and $\star' = w_\tau^{t,i} t' + (w_\tau^i)^T x' + \eta_\tau^i$. Denoting now by $\mu_\tau^N = \frac{1}{N} \sum_{i=1}^N \delta_{c_\tau^i, w_\tau^{t,i}, w_\tau^i, \eta_\tau^i}$ the empirical measure at training time τ of our fully-connected NN (3) with a single hidden layer with N neuron and their parameters $\theta_\tau =$

$(c_\tau^i, w_\tau^{t,i}, w_\tau^i, \eta_\tau^i)_{i=1,\dots,N}$, and using the definition of the NN kernel B from (13), we can rewrite the formula for $\frac{d}{d\tau}g_{\theta_\tau}^N$ in (41) as

$$\frac{d}{d\tau}g_{\theta_\tau}^N(t, x) = -\alpha_\tau \int_0^T \int_D B(t, x, t', x'; \mu_\tau^N) \widehat{u}_{\theta_\tau}^N(t', x') dx' dt' = -\alpha_\tau T_{B(\mu_\tau^N)} \widehat{u}_{\theta_\tau}^N, \quad (42)$$

where we used the definition of the NN integral operator T_B from (12) in the last step. Simple integration in the training time τ yields by the fundamental theorem of calculus

$$g_{\theta_\tau}^N(t, x) = g_{\theta_0}^N(t, x) + \int_0^\tau \frac{d}{ds} g_{\theta_s}^N(t, x) ds = g_{\theta_0}^N(t, x) - \int_0^\tau \alpha_s T_{B(\mu_s^N)} \widehat{u}_{\theta_s}^N ds. \quad (43)$$

It is straightforward to see from the definition of the kernel B in (13) that we can bound

$$\|B(\mu_\tau^N)\|_{L_2(D_T \times D_T)} \leq C \left(1 + \frac{1}{N} \sum_{i=1}^N (c_\tau^i)^2 \right) \leq C \quad (44)$$

for a constant $C = C(T, D, \sigma, C_c) < \infty$ due to the boundedness Assumptions B1, B2, and A2, and using (34) in the last step. We can use this in (43) to bound with Young's inequality, Cauchy-Schwarz inequality and the second part of (6),

$$\begin{aligned} \|g_{\theta_\tau}^N\|_{L_2(D_T)}^2 &\leq 2 \|g_{\theta_0}^N\|_{L_2(D_T)}^2 + 2 \left\| \int_0^\tau \alpha_s T_{B(\mu_s^N)} \widehat{u}_{\theta_s}^N ds \right\|_{L_2(D_T)}^2 \\ &\leq 2 \|g_{\theta_0}^N\|_{L_2(D_T)}^2 + 2 \int_0^\tau \alpha_s^2 ds \int_0^\tau \|T_{B(\mu_s^N)} \widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^2 ds \\ &\leq 2 \|g_{\theta_0}^N\|_{L_2(D_T)}^2 + C \int_0^\tau \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^2 ds \end{aligned} \quad (45)$$

for a constant $C = C(\alpha, T, D, \sigma, C_c) < \infty$ which is in particular independent of N . Using the explicit form (3) we notice further that for any $p \geq 2$ it holds

$$\begin{aligned} \mathbb{E} \|g_{\theta_0}^N\|_{L_p(D_T)}^p &\leq \frac{C(p)}{N^{p\beta}} \mathbb{E} \int_0^T \int_D \left(\sum_{i=1}^N |c_0^i \sigma(w_0^{t,i} t + (w_0^i)^T x + \eta_0^i)|^2 \right)^{p/2} dx dt \\ &\leq \frac{1}{N^{p\beta-p/2}} C \leq C \end{aligned} \quad (46)$$

for a constant $C = C(p, T, D, \sigma, C_c) < \infty$. To obtain the first inequality in (46) we used the Marcinkiewicz-Zygmund inequality with random variables $z^i(t, x) = c_0^i \sigma(w_0^{t,i} t + (w_0^i)^T x + \eta_0^i)$, which are independent thanks to the initial independence of the parameters $\theta_0^i = (c_0^i, w_0^{t,i}, w_0^i, \eta_0^i)$ as of Assumption B3, mean-zero due to the c_0^i having zero mean and being drawn independently from the other parameters as of Assumptions B3(ii) and (i), and have finite p th moments due to the c_0^i 's being compactly supported as of Assumption B3(ii) together with the boundedness of σ from Assumption B1. The last two reasons also justify the second inequality in (46), while the third inequality holds since $\beta \in (1/2, 1)$. With $p = 2$, this allows to conclude (45) after taking the expectation with the estimate

$$\sup_{s \in [0, \tau]} \mathbb{E} \|g_{\theta_s}^N\|_{L_2(D_T)}^2 \leq C \left(1 + \int_0^\tau \mathbb{E} \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^2 ds \right) \quad (47)$$

for a constant $C = C(\alpha, T, D, \sigma, C_c) < \infty$ which is in particular independent of N .

Leveraging this estimate we can bound the norm of the adjoint in expectation as

$$\begin{aligned}
 & \sup_{s \in [0, \tau]} \mathbb{E} \left[\|\widehat{u}_{\theta_s}^N\|_{L_2([0, T], H^1(D))}^2 + \|\widehat{u}_{\theta_s}^N\|_{L_\infty([0, T], L_2(D))}^2 \right] \\
 & \leq C \left(\sup_{s \in [0, \tau]} \mathbb{E} \left[\|u_{\theta_s}^N\|_{L_2(D_T)}^2 \right] + \|h\|_{L_2(D_T)}^2 \right) \\
 & \leq C \left(\sup_{s \in [0, \tau]} \mathbb{E} \left[\|g_{\theta_s}^N\|_{L_2(D_T)}^2 \right] + \|f\|_{L_2(D)}^2 + \|h\|_{L_2(D_T)}^2 + 1 \right) \\
 & \leq C \left(\int_0^\tau \mathbb{E} \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^2 ds + \|f\|_{L_2(D)}^2 + \|h\|_{L_2(D_T)}^2 + 1 \right) \\
 & \leq C \left(\int_0^\tau \mathbb{E} \left[\|\widehat{u}_{\theta_s}^N\|_{L_2([0, T], H^1(D))}^2 + \|\widehat{u}_{\theta_s}^N\|_{L_\infty([0, T], L_2(D))}^2 \right] ds + \|f\|_{L_2(D)}^2 + \|h\|_{L_2(D_T)}^2 + 1 \right), \tag{48}
 \end{aligned}$$

where we used the estimates (27) and (26) in the first and second inequality, respectively. After employing (47) to obtain the next-to-last line, an application of Grönwall's inequality yields the uniform bound

$$\sup_{s \in [0, \tau]} \mathbb{E} \left[\|\widehat{u}_{\theta_s}^N\|_{L_2([0, T], H^1(D))}^2 + \|\widehat{u}_{\theta_s}^N\|_{L_\infty([0, T], L_2(D))}^2 \right] \leq C \tag{49}$$

for a constant $C = C(\alpha, \mathcal{T}, T, D, \mathcal{L}, q, \sigma, C_c) < \infty$ which is in particular independent of N . Herefore, note that $\mathbb{E}[\|\widehat{u}_{\theta_0}^N\|_{L_2([0, T], H^1(D))}^2 + \|\widehat{u}_{\theta_0}^N\|_{L_\infty([0, T], L_2(D))}^2] \leq C$ for a constant $C = C(T, D, \mathcal{L}, q, \sigma, C_c) < \infty$ as of (27) and (26) together with the fact that $\mathbb{E}\|g_{\theta_0}^N\|_{L_2(D_T)}^2 \leq C$ according to (46).

Step 3: Convergence as $N \rightarrow \infty$. With the adjoint $\widehat{u}_{\theta_\tau}^N$ being bounded uniformly (in the number N of NN parameters) in expectation as of (49), we immediately derive from (25) and (35) after taking the expectation that

$$\sup_{i=1, \dots, N, \tau \in [0, T]} \mathbb{E} \left[|c_\tau^i - c_0^i| + |w_\tau^{t,i} - w_0^{t,i}| + \|w_\tau^i - w_0^i\| + |\eta_\tau^i - \eta_0^i| \right] \leq \frac{C}{N^{1-\beta}} \tag{50}$$

for a constant $C = C(\alpha, \mathcal{T}, T, D, \mathcal{L}, q, \sigma, \mu_0) < \infty$ which is in particular independent of N .

Let us now prove (18). Recall that $(u_{\theta_\tau}^N, \widehat{u}_{\theta_\tau}^N)$ and $(u_\tau^*, \widehat{u}_\tau^*)$ are solutions to the PDE system (1) & (8) coupled with the gradient descent update (5) and the PDE system (16)–(17) coupled with the integro-differential equation (15), respectively. Following the computations of *Step 2c* in the proof of Lemma 13 in Appendix A that lead to in particular (202) and (207) as well as (219) we obtain the bounds

$$\|u_{\theta_\tau}^N - u_\tau^*\|_{L_2([0, T], H^1(D))} + \|u_{\theta_\tau}^N - u_\tau^*\|_{L_\infty([0, T], L_2(D))} \leq C \|g_{\theta_\tau}^N - g_\tau^*\|_{L_2(D_T)}, \tag{51}$$

and, with $p = d + 2$ which satisfies $p > d + 1$, by employing Morrey's inequality

$$\|u_{\theta_\tau}^N - u_\tau^*\|_{L_\infty(D_T)} \leq C \|g_{\theta_\tau}^N - g_\tau^*\|_{L_p(D_T)} \tag{52}$$

as well as,

$$\begin{aligned}
 & \left\| \widehat{u}_{\theta_\tau}^N - \widehat{u}_\tau^* \right\|_{L_2([0,T],H^1(D))} + \left\| \widehat{u}_{\theta_\tau}^N - \widehat{u}_\tau^* \right\|_{L_\infty([0,T],L_2(D))} \\
 & \leq C \left(\left\| u_{\theta_\tau}^N - u_\tau^* \right\|_{L_2(D_T)} + \left\| \widehat{u}_\tau^* \right\|_{L_\infty([0,T],L_2(D))} \left\| u_{\theta_\tau}^N - u_\tau^* \right\|_{L_\infty(D_T)} \right) \\
 & \leq C \left(\left\| u_{\theta_\tau}^N - u_\tau^* \right\|_{L_2(D_T)} + \left\| u_{\theta_\tau}^N - u_\tau^* \right\|_{L_\infty(D_T)} \right)
 \end{aligned} \tag{53}$$

at the cost of some other, potentially larger, constant $C = C(\alpha, \mathcal{T}, T, D, \mathcal{L}, q, \sigma, \mu_0, C_2^B) < \infty$. Here, we used in the last step the fact that \widehat{u}_τ^* obeys a deterministic bound, which follows after an application of Grönwall's inequality from the chain of inequalities

$$\begin{aligned}
 \left\| \widehat{u}_\tau^* \right\|_{L_2([0,T],H^1(D))}^2 + \left\| \widehat{u}_\tau^* \right\|_{L_\infty([0,T],L_2(D))}^2 & \leq C \left(\left\| u_\tau^* \right\|_{L_2(D_T)}^2 + \left\| h \right\|_{L_2(D_T)}^2 \right) \\
 & \leq C \left(\left\| g_\tau^* \right\|_{L_2(D_T)}^2 + \left\| f \right\|_{L_2(D)}^2 + \left\| h \right\|_{L_2(D_T)}^2 + 1 \right) \\
 & \leq C \left(\int_0^\tau \left\| \widehat{u}_s^* \right\|_{L_2(D_T)}^2 ds + \left\| f \right\|_{L_2(D)}^2 + \left\| h \right\|_{L_2(D_T)}^2 + 1 \right),
 \end{aligned} \tag{54}$$

where the first two inequalities are obtained by following the computations of *Step 1c* in the proof of Lemma 13 in Appendix A that lead to in particular (186) and (179), while the last step is true since by Cauchy-Schwarz inequality, Lemma 22 and the second part of (6) we have $\left\| g_\tau^* \right\|_{L_2(D_T)}^2 = \left\| \int_0^\tau \alpha_s T_{B_0} \widehat{u}_s^* ds \right\|_{L_2(D_T)}^2 \leq \int_0^\tau \alpha_s^2 ds \int_0^\tau \left\| T_{B_0} \widehat{u}_s^* \right\|_{L_2(D_T)}^2 ds \leq C \int_0^\tau \left\| \widehat{u}_s^* \right\|_{L_2(D_T)}^2 ds$ for a constant $C = C(\alpha, C_2^B)$.

Since, as we established in (43), the NN $g_{\theta_\tau}^N$ in the source term of the PDE (8) can be represented during training by an integro-differential equation similarly to the representation of g_τ^* in (15), we can estimate with triangle inequality

$$\begin{aligned}
 \left\| g_{\theta_\tau}^N - g_\tau^* \right\|_{L_p(D_T)} & = \left\| g_{\theta_0}^N - \int_0^\tau \alpha_s \left(T_{B(\mu_s^N)} \widehat{u}_{\theta_s}^N - T_{B_0} \widehat{u}_s^* \right) ds \right\|_{L_p(D_T)} \\
 & \leq \left\| g_{\theta_0}^N \right\|_{L_p(D_T)} + \int_0^\tau \alpha_s \left\| \left(T_{B(\mu_s^N)} - T_{B_0} \right) \widehat{u}_{\theta_s}^N \right\|_{L_p(D_T)} ds \\
 & \quad + \int_0^\tau \alpha_s \left\| T_{B_0} \left(\widehat{u}_{\theta_s}^N - \widehat{u}_s^* \right) \right\|_{L_p(D_T)} ds
 \end{aligned} \tag{55}$$

and it remains to control each term individually in expected value. To estimate the expectation of the first term on the right-hand side of (55) we use Jensen's inequality recalling that $p \geq 1$ since $p = d + 2$ and use (46) which yields

$$\mathbb{E} \left\| g_{\theta_0}^N \right\|_{L_p(D_T)} \leq \left(\mathbb{E} \left\| g_{\theta_0}^N \right\|_{L_p(D_T)}^p \right)^{1/p} \leq \frac{1}{N^{\beta-1/2}} C \tag{56}$$

for a constant $C = C(p, T, D, \sigma, C_c) < \infty$ which is in particular independent of N . To bound the second term of (55) in expectation we first note that with triangle inequality it holds

$$\begin{aligned}
 & \mathbb{E} \left\| \left(T_{B(\mu_s^N)} - T_{B_0} \right) \widehat{u}_{\theta_s}^N \right\|_{L_p(D_T)} \\
 & \leq \mathbb{E} \left\| \left(T_{B(\mu_s^N)} - T_{B(\mu_0^N)} \right) \widehat{u}_{\theta_s}^N \right\|_{L_p(D_T)} + \mathbb{E} \left\| \left(T_{B(\mu_0^N)} - T_{B_0} \right) \widehat{u}_{\theta_s}^N \right\|_{L_p(D_T)}.
 \end{aligned} \tag{57}$$

For the first term in (57) we bound with two applications of Cauchy-Schwarz inequality

$$\begin{aligned}
 & \mathbb{E} \left\| \left(T_{B(\mu_s^N)} - T_{B(\mu_0^N)} \right) \widehat{u}_{\theta_s}^N \right\|_{L_p(D_T)} \\
 &= \mathbb{E} \left(\int_0^T \int_D \left| \int_0^T \int_D B(t, x, t', x'; \mu_s^N - \mu_0^N) \widehat{u}_{\theta_s}^N(t', x') dx' dt' \right|^p dx dt \right)^{1/p} \\
 &\leq \mathbb{E} \left[\left(\int_0^T \int_D \left(\int_0^T \int_D (B(t, x, t', x'; \mu_s^N - \mu_0^N))^2 dx' dt' \right)^{p/2} dx dt \right)^{1/p} \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)} \right] \\
 &\leq \left(\mathbb{E} \left(\int_0^T \int_D \left(\int_0^T \int_D (B(t, x, t', x'; \mu_s^N - \mu_0^N))^2 dx' dt' \right)^{p/2} dx dt \right)^{2/p} \right)^{1/2} \\
 &\quad \cdot \left(\mathbb{E} \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^2 \right)^{1/2} \\
 &\leq C \left(\mathbb{E} \left(\int_0^T \int_D \left(\int_0^T \int_D (B(t, x, t', x'; \mu_s^N - \mu_0^N))^2 dx' dt' \right)^{p/2} dx dt \right)^{2/p} \right)^{1/2}
 \end{aligned} \tag{58}$$

for a constant $C = C(\alpha, \mathcal{T}, T, D, \mathcal{L}, q, \sigma, C_c)$ after using (49) in the last step to bound the expected value of the norm of the adjoint. Observing further after recalling the definition of B from (13) that it holds with Lemma 20

$$\begin{aligned}
 B(t, x, t', x'; \mu_s^N - \mu_0^N) &\leq \frac{1}{N} \sum_{i=1}^N L_k(c_s^i, c_0^i) \left(|c_s^i - c_0^i| + |w_s^{t,i} - w_0^{t,i}| + \|w_s^i - w_0^i\| + |\eta_s^i - \eta_0^i| \right) \\
 &\leq C \frac{1}{N} \sum_{i=1}^N \left(|c_s^i - c_0^i| + |w_s^{t,i} - w_0^{t,i}| + \|w_s^i - w_0^i\| + |\eta_s^i - \eta_0^i| \right)
 \end{aligned} \tag{59}$$

for a constant $C = C(T, D, \sigma, C_c)$ after using (34) in the last step to bound $L_k(c_s^i, c_0^i)$, which is quadratic in c_s^i and c_0^i (see Lemma 20), we are left with

$$\begin{aligned}
 & \mathbb{E} \left\| \left(T_{B(\mu_s^N)} - T_{B(\mu_0^N)} \right) \widehat{u}_{\theta_s}^N \right\|_{L_p(D_T)} \\
 &\leq C \left(\mathbb{E} \left(\frac{1}{N} \sum_{i=1}^N \left(|c_s^i - c_0^i| + |w_s^{t,i} - w_0^{t,i}| + \|w_s^i - w_0^i\| + |\eta_s^i - \eta_0^i| \right) \right)^2 \right)^{1/2} \\
 &\leq C \left(\mathbb{E} \left(\frac{1}{N} \sum_{i=1}^N \left(\frac{1}{N^{1-\beta}} \int_0^\tau \alpha_s \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)} ds \right) \right)^2 \right)^{1/2} \\
 &= \frac{C}{N^{1-\beta}} \left(\mathbb{E} \left(\int_0^\tau \alpha_s \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)} ds \right)^2 \right)^{1/2}
 \end{aligned}$$

$$\begin{aligned}
 &\leq \frac{C}{N^{1-\beta}} \left(\int_0^\tau \alpha_s^2 ds \mathbb{E} \int_0^\tau \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^2 ds \right)^{1/2} \\
 &\leq \frac{C}{N^{1-\beta}}
 \end{aligned} \tag{60}$$

for a constant $C = C(\alpha, \mathcal{T}, T, D, \mathcal{L}, q, \sigma, \mu_0, C_c)$ which is in particular independent of N , where we used (25) and (35) in the second line, Cauchy-Schwarz inequality in the next-to-last step, and the second part of (6) together with the bound (49) on the expected value of the norm of the adjoint to obtain the last inequality. In order to tackle the second term in (57), let us first introduce the random variables $Z^i(t, x, t', x') = k(t, x, t', x'; c_0^i, w_0^{t,i}, w_0^i, \eta_0^i) - \int k(t, x, t', x'; c, w^t, w, \eta) d\mu_0(c, w^t, w, \eta)$, which are independent thanks to the initial independence of the parameters $\theta_0^i = (c_0^i, w_0^{t,i}, w_0^i, \eta_0^i)$ as of Assumption B3, mean-zero, and have finite p th moments due to the c_0^i 's being compactly supported as of Assumption B3(ii) together with the boundedness of σ and σ' from Assumptions B1 and B2 and the boundedness of the domain D as of Assumption A2. After taking the expectation we can estimate with two applications of Cauchy-Schwarz inequality in the second and third step, (49) in the fourth step to bound the expected value of the norm of the adjoint, and two applications of Jensen's inequality in the last step (once for the expectation in the setting of a concave function and once for the time-space integral in the setting of a convex function at the cost of a constant depending only on $T \text{vol}(D)$ and p ; herefore, recall that $p/2 \geq 1$ and $2/p \leq 1$ since $p = d + 2$) that

$$\begin{aligned}
 &\mathbb{E} \left\| \left(T_{B(\mu_0^N)} - T_{B_0} \right) \widehat{u}_{\theta_s}^N \right\|_{L_p(D_T)} \\
 &= \mathbb{E} \left(\int_0^T \int_D \left| \int_0^T \int_D B(t, x, t', x'; \mu_0^N - \mu_0) \widehat{u}_{\theta_s}^N(t', x') dx' dt' \right|^p dx dt \right)^{1/p} \\
 &\leq \mathbb{E} \left[\left(\int_0^T \int_D \left(\int_0^T \int_D (B(t, x, t', x'; \mu_0^N - \mu_0))^2 dx' dt' \right)^{p/2} dx dt \right)^{1/p} \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)} \right] \\
 &\leq \left(\mathbb{E} \left(\int_0^T \int_D \left(\int_0^T \int_D (B(t, x, t', x'; \mu_0^N - \mu_0))^2 dx' dt' \right)^{p/2} dx dt \right)^{2/p} \right)^{1/2} \\
 &\quad \cdot \left(\mathbb{E} \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^2 \right)^{1/2} \\
 &\leq C \left(\mathbb{E} \left(\int_0^T \int_D \left(\int_0^T \int_D (B(t, x, t', x'; \mu_0^N - \mu_0))^2 dx' dt' \right)^{p/2} dx dt \right)^{2/p} \right)^{1/2} \\
 &= \frac{C}{N} \left(\mathbb{E} \left(\int_0^T \int_D \left(\int_0^T \int_D \left(\sum_{i=1}^N Z^i(t, x, t', x') \right)^2 dx' dt' \right)^{p/2} dx dt \right)^{2/p} \right)^{1/2} \\
 &\leq \frac{C}{N} \left(\int_0^T \int_D \int_0^T \int_D \mathbb{E} \left| \sum_{i=1}^N Z^i(t, x, t', x') \right|^p dx' dt' dx dt \right)^{1/p}
 \end{aligned} \tag{61}$$

for a constant $C = C(p, \mathcal{T}, T, D, \mathcal{L}, q, \mu_0)$. We can now employ the Marcinkiewicz-Zygmund inequality with random variables Z^i to obtain

$$\begin{aligned} \mathbb{E} \left\| \left(T_{B(\mu_0^N)} - T_{B_0} \right) \widehat{u}_{\theta_s}^N \right\|_{L_p(D_T)} &\leq \frac{C}{N} \left(\int_0^T \int_D \int_0^T \int_D \mathbb{E} \left(\sum_{i=1}^N |Z^i(t, x, t', x')|^2 \right)^{p/2} dx' dt' dx dt \right)^{1/p} \\ &\leq \frac{C}{N} \left(\int_0^T \int_D \int_0^T \int_D \mathbb{E} \left[N^{p/2-1} \sum_{i=1}^N |Z^i(t, x, t', x')|^p \right] dx' dt' dx dt \right)^{1/p} \\ &\leq \frac{C}{N^{1/2}} \left(\int_0^T \int_D \int_0^T \int_D \mathbb{E} |Z^1(t, x, t', x')|^p dx' dt' dx dt \right)^{1/p} \end{aligned} \quad (62)$$

for some other, potentially larger, constant C after using Hölder's inequality in the next-to-last and the fact that the random variables $Z^i(t, x, t', x')$ are identically distributed in the last step. It further holds after recalling the definition of the random variables $Z^i(t, x, t', x')$ with Lemma 20 that

$$\begin{aligned} \mathbb{E} \left\| \left(T_{B(\mu_0^N)} - T_{B_0} \right) \widehat{u}_{\theta_s}^N \right\|_{L_p(D_T)} &\leq \frac{C}{N^{1/2}} \left(\mathbb{E} \left| \int L_k(c_0^i, c) (|c_0^i - c| + |w_0^{t,i} - w^t| + \|w_0^i - w\| + |\eta_0^i - \eta|) d\mu_0(c, w^t, w, \eta) \right|^p \right)^{1/p} \\ &\leq \frac{C}{N^{1/2}} \end{aligned} \quad (63)$$

for a constant $C = C(p, \mathcal{T}, T, D, \mathcal{L}, q, \sigma, \mu_0)$ which is in particular independent of N . In the last step, recalling that $L_k(c_0^i, c)$ is quadratic in c_0^i and c (see Lemma 20), we firstly used that the initial condition μ_0 has a compactly supported marginal distribution $\mu_{0,c}$ as of Assumptions B3(ii) and that the parameters $c_0^i \sim \mu_{0,c}$, and secondly that the marginal distribution $\mu_{0,(w^t, w, \eta)}$ has bounded p th moments and that the parameters $(w_0^{t,i}, w_0^i, \eta_0^i) \sim \mu_{0,(w^t, w, \eta)}$. Employing (60) and (63) in (57) we eventually obtain the bound

$$\mathbb{E} \left\| \left(T_{B(\mu_s^N)} - T_{B_0} \right) \widehat{u}_{\theta_s}^N \right\|_{L_p(D_T)} \leq \frac{C}{N^{1-\beta}} + \frac{C}{N^{1/2}} \quad (64)$$

for a constant $C = C(p, \alpha, \mathcal{T}, T, D, \mathcal{L}, q, \sigma, \mu_0, C_c)$ which is in particular independent of N .

To estimate the last term of (55) we can directly employ Lemma 24 to obtain

$$\begin{aligned} \|T_{B_0}(\widehat{u}_{\theta_s}^N - \widehat{u}_s^*)\|_{L_p(D_T)} &= \left(\int_0^T \int_D |[T_{B_0}(\widehat{u}_{\theta_s}^N - \widehat{u}_s^*)](t, x)|^p dx dt \right)^{1/p} \\ &\leq C \|\widehat{u}_{\theta_s}^N - \widehat{u}_s^*\|_{L_2(D_T)} \end{aligned} \quad (65)$$

for a constant $C = C(p, T, D, C_\infty^T)$.

Combining the estimates (56), (64) and (65), and plugging them into (55) after taking the expectation we eventually arrive at

$$\mathbb{E} \|g_{\theta_\tau}^N - g_\tau^*\|_{L_p(D_T)} \leq C \left(\frac{1}{N^{\beta-1/2}} + \frac{1}{N^{1-\beta}} + \frac{1}{N^{1/2}} + \int_0^\tau \alpha_s \mathbb{E} \|\widehat{u}_{\theta_s}^N - \widehat{u}_s^*\|_{L_2(D_T)} ds \right). \quad (66)$$

Inserting this now into (51) and (52) and, consecutively, the results into (53), we get

$$\begin{aligned} & \mathbb{E} \left[\left\| \widehat{u}_{\theta_\tau}^N - \widehat{u}_\tau^* \right\|_{L_2([0,T],H^1(D))} + \left\| \widehat{u}_{\theta_\tau}^N - \widehat{u}_\tau^* \right\|_{L_\infty([0,T],L_2(D))} \right] \\ & \leq C \left(\frac{1}{N^{\beta-1/2}} + \frac{1}{N^{1-\beta}} + \frac{1}{N^{1/2}} + \int_0^\tau \alpha_s \mathbb{E} \left\| \widehat{u}_{\theta_s}^N - \widehat{u}_s^* \right\|_{L_2(D_T)} ds \right), \end{aligned} \quad (67)$$

which yields after an application of Grönwall's inequality

$$\begin{aligned} & \mathbb{E} \left[\left\| \widehat{u}_{\theta_\tau}^N - \widehat{u}_\tau^* \right\|_{L_2([0,T],H^1(D))} + \left\| \widehat{u}_{\theta_\tau}^N - \widehat{u}_\tau^* \right\|_{L_\infty([0,T],L_2(D))} \right] \\ & \leq C \left(\frac{1}{N^{\beta-1/2}} + \frac{1}{N^{1-\beta}} + \frac{1}{N^{1/2}} \right) \exp \left(\int_0^\tau \alpha_s ds \right). \end{aligned} \quad (68)$$

Since $\beta \in (1/2, 1)$ and $\mathcal{T} < \infty$, (18b) follows. Utilizing this, (18c) follows from (66), and (18a) eventually follows therefrom with (51) after taking the expected value on both sides of (51). This concludes the proof. \blacksquare

2.6 Main Convergence Result for the NN-PDE

We are now ready to discuss our second main theoretical result, Theorem 3, which is about the convergence of the NN-PDE solution u_τ^* to the target data h , i.e., a global minimizer of the loss \mathcal{J}^* defined in (19), as the training time $\tau \rightarrow \infty$.

A few comments about Theorem 3 are in order. Sufficient conditions for the well-posedness (i.e., uniqueness and existence) of a solution (u^*, \widehat{u}^*) to the PDE system (16)–(17) coupled with the integro-differential equation (15) are provided by Lemma 13 and Remark 14. While they, and in particular the additional Assumption 7, are sufficient, they may not be necessary and the well-posedness of the system could be guaranteed under another set of assumptions, see Remark 8.

Theorem 3 proves the global convergence of the adjoint gradient descent optimization method (5) in the infinite-width NN hidden layer limit as the training time $\tau \rightarrow \infty$. While it is, to the best of our knowledge, a first-of-its-kind convergence result in the setting of semi-linear (and therefore strictly nonlinear) parabolic PDEs, we substantially strengthen beyond that the notion of convergence compared to prior results (Sirignano et al., 2023), which considered the setting of linear PDEs, see Remark 36 for more technical details.

Let us now provide a proof sketch of the statement, which gives an outline of the subsequent Sections 4 to 9 comprising the central steps involved in the proof of Theorem 3.

Proof [Proof sketch of Theorem 3] Properties of the infinite-width NN (Section 4).

The training time derivative of the PDE right-hand side $g_\tau^* = -\int_0^\tau \alpha_s T_{B_0} \widehat{u}_s^* ds$ given in (15) is $\frac{d}{d\tau} g_\tau^* = -\alpha_\tau T_{B_0} \widehat{u}_\tau^*$. Due to the NN kernel operator T_{B_0} being a Hilbert-Schmidt operator as of Remark 21 and Lemma 22, $T_{B_0} \widehat{u}_\tau^* \in L_2(D_T)$ for every τ and $\|T_{B_0} \widehat{u}_\tau^*\|_{L_2(D_T)} \leq C_2^B \|\widehat{u}_\tau^*\|_{L_2(D_T)}$. Leveraging that the NN kernel B_0 is uniformly bounded in L_∞ as a consequence of Assumption 9 on the NN architecture, we further show in Lemma 24 that $T_{B_0} \widehat{u}_\tau^* \in L_\infty(D_T)$ for every τ and $\|T_{B_0} \widehat{u}_\tau^*\|_{L_\infty(D_T)} \leq C_\infty^{T_B} \|\widehat{u}_\tau^*\|_{L_2(D_T)}$. Furthermore, T_{B_0} is positive definite according to Lemma 23 and its eigenfunctions form an orthonormal basis of $L_2(D_T)$ as of Lemma 22.

Step 1: Decay of the loss \mathcal{J}^* (Section 5). With chain rule and by leveraging the adjoint PDE (17), we obtain in Proposition 26 for the training time derivative $\frac{d}{d\tau}\mathcal{J}_\tau^*$ of the loss \mathcal{J}^* defined in (19) with partial integration that

$$\begin{aligned} \frac{d}{d\tau}\mathcal{J}_\tau^* &= \int_0^T \int_D (u_\tau^*(t, x) - h(t, x)) \frac{d}{d\tau} u_\tau^*(t, x) dx dt \\ &= \int_0^T \int_D \widehat{u}_\tau^*(t, x) \frac{d}{d\tau} g_\tau^*(t, x) dx dt = -\alpha_\tau (\widehat{u}_\tau^*, T_{B_0} \widehat{u}_\tau^*)_{L_2} = -\alpha_\tau \mathcal{Q}_\tau^*, \end{aligned} \quad (69)$$

where we used that $\frac{d}{d\tau} g_\tau^* = -\alpha_\tau T_{B_0} \widehat{u}_\tau^*$ in the next-to-last step and the definition (112) of the functional $\mathcal{Q}_\tau^* = (\widehat{u}_\tau^*, T_{B_0} \widehat{u}_\tau^*)_{L_2}$ in the last step.

Step 2: Cycle of stopping times analysis (Section 8). Following the frameworks of (Bertsekas and Tsitsiklis, 2000; Sirignano and Spiliopoulos, 2022) on gradient convergence in gradient methods, we prove in Proposition 32 that

$$\lim_{\tau \rightarrow \infty} \mathcal{Q}_\tau^* = 0. \quad (70)$$

It is immediate to observe that $\liminf_{\tau \rightarrow \infty} \mathcal{Q}_\tau^* = 0$. Namely, if there existed an $\varepsilon > 0$ such that $\mathcal{Q}_\tau^* \geq \varepsilon$ for all $\tau \geq \bar{\tau}$, we would have had by (69) and the fundamental theorem of calculus that $\mathcal{J}_\tau^* = \mathcal{J}_{\bar{\tau}}^* - \int_{\bar{\tau}}^\tau \alpha_s \mathcal{Q}_s^* ds \leq \mathcal{J}_{\bar{\tau}}^* - \varepsilon \int_{\bar{\tau}}^\tau \alpha_s ds \rightarrow -\infty$ as $\tau \rightarrow \infty$ due to condition (6) on the learning rate α_τ . This contradicts the positivity of the loss \mathcal{J}^* . Thus, the case that the functional \mathcal{Q}_τ^* is larger than some ε for all but a finite amount of time cannot occur. However, it remains to outrule the case that the functional \mathcal{Q}_τ^* spikes above ε forever, while being small most of the time. Let us therefore bring $\limsup_{\tau \rightarrow \infty} \mathcal{Q}_\tau^* > 0$ to the contradiction. To this end, assume that there exists an $\varepsilon > 0$ such that $\mathcal{Q}_\tau^* < \varepsilon/2$ for infinitely many τ 's as well as $\mathcal{Q}_\tau^* > \varepsilon$ for infinitely many τ 's. Then there exists an infinite cycle of stopping times

$$0 = \sigma_0 \leq \tau_1 \leq \sigma_1 \leq \tau_2 \leq \sigma_2 \leq \tau_3 \leq \dots, \quad (71)$$

with τ_k and σ_k being defined for $k = 1, 2, \dots$ according to

$$\begin{aligned} \tau_k &= \inf \{ \tau > \sigma_{k-1} : \mathcal{Q}_\tau^* \geq \varepsilon \} \\ \sigma_k &= \sup \left\{ \tau \geq \tau_k : \frac{1}{2} \mathcal{Q}_{\tau_k}^* \leq \mathcal{Q}_s^* \leq 2 \mathcal{Q}_{\tau_k}^* \text{ for all } s \in [\tau_k, \tau] \text{ and } \int_{\tau_k}^\tau \alpha_s ds \leq \frac{\varepsilon}{2L_Q} \right\}, \end{aligned} \quad (72)$$

where $L_Q > 0$ will be defined in *Step 3*. By a telescopic sum argument, we have for sufficiently large \tilde{n} and for all $n \geq \tilde{n}$ that

$$\begin{aligned} \mathcal{J}_{\tau_{n+1}}^* &= \mathcal{J}_{\tau_{\tilde{n}}}^* + \sum_{k=\tilde{n}}^n (\mathcal{J}_{\tau_{k+1}}^* - \mathcal{J}_{\tau_k}^*) \\ &= \mathcal{J}_{\tau_{\tilde{n}}}^* + \sum_{k=\tilde{n}}^n \left[\underbrace{(\mathcal{J}_{\tau_{k+1}}^* - \mathcal{J}_{\sigma_k}^*)}_{\leq 0} + \underbrace{(\mathcal{J}_{\sigma_k}^* - \mathcal{J}_{\tau_k}^*)}_{\leq -(1-\vartheta)\varepsilon^2/(4L_Q)} \right] \rightarrow -\infty \end{aligned} \quad (73)$$

as $n \rightarrow \infty$, which is again a contradiction as the loss \mathcal{J}^* is positive. Thus, (70) holds. The properties under the brackets are derived as follows.

- On the intervals $I_{k+1}^1 = [\sigma_k, \tau_{k+1})$, where $\mathcal{Q}_\tau^* \leq \varepsilon$, i.e., where \mathcal{Q}_τ^* is negligibly small, we just show $\mathcal{J}_{\tau_{k+1}}^* - \mathcal{J}_{\sigma_k}^* \leq 0$ by using (69), the fundamental theorem of calculus and the positive definiteness of T_{B_0} .
- On the intervals $I_k^2 = [\tau_k, \sigma_k)$, on the other hand, where $\mathcal{Q}_{\tau_k}^*/2 \leq \mathcal{Q}_\tau^* \leq 2\mathcal{Q}_{\tau_k}^*$, it holds $\mathcal{J}_{\sigma_k}^* - \mathcal{J}_{\tau_k}^* \leq -(1-\vartheta)\varepsilon^2/(4L_{\mathcal{Q}})$ for any $\vartheta \in (0, 1)$ as we prove in detail in Proposition 32. The intuition behind this bound is that on those intervals $\mathcal{Q}_\tau^* \geq \mathcal{Q}_{\tau_k}^*/2 \geq \varepsilon/2$, i.e., \mathcal{Q}_τ^* is non-negligibly large, while, at the same time, enough training progress is made in the sense that $\int_{\tau_k}^{\sigma_k} \alpha_\tau d\tau \geq (1-\vartheta)\varepsilon/(2L_{\mathcal{Q}})$. The former is by definition of the stopping times. The latter is either ensured by the definition of the stopping time σ_k or guaranteed, as proven in Lemma 33, by a regularity bound (74) for the functional \mathcal{Q}_τ^* in terms of the learning rate α_τ , which we derive in *Step 3*. It allows to lower bound in this case $\int_{\tau_k}^{\sigma_k} \alpha_\tau d\tau$ by the change of the functional $\mathcal{Q}_{\tau_k}^*$, which is (up to an arbitrarily small factor ϑ) at least $\varepsilon/2$ on the intervals I_k^2 .

Step 3: PDE considerations and a regularity bound for the functional \mathcal{Q}_τ^* in terms of the learning rate (Sections 6 and 7). A crucial property of the functional \mathcal{Q}_τ^* in the preceding argument is the regularity bound

$$|\mathcal{Q}_{\tau_2}^* - \mathcal{Q}_{\tau_1}^*| \leq L_{\mathcal{Q}} \int_{\tau_1}^{\tau_2} \alpha_\tau d\tau, \quad (74)$$

which holds for all $0 \leq \tau_1 \leq \tau_2$. To prove (74), we develop a novel approach in Proposition 31. To this end, let us introduce for \mathcal{Q}_τ^* and the coupled PDE system (16)–(17) the second-level adjoint system (114)–(115) with variables $(\widehat{v}_\tau^*, \widehat{w}_\tau^*)$. With the fundamental theorem of calculus and by leveraging those adjoint PDEs, we can represent

$$\mathcal{Q}_{\tau_2}^* - \mathcal{Q}_{\tau_1}^* = \int_{\tau_1}^{\tau_2} \frac{d}{d\tau} \mathcal{Q}_\tau^* d\tau = \int_{\tau_1}^{\tau_2} \left(\frac{d}{d\tau} g_\tau^*, \widehat{v}_\tau^* \right)_{L_2(D_T)} d\tau = - \int_{\tau_1}^{\tau_2} \alpha_\tau (T_{B_0} \widehat{u}_\tau^*, \widehat{v}_\tau^*)_{L_2(D_T)} d\tau \quad (75)$$

and it remains to employ Cauchy-Schwarz inequality and to derive uniform (in the training time τ) L_2 bounds on \widehat{u}_τ^* and \widehat{v}_τ^* . In Lemma 28 we establish such bound for \widehat{u}_τ^* , which is a consequence of an energy estimate and \mathcal{J}_τ^* being monotonically non-increasing as of Proposition 26. In Lemma 30 we prove that also \widehat{v}_τ^* is uniformly (in the training time τ) bounded in L_2 . While the proof is again based on an energy estimate, the technical difficulty arises from the structure of the source term in (114) which is of the form $\widehat{w}_\tau^* + q_{uu}(u_\tau^*) \widehat{u}_\tau^* \widehat{w}_\tau^*$. In order to control this term in L_2 , we establish in Lemma 29 uniform (in the training time τ) estimates on the L_2 and L_∞ norms of the second-level adjoint \widehat{w}_τ^* . That we are in particular able to obtain an L_∞ bound is a consequence of the source term $2T_{B_0} \widehat{u}_\tau^*$ in (115) being in L_∞ by Lemma 24. Cauchy-Schwarz inequality now yields

$$\begin{aligned} |\mathcal{Q}_{\tau_2}^* - \mathcal{Q}_{\tau_1}^*| &\leq \int_{\tau_1}^{\tau_2} \alpha_\tau \|T_{B_0} \widehat{u}_\tau^*\|_{L_2(D_T)} \|\widehat{v}_\tau^*\|_{L_2(D_T)} d\tau \\ &\leq L_{\mathcal{Q}} \int_{\tau_1}^{\tau_2} \alpha_\tau d\tau \end{aligned} \quad (76)$$

with $L_{\mathcal{Q}} = C_2^B C^{\widehat{u}} C^{\widehat{v}}$ after employing Lemmas 22 and 28 and Lemma 30 in the last step.

Step 4: Convergence of the adjoint \widehat{u}_τ^* and the solution u_τ^* (Section 9). Since the functional $\mathcal{Q}_\tau^* = (\widehat{u}_\tau^*, T_{B_0}\widehat{u}_\tau^*)_{L_2(D_T)}$ converges to zero as $\tau \rightarrow \infty$ according to (70) and since the NN kernel operator T_{B_0} is positive definite as of Lemma 23, the adjoint \widehat{u}_τ^* in (17) converges to zero weakly in L_2 as $\tau \rightarrow \infty$ as we prove in Proposition 34. Leveraging the structure of the adjoint PDE (8), we infer therefrom in Proposition 35 that the solution u_τ^* in (16) converges to the target data h weakly in L_2 , which concludes the proof. \blacksquare

Before turning the focus of the manuscript to the numerical experiments and the proof details thereafter, let us draw in the following remark an analogy to the convergence analysis of gradient descent methods in the setting of finite-dimensional optimization problems, which highlights the challenges arising from our infinite-dimensional PDE-based setting.

Remark 15 *To minimize a finite-dimensional objective function or loss $J : \mathbb{R}^d \rightarrow \mathbb{R}$, we run gradient descent $\frac{d}{d\tau}\theta_\tau = -\alpha_\tau \nabla_\theta J(\theta_\tau)$ with learning rate α_τ .*

We can compute with chain rule that $\frac{d}{d\tau}J(\theta_\tau) = \nabla_\theta J(\theta_\tau) \frac{d}{d\tau}\theta_\tau = -\alpha_\tau \|\nabla_\theta J(\theta_\tau)\|^2$, cf. (69), showing the decay of the loss J . Assuming that the loss J is L_J -smooth, i.e., has a L_J -Lipschitz continuous gradient, and that $\nabla_\theta J$ is bounded by $C_{\nabla J}$, we can easily verify that $|\|\nabla_\theta J(\theta_{\tau_2})\|^2 - \|\nabla_\theta J(\theta_{\tau_1})\|^2| \leq 2C_{\nabla J} \|\nabla_\theta J(\theta_{\tau_2}) - \nabla_\theta J(\theta_{\tau_1})\| \leq 2C_{\nabla J} L_J \|\theta_{\tau_2} - \theta_{\tau_1}\| \leq 2C_{\nabla J} L_J \int_{\tau_1}^{\tau_2} \alpha_\tau \|\nabla_\theta J(\theta_\tau)\| d\tau \leq 2C_{\nabla J}^2 L_J \int_{\tau_1}^{\tau_2} \alpha_\tau d\tau$, which matches our regularity bound in terms of the learning rate, cf. (74). Note that in that case, one arrives at the same statement albeit in a considerably more straightforward way compared to how we derived (74) in the infinite dimensional setting studied in our paper. Then, following the same steps as in the cycle of stopping times analysis from above, we can infer that $\lim_{\tau \rightarrow \infty} \|\nabla_\theta J(\theta_\tau)\| = 0$, cf. (70), i.e., that gradient descent converges to a stationary point.

3. Numerical Experiments

Several numerical studies are presented below that illustrate and support our theoretical findings. We consider the following second-order parabolic partial differential equation

$$\begin{aligned} \partial_t u - 0.01\Delta u - q(u) &= g_{\text{target}} && \text{in } [0, T] \times D, \\ u &= 0 && \text{on } [0, T] \times \partial D, \\ u &= 0.2 \sin(4\pi x) \sin(2\pi y) && \text{on } \{0\} \times D, \end{aligned} \tag{77}$$

on the spatial domain $D = [0, 0.5] \times [0, 1]$ and with time horizon $T = 1$. Two scenarios, namely

- (i) the (linear) heat equation, i.e., the case where $q \equiv 0$, and
- (ii) the (nonlinear) Allen-Cahn equation with nonlinear term $q(u) = u^3 - u$

are investigated. In either case, the target source term, which is unknown to the practitioner, is given by $g_{\text{target}}(t, x, y) = 1600x(1 - 2x)y^2(0.2 + 0.6t - y)^2(1 - y)^2$. To learn it from data by running the adjoint gradient descent optimization method (5) with the gradient being computed according to (7), we model g_{target} by an NN g_θ^N with N neurons of the form (3). As hyperparameter we choose $\beta = 2/3$, as activation function σ we choose the tanh, and the

NN parameters are initialized according to $c_0^i \sim \mathcal{U}([-1, 1])$, $w_0^{t,i} \sim \mathcal{N}(0, 1)$, $w_0^i \sim \mathcal{N}(0, \text{Id}_{2 \times 2})$ and $\eta_0^i \sim \mathcal{N}(0, 1)$ for all $i \in \{1, \dots, N\}$, which is in accordance with Assumption 9.

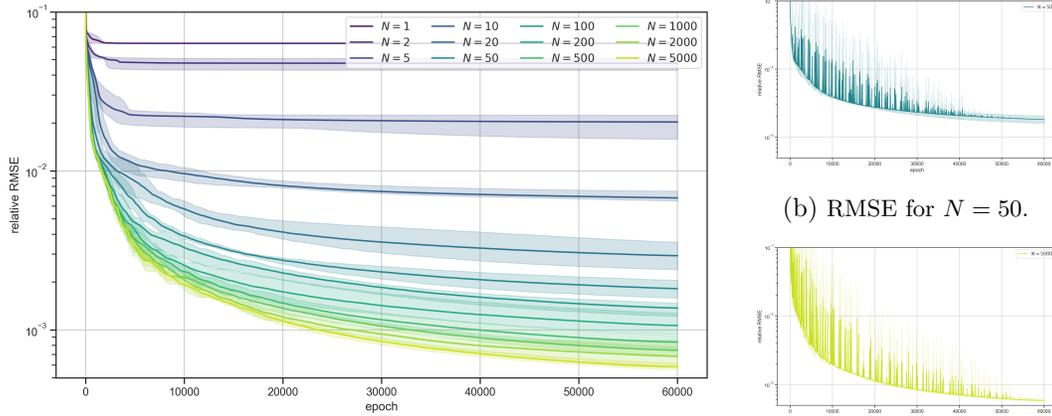
We use the **Adam** optimizer with hyperparameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The learning rate is initially set to $\alpha_0^N = \frac{0.01}{N^{1-2\beta}}$ and decreased adaptively by the **ReduceLRonPlateau** scheduler with factor 0.95 and patience 100, which we gradually reduce during training. For numerical stability, the gradients are clipped using **ZClip** (Kumar et al., 2025), an algorithm for robust gradient norm statistics estimation, which employs z-score-based anomaly detection and leverages exponential moving averages. For its hyperparameters, we choose a smoothing factor of $\alpha = 0.98$ and a z-score threshold of 0.4. Qualitatively similar results can be obtained with other optimizers such as **RMSprop** or **SGD**.

In Figures 1 and 2, respectively, we depict for the linear heat equation, i.e., scenario (i), and the nonlinear Allen-Cahn equation with nonlinear term $q(u) = u^3 - u$, i.e., scenario (ii), the relative root mean square error (RMSE)

$$\begin{aligned} \text{RMSE}(\theta) &= \frac{1}{\|h\|_{L_\infty(D_T)}} \sqrt{2\mathcal{J}_\theta^N} \\ &= \frac{1}{\|h\|_{L_\infty(D_T)}} \sqrt{\int_0^T \int_D (u_\theta^N(t, x) - h(t, x))^2 dx dt} \end{aligned} \tag{78}$$

during training of the NN g_θ^N for a range of different numbers of neurons N . Our results are averaged across five runs with different seeds and visualized as described in the captions.

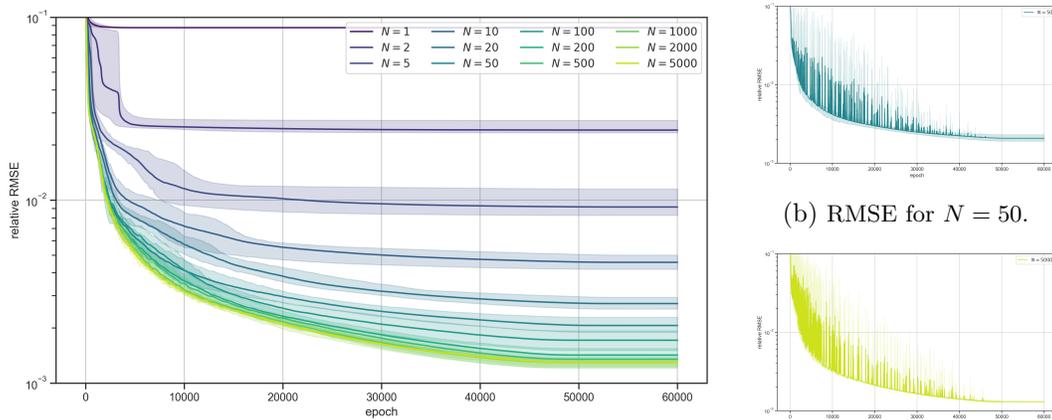
We observe that the NN-PDE solution \hat{u}_θ^N converges to the target data h as the number of neurons N in the NN g_θ^N increases (see Figures 1a and 2a), which confirms our theoretical expectations, cf. Theorems 1 and 3. For a visualization of the target source term g_{target} , the NN g_θ^N as well as the target data h and the NN-PDE solution u_θ^N , we refer the reader to the GitHub repository <https://github.com/KonstantinRiedl/MNPDEs>. In both experimental scenarios (i) and (ii), the loss decreases quickly from the beginning. For small values of N , convergence saturates earlier at higher loss plateaus, while the error continues to decrease to very low loss plateaus even after 60,000 epochs for larger N . The spikes and oscillations present in the trajectory of the raw losses (see Figures 1b, 1c, 2b and 2c) are reminiscent of the edge of stability phenomenon observed in (Cohen et al., 2021, 2022) when using gradient methods to minimize neural network training objectives. They observe for a wide range of machine learning tasks that the training loss does not behave monotonically over short timescales, but consistently decreases over long ones, due to exhibiting a self-stabilization property (Damian et al., 2023). This is in accordance with our observations. These artifacts are intensified by the strong non-convexity of our underlying optimization problem in both the linear and nonlinear PDE scenario. Although this non-convexity is expected to vanish in the infinite-width hidden layer limit in the case of the linear heat equation due to a convexification of the optimization problem, for moderately sized finite-dimensional neural networks, this non-convexity appears to still have an effect on the training. In the nonlinear case, where the problem does not convexify, this may lead to the observed more frequent oscillations in the loss.



(a) Best RMSE for $N \in \{1, 2, 5, 10, 20, 50, 100, 200, 500, 1000, 2000, 5000\}$. (c) RMSE for $N = 5000$.

Figure 1: Decay of the relative $\text{RMSE}(\theta)$ during training of the NN g_θ^N for different numbers of neurons N (colored in blue to green as N increases) in case of the linear heat equation, i.e., scenario (i). In (a), we depict for a range of different numbers of neurons $N \in \{1, \dots, 5000\}$ the relative RMSE of the best model observed during training up until the current epoch. That is, if the RMSE of a model θ_k at epoch k is $\text{RMSE}(\theta_k)$, then the plot displays $\min_{\ell \leq k} \text{RMSE}(\theta_\ell)$ at each epoch k . As we increase the number of neurons N , we observe an improvement in the respective RMSE. In (b) and (c), we plot the instantaneous relative RMSE of the current model at each epoch for $N = 50$ and $N = 5000$ neurons. I.e., these plots display $\text{RMSE}(\theta_k)$ at each epoch k .

In all three plots, we display the mean across five individual training runs with different initializations as a solid line together with the maximal deviation therefrom by a shaded area.



(a) Best RMSE for $N \in \{1, 2, 5, 10, 20, 50, 100, 200, 500, 1000, 2000, 5000\}$. (c) RMSE for $N = 5000$.

Figure 2: We repeat the experiment of Figure 1 for the nonlinear Allen-Cahn equation, i.e., scenario (ii).

4. Neural Network in the PDE Source Term

This section is about the mathematical tools related to the neural network (NN) g_θ^N defined in (3), which constitutes the source term of the PDE (1).

4.1 Properties of the NN Kernel B

The NN kernel B , a.k.a. the neural tangent kernel (NTK) is given as in (13). In Lemmas 16, 17 and 19 we establish some properties of the NN kernel B_0 at initialization that will be useful throughout the manuscript. First, we show that the kernel B_0 is uniformly bounded, which is a direct consequence of Assumptions B1, B2 and B3(ii).

Lemma 16 (L_∞ -boundedness of B_0) *The kernel $B_0 = B(\mu_0)$ defined in (13) is uniformly bounded in L_∞ , i.e., it holds*

$$|B_0(t, x, t', x')| = |B(t, x, t', x'; \mu_0)| \leq C_\infty^B \quad (79)$$

for all $(t, x), (t', x') \in D_T$ for a constant $C_\infty^B = C_\infty^B(T, D, \sigma, \mu_0)$.

Lemma 16 directly implies that the kernel B_0 is bounded in L_2 .

Lemma 17 (L_2 -boundedness of B_0) *The kernel $B_0 = B(\mu_0)$ defined in (13) is bounded in L_2 , i.e., $B_0 \in L_2(D_T \times D_T)$. We abbreviate $C_2^B = \|B_0\|_{L_2(D_T \times D_T)}$.*

Remark 18 *Since the NN kernel B is symmetric, as easily verifiable by noting that $B(t, x, t', x'; \mu) = B(t', x', t, x; \mu)$, and since $B_0 \in L_2(D_T \times D_T)$ as establish in Lemma 17, B_0 is a Hilbert-Schmidt kernel.*

We further show that the kernel B_0 is Lipschitz continuous in the time and space variables.

Lemma 19 (Lipschitz continuity of B_0) *The kernel $B_0 = B(\mu_0)$ defined in (13) is L_B -Lipschitz continuous, i.e., it holds*

$$|B(t^1, x^1, t', x'; \mu_0) - B(t^2, x^2, t', x'; \mu_0)| \leq L_B (|t^1 - t^2| + \|x^1 - x^2\|) \quad (80)$$

for all $(t^1, x^2), (t^2, x^2), (t', x') \in D_T$ for a constant $L_B = L_B(T, D, \sigma, \mu_0)$.

Proof The Lipschitzness and boundedness of σ and σ' via Assumptions B1 and B2 together with Jensen's inequality give

$$\begin{aligned} & |B(t^1, x^1, t', x'; \mu_0) - B(t^2, x^2, t', x'; \mu_0)| \\ & \leq \int C(1 + c^2) ((1 + |w^t|) |t^1 - t^2| + (1 + \|w\|) \|x^1 - x^2\|) d\mu_0(w^t, w, \eta, c) \end{aligned} \quad (81)$$

for a constant $C = C(T, D, \sigma)$. Since μ_0 is such that the marginal distribution $\mu_{0,c}$ is compactly supported and the marginal distribution $\mu_{0,(w^t,w,\eta)}$ has bounded moments according to Assumption B3, the statement follows. \blacksquare

To wrap up this section, let us show that k is (locally) Lipschitz continuous in the NN parameters.

Lemma 20 (Lipschitz continuity of k) *The function k defined in (14) is L_k -Lipschitz continuous, i.e., it holds*

$$\begin{aligned} & |k(t, x, t', x'; c^1, w^{t,1}, w^1, \eta^1) - k(t, x, t', x'; c^2, w^{t,2}, w^2, \eta^2)| \\ & \leq L_k(c^1, c^2) (|c^1 - c^2| + |w^{t,1} - w^{t,2}| + \|w^1 - w^2\| + |\eta^1 - \eta^2|) \end{aligned} \quad (82)$$

for all $(c^1, c^2), (w^{t,1}, w^{t,2}), (w^1, w^2), (\eta^1, \eta^2)$ and for all $(t, x), (t', x') \in D_T$ for a constant $L_k(c^1, c^2) = L_k(T, D, \sigma, c^1, c^2)$ that is quadratic in c^1 and c^2 .

Proof The proof follows directly using the Lipschitzness and boundedness of σ and σ' via Assumptions B1, B2. Details are omitted. \blacksquare

4.2 Properties of the NN Integral Operator T_{B_0}

The NN integral operator T_{B_0} is given as in (12).

Remark 21 *With the kernel B_0 being, as discussed in Remark 18, symmetric and in L_2 , i.e., a Hilbert-Schmidt kernel, the associated operator $T_{B_0} : L_2(\overline{D_T}) \rightarrow L_2(\overline{D_T})$ is a Hilbert-Schmidt integral operator.*

In Lemmas 22 to 25 we establish some properties of the NN integral operator T_{B_0} that will be useful throughout the manuscript. First, we show that the eigenfunctions of the NN integral operator T_{B_0} come with real eigenvalues and form an orthonormal basis of $L_2(D_T)$.

Lemma 22 (Properties of T_{B_0}) *The operator T_{B_0} defined in (12) with $B_0 = B(\mu_0)$ is a self-adjoint compact linear operator with operator norm $\|T_{B_0}\| \leq \|T_{B_0}\|_{\text{HS}} = \|B_0\|_{L_2(D_T \times D_T)} = C_2^B$, where $\|\cdot\|_{\text{HS}}$ denotes the Hilbert-Schmidt norm. Furthermore, the eigenfunctions $\{e_k(t, x)\}_{k=1}^\infty$ of T_{B_0} have real eigenvalues $\{\lambda_k\}_{k=1}^\infty$ and form an orthonormal basis of $L_2(D_T)$.*

Proof Since the kernel B_0 is symmetric and of Hilbert-Schmidt type as verified in Lemma 17, see also Remark 18, the operator T_{B_0} is a self-adjoint, compact linear operator. The Hilbert-Schmidt norm is $\|T_{B_0}\|_{\text{HS}} = \|B_0\|_{L_2(D_T \times D_T)}$ and provides an upper bound to the operator norm. Furthermore, the spectral theorem ensures the existence of an orthonormal basis of $L_2(D_T)$ consisting of eigenvectors of T_{B_0} with real eigenvalue, see (Brezis, 2011, Theorem 6.12). \blacksquare

In fact, as we show next, the eigenvalues of the NN integral operator T_{B_0} can be shown to be strictly positive (Sirignano et al., 2023, pages 27–28).

Lemma 23 (Positive definiteness of T_{B_0}) *The eigenvalues $\{\lambda_k\}_{k=1}^\infty$ of the operator T_{B_0} defined in (12) with $B_0 = B(\mu_0)$ are strictly positive, i.e., it holds $\lambda_k > 0$ for all k . Moreover, it holds $\lambda_k \leq \|B_0\|_{L_2(D_T \times D_T)}$ for all k .*

Proof We first prove that the eigenvalues are strictly positive, i.e., $\lambda_k > 0$ for all k . Using the definition of the kernel $B_0 = B(\mu_0)$ in (13) we can directly compute that for any function

$\hat{u} = \hat{u}(t, x)$ it holds

$$\begin{aligned}
 (\hat{u}, T_{B_0} \hat{u})_{L_2(D_T)} &= \int_0^T \int_D \hat{u}(t, x) \int_0^T \int_D B_0(t, x, t', x') \hat{u}(t', x') dx' dt' dx dt \\
 &= \int \left(\int_0^T \int_D \sigma(w^t t + w^T x + \eta) \hat{u}(t, x) dx dt \right)^2 \\
 &\quad + \left(\int_0^T \int_D c \sigma'(w^t t + w^T x + \eta) t \hat{u}(t, x) dx dt \right)^2 \\
 &\quad + \left\| \int_0^T \int_D c \sigma'(w^t t + w^T x + \eta) x \hat{u}(t, x) dx dt \right\|^2 \\
 &\quad + \left(\int_0^T \int_D c \sigma'(w^t t + w^T x + \eta) \hat{u}(t, x) dx dt \right)^2 d\mu_0(c, w^t, w, \eta) \\
 &\geq \int \left(\int_0^T \int_D \sigma(w^t t + w^T x + \eta) \hat{u}(t, x) dx dt \right)^2 d\mu_{0, (w^t, w, \eta)}(w^t, w, \eta) \\
 &\geq 0,
 \end{aligned} \tag{83}$$

where the inequality in the next-to-last step holds due to the non-negativity of the individual summands in lines 3–5. Since also the summand in line 2 is non-negative, the last inequality holds, which verifies $\lambda_k = (e_k, T_{B_0} e_k)_{L_2(D_T)} \geq 0$.

Let us now show further that $(\hat{u}, T_{B_0} \hat{u})_{L_2(D_T)} = 0$ if and only if $\hat{u} = 0$ everywhere. The “if” direction is immediate. For the “only if” direction, we proceed by contradiction. Consider a function \hat{u} which is not everywhere 0 but suppose that $(\hat{u}, T_{B_0} \hat{u})_{L_2(D_T)} = 0$. The latter implies due to the inequality (83) that

$$\int_0^T \int_D \sigma(w^t t + w^T x + \eta) \hat{u}(t, x) dx dt = 0 \quad \text{for all } w^t \in \mathbb{R}, w \in \mathbb{R}^d, \eta \in \mathbb{R}, \tag{84}$$

since the marginal distribution $\mu_{0, (w^t, w, \eta)}$ assigns positive probability to every set with positive Lebesgue measure as of Assumption B3(iv) and continuity of the integrand w.r.t. the NN parameters w^t, w, η . Since σ is non-constant and bounded as of Assumption B1, it is, according to (Hornik, 1991, Theorem 5), discriminatory in the sense of (Cybenko, 1989; Hornik, 1991). This ensures (note that $\hat{u}(t, x) dx dt$ is a finite signed measure since $\hat{u} \in L_1(D_T)$ by Jensen’s inequality and the fact that $\hat{u} \in L_2(D_T)$ and D_T being bounded as of Assumption A2) that (84) implies that $\hat{u} = 0$ by the definition of σ being discriminatory, see (Hornik, 1991). Since this is a contradiction, $(\hat{u}, T_{B_0} \hat{u})_{L_2(D_T)} > 0$ if \hat{u} is not everywhere 0. In particular, for the eigenfunctions e_k it thus holds $\lambda_k = (e_k, T_{B_0} e_k)_{L_2(D_T)} > 0$, which proves the first part of the statement.

It remains to show that the eigenvalues are bounded from above, i.e., $\lambda_k < \infty$. For this note that by Cauchy-Schwarz inequality it holds

$$\begin{aligned}
 \lambda_k &= (e_k, T_{B_0} e_k)_{L_2(D_T)} \leq \|e_k\|_{L_2(D_T)} \|T_{B_0} e_k\|_{L_2(D_T)} \\
 &\leq \|B_0\|_{L_2(D_T \times D_T)} \|e_k\|_{L_2(D_T)}^2 = \|B_0\|_{L_2(D_T \times D_T)},
 \end{aligned} \tag{85}$$

where the last inequality is due to Lemma 22. This concludes the statement. \blacksquare

Next, let us show that the NN integral operator T_{B_0} maps L_2 to L_∞ .

Lemma 24 (L_∞ -Boundedness of $T_{B_0}\hat{u}$) *Let $\hat{u} \in L_2(D_T)$. Then $T_{B_0}\hat{u}$ defined in (12) with $B_0 = B(\mu_0)$ is uniformly bounded in L_∞ , i.e., it holds*

$$|[T_{B_0}\hat{u}](t, x)| \leq C_\infty^{T_B} \|\hat{u}\|_{L_2(D_T)} \quad (86)$$

for all $(t, x) \in D_T$ for a constant $C_\infty^{T_B} = C_\infty^{T_B}(T, D, C_\infty^B)$.

Proof Using Cauchy-Schwarz inequality and employing Lemma 16 afterwards, we bound

$$\begin{aligned} |[T_{B_0}\hat{u}](t, x)| &\leq \|\hat{u}\|_{L_2(D_T)} \sqrt{\int_0^T \int_D (B_0(t, x, t', x'))^2 dx' dt'} \\ &\leq \|\hat{u}\|_{L_2(D_T)} \sqrt{T \text{vol}(D)} C_\infty^B \end{aligned} \quad (87)$$

for each $(t, x) \in D_T$, which proves the assertion with $C_\infty^{T_B} = \sqrt{T \text{vol}(D)} C_\infty^B$. \blacksquare

To wrap up this section, we furthermore show that $T_{B_0}\hat{u}$ is Lipschitz continuous.

Lemma 25 (Lipschitz continuity of $T_{B_0}\hat{u}$) *Let $\hat{u} \in L_2(D_T)$. Then $T_{B_0}\hat{u}$ defined in (12) with $B_0 = B(\mu_0)$ is L_{T_B} -Lipschitz continuous, i.e., it holds*

$$|[T_{B_0}\hat{u}](t^1, x^1) - [T_{B_0}\hat{u}](t^2, x^2)| \leq L_{T_B} (|t^1 - t^2| + \|x^1 - x^2\|) \quad (88)$$

for all $(t^1, x^2), (t^2, x^2) \in D_T$ for a constant $L_{T_B} = L_{T_B}(T, D, L_B)$.

Proof Using Cauchy-Schwarz inequality and employing Lemma 19 afterwards, we bound

$$\begin{aligned} &|[T_{B_0}\hat{u}](t^1, x^1) - [T_{B_0}\hat{u}](t^2, x^2)| \\ &\leq \|\hat{u}\|_{L_2(D_T)} \sqrt{\int_0^T \int_D (B(t^1, x^1, t', x'; \mu_0) - B(t^2, x^2, t', x'; \mu_0))^2 dx' dt'} \\ &\leq \|\hat{u}\|_{L_2(D_T)} L_B \sqrt{T \text{vol}(D)} (|t^1 - t^2| + \|x^1 - x^2\|) \end{aligned} \quad (89)$$

for $(t^1, x^2), (t^2, x^2) \in D_T$. \blacksquare

5. Decay of the Loss \mathcal{J}_τ^*

Proposition 26 in this section establishes that the loss \mathcal{J}_τ^* defined in (19) is monotonically non-increasing in the training time τ .

We state the result for a training time interval I , which may be either $[0, \mathcal{T}]$ or $[0, \infty)$.

Proposition 26 (Decay of the loss \mathcal{J}_τ^*) *Let $((u_\tau^*, \hat{u}_\tau^*))_{\tau \in I} \in \mathcal{C}(I, \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15)*

in the sense of Lemma 13 and Remark 14 on the training time interval I . Define the loss \mathcal{J}_τ^* as in (19). Then, for the training time derivative $\frac{d}{d\tau}\mathcal{J}_\tau^*$ it holds

$$\begin{aligned} \frac{d}{d\tau}\mathcal{J}_\tau^* &= -\alpha_\tau(\widehat{u}_\tau^*, T_{B_0}\widehat{u}_\tau^*)_{L_2(D_T)} \\ &= -\alpha_\tau \int_0^T \int_D \widehat{u}_\tau^*(t, x) \int_0^T \int_D B(t, x, t', x'; \mu_0) \widehat{u}_\tau^*(t', x') dx' dt' dx dt \end{aligned} \quad (90)$$

for all $\tau \in I$ with the operator T_{B_0} defined in (12) and where the kernel $B_0 = B(\mu_0)$ is as in (13). In particular, we have

$$\frac{d}{d\tau}\mathcal{J}_\tau^* \leq 0 \quad (91)$$

for all $\tau \in I$.

Proof Taking the training time derivative of our loss \mathcal{J}_τ^* , i.e., the derivative w.r.t. the training time τ , we obtain by chain rule and by using that \widehat{u}_τ^* is a weak solution to the adjoint PDE (17) in the sense of Definition 5 with right-hand side $(u_\tau^* - h)$ that

$$\begin{aligned} \frac{d}{d\tau}\mathcal{J}_\tau^* &= \frac{d}{d\tau} \frac{1}{2} \int_0^T \int_D (u_\tau^*(t, x) - h(t, x))^2 dx dt \\ &= \int_0^T \int_D (u_\tau^*(t, x) - h(t, x)) \frac{d}{d\tau} u_\tau^*(t, x) dx dt = \int_0^T \left(u_\tau^*(t, \bullet) - h(t, \bullet), \frac{d}{d\tau} u_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\ &= \int_0^T \left\langle -\partial_t \widehat{u}_\tau^*(t, \bullet), \frac{d}{d\tau} u_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B}^\dagger \left[\widehat{u}_\tau^*(t, \bullet), \frac{d}{d\tau} u_\tau^*(t, \bullet); t \right] \\ &\quad - \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \widehat{u}_\tau^*(t, \bullet), \frac{d}{d\tau} u_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\ &= \int_0^T \left\langle \partial_t \frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{u}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B} \left[\frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{u}_\tau^*(t, \bullet); t \right] \\ &\quad - \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{u}_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\ &= \int_0^T \left(\frac{d}{d\tau} g_\tau^*(t, \bullet), \widehat{u}_\tau^*(t, \bullet) \right)_{L_2(D)} dt = \int_0^T \int_D \left(\frac{d}{d\tau} g_\tau^*(t, x) \right) \widehat{u}_\tau^*(t, x) dx dt. \end{aligned} \quad (92)$$

For the weak solution property in the third line of (92), we note that $\frac{d}{d\tau} u_\tau^*$, the weak solution to the linear parabolic PDE

$$\begin{aligned} \partial_t \frac{d}{d\tau} u_\tau^* + \mathcal{L} \frac{d}{d\tau} u_\tau^* - q_u(u_\tau^*) \frac{d}{d\tau} u_\tau^* &= \frac{d}{d\tau} g_\tau^* && \text{in } D_T, \\ \frac{d}{d\tau} u_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\ \frac{d}{d\tau} u_\tau^* &= \frac{d}{d\tau} f = 0 && \text{on } \{0\} \times D, \end{aligned} \quad (93)$$

which is obtained by taking in (16) the derivative w.r.t. the training time τ , can be used as a test function in the weak formulation of (17), see Definition 5, since $\frac{d}{d\tau}u_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$. Existence, uniqueness and regularity of a weak solution to (93) in a sense analogous to Definition 5 follow from classical results, see, e.g., (Evans, 2010, Chapter 7.1, Theorem 3) and (Evans, 2010, Chapter 7.1, Theorem 4), as $\frac{d}{d\tau}g_\tau^* = -\alpha_\tau T_{B_0}\widehat{u}_\tau^* \in L_2(D_T)$ by Lemma 22.

For the step in the fourth line of (92) we first recall that since $\frac{d}{d\tau}u_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$ and since $\partial_t \widehat{u}_\tau^*(t, \bullet) \in L_2(D)$ for a.e. $t \in [0, T]$, the dual pairing between $H^{-1}(D)$ and $H_0^1(D)$ coincides with the $L_2(D)$ scalar product (Evans, 2010, Chapter 5.9, Theorem 1(iii)). This allows to compute with partial integration, which applies since $\widehat{u}_\tau^*(t, \bullet), \frac{d}{d\tau}u_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$, that

$$\begin{aligned}
 & \int_0^T \left\langle \partial_t \widehat{u}_\tau^*(t, \bullet), \frac{d}{d\tau}u_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} dt \\
 &= \int_0^T \left(\partial_t \widehat{u}_\tau^*(t, \bullet), \frac{d}{d\tau}u_\tau^*(t, \bullet) \right)_{L_2(D)} dt = \int_0^T \int_D (\partial_t \widehat{u}_\tau^*(t, x)) \frac{d}{d\tau}u_\tau^*(t, x) dx dt \\
 &= \int_D \underbrace{\widehat{u}_\tau^*(t, x) \frac{d}{d\tau}u_\tau^*(t, x)}_{=0} \Big|_0^T dx - \int_0^T \int_D \widehat{u}_\tau^*(t, x) \partial_t \frac{d}{d\tau}u_\tau^*(t, x) dx dt \\
 &\quad \text{since } \widehat{u}_\tau^*=0 \text{ on } \{T\} \times D \text{ and} \\
 &\quad \text{since } \frac{d}{d\tau}u_\tau^*=0 \text{ on } \{0\} \times D \\
 &= - \int_0^T \int_D \left(\partial_t \frac{d}{d\tau}u_\tau^*(t, x) \right) \widehat{u}_\tau^*(t, x) dx dt = - \int_0^T \left(\partial_t \frac{d}{d\tau}u_\tau^*(t, \bullet), \widehat{u}_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\
 &= - \int_0^T \left\langle \partial_t \frac{d}{d\tau}u_\tau^*(t, \bullet), \widehat{u}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} dt,
 \end{aligned} \tag{94}$$

where the last step holds again since now $\widehat{u}_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$ and $\partial_t \frac{d}{d\tau}\widehat{u}_\tau^*$ is in $L_2(D)$ for a.e. $t \in [0, T]$ by (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1) with $p = 2$. Those computations are analogous to the ones of Lemma 13 for the PDE (93) due to its with (17) identical structure and since $\frac{d}{d\tau}g_\tau^* \in L_2(D_T)$. Secondly, by definition of the adjoint bilinear form \mathcal{B}^\dagger (see Definition 5) it holds

$$\mathcal{B}^\dagger \left[\widehat{u}_\tau^*(t, \bullet), \frac{d}{d\tau}u_\tau^*(t, \bullet); t \right] = \mathcal{B} \left[\frac{d}{d\tau}u_\tau^*(t, \bullet), \widehat{u}_\tau^*(t, \bullet); t \right] \tag{95}$$

for a.e. $t \in [0, T]$ since $\widehat{u}_\tau^*(t, \bullet), \frac{d}{d\tau}u_\tau^*(t, \bullet) \in H_0^1(D)$.

The penultimate step of (92) holds since $\frac{d}{d\tau}\widehat{u}_\tau^*$ is a weak solution to the PDE (93) and since $\widehat{u}_\tau^*(t, \bullet)$ is a suitable test function as it is in $H_0^1(D)$ for a.e. $t \in [0, T]$.

Now, recalling the definition of the right-hand side g_τ^* from (15) and taking its training time derivative to obtain $\frac{d}{d\tau}g_\tau^* = -\alpha_\tau T_{B_0}\widehat{u}_\tau^*$, as well as recalling the definition of the operator T_{B_0} from (12), we can continue (92) to obtain

$$\begin{aligned}
 \frac{d}{d\tau}\mathcal{J}_\tau^* &= -\alpha_\tau \int_0^T \int_D \widehat{u}_\tau^*(t, x) \int_0^T \int_D B(t, x, t', x'; \mu_0) \widehat{u}_\tau^*(t', x') dx' dt' dx dt \\
 &= -\alpha_\tau (\widehat{u}_\tau^*, T_{B_0}\widehat{u}_\tau^*)_{L_2(D_T)},
 \end{aligned} \tag{96}$$

which concludes the first part of the proof.

The second part now follows immediately thanks to the operator T_{B_0} being positive definite as of Lemma 23. \blacksquare

Following analogous steps we can prove Lemma 10.

Proof [Proof of Lemma 10] Taking the gradient of the loss \mathcal{J}_θ^N w.r.t. the NN parameters θ , we obtain by chain rule and by using that \widehat{u}_θ^N is a weak solution to the adjoint PDE (8) in the sense of Definition 5 that

$$\begin{aligned}
 \nabla_\theta \mathcal{J}_\theta^N &= \nabla_\theta \frac{1}{2} \int_0^T \int_D (u_\theta^N(t, x) - h(t, x))^2 dx dt \\
 &= \int_0^T \int_D (u_\theta^N(t, x) - h(t, x)) \nabla_\theta u_\theta^N(t, x) dx dt = \int_0^T (u_\theta^N(t, \bullet) - h(t, \bullet), \nabla_\theta u_\theta^N(t, \bullet))_{L_2(D)} dt \\
 &= \int_0^T \langle -\partial_t \widehat{u}_\theta^N(t, \bullet), \nabla_\theta u_\theta^N(t, \bullet) \rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B}^\dagger[\widehat{u}_\theta^N(t, \bullet), \nabla_\theta u_\theta^N(t, \bullet); t] \\
 &\quad - (q_u(t, \bullet, u_\theta^N(t, \bullet)) \widehat{u}_\theta^N(t, \bullet), \nabla_\theta u_\theta^N(t, \bullet))_{L_2(D)} dt \\
 &= \int_0^T \langle \partial_t \nabla_\theta u_\theta^N(t, \bullet), \widehat{u}_\theta^N(t, \bullet) \rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B}[\nabla_\theta u_\theta^N(t, \bullet), \widehat{u}_\theta^N(t, \bullet); t] \\
 &\quad - (q_u(t, \bullet, u_\theta^N(t, \bullet)) \nabla_\theta u_\theta^N(t, \bullet), \widehat{u}_\theta^N(t, \bullet))_{L_2(D)} dt \\
 &= \int_0^T (\nabla_\theta g_\theta^N(t, \bullet), \widehat{u}_\theta^N(t, \bullet))_{L_2(D)} dt = \int_0^T \int_D (\nabla_\theta g_\theta^N(t, x)) \widehat{u}_\theta^N(t, x) dx dt.
 \end{aligned} \tag{97}$$

For the weak solution property in the third line of (97), we note that $\nabla_\theta u_\theta^N$, the weak solution to the linear parabolic PDE

$$\begin{aligned}
 \partial_t \nabla_\theta u_\theta^N + \mathcal{L} \nabla_\theta u_\theta^N - q_u(u_\theta^N) \nabla_\theta u_\theta^N &= \nabla_\theta g_\theta^N && \text{in } D_T, \\
 \nabla_\theta u_\theta^N &= 0 && \text{on } [0, T] \times \partial D, \\
 \nabla_\theta u_\theta^N &= \nabla_\theta f = 0 && \text{on } \{0\} \times D,
 \end{aligned} \tag{98}$$

which is obtained by taking in (1) the gradient w.r.t. the NN parameters θ , can be used as a test function in the weak formulation of (8), see Definition 5, since $\nabla_\theta u_\theta^N(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$. Existence, uniqueness and regularity of a weak solution to (98) in a sense analogous to Definition 5 follow from classical results as $\nabla_\theta g_\theta^N \in L_2(D_T)$, see, e.g., (Evans, 2010, Chapter 7.1, Theorem 3) and (Evans, 2010, Chapter 7.1, Theorem 4).

For the step in the fourth line of (97) we use partial integration and the definition of the adjoint bilinear form \mathcal{B}^\dagger with the same argumentation as in the proof of Proposition 26.

The penultimate step of (97) holds since $\nabla_\theta u_\theta^N$ is a weak solution to the PDE (98) and since $\widehat{u}_\theta^N(t, \bullet)$ is a suitable test function as it is in $H_0^1(D)$ for a.e. $t \in [0, T]$. \blacksquare

6. PDE Considerations

Leveraging that the loss \mathcal{J}_τ^* defined in (19) is non-increasing in the training time τ as established in Proposition 26, we provide in Sections 6.1 and 6.2 uniform (in the training time τ) estimates for the norms of the PDE solution u_τ^* to (16) and its adjoint \widehat{u}_τ^* in (17). Those bounds are in particular independent of and thus uniform in the training time τ , depending only on properties of the PDE and the NN initialization at training time $\tau = 0$.

We state the results for a training time interval I , which may be either $[0, \mathcal{T}]$ or $[0, \infty)$.

6.1 Boundedness of the PDE Solution u_τ^* Uniformly in the Training Time

The following uniform (in the training time τ) bound on the L_2 norm of the PDE solution u_τ^* to (16) is an immediate consequence of the loss \mathcal{J}_τ^* being monotonically non-increasing.

Lemma 27 *Let $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in I} \in \mathcal{C}(I, \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) in the sense of Lemma 13 and Remark 14 on the training time interval I . Then the solution u_τ^* is uniformly (in the training time τ) bounded in $L_2(D_T)$ on that interval I , i.e., it holds*

$$\sup_{\tau \in I} \|u_\tau^*\|_{L_2(D_T)} \leq C^u \quad (99)$$

for the constant $C^u = 4\mathcal{J}_0^* + 2\|h\|_{L_2(D_T)}^2$.

Proof For the solution u_τ^* to (16) we can compute with Young's inequality

$$\begin{aligned} \|u_\tau^*\|_{L_2(D_T)}^2 &= \int_0^T \int_D (u_\tau^*(t, x))^2 dx dt = \int_0^T \int_D (u_\tau^*(t, x) - h(t, x) + h(t, x))^2 dx dt \\ &\leq \int_0^T \int_D 2(u_\tau^*(t, x) - h(t, x))^2 + 2(h(t, x))^2 dx dt = 4\mathcal{J}_\tau^* + 2\|h\|_{L_2(D_T)}^2 \\ &\leq 4\mathcal{J}_0^* + 2\|h\|_{L_2(D_T)}^2, \end{aligned} \quad (100)$$

where the last step is a consequence of \mathcal{J}_τ^* being monotonically non-increasing on the training time interval I according to Proposition 26. \blacksquare

6.2 Boundedness of the Adjoint \widehat{u}_τ^* Uniformly in the Training Time

Uniform (in the training time τ) bounds on the $L_2([0, T], H^1(D))$ - and $L_\infty([0, T], L_2(D))$ -norms of the adjoint \widehat{u}_τ^* in (17) are obtained via an energy estimate for the linear parabolic PDE (17) leveraging that the loss \mathcal{J}_τ^* is monotonically non-increasing in the training time τ .

Lemma 28 *Let $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in I} \in \mathcal{C}(I, \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) in the sense of Lemma 13 and Remark 14 on the training time interval I . Then the adjoint \widehat{u}_τ^* in (17) is uniformly (in the training time τ) bounded in $L_2([0, T], H^1(D))$ and $L_\infty([0, T], L_2(D))$ on that interval I , i.e., it holds*

$$\sup_{\tau \in I} \left(\|\widehat{u}_\tau^*\|_{L_2([0, T], H^1(D))} + \|\widehat{u}_\tau^*\|_{L_\infty([0, T], L_2(D))} \right) \leq C^{\widehat{u}} \quad (101)$$

for a constant $C^{\hat{u}} = C^{\hat{u}}(T, \mathcal{L}, \mathcal{J}_0^*)$.

Proof Let us first reverse the adjoint parabolic backward PDE (17) in time to obtain with a time transformation for $\hat{u}_\tau^* = \hat{u}_\tau^*(t, x) = \hat{u}_\tau^*(T - t, x)$ the parabolic forward PDE

$$\begin{aligned} \partial_t \hat{u}_\tau^* + \underline{\mathcal{L}}^* \hat{u}_\tau^* - \underline{q}_u(u_\tau^*(T - \cdot, \cdot)) \hat{u}_\tau^* &= (u_\tau^*(T - \cdot, \cdot) - h(T - \cdot, \cdot)) && \text{in } D_T, \\ \hat{u}_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\ \hat{u}_\tau^* &= 0 && \text{on } \{0\} \times D, \end{aligned} \quad (102)$$

where $\underline{\mathcal{L}}^* = \underline{\mathcal{L}}^*(t, x) = \mathcal{L}^\dagger(T - t, x)$ and $\underline{q} = \underline{q}(t, x, u) = \underline{q}(T - t, x, u)$.

Let us now start by estimating $\|\hat{u}_\tau^*(t, \cdot)\|_{L_2(D)}^2 = \int_D (\hat{u}_\tau^*(t, x))^2 dx$. With chain rule and by using that \hat{u}_τ^* is a weak solution to the time-reversed adjoint PDE (102) we have

$$\begin{aligned} \partial_t \|\hat{u}_\tau^*(t, \cdot)\|_{L_2(D)}^2 &= 2 (\hat{u}_\tau^*(t, \cdot), \partial_t \hat{u}_\tau^*(t, \cdot))_{L_2(D)} = 2 \langle \partial_t \hat{u}_\tau^*(t, \cdot), \hat{u}_\tau^*(t, \cdot) \rangle_{H^{-1}(D), H_0^1(D)} \\ &= -2 \underline{\mathcal{B}}^*[\hat{u}_\tau^*(t, \cdot), \hat{u}_\tau^*(t, \cdot); t] + 2(\underline{q}_u(t, \cdot, u_\tau^*(T - t, \cdot)) \hat{u}_\tau^*(t, \cdot), \hat{u}_\tau^*(t, \cdot))_{L_2(D)} \\ &\quad + 2(u_\tau^*(T - t, \cdot) - h(T - t, \cdot), \hat{u}_\tau^*(t, \cdot))_{L_2(D)}, \end{aligned} \quad (103)$$

where $\underline{\mathcal{B}}^*[\hat{u}, u; t] = \mathcal{B}^\dagger[\hat{u}, u; T - t]$. For the second step recall that since $\hat{u}_\tau^*(t, \cdot) \in H_0^1(D)$ for a.e. $t \in [0, T]$ and since $\partial_t \hat{u}_\tau^*(t, \cdot) \in L_2(D)$ for a.e. $t \in [0, T]$, the dual pairing between $H^{-1}(D)$ and $H_0^1(D)$ coincides with the $L_2(D)$ scalar product (Evans, 2010, Chapter 5.9, Theorem 1(iii)). For the third step, i.e., the weak solution property, note that $\hat{u}_\tau^*(t, \cdot)$ is a valid test function since $\hat{u}_\tau^*(t, \cdot) \in H_0^1(D)$ for a.e. $t \in [0, T]$. To upper bound the right-hand side of (103), we consider each of the three terms separately. For the first term, by using the definition of the bilinear form \mathcal{B} as well as that by Assumption A3 the PDE operator is uniformly parabolic and that by Assumption A4 the coefficients are in L_∞ , we can estimate with Cauchy-Schwarz and Young's inequality

$$\begin{aligned} -\underline{\mathcal{B}}^*[\hat{u}_\tau^*(t, \cdot), \hat{u}_\tau^*(t, \cdot); t] &= -\mathcal{B}^\dagger[\hat{u}_\tau^*(t, \cdot), \hat{u}_\tau^*(t, \cdot); T - t] = -\mathcal{B}[\hat{u}_\tau^*(t, \cdot), \hat{u}_\tau^*(t, \cdot); T - t] \\ &= -\int_U \sum_{i,j=1}^d a^{ij}(T - t, x) \partial_{x_i} \hat{u}_\tau^*(t, x) \partial_{x_j} \hat{u}_\tau^*(t, x) \\ &\quad + \sum_{i=1}^d b^i(T - t, x) \partial_{x_i} \hat{u}_\tau^*(t, x) \hat{u}_\tau^*(t, x) + c(T - t, x) \hat{u}_\tau^*(t, x) \hat{u}_\tau^*(t, x) dx \\ &\leq \int_U -\nu \|\nabla_x \hat{u}_\tau^*(t, x)\|^2 + \frac{\nu}{2} \|\nabla_x \hat{u}_\tau^*(t, x)\|^2 \\ &\quad + \frac{1}{2\nu} \sum_{i=1}^d \|b^i\|_{L_\infty(D_T)} (\hat{u}_\tau^*(t, x))^2 + \|c\|_{L_\infty(D_T)} (\hat{u}_\tau^*(t, x))^2 dx \\ &\leq \frac{\nu}{2} \|\hat{u}_\tau^*(t, \cdot)\|_{H^1(D)}^2 + \left(\frac{1}{2\nu} \sum_{i=1}^d \|b^i\|_{L_\infty(D_T)} + \|c\|_{L_\infty(D_T)} \right) \|\hat{u}_\tau^*(t, \cdot)\|_{L_2(D)}^2, \end{aligned} \quad (104)$$

where for the middle term in the next-to-last step we note that with Young's inequality it holds

$$\begin{aligned} \sum_{i=1}^d b^i(T-t, x) \partial_{x_i} \widehat{u}_\tau^*(t, x) \widehat{u}_\tau^*(t, x) &\leq \sum_{i=1}^d \left(\frac{\nu}{2} (\partial_{x_i} \widehat{u}_\tau^*(t, x))^2 + \frac{1}{2\nu} (b^i(T-t, x) \widehat{u}_\tau^*(t, x))^2 \right) \\ &\leq \frac{\nu}{2} \|\nabla_x \widehat{u}_\tau^*(t, x)\|^2 + \frac{1}{2\nu} \sum_{i=1}^d \|b^i\|_{L^\infty(D_T)} (\widehat{u}_\tau^*(t, x))^2. \end{aligned} \quad (105)$$

For the second term, by using that by Assumption A5 q_u is bounded, we can estimate

$$\begin{aligned} (q_u(t, \bullet, u_\tau^*(T-t, \bullet)) \widehat{u}_\tau^*(t, \bullet), \widehat{u}_\tau^*(t, \bullet))_{L_2(D)} &= (q_u(T-t, \bullet, u_\tau^*(T-t, \bullet)) \widehat{u}_\tau^*(t, \bullet), \widehat{u}_\tau^*(t, \bullet))_{L_2(D)} \\ &\leq c_q \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2. \end{aligned} \quad (106)$$

For the third and last term, by Cauchy-Schwarz and Young's inequality we can derive the upper bound

$$\begin{aligned} (u_\tau^*(T-t, \bullet) - h(T-t, \bullet), \widehat{u}_\tau^*(t, \bullet))_{L_2(D)} &\leq \|u_\tau^*(T-t, \bullet) - h(T-t, \bullet)\|_{L_2(D)} \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)} \\ &\leq \frac{1}{2} \left(\|u_\tau^*(T-t, \bullet) - h(T-t, \bullet)\|_{L_2(D)}^2 + \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 \right). \end{aligned} \quad (107)$$

Combining the bounds established in (104)–(107) and inserting them into (103), we can continue bounding (103) as

$$\begin{aligned} \partial_t \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 + \frac{\nu}{2} |\widehat{u}_\tau^*(t, \bullet)|_{H^1(D)}^2 &\leq \left(\frac{1}{\nu} \sum_{i=1}^d \|b^i\|_{L^\infty(D_T)} + 2\|c\|_{L^\infty(D_T)} \right) \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 + 2c_q \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 \\ &\quad + \left(\|u_\tau^*(T-t, \bullet) - h(T-t, \bullet)\|_{L_2(D)}^2 + \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 \right) \\ &\leq C \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 + \|u_\tau^*(T-t, \bullet) - h(T-t, \bullet)\|_{L_2(D)}^2 \end{aligned} \quad (108)$$

for a constant $C = C(\mathcal{L}, q)$. Defining $\widehat{N}_\tau(t) = \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 + \frac{\nu}{2} \int_0^t |\widehat{u}_\tau^*(s, \bullet)|_{H^1(D)}^2 ds$, (108) translates to

$$\partial_t \widehat{N}_\tau(t) \leq C \widehat{N}_\tau(t) + \|u_\tau^*(T-t, \bullet) - h(T-t, \bullet)\|_{L_2(D)}^2. \quad (109)$$

We can now employ Grönwall's inequality to obtain

$$\widehat{N}_\tau(t) \leq \left(\widehat{N}_\tau(0) + \|u_\tau^* - h\|_{L_2(D_T)}^2 \right) e^{Ct} \leq \left(\widehat{N}_\tau(0) + \|u_\tau^* - h\|_{L_2(D_T)}^2 \right) e^{CT}. \quad (110)$$

Recalling that $\widehat{N}_\tau(0) = 0$ by the initial condition in (102) shows

$$\|\widehat{u}_\tau^*\|_{L_2([0, T], H^1(D))} + \|\widehat{u}_\tau^*\|_{L^\infty([0, T], L_2(D))} \leq 2\mathcal{J}_\tau^* e^{CT} \leq 2\mathcal{J}_0^* e^{CT} \quad (111)$$

where the last step is a consequence of \mathcal{J}_τ^* being monotonically non-increasing on the training time interval I according to Proposition 26, which concludes the proof. \blacksquare

7. The Functional $\mathcal{Q}_\tau^* = (\widehat{u}_\tau^*, T_{B_0}\widehat{u}_\tau^*)_{L_2(D_T)}$

This section is dedicated to proving in Proposition 31 in Section 7.3 a regularity bound for the functional

$$\mathcal{Q}_\tau^* = (\widehat{u}_\tau^*, T_{B_0}\widehat{u}_\tau^*)_{L_2(D_T)} \quad (112)$$

of the form

$$|\mathcal{Q}_{\tau_2}^* - \mathcal{Q}_{\tau_1}^*| \leq L_Q \int_{\tau_1}^{\tau_2} \alpha_\tau d\tau, \quad (113)$$

which holds for all $0 \leq \tau_1 \leq \tau_2$, for a constant $L_Q > 0$ as specified after (130). Here, the operator T_{B_0} is defined in (12) and the kernel $B_0 = B(\mu_0)$ is as in (13).

In order to derive this bound, let us introduce for the functional \mathcal{Q}_τ^* defined in (112) and the PDE system (16)–(17) the second-level adjoint system with variables $(\widehat{v}^*, \widehat{w}^*)$ given by

$$\begin{aligned} -\partial_t \widehat{v}_\tau^* + \mathcal{L}^\dagger \widehat{v}_\tau^* - q_u(u_\tau^*) \widehat{v}_\tau^* &= \widehat{w}_\tau^* + q_{uu}(u_\tau^*) \widehat{u}_\tau^* \widehat{w}_\tau^* && \text{in } D_T, \\ \widehat{v}_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\ \widehat{v}_\tau^* &= 0 && \text{on } \{T\} \times D, \end{aligned} \quad (114)$$

and

$$\begin{aligned} \partial_t \widehat{w}_\tau^* + \mathcal{L} \widehat{w}_\tau^* - q_u(u_\tau^*) \widehat{w}_\tau^* &= 2T_{B_0} \widehat{u}_\tau^* && \text{in } D_T, \\ \widehat{w}_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\ \widehat{w}_\tau^* &= 0 && \text{on } \{0\} \times D. \end{aligned} \quad (115)$$

Before discussing the main statement of this section, Proposition 31, we establish in Sections 7.1 and 7.2 uniform (in the training time τ) estimates for several norms of the second-level adjoints \widehat{w}_τ^* in (114) and \widehat{v}_τ^* in (115), respectively.

7.1 Boundedness of the Second-Level Adjoint \widehat{w}_τ^* Uniformly in the Training Time

We show well-posedness of the second-level adjoint \widehat{w}_τ^* in (115) and derive uniform (in the training time τ) bounds on its $L_\infty(D_T)$ -, $L_2(D_T)$ - and $L_\infty([0, T], L_2(D))$ -norms. The uniformity of the bound in time and space is a consequence of the right-hand side of (115) being in L_∞ as of Lemma 24, while the uniformity in the training time τ follows from the uniformity of the bound on the adjoint \widehat{u}_τ^* as of Lemma 28.

Lemma 29 *Let $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in I} \in \mathcal{C}(I, \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) in the sense of Lemma 13 and Remark 14 on the training time interval I . Then the linear parabolic PDE (115) admits for every $\tau \in I$ a unique weak solution \widehat{w}_τ^* in a sense analogous to Definition 5 with right-hand side $2T_{B_0}\widehat{u}_\tau^*$, which satisfies $\partial_t \widehat{w}_\tau^*(t, \bullet) \in L_2(D)$ for a.e. $t \in [0, T]$.*

In addition, the adjoint \widehat{w}_τ^ in (115) is uniformly (in the training time τ) bounded in $L_\infty(D_T)$ on that interval I , i.e., it holds*

$$\sup_{\tau \in I} \|\widehat{w}_\tau^*\|_{L_\infty(D_T)} \leq C_\infty^{\widehat{w}} \quad (116)$$

for a constant $C_{\infty}^{\widehat{w}} = C_{\infty}^{\widehat{w}}(T, \mathcal{L}, q, C_2^B, C^{\widehat{u}})$. Furthermore, the adjoint \widehat{w}_{τ}^* is uniformly (in the training time τ) bounded in $L_2(D_T)$ and $L_{\infty}([0, T], L_2(D))$ on that interval I , i.e., it holds

$$\sup_{\tau \in I} \left(\|\widehat{w}_{\tau}^*\|_{L_2(D_T)} + \|\widehat{w}_{\tau}^*\|_{L_{\infty}([0, T], L_2(D))} \right) \leq C^{\widehat{w}} \quad (117)$$

for a constant $C^{\widehat{w}} = C^{\widehat{w}}(D, C_{\infty}^{\widehat{u}})$.

Proof Step 1a: *Existence of a unique weak solution \widehat{w}_{τ}^* .* Existence and uniqueness of a weak solution to (115) in a sense analogous to Definition 5 follow from classical results, see, e.g., (Evans, 2010, Chapter 7.1, Theorem 3) and (Evans, 2010, Chapter 7.1, Theorem 4), as $2T_{B_0}\widehat{u}_{\tau}^* \in L_2(D_T)$ according to Lemmas 22 and 28.

The remainder of the statement follows from an application of Morrey's inequality after leveraging (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1) for any $p \geq 2$.

Step 1b: Existence of a unique solution $\widehat{w}_{\tau}^ \in W_p^{1,2}(D_T)$ for any $p \geq 2$.* We first notice that, in the notation of (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1), the coefficients $a_{ij}(t, x) = a^{ij}(t, x)$ of the linear PDE operator of the parabolic PDE (115) are bounded continuous functions in D_T for all $i, j = 1, \dots, d$, while the coefficients $a_i(t, x) = b^i(t, x) - \sum_{j=1}^d \partial_{x_j} a^{ji}(t, x)$ and $a(t, x) = c(t, x) - \sum_{i=1}^d \partial_{x_i} b^i(t, x) - q_u(t, x, u(t, x))$ have finite norms $\|a_i\|_{L_r(D_T)}$ and $\|a\|_{L_s(D_T)}$ for any $r, s > 0$. This is due to the uniform boundedness of the coefficients per Assumptions A4 and A5 combined with the boundedness of the domain per Assumption A2, see the subsequent computations with $T' = 0$ and $\Delta T' = T$. Moreover, since it hold $\|a_i\|_{L_r(D_{T', T'+\Delta T'})} \leq (\|b^i\|_{L_{\infty}(D_T)} + \sum_{j=1}^d \|\partial_{x_j} a^{ji}\|_{L_{\infty}(D_T)}) (\Delta T' \text{vol}(D))^{1/r}$ for all $i = 1, \dots, d$ and $\|a\|_{L_s(D_{T', T'+\Delta T'})} \leq (\|c\|_{L_{\infty}(D_T)} + \sum_{i=1}^d \|\partial_{x_i} b^i\|_{L_{\infty}(D_T)} + c_q) (\Delta T' \text{vol}(D))^{1/s}$, $\|a_i\|_{L_r(D_{T', T'+\Delta T'})}$ and $\|a\|_{L_s(D_{T', T'+\Delta T'})}$ tend to zero as $\Delta T' \rightarrow 0$. Furthermore, ∂D is sufficiently smooth as of Assumption A1. The right-hand side $f = 2T_{B_0}\widehat{u}_{\tau}^* \in L_p(D_T)$ for any $p \geq 2$ due to being uniformly (in the training time τ) bounded in L_{∞} as of Lemma 24, Lemma 28 and the domain D_T being bounded as of Assumption A2, which ensures

$$\begin{aligned} \|2T_{B_0}\widehat{u}_{\tau}^*\|_{L_p(D_T)} &= \left(\int_0^T \int_D |[2T_{B_0}\widehat{u}_{\tau}^*](t, x)|^p dx dt \right)^{1/p} \\ &\leq 2C_{\infty}^{T_{B_0}} \|\widehat{u}_{\tau}^*\|_{L_2(D_T)} (T \text{vol}(D))^{1/p} \leq 2C_{\infty}^{T_{B_0}} C^{\widehat{u}} (T \text{vol}(D))^{1/p}. \end{aligned} \quad (118)$$

Moreover, both the initial and the boundary conditions $\phi = 0 \in W_p^{2-2/p}(D)$ and $\Phi = 0 \in W_p^{1-1/(2p), 2-1/p}(\partial D_T)$ trivially satisfy the compatibility condition $\phi|_{\partial D} = \Phi|_{t=0}$. Thus, (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1) ensures that $\widehat{w}_{\tau}^* \in W_p^{1,2}(D_T)$ obeys the bound

$$\|\widehat{w}_{\tau}^*\|_{W_p^{1,2}(D_T)} \leq C \|2T_{B_0}\widehat{u}_{\tau}^*\|_{L_p(D_T)} \quad (119)$$

for a constant $C = C(T, \mathcal{L}, q)$. With the uniform (in the training time τ) bound (118) at our disposal, $\|\widehat{w}_{\tau}^*\|_{W_p^{1,2}(D_T)}$ can be controlled uniformly (in the training time τ) as

$$\|\widehat{w}_{\tau}^*\|_{W_p^{1,2}(D_T)} \leq 2CC_{\infty}^{T_{B_0}} C^{\widehat{u}} (T \text{vol}(D))^{1/p}. \quad (120)$$

This in particular proves that $\widehat{w}_\tau^* \in W_2^{1,2}(D_T)$ obeying (120) with $p = 2$, concluding the first part of the statement since $\partial_t \widehat{w}_\tau^*(t, \bullet) \in L_2(D)$ has to necessarily hold for a.e. $t \in [0, T]$.

Step 2a: Boundedness of the $L_\infty(D_T)$ norm of \widehat{w}_τ^ .* With the conditions of (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1) being fulfilled for any $p \geq 2$ as we verified before, they are in particular fulfilled for $p > d + 1$. Since we have for such p the continuous embedding $W_p^{1,2}(D_T) \hookrightarrow W_p^{1,1}(D_T) \hookrightarrow L_\infty(\overline{D}_T)$ by Morrey's inequality (Brezis, 2011, Theorem 9.12), we have the first inequality in

$$\|\widehat{w}_\tau^*\|_{L_\infty(D_T)} \leq c(d, p) \|\widehat{w}_\tau^*\|_{W_p^{1,2}(D_T)} \leq 2c(d, p) C C_\infty^{T B_0} C^{\widehat{u}} (T \text{vol}(D))^{1/p}, \quad (121)$$

with the second one being due to (120). As the right-hand side is bounded uniformly (in the training time τ), and since $\widehat{w}_\tau^* \in W_p^{1,2}(D_T)$ has a continuous version (Evans, 2010, Chapter 5.6, Theorem 5), this concludes the second part of the statement.

Step 2b: Boundedness of the $L_2(D_T)$ and $L_\infty([0, T], L_2(D))$ norms of \widehat{w}_τ^ .* The last part of the statement follows since $\|\widehat{w}_\tau^*\|_{L_\infty([0, T], L_2(D))} \leq \sqrt{\text{vol}(D)} \|\widehat{w}_\tau^*\|_{L_\infty(D_T)}$ and $\|\widehat{w}_\tau^*\|_{L_2(D_T)} \leq \sqrt{T \text{vol}(D)} \|\widehat{w}_\tau^*\|_{L_\infty(D_T)}$. \blacksquare

7.2 Boundedness of the Second-Level Adjoint \widehat{v}_τ^* Uniformly in the Training Time

We now show well-posedness of the second-level adjoint \widehat{v}_τ^* in (114) and derive uniform (in the training time τ) bounds on its $L_2([0, T], H^1(D))$ - and $L_\infty([0, T], L_2(D))$ -norms.

Lemma 30 *Let $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in I} \in \mathcal{C}(I, \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) in the sense of Lemma 13 and Remark 14 on the training time interval I . Then the linear parabolic PDE (114) admits for every $\tau \in I$ a unique weak solution \widehat{v}_τ^* in a sense analogous to Definition 5 with right-hand side $\widehat{w}_\tau^* + q_{uu}(u_\tau^*) \widehat{u}_\tau^* \widehat{w}_\tau^*$, which satisfies $\partial_t \widehat{v}_\tau^*(t, \bullet) \in L_2(D)$ for a.e. $t \in [0, T]$.*

In addition, the adjoint \widehat{v}_τ^ in (114) is uniformly (in the training time τ) bounded in $L_2([0, T], H_1(D))$ and $L_\infty([0, T], L_2(D))$ on that interval I , i.e., it holds*

$$\sup_{\tau \in I} \left(\|\widehat{v}_\tau^*\|_{L_2([0, T], H^1(D))} + \|\widehat{v}_\tau^*\|_{L_\infty([0, T], L_2(D))} \right) \leq C^{\widehat{v}} \quad (122)$$

for a constant $C^{\widehat{v}} = C^{\widehat{v}}(T, \mathcal{L}, q, C^{\widehat{u}}, C^{\widehat{w}}, C_\infty^{\widehat{w}})$.

Proof Let us first reverse the adjoint parabolic backward PDE (114) in time to obtain with a time transformation for $\widehat{v}_\tau^* = \underline{\widehat{v}}_\tau^*(t, x) = \widehat{v}_\tau^*(T - t, x)$ the parabolic forward PDE

$$\begin{aligned} \partial_t \widehat{v}_\tau^* + \underline{\mathcal{L}}^* \widehat{v}_\tau^* - \underline{q}_u(u_\tau^*(T - \bullet, \bullet)) \widehat{v}_\tau^* &= \widehat{w}_\tau^*(T - \bullet, \bullet) + q_{uu}(u_\tau^*(T - \bullet, \bullet)) \\ &\quad \cdot \widehat{u}_\tau^*(T - \bullet, \bullet) \widehat{w}_\tau^*(T - \bullet, \bullet) \quad \text{in } D_T, \\ \widehat{v}_\tau^* &= 0 \quad \text{on } [0, T] \times \partial D, \\ \widehat{v}_\tau^* &= 0 \quad \text{on } \{0\} \times D, \end{aligned} \quad (123)$$

where $\underline{\mathcal{L}}^* = \underline{\mathcal{L}}^*(t, x) = \mathcal{L}^\dagger(T - t, x)$ and $\underline{q} = \underline{q}(t, x, u) = \underline{q}(T - t, x, u)$.

Step 1: Existence of a unique solution \widehat{v}_τ^ .* Existence, uniqueness and regularity of a weak solution to (114) in a sense analogous to Definition 5 follow analogously to *Steps 1a* and *b* of the proof of Lemma 29 from classical results, namely (Evans, 2010, Chapter 7.1, Theorem 3) and (Evans, 2010, Chapter 7.1, Theorem 4) as well as (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1) for $p = 2$. Herefore note that $\widehat{w}_\tau^* + q_{uu}(u_\tau^*)\widehat{u}_\tau^*\widehat{w}_\tau^* \in L_2(D_T)$ by combining Lemmas 28 and 29 with Assumption A6.

Step 2: Boundedness of the $L_2([0, T], H^1(D))$ and $L_\infty([0, T], L_2(D))$ norms of \widehat{v}_τ^ .* Let us now estimate $\|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)}^2 = \int_D (\widehat{v}_\tau^*(t, x))^2 dx$. With chain rule and by using that \widehat{v}_τ^* is a weak solution to the time-reversed adjoint PDE (123) we have

$$\begin{aligned} \partial_t \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)}^2 &= 2 \langle \widehat{v}_\tau^*(t, \bullet), \partial_t \widehat{v}_\tau^*(t, \bullet) \rangle_{L_2(D)} = 2 \langle \partial_t \widehat{v}_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet) \rangle_{H^{-1}(D), H_0^1(D)} \\ &= -2\mathcal{B}^*[\widehat{v}_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet); t] + 2(q_u(t, \bullet, u_\tau^*(T-t, \bullet))\widehat{v}_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet))_{L_2(D)} \\ &\quad + 2(\widehat{w}_\tau^*(T-t, \bullet) + q_{uu}(t, \bullet, u_\tau^*(T-t, \bullet))\widehat{u}_\tau^*(T-t, \bullet)\widehat{w}_\tau^*(T-t, \bullet), \widehat{v}_\tau^*(t, \bullet))_{L_2(D)}, \end{aligned} \quad (124)$$

where $\mathcal{B}^*[\widehat{u}, u; t] = \mathcal{B}^\dagger[\widehat{u}, u; T-t]$ and where we recall for the second step that since $\widehat{v}_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$ and since $\partial_t \widehat{v}_\tau^*(t, \bullet) \in L_2(D)$ for a.e. $t \in [0, T]$, the dual pairing between $H^{-1}(D)$ and $H_0^1(D)$ coincides with the $L_2(D)$ scalar product (Evans, 2010, Chapter 5.9, Theorem 1(iii)). For the third step, i.e., the weak solution property, note that $\widehat{v}_\tau^*(t, \bullet)$ is a valid test function since $\widehat{v}_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$. To upper bound the right-hand side of (124), we consider again each of the three terms separately. Analogously to (104) we have for the first term

$$\begin{aligned} &- \mathcal{B}^*[\widehat{v}_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet); t] \\ &\leq -\frac{\nu}{2} \|\widehat{v}_\tau^*(t, \bullet)\|_{H^1(D)}^2 + \left(\frac{1}{2\nu} \sum_{i=1}^d \|b^i\|_{L_\infty(D_T)} + \|c\|_{L_\infty(D_T)} \right) \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)}^2, \end{aligned} \quad (125)$$

where we used the definition of the bilinear form \mathcal{B} as well as that by Assumption A3 the PDE operator is uniformly parabolic and that by Assumption A4 the coefficients are in L_∞ . For the second term we have as in (106) with Assumption A5 that

$$(q_u(t, \bullet, u_\tau^*(T-t, \bullet))\widehat{v}_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet))_{L_2(D)} \leq c_q \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)}^2. \quad (126)$$

For the third and last term, using Assumption A6, by Cauchy-Schwarz, Hölder's and Young's inequality we upper bound

$$\begin{aligned} &(\widehat{w}_\tau^*(T-t, \bullet) + q_{uu}(T-t, \bullet, u_\tau^*(T-t, \bullet))\widehat{u}_\tau^*(T-t, \bullet)\widehat{w}_\tau^*(T-t, \bullet), \widehat{v}_\tau^*(t, \bullet))_{L_2(D)} \\ &\leq \left(\|\widehat{w}_\tau^*(T-t, \bullet)\|_{L_2(D)} + c'_q \|\widehat{u}_\tau^*(T-t, \bullet)\widehat{w}_\tau^*(T-t, \bullet)\|_{L_2(D)} \right) \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)} \\ &\leq \left(\|\widehat{w}_\tau^*(T-t, \bullet)\|_{L_2(D)} + c'_q \|\widehat{u}_\tau^*(T-t, \bullet)\|_{L_2(D)} \|\widehat{w}_\tau^*(T-t, \bullet)\|_{L_\infty(D)} \right) \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)} \\ &\leq \frac{1}{2} \left(\left(\|\widehat{w}_\tau^*(T-t, \bullet)\|_{L_2(D)} + c'_q \|\widehat{u}_\tau^*(T-t, \bullet)\|_{L_2(D)} \|\widehat{w}_\tau^*(T-t, \bullet)\|_{L_\infty(D)} \right)^2 + \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)}^2 \right) \\ &\leq \frac{1}{2} \left(\left(C^{\widehat{w}} + c'_q C^{\widehat{u}} C^{\widehat{w}} \right)^2 + \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)}^2 \right), \end{aligned} \quad (127)$$

where we employed Lemma 28 to bound the L_2 norm of \widehat{u}_τ^* and Lemma 29 to control the L_2 and L_∞ norms of \widehat{w}_τ^* . Combining the bounds established in (125)–(127) and inserting them into (124), we can continue bounding (124) as

$$\begin{aligned} \partial_t \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)}^2 + \frac{\nu}{2} |\widehat{v}_\tau^*(t, \bullet)|_{H^1(D)}^2 &\leq \left(\frac{1}{\nu} \sum_{i=1}^d \|b^i\|_{L_\infty(D_T)} + 2 \|c\|_{L_\infty(D_T)} \right) \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)}^2 \\ &\quad + 2c_q \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)}^2 + \left((C^{\widehat{w}} + c'_q C^{\widehat{u}} C^{\widehat{w}})^2 + \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)}^2 \right) \\ &\leq C \|\widehat{v}_\tau^*(t, \bullet)\|_{L_2(D)}^2 + (C^{\widehat{w}} + c'_q C^{\widehat{u}} C^{\widehat{w}})^2 \end{aligned} \quad (128)$$

for a constant $C = C(\mathcal{L}, q)$. Recalling that $\widehat{v}_\tau^*(0, \bullet) = 0$ by the initial condition in (123), an application of Grönwall's inequality shows

$$\|\widehat{v}_\tau^*\|_{L_2([0, T], H^1(D))} + \|\widehat{v}_\tau^*\|_{L_\infty([0, T], L_2(D))} \leq (C^{\widehat{w}} + c'_q C^{\widehat{u}} C^{\widehat{w}})^2 T e^{CT}, \quad (129)$$

which concludes the proof. \blacksquare

7.3 Regularity Bound for the Functional \mathcal{Q}_τ^* in Terms of the Learning Rate

We now have all technical tools at hand to derive a regularity bound for the functional \mathcal{Q}_τ^* in terms of the learning rate α_τ , which is the main result of this section.

Proposition 31 *Let $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in I} \in \mathcal{C}(I, \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) in the sense of Lemma 13 and Remark 14 on the training time interval I . Then the functional \mathcal{Q}_τ^* as defined in (112) obeys the regularity bound*

$$|\mathcal{Q}_{\tau_2}^* - \mathcal{Q}_{\tau_1}^*| \leq L_{\mathcal{Q}} \int_{\tau_1}^{\tau_2} \alpha_\tau d\tau \quad (130)$$

for all $\tau_1, \tau_2 \in I$ with $0 \leq \tau_1 \leq \tau_2$ for a constant $L_{\mathcal{Q}} = L_{\mathcal{Q}}(C^{\widehat{u}}, C^{\widehat{v}}, C_2^B)$.

Proof By the fundamental theorem of calculus, it holds for all $0 \leq \tau_1 \leq \tau_2$ that

$$\mathcal{Q}_{\tau_2}^* - \mathcal{Q}_{\tau_1}^* = \int_{\tau_1}^{\tau_2} \frac{d}{d\tau} \mathcal{Q}_\tau^* d\tau \quad (131)$$

and it thus remains to compute and estimate $\frac{d}{d\tau} \mathcal{Q}_\tau^*$. Recalling that $\mathcal{Q}_\tau^* = (\widehat{u}_\tau^*, T_{B_0} \widehat{u}_\tau^*)_{L_2(D_T)}$ as defined in (112), we obtain for its training time derivative by chain rule and by using that \widehat{w}_τ^* and \widehat{v}_τ^* are weak solutions (in a sense analogous to Definition 5) to the second-level

adjoint system consisting of the PDEs (115) and (114) that

$$\begin{aligned}
 \frac{d}{d\tau} \mathcal{Q}_\tau^* &= \frac{d}{d\tau} \int_0^T \int_D \widehat{u}_\tau^*(t, x) [T_{B_0} \widehat{u}_\tau^*](t, x) \, dx dt \\
 &= \int_0^T \int_D 2[T_{B_0} \widehat{u}_\tau^*](t, x) \frac{d}{d\tau} \widehat{u}_\tau^*(t, x) \, dx dt = \int_0^T \left(2[T_{B_0} \widehat{u}_\tau^*](t, \bullet), \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\
 &= \int_0^T \left\langle \partial_t \widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B} \left[\widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet); t \right] \\
 &\quad - \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\
 &= \int_0^T \left\langle \partial_t \widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B} \left[\widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet); t \right] \\
 &\quad - \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\
 &\quad + \int_0^T \left\langle -\partial_t \widehat{v}_\tau^*(t, \bullet), \frac{d}{d\tau} u_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B}^\dagger \left[\widehat{v}_\tau^*(t, \bullet), \frac{d}{d\tau} u_\tau^*(t, \bullet); t \right] \\
 &\quad - \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \widehat{v}_\tau^*(t, \bullet), \frac{d}{d\tau} u_\tau^*(t, \bullet) \right)_{L_2(D)} \\
 &\quad - \left(\widehat{w}_\tau^*(t, \bullet) + q_{uu}(t, \bullet, u_\tau^*(t, \bullet)) \widehat{u}_\tau^*(t, \bullet) \widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau} u_\tau^*(t, \bullet) \right)_{L_2(D)} dt.
 \end{aligned} \tag{132}$$

For the weak solution property in the third line of (132), we note that $\frac{d}{d\tau} \widehat{u}_\tau^*$, the weak solution to the linear parabolic PDE

$$\begin{aligned}
 -\partial_t \frac{d}{d\tau} \widehat{u}_\tau^* + \mathcal{L}^\dagger \frac{d}{d\tau} \widehat{u}_\tau^* - q_u(u_\tau^*) \frac{d}{d\tau} \widehat{u}_\tau^* &= \frac{d}{d\tau} u_\tau^* + q_{uu}(u_\tau^*) \widehat{u}_\tau^* \frac{d}{d\tau} u_\tau^* && \text{in } D_T, \\
 \frac{d}{d\tau} \widehat{u}_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\
 \frac{d}{d\tau} \widehat{u}_\tau^* &= 0 && \text{on } \{T\} \times D,
 \end{aligned} \tag{133}$$

which is obtained by taking in (16) the derivative w.r.t. the training time τ , can be used as a test function in the weak formulation of (17), see Definition 5, since $\frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$. Existence and uniqueness of a weak solution to (133) in a sense analogous to Definition 5 follow from classical results, see, e.g., (Evans, 2010, Chapter 7.1, Theorem 3) and (Evans, 2010, Chapter 7.1, Theorem 4), as $\frac{d}{d\tau} u_\tau^* + q_{uu}(u_\tau^*) \widehat{u}_\tau^* \frac{d}{d\tau} u_\tau^* \in L_2(D_T)$. That the right-hand side is indeed in L_2 follows directly after noting that the PDE (93) for $\frac{d}{d\tau} u_\tau^*$ has a structure identical to (115) with right-hand side $\frac{d}{d\tau} g_\tau^* = -\alpha_\tau T_{B_0} \widehat{u}_\tau^*$, i.e., the same up to a constant factor. Following the lines of the proof of Lemma 29 this ensures that $\frac{d}{d\tau} u_\tau^* \in L_\infty(D_T)$. For the weak solution property in the fifth step of (132), we note that $\frac{d}{d\tau} u_\tau^*$, the weak solution to the linear parabolic PDE (93) can be used as a test function in

the weak formulation of (114), since $\frac{d}{d\tau}u_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$, see the discussion after (93).

We now perform partial integration. For this purpose, first recall that since $\frac{d}{d\tau}\widehat{u}_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$ and since $\partial_t\widehat{w}_\tau^*(t, \bullet) \in L_2(D)$ for a.e. $t \in [0, T]$ according to Lemma 29, the dual pairing between $H^{-1}(D)$ and $H_0^1(D)$ coincides with the $L_2(D)$ scalar product (Evans, 2010, Chapter 5.9, Theorem 1(iii)). This allows to compute with partial integration, which applies since $\widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau}\widehat{u}_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$, that

$$\begin{aligned}
 & \int_0^T \left\langle \partial_t \widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} dt \\
 &= \int_0^T \left(\partial_t \widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet) \right)_{L_2(D)} dt = \int_0^T \int_D (\partial_t \widehat{w}_\tau^*(t, x)) \frac{d}{d\tau} \widehat{u}_\tau^*(t, x) dx dt \\
 &= \int_D \underbrace{\widehat{w}_\tau^*(t, x) \frac{d}{d\tau} \widehat{u}_\tau^*(t, x)}_{=0} \Big|_0^T dx - \int_0^T \int_D \widehat{w}_\tau^*(t, x) \partial_t \frac{d}{d\tau} \widehat{u}_\tau^*(t, x) dx dt \\
 &\quad \text{since } \widehat{w}_\tau^* = 0 \text{ on } \{0\} \times D \text{ and} \\
 &\quad \text{since } \frac{d}{d\tau} \widehat{u}_\tau^* = 0 \text{ on } \{T\} \times D \\
 &= - \int_0^T \int_D \left(\partial_t \frac{d}{d\tau} \widehat{u}_\tau^*(t, x) \right) \widehat{w}_\tau^*(t, x) dx dt = - \int_0^T \left(\partial_t \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\
 &= - \int_0^T \left\langle \partial_t \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} dt,
 \end{aligned} \tag{134}$$

where the last step holds again since now $\widehat{w}_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$ and $\partial_t \frac{d}{d\tau} \widehat{u}_\tau^*$ is in $L_2(D)$ for a.e. $t \in [0, T]$, which follows again analogously to Lemma 30 for the PDE (133) due to its with (114) identical structure and $\frac{d}{d\tau}u_\tau^* + q_{uu}(u_\tau^*)\widehat{u}_\tau^* \frac{d}{d\tau}u_\tau^* \in L_2(D_T)$. Similarly, since $\frac{d}{d\tau}u_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$ (see the discussion after (93)) and since $\partial_t \widehat{v}_\tau^*(t, \bullet) \in L_2(D)$ for a.e. $t \in [0, T]$ according to Lemma 30, we may compute analogously to (94) that

$$\int_0^T \left\langle \partial_t \widehat{v}_\tau^*(t, \bullet), \frac{d}{d\tau} u_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} dt = - \int_0^T \left\langle \partial_t \frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} dt \tag{135}$$

since $\widehat{v}_\tau^*(t, \bullet), \frac{d}{d\tau}u_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$. Secondly, by definition of the adjoint bilinear form \mathcal{B}^\dagger (see Definition 5) it hold $\mathcal{B}[\widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau}\widehat{u}_\tau^*(t, \bullet); t] = \mathcal{B}^\dagger[\frac{d}{d\tau}\widehat{u}_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet); t]$ and $\mathcal{B}^\dagger[\widehat{v}_\tau^*(t, \bullet), \frac{d}{d\tau}u_\tau^*(t, \bullet); t] = \mathcal{B}[\frac{d}{d\tau}u_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet); t]$ for a.e. $t \in [0, T]$ since $\widehat{w}_\tau^*(t, \bullet), \frac{d}{d\tau}\widehat{u}_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet), \frac{d}{d\tau}u_\tau^*(t, \bullet) \in H_0^1(D)$. With (134), (135) and the former, we can continue (132) as

$$\begin{aligned}
 \frac{d}{d\tau} \mathcal{Q}_\tau^* &= \int_0^T \left\langle -\partial_t \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B}^\dagger \left[\frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet); t \right] \\
 &\quad - \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet) \right)_{L_2(D)} dt
 \end{aligned}$$

$$\begin{aligned}
 & + \int_0^T \left\langle \partial_t \frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B} \left[\frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet); t \right] \\
 & - \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet) \right)_{L_2(D)} \\
 & - \left(\frac{d}{d\tau} u_\tau^*(t, \bullet) + q_{uu}(t, \bullet, u_\tau^*(t, \bullet)) \widehat{u}_\tau^*(t, \bullet) \frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet) \right)_{L_2(D)} dt.
 \end{aligned} \tag{136}$$

A simple reordering of the terms for later convenience gives

$$\begin{aligned}
 \frac{d}{d\tau} \mathcal{Q}_\tau^* & = \int_0^T \left\langle -\partial_t \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B}^\dagger \left[\frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet); t \right] \\
 & - \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \frac{d}{d\tau} \widehat{u}_\tau^*(t, \bullet) \right. \\
 & \quad \left. + q_{uu}(t, \bullet, u_\tau^*(t, \bullet)) \widehat{u}_\tau^*(t, \bullet) \frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\
 & + \int_0^T \left\langle \partial_t \frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B} \left[\frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet); t \right] \\
 & - \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\
 & - \int_0^T \left(\frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet) \right)_{L_2(D)} dt.
 \end{aligned} \tag{137}$$

Leveraging now in the first and second line of (137) the weak formulation of (133), with test function $\widehat{w}_\tau^*(t, \bullet)$ (suitable due to Lemma 29), and in the third and fourth line the weak formulation of (93), with test function $\widehat{v}_\tau^*(t, \bullet)$ (suitable due to Lemma 30), we arrive at

$$\begin{aligned}
 \frac{d}{d\tau} \mathcal{Q}_\tau^* & = \int_0^T \left(\frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet) \right)_{L_2(D)} dt + \int_0^T \left(\frac{d}{d\tau} g_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\
 & - \int_0^T \left(\frac{d}{d\tau} u_\tau^*(t, \bullet), \widehat{w}_\tau^*(t, \bullet) \right)_{L_2(D)} dt \\
 & = \int_0^T \left(\frac{d}{d\tau} g_\tau^*(t, \bullet), \widehat{v}_\tau^*(t, \bullet) \right)_{L_2(D)} dt = \left(\frac{d}{d\tau} g_\tau^*, \widehat{v}_\tau^* \right)_{L_2(D_T)}.
 \end{aligned} \tag{138}$$

With the expression derived in (138) for $\frac{d}{d\tau} \mathcal{Q}_\tau^*$, we can now obtain a bound on (131). Recalling that $\frac{d}{d\tau} g_\tau^* = -\alpha_\tau T_{B_0} \widehat{u}_\tau^*$ by taking the training time derivative of g_τ as defined in (15) and employing Cauchy-Schwarz inequality yields

$$\begin{aligned}
 |\mathcal{Q}_{\tau_2}^* - \mathcal{Q}_{\tau_1}^*| & = \left| \int_{\tau_1}^{\tau_2} \frac{d}{d\tau} \mathcal{Q}_\tau^* d\tau \right| = \left| \int_{\tau_1}^{\tau_2} \left(\frac{d}{d\tau} g_\tau^*, \widehat{v}_\tau^* \right)_{L_2(D_T)} d\tau \right| = \left| \int_{\tau_1}^{\tau_2} (\alpha_\tau T_{B_0} \widehat{u}_\tau^*, \widehat{v}_\tau^*)_{L_2(D_T)} d\tau \right| \\
 & \leq \int_{\tau_1}^{\tau_2} \alpha_\tau \|T_{B_0}\| \|\widehat{u}_\tau^*\|_{L_2(D_T)} \|\widehat{v}_\tau^*\|_{L_2(D_T)} d\tau \leq C_2^B C^{\widehat{u}} C^{\widehat{v}} \int_{\tau_1}^{\tau_2} \alpha_\tau d\tau,
 \end{aligned} \tag{139}$$

where the last inequality is due to the operator norm of T_{B_0} being bounded by the L_2 -norm of the kernel B as of Lemma 22, $\sup_{\tau \in I} \|\widehat{u}_\tau^*\|_{L_2(D_T)} \leq C^{\widehat{u}}$ according to Lemma 28, and $\sup_{\tau \in I} \|\widehat{v}_\tau^*\|_{L_2(D_T)} \leq C^{\widehat{v}}$ according to Lemma 30. \blacksquare

8. Cycle of Stopping Times Analysis

Exploiting the regularity bound for the functional \mathcal{Q}_τ^* in terms of the learning rate α_τ established in Proposition 31 together with the fact that $\frac{d}{d\tau} \mathcal{J}_\tau^* = -\alpha_\tau \mathcal{Q}_\tau^*$ as shown in Proposition 26, we prove in Proposition 32 of this section by using a cycle of stopping times analysis as conducted in (Bertsekas and Tsitsiklis, 2000; Sirignano and Spiliopoulos, 2022) that this entails $\mathcal{Q}_\tau^* \rightarrow 0$ as $\tau \rightarrow \infty$ provided that the learning rate $(\alpha_\tau)_{\tau \geq 0}$ is decreasing and such that $\int_0^\infty \alpha_\tau d\tau = \infty$.

Proposition 32 *Let $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in [0, \infty)} \in \mathcal{C}([0, \infty), \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) in the sense of Remark 14 on the training time interval $[0, \infty)$. Then,*

$$\lim_{\tau \rightarrow \infty} \mathcal{Q}_\tau^* = 0 \quad (140)$$

and thus also $\lim_{\tau \rightarrow \infty} \frac{d}{d\tau} \mathcal{J}_\tau^* = 0$.

Proof The proof borrows the cycle of stopping times argument from (Bertsekas and Tsitsiklis, 2000, Proposition 1) and (Sirignano and Spiliopoulos, 2022, Theorem 3.1), which crucially depends on the regularity bound (130) for the functional \mathcal{Q}_τ^* in terms of the learning rate as apparent in the proof of Lemma 33.

Setup. Let $\varepsilon > 0$ and set $A = \varepsilon / (2L_Q) > 0$. We define the cycle of stopping times

$$0 = \sigma_0 \leq \tau_1 \leq \sigma_1 \leq \tau_2 \leq \sigma_2 \leq \tau_3 \leq \dots, \quad (141)$$

where τ_k and σ_k are defined for $k = 1, 2, \dots$ according to

$$\begin{aligned} \tau_k &= \inf \{ \tau > \sigma_{k-1} : \mathcal{Q}_\tau^* \geq \varepsilon \} \\ \sigma_k &= \sup \left\{ \tau \geq \tau_k : \frac{1}{2} \mathcal{Q}_{\tau_k}^* \leq \mathcal{Q}_s^* \leq 2 \mathcal{Q}_{\tau_k}^* \text{ for all } s \in [\tau_k, \tau] \text{ and } \int_{\tau_k}^\tau \alpha_s ds \leq A \right\}. \end{aligned} \quad (142)$$

We further introduce the intervals $I_k^1 = [\sigma_{k-1}, \tau_k)$ and $I_k^2 = [\tau_k, \sigma_k)$. It is easy to convince ourselves that by continuity (in the training time τ) it holds $\mathcal{Q}_\tau^* < \varepsilon$ for $\tau \in I_k^1$ as well as $\mathcal{Q}_{\tau_k}^* / 2 \leq \mathcal{Q}_\tau^* \leq 2 \mathcal{Q}_{\tau_k}^*$ for $\tau \in I_k^2$ according to the definitions of the stopping times.

Main Proof. We wish to show that there exists a finite time \mathcal{T}^* such that it holds $\mathcal{Q}_\tau^* \leq \varepsilon$ for all $\tau > \mathcal{T}^*$. Since ε was arbitrary, the statement then follows.

Case 1a: Finitely many τ_k 's, $\tau_K = \infty$. In this case, since there are only finitely many τ_k 's with $\tau_K = \infty$, there indeed exists \mathcal{T}^* such that $\mathcal{Q}_\tau^* \leq \varepsilon$ for all $\tau > \mathcal{T}^*$.

Case 1b: Finitely many τ_k 's, $\sigma_K = \infty$. This case cannot occur, since it would necessitate $\int_{\tau_k}^\infty \alpha_\tau d\tau \leq A$, which contradicts that by assumption on the learning rate $\int_0^\infty \alpha_\tau d\tau = \infty$.

It thus remains to show that the case of infinitely many τ_k 's cannot occur either.

Case 2: Infinitely many τ_k 's. In this case, we have for sufficiently large \tilde{n} and for all $n \geq \tilde{n}$ by a telescopic sum argument that

$$\mathcal{J}_{\tau_{n+1}}^* - \mathcal{J}_{\tau_{\tilde{n}}}^* = \sum_{k=\tilde{n}}^n (\mathcal{J}_{\tau_{k+1}}^* - \mathcal{J}_{\tau_k}^*) = \sum_{k=\tilde{n}}^n \left[(\mathcal{J}_{\tau_{k+1}}^* - \mathcal{J}_{\sigma_k}^*) + (\mathcal{J}_{\sigma_k}^* - \mathcal{J}_{\tau_k}^*) \right], \quad (143)$$

where, in the last line, the respective first term captures the behavior on the intervals $I_{k+1}^1 = [\sigma_k, \tau_{k+1})$, while the second term captures the behavior on the intervals $I_k^2 = [\tau_k, \sigma_k)$.

On the intervals $I_{k+1}^1 = [\sigma_k, \tau_{k+1})$ we have $\mathcal{Q}_\tau^* \leq \varepsilon$ for $\tau \in I_{k+1}^1$. By the fundamental theorem of calculus it holds

$$\mathcal{J}_{\tau_{k+1}}^* - \mathcal{J}_{\sigma_k}^* = \int_{\sigma_k}^{\tau_{k+1}} \frac{d}{d\tau} \mathcal{J}_\tau^* d\tau = - \int_{\sigma_k}^{\tau_{k+1}} \alpha_\tau \mathcal{Q}_\tau^* d\tau \leq 0, \quad (144)$$

where we used Proposition 26 to obtain the second equality and the positivity of \mathcal{Q}_τ^* , a consequence of the positive definiteness of T_{B_0} from Lemma 23, for the last inequality.

On the other hand, on the intervals $I_k^2 = [\tau_k, \sigma_k)$ we have

$$\frac{1}{2} \mathcal{Q}_{\tau_k}^* \leq \mathcal{Q}_\tau^* \leq 2 \mathcal{Q}_{\tau_k}^* \quad \text{and} \quad \int_{\tau_k}^\tau \alpha_s ds \leq A \quad (145)$$

for $\tau \in I_k^2$. Thus, again by the fundamental theorem of calculus and using Proposition 26 in the second equality, it holds

$$\mathcal{J}_{\sigma_k}^* - \mathcal{J}_{\tau_k}^* = \int_{\tau_k}^{\sigma_k} \frac{d}{d\tau} \mathcal{J}_\tau^* d\tau = - \int_{\tau_k}^{\sigma_k} \alpha_\tau \mathcal{Q}_\tau^* d\tau \leq -\frac{1}{2} \mathcal{Q}_{\tau_k}^* \int_{\tau_k}^{\sigma_k} \alpha_\tau d\tau \leq -\frac{(1-\vartheta)}{2} \varepsilon A \quad (146)$$

for any $\vartheta \in (0, 1)$, where the third inequality is due to the property of the interval I_k^2 , while the fourth inequality is firstly since by continuity and by definition of the stopping time τ_k it holds $\mathcal{Q}_{\tau_k}^* \geq \varepsilon$ and secondly since as of Lemma 33 it holds $(1-\vartheta)A \leq \int_{\tau_k}^{\sigma_k} \alpha_\tau d\tau$.

Inserting (144) and (146) into (143) yields

$$\mathcal{J}_{\tau_{n+1}}^* \leq \mathcal{J}_{\tau_{\tilde{n}}}^* - \sum_{k=\tilde{n}}^n \frac{(1-\vartheta)}{2} \varepsilon A = \mathcal{J}_{\tau_{\tilde{n}}}^* - \sum_{k=\tilde{n}}^n \frac{(1-\vartheta)\varepsilon^2}{4L_{\mathcal{Q}}}. \quad (147)$$

Letting $n \rightarrow \infty$, we would obtain that $\mathcal{J}_{\tau_{n+1}}^* \rightarrow -\infty$, which contradicts the fact that $\mathcal{J}_\tau^* \geq 0$ by definition. By excluding that this case can occur, the proof is concluded. \blacksquare

In the proof of Proposition 32 we made use of the following auxiliary result.

Lemma 33 *Let $((u_\tau^*, \hat{u}_\tau^*))_{\tau \in [0, \infty)} \in \mathcal{C}([0, \infty), \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) in the sense of Remark 14 on the training time interval $[0, \infty)$. For given $\varepsilon > 0$, let $A = \varepsilon/(2L_{\mathcal{Q}})$. Then, for k large enough and for $\eta > 0$ small enough (potentially depending on k), one has $\int_{\tau_k}^{\sigma_k + \eta} \alpha_\tau d\tau > A$. Moreover, we also have $(1-\vartheta)A \leq \int_{\tau_k}^{\sigma_k} \alpha_\tau d\tau \leq A$ for any $\vartheta \in (0, 1)$.*

Proof The proof of the first part of the statement proceeds by contradiction. Let us therefore assume that $\int_{\tau_k}^{\sigma_k+\eta} \alpha_\tau d\tau \leq A$. Leveraging the regularity bound for the functional \mathcal{Q}_τ^* in terms of the learning rate established in Proposition 31 with $\tau_1 = \tau_k$ and $\tau_2 = \sigma_k + \eta$, we have

$$\mathcal{Q}_{\sigma_k+\eta}^* - \mathcal{Q}_{\tau_k}^* \leq |\mathcal{Q}_{\sigma_k+\eta}^* - \mathcal{Q}_{\tau_k}^*| \leq L_{\mathcal{Q}} \int_{\tau_k}^{\sigma_k+\eta} \alpha_\tau d\tau \leq L_{\mathcal{Q}} A = \frac{1}{2}\varepsilon \leq \frac{1}{2}\mathcal{Q}_{\tau_k}^*, \quad (148)$$

where we used the contradiction assumption in the third step, the definition of $A = \varepsilon/(2L_{\mathcal{Q}})$ in the fourth step and that by definition of the stopping time τ_k it holds $\mathcal{Q}_{\tau_k}^* \geq \varepsilon$ in the last. The computation (148) implies $\mathcal{Q}_{\sigma_k+\eta}^* \leq \mathcal{Q}_{\tau_k}^* + \mathcal{Q}_{\tau_k}^*/2 \leq 2\mathcal{Q}_{\tau_k}^*$ by simple reordering as well as $\mathcal{Q}_{\tau_k}^* - \mathcal{Q}_{\sigma_k+\eta}^* \leq |\mathcal{Q}_{\sigma_k+\eta}^* - \mathcal{Q}_{\tau_k}^*| \leq \mathcal{Q}_{\tau_k}^*/2$, or rearranged $\mathcal{Q}_{\tau_k}^*/2 \leq \mathcal{Q}_{\sigma_k+\eta}^*$. In summary, $\frac{1}{2}\mathcal{Q}_{\tau_k}^* \leq \mathcal{Q}_{\sigma_k+\eta}^* \leq 2\mathcal{Q}_{\tau_k}^*$. Since the same reasoning holds for any $0 < \tilde{\eta} \leq \eta$, this yields a contradiction, as this would imply that $\sigma_k = \sigma_k + \eta$, contradicting $\eta > 0$. Thus, $\int_{\tau_k}^{\sigma_k+\eta} \alpha_\tau d\tau > A$ holds proving the first part of the statement.

What concerns the second part, since the learning rate α_τ is decreasing in τ , for large enough k and small enough η we can ensure $\int_{\sigma_k}^{\sigma_k+\eta} \alpha_\tau d\tau \leq \vartheta A$. Thus,

$$\int_{\tau_k}^{\sigma_k} \alpha_\tau d\tau = \int_{\tau_k}^{\sigma_k+\eta} \alpha_\tau d\tau - \int_{\sigma_k}^{\sigma_k+\eta} \alpha_\tau d\tau \geq A - \vartheta A = (1 - \vartheta)A. \quad (149)$$

Since by definition $\int_{\tau_k}^{\sigma_k} \alpha_\tau d\tau \leq A$, this concludes the proof. \blacksquare

9. Convergence of the Adjoint \widehat{u}_τ^* and the Solution u_τ^*

Since the functional $\mathcal{Q}_\tau^* = (\widehat{u}_\tau^*, T_{B_0}\widehat{u}_\tau^*)_{L_2(D_T)}$ converges to zero as $\tau \rightarrow \infty$ according to Proposition 32 and since the NN kernel operator T_{B_0} is positive definite as of Lemma 23, we can derive in Proposition 34 in Section 9.1 the weak L_2 convergence of the adjoint \widehat{u}_τ^* in (17) to zero as $\tau \rightarrow \infty$. Noticing that this entails that the left-hand side of the adjoint PDE (8) converges to zero when evaluated against any test function, we infer therefrom in Proposition 35 in Section 9.2 the weak L_2 convergence of the solution u_τ^* in (16) to the target data h by definition of the adjoint PDE. In Section 9.3, we provide a result of independent interest showing that (strong) limit points of the trained NN-PDE solution are global minimizers of the loss \mathcal{J}^* for an even more general class of second-order parabolic NN-PDEs.

9.1 Convergence of the Adjoint \widehat{u}_τ^* as $\tau \rightarrow \infty$

Let us first infer the weak L_2 convergence of the adjoint \widehat{u}_τ^* in (17) to zero.

Proposition 34 *Let $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in [0, \infty)} \in \mathcal{C}([0, \infty), \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) in the sense of Remark 14 on the training time interval $[0, \infty)$. Then,*

$$\widehat{u}_\tau^* \rightharpoonup 0 \text{ in } L_2 \quad \text{as } \tau \rightarrow \infty, \quad (150)$$

i.e., for each test function $\phi \in L_2(D_T)$ it holds $\lim_{\tau \rightarrow \infty} (\widehat{u}_\tau^, \phi)_{L_2(D_T)} = 0$.*

Proof Since the eigenfunctions $\{e_k(t, x)\}_{k=1}^\infty$ of T_{B_0} form an orthonormal basis of $L_2(D_T)$ according to Lemma 22, we have for \widehat{u}_τ^* the expansion $\widehat{u}_\tau^*(t, x) = \sum_{k=1}^\infty c_k(\tau) e_k(t, x)$. Using this, we can express $\mathcal{Q}_\tau^* = (\widehat{u}_\tau^*, T_{B_0} \widehat{u}_\tau^*)_{L_2(D_T)} = \sum_{k=1}^\infty \lambda_k c_k^2(\tau)$, where the last equality holds as T_{B_0} is a continuous operator. Taking the limit $\tau \rightarrow \infty$ and leveraging Proposition 32 in the last step of the following display, this shows

$$\lim_{\tau \rightarrow \infty} \sum_{k=1}^\infty \lambda_k c_k^2(\tau) = \lim_{\tau \rightarrow \infty} (\widehat{u}_\tau^*, T_{B_0} \widehat{u}_\tau^*)_{L_2(D_T)} = \lim_{\tau \rightarrow \infty} \mathcal{Q}_\tau^* = 0. \quad (151)$$

Consequently, for k fixed, it holds $\lim_{\tau \rightarrow \infty} \lambda_k c_k^2(\tau) = 0$. Furthermore, with $\lambda_k > 0$ according to Lemma 23, for k fixed, it also holds $\lim_{\tau \rightarrow \infty} c_k(\tau) = 0$.

Let $\varepsilon > 0$ and let $\phi \in L_2(D_T)$ denote a test function, which we can represent as $\phi(t, x) = \sum_{k=1}^\infty \varphi_k e_k(t, x)$ with $\sum_{k=1}^\infty \varphi_k^2 < \infty$. Thus, there exists $K > 0$ such that $\sum_{k=K+1}^\infty \varphi_k^2 \leq \varepsilon^2 / (2C^{\widehat{u}})^2$.

If $\varphi_k = 0$ for all $k = 1, \dots, K$, then it holds $|\sum_{k=1}^K \varphi_k c_k(\tau)| = 0$. Otherwise, recalling that $\lim_{\tau \rightarrow \infty} c_k(\tau) = 0$ for any fixed k , there exists $\bar{\tau} > 0$ such that we have $|c_k(\tau)| \leq \varepsilon / (2K \max_{k=1, \dots, K} |\varphi_k|)$ (uniformly for $k = 1, \dots, K$) for all $\tau \geq \bar{\tau}$. This shows in particular that $|\sum_{k=1}^K \varphi_k c_k(\tau)| \leq \sum_{k=1}^K |\varphi_k| |c_k(\tau)| \leq \varepsilon/2$. We can now estimate with triangle inequality for all such $\tau \geq \bar{\tau}$ that

$$\begin{aligned} |(\phi, \widehat{u}_\tau^*)_{L_2(D_T)}| &= \left| \sum_{k=1}^\infty \varphi_k c_k(\tau) \right| \leq \left| \sum_{k=1}^K \varphi_k c_k(\tau) \right| + \left| \sum_{k=K+1}^\infty \varphi_k c_k(\tau) \right| \\ &\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2C^{\widehat{u}}} \|\widehat{u}_\tau^*\|_{L_2(D_T)} \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2C^{\widehat{u}}} C^{\widehat{u}} \leq \varepsilon, \end{aligned} \quad (152)$$

where we used Cauchy-Schwarz inequality and the former estimates together with Lemma 28 to obtain the bound on the tail of the series in the inequalities in the second line. Thus $|(\phi, \widehat{u}_\tau^*)_{L_2(D_T)}| \leq \varepsilon$ for all $\tau \geq \bar{\tau}$. Since $\varepsilon > 0$ was arbitrary, this shows that it holds $\lim_{\tau \rightarrow \infty} (\phi, \widehat{u}_\tau^*)_{L_2(D_T)} = 0$ for all test functions $\phi \in L_2(D_T)$, proving the weak convergence of \widehat{u}_τ^* to zero in L_2 as $\tau \rightarrow \infty$. \blacksquare

9.2 Convergence of the Solution u_τ^* as $\tau \rightarrow \infty$

It remains to infer the weak L_2 convergence of the solution u_τ^* to (16) to the target data h .

Proposition 35 *Let $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in [0, \infty)} \in \mathcal{C}([0, \infty), \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) in the sense of Remark 14 on the training time interval $[0, \infty)$. Then,*

$$u_\tau^* \rightharpoonup h \text{ in } L_2 \quad \text{as } \tau \rightarrow \infty, \quad (153)$$

i.e., for each test function $\phi \in L_2(D_T)$ it holds $\lim_{\tau \rightarrow \infty} (u_\tau^ - h, \phi)_{L_2(D_T)} = 0$.*

Proof Let us first show that $\lim_{\tau \rightarrow \infty} (u_\tau^* - h, \widetilde{\phi})_{L_2(D_T)} = 0$ for each test function $\widetilde{\phi} \in C_c^\infty(D_T)$ that vanishes on the boundary. By using that \widehat{u}_τ^* is a weak solution to the adjoint PDE (17)

in the sense of Definition 5 with right-hand side $(u_\tau^* - h)$ we compute for the test function $\tilde{\phi} \in C_c^\infty(D_T)$ that

$$\begin{aligned}
 (u_\tau^* - h, \tilde{\phi})_{L_2(D_T)} &= \int_0^T \left(u_\tau^*(t, \bullet) - h(t, \bullet), \tilde{\phi}(t, \bullet) \right)_{L_2(D)} dt \\
 &= \int_0^T \left\langle -\partial_t \hat{u}_\tau^*(t, \bullet), \tilde{\phi}(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B}^\dagger \left[\hat{u}_\tau^*(t, \bullet), \tilde{\phi}(t, \bullet); t \right] \\
 &\quad - \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \hat{u}_\tau^*(t, \bullet), \tilde{\phi}(t, \bullet) \right)_{L_2(D)} dt \\
 &= \int_0^T \left\langle \partial_t \tilde{\phi}(t, \bullet), \hat{u}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B} \left[\tilde{\phi}(t, \bullet), \hat{u}_\tau^*(t, \bullet); t \right] \\
 &\quad - \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \tilde{\phi}(t, \bullet), \hat{u}_\tau^*(t, \bullet) \right)_{L_2(D)} dt
 \end{aligned} \tag{154}$$

with the last step following analogously to (92), where we justified the individual steps in detail, see (94)–(95). Therefore, note that in the case here, even $\tilde{\phi} \in C_c^\infty(D_T)$.

As a consequence of the convergence $\hat{u}_\tau^* \rightarrow 0$ in L_2 as $\tau \rightarrow \infty$, which we established in Proposition 34, the right-hand side of (154) converges to zero as $\tau \rightarrow \infty$. To be precise, let us discuss each of the three terms. Firstly, since $\partial_t \tilde{\phi} \in C_c^\infty(D_T) \subset L_2(D_T)$ and $\hat{u}_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$, the dual pairing between $H^{-1}(D)$ and $H_0^1(D)$ coincides with the $L_2(D)$ scalar product (Evans, 2010, Chapter 5.9, Theorem 1(iii)) and thus

$$\begin{aligned}
 \int_0^T \left\langle \partial_t \tilde{\phi}(t, \bullet), \hat{u}_\tau^*(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} dt &= \int_0^T \left(\hat{u}_\tau^*(t, \bullet), \partial_t \tilde{\phi}(t, \bullet) \right)_{L_2(D)} dt \\
 &= \left(\hat{u}_\tau^*, \partial_t \tilde{\phi} \right)_{L_2(D_T)},
 \end{aligned} \tag{155}$$

which converges to zero as $\tau \rightarrow \infty$ since $\hat{u}_\tau^* \rightarrow 0$ in L_2 according to Proposition 34 with test function $\partial_t \tilde{\phi} \in C_c^\infty(D_T) \subset L_2(D_T)$. Secondly, by definition of the bilinear form \mathcal{B} in (21) we have

$$\begin{aligned}
 \int_0^T \mathcal{B} \left[\tilde{\phi}(t, \bullet), \hat{u}_\tau^*(t, \bullet); t \right] dt &= \int_0^T \int_U \sum_{i,j=1}^d a^{ij}(t, x) \partial_{x_i} \tilde{\phi}(t, x) \partial_{x_j} \hat{u}_\tau^*(t, x) \\
 &\quad + \sum_{i=1}^d b^i(t, x) \partial_{x_i} \tilde{\phi}(t, x) \hat{u}_\tau^*(t, x) + c(t, x) \tilde{\phi}(t, x) \hat{u}_\tau^*(t, x) dx dt \\
 &= \int_0^T \int_U - \sum_{i,j=1}^d \partial_{x_j} \left(a^{ij}(t, x) \partial_{x_i} \tilde{\phi}(t, x) \right) \hat{u}_\tau^*(t, x) \\
 &\quad + \sum_{i=1}^d b^i(t, x) \partial_{x_i} \tilde{\phi}(t, x) \hat{u}_\tau^*(t, x) + c(t, x) \tilde{\phi}(t, x) \hat{u}_\tau^*(t, x) dx dt \\
 &= \int_0^T \int_U - \sum_{i,j=1}^d a^{ij}(t, x) \partial_{x_i x_j}^2 \tilde{\phi}(t, x) \hat{u}_\tau^*(t, x) \\
 &\quad - \sum_{i,j=1}^d \partial_{x_j} a^{ij}(t, x) \partial_{x_i} \tilde{\phi}(t, x) \hat{u}_\tau^*(t, x)
 \end{aligned}$$

$$\begin{aligned}
 & + \sum_{i=1}^d b^i(t, x) \partial_{x_i} \tilde{\phi}(t, x) \widehat{u}_\tau^*(t, x) + c(t, x) \tilde{\phi}(t, x) \widehat{u}_\tau^*(t, x) dx dt \\
 & = \left(\widehat{u}_\tau^*, - \sum_{i,j=1}^d a^{ij} \partial_{x_i x_j}^2 \tilde{\phi} - \sum_{i,j=1}^d \partial_{x_j} a^{ij} \partial_{x_i} \tilde{\phi} + \sum_{i=1}^d b^i \partial_{x_i} \tilde{\phi} + c \tilde{\phi} \right)_{L_2(D_T)}, \quad (156)
 \end{aligned}$$

where the second step is just partial integration with all boundary terms vanishing since also $\partial_{x_i} \tilde{\phi} \in C_c^\infty(D_T)$. Since the coefficients $a^{ij}, \partial_{x_j} a^{ij}, b^i, c \in L_\infty(D_T)$ as of Assumption A4, the test function in the scalar-product in the last line of (156) is in $L_2(D_T)$ and thus the right-hand side of (156) converges to zero as $\tau \rightarrow \infty$ since $\widehat{u}_\tau^* \rightharpoonup 0$ in L_2 according to Proposition 34. Thirdly and lastly, since with $\tilde{\phi} \in C_c^\infty(D_T) \subset L_2(D_T)$ and q_u being uniformly bounded as of Assumption A5, also $q_u(\bullet, \bullet, u_\tau^*) \tilde{\phi} \in L_2(D_T)$,

$$\int_0^T \left(q_u(t, \bullet, u_\tau^*(t, \bullet)) \tilde{\phi}(t, \bullet), \widehat{u}_\tau^*(t, \bullet) \right)_{L_2(D)} dt = (\widehat{u}_\tau^*, q_u(u_\tau^*) \tilde{\phi})_{L_2(D_T)} \quad (157)$$

converges to zero as $\tau \rightarrow \infty$ since $\widehat{u}_\tau^* \rightharpoonup 0$ in L_2 according to Proposition 34. With this we have shown that $(u_\tau^* - h, \tilde{\phi})_{L_2(D_T)} \rightarrow 0$ for all $\tilde{\phi} \in C_c^\infty(D_T)$.

Let now $\phi \in L_2(D_T)$ and $\varepsilon > 0$. Since $C_c^\infty(D_T)$ is dense in $L_2(D_T)$ (Brezis, 2011, Corollary 4.23), there exists $\tilde{\phi} \in C_c^\infty(D_T)$ such that $\|\phi - \tilde{\phi}\|_{L_2(D_T)} \leq \varepsilon / \sqrt{2\mathcal{J}_0^*}$. We can thus estimate

$$\begin{aligned}
 |(u_\tau^* - h, \phi)_{L_2(D_T)}| & \leq |(u_\tau^* - h, \phi - \tilde{\phi})_{L_2(D_T)}| + |(u_\tau^* - h, \tilde{\phi})_{L_2(D_T)}| \\
 & \leq \|u_\tau^* - h\|_{L_2(D_T)} \|\phi - \tilde{\phi}\|_{L_2(D_T)} + |(u_\tau^* - h, \tilde{\phi})_{L_2(D_T)}| \\
 & = \sqrt{2\mathcal{J}_\tau^*} \|\phi - \tilde{\phi}\|_{L_2(D_T)} + |(u_\tau^* - h, \tilde{\phi})_{L_2(D_T)}| \quad (158) \\
 & \leq \sqrt{2\mathcal{J}_0^*} \|\phi - \tilde{\phi}\|_{L_2(D_T)} + |(u_\tau^* - h, \tilde{\phi})_{L_2(D_T)}| \\
 & \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon
 \end{aligned}$$

for sufficiently large τ . In the next-to-last step we used that by Proposition 26 the loss \mathcal{J}_τ^* is non-increasing. The last step holds since $|(u_\tau^* - h, \tilde{\phi})_{L_2(D_T)}| \rightarrow 0$ for $\tilde{\phi} \in C_c^\infty(D_T)$, thus $|(u_\tau^* - h, \tilde{\phi})_{L_2(D_T)}| \leq \varepsilon/2$ for sufficiently large τ . Consequently, $|(u_\tau^* - h, \phi)_{L_2(D_T)}| \rightarrow 0$ for all $\phi \in L_2(D_T)$, which concludes the proof. \blacksquare

Before closing this section, let us compare Proposition 35 to prior work to indicate that we substantially strengthen the notion of convergence for a significantly wider class of PDEs and a more general loss.

Remark 36 *The weak convergences $\widehat{u}_\tau^* \rightharpoonup 0$ in L_2 and $u_\tau^* \rightharpoonup h$ in L_2 as $\tau \rightarrow \infty$ established in Propositions 34 and 35, respectively, significantly improve prior work (Sirignano et al., 2023), where only convergence of the time averages has been established, cf. (Sirignano et al., 2023, Theorem 9.3). In the elliptic linear PDE setting, the authors of (Sirignano et al., 2023) prove $\lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_0^\tau (\phi, \widehat{u}_s^*)_{L_2}^2 ds = 0$ for all $\phi \in L_2$ and $\lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_0^\tau (\psi, \widehat{u}_s^* - h)_{L_2}^2 ds = 0$ for all $\psi \in \mathcal{A} := \{\psi \in H_0^1 : \mathcal{L}\psi \in L_2\} \subset L_2$.*

To see that Propositions 34 and 35 are stronger, simply observe that the time average $\frac{1}{\tau} \int_0^\tau f_s^2 ds \rightarrow 0$ might converge while $f_\tau \not\rightarrow 0$. (f_τ corresponds here to either $(\phi, \widehat{u}_\tau^*)_{L_2}^2$ or $(\psi, \widehat{u}_\tau^* - h)_{L_2}^2$.) A straightforward smooth example is given by

$$f_\tau = \begin{cases} \exp\left(1 - \frac{1}{1 - (\tau - 2^\ell)^2}\right), & \text{for } \tau \in [2^\ell - 1, 2^\ell + 1] \text{ for } \ell = 1, 2, \dots, \\ 0, & \text{else.} \end{cases} \quad (159)$$

The function $\tau \mapsto f_\tau$ concatenates infinitely many bump functions centered around 2^ℓ , $\ell = 1, 2, \dots$, with width 2 and maximal height 1. Therefore, clearly, $u_\tau \not\rightarrow 0$. However, since there are $\lceil \log_2(\tau) \rceil$ such bumps before time τ ,

$$\frac{1}{\tau} \int_0^\tau f_s^2 ds \leq \frac{1}{\tau} \sum_{\ell=1}^{\lceil \log_2(\tau) \rceil} 2 \leq \frac{2}{\tau} \log_2(\tau) \rightarrow 0 \quad \text{as } \tau \rightarrow \infty. \quad (160)$$

Conversely, it is immediate to see that $f_\tau \rightarrow 0$ implies $\frac{1}{\tau} \int_0^\tau f_s^2 ds \rightarrow 0$.

Secondly, unlike (Sirignano et al., 2023), where the considered loss is given by

$$\widetilde{\mathcal{J}}_\tau^* = \frac{1}{2} \sum_{\ell=1}^L (u_\tau^* - h, m_\ell)_{L_2}^2 \quad (161)$$

for given functions $\{m_\ell\}_{\ell=1}^L$, we consider the stronger loss $\mathcal{J}_\tau^* = \|u_\tau^* - h\|_{L_2}^2$ as in (19).

We therefore generalize in this paper not just the class of considered PDEs substantially by allowing for nonlinear PDEs, but significantly improve the notion of convergence.

9.3 Limit Points of the Trained NN-PDE Solution are Global Minimizers of the Loss \mathcal{J}^*

To conclude the theoretical contributions of this work, let us provide a result about the limit points of the trained NN-PDE solution u_τ^* , which holds for the even more general class of fully nonlinear second-order parabolic NN-PDEs

$$\begin{aligned} \partial_t u_\tau^* + \mathcal{L}u_\tau^* - q(u_\tau^*, \nabla_x u_\tau^*, \mathbf{H}_{xx} u_\tau^*) &= g_\tau^* && \text{in } D_T, \\ u_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\ u_\tau^* &= f && \text{on } \{0\} \times D, \end{aligned} \quad (162)$$

with associated adjoint PDE

$$\begin{aligned} -\partial_t \widehat{u}_\tau^* + \mathcal{L}^\dagger \widehat{u}_\tau^* - q_u(u_\tau^*, \nabla_x u_\tau^*, \mathbf{H}_{xx} u_\tau^*) \widehat{u}_\tau^* \\ + \sum_{i=1}^d \partial_{x_i} (q_{p_i}(u_\tau^*, \nabla_x u_\tau^*, \mathbf{H}_{xx} u_\tau^*) \widehat{u}_\tau^*) \\ - \sum_{i,j=1}^d \partial_{x_i x_j}^2 (q_{\mathbf{H}_{ij}}(u_\tau^*, \nabla_x u_\tau^*, \mathbf{H}_{xx} u_\tau^*) \widehat{u}_\tau^*) &= (u_\tau^* - h) && \text{in } D_T, \\ \widehat{u}_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\ \widehat{u}_\tau^* &= 0 && \text{on } \{T\} \times D, \end{aligned} \quad (163)$$

and coupled with the integro-differential equation (15) for g_τ^* .

We show that any (strong) limit point of the solution of the trained NN-PDE, when using the adjoint gradient descent optimization method (5) with the gradient being computed according to (7), is a global minimizer of the loss \mathcal{J}^* .

Theorem 37 Let $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in [0, \infty)} \in \mathcal{C}([0, \infty), \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the more general PDE system (162)–(163) coupled with the integro-differential equation (15) in a sense analogous to Lemma 13 and Remark 14 on the training time interval $[0, \infty)$. Assume that $(u_\tau^*, \widehat{u}_\tau^*)$ converges to some $(u_\infty^*, \widehat{u}_\infty^*)$ in $L_2(D_T)$ as $\tau \rightarrow \infty$. Then $\widehat{u}_\infty^* \equiv 0$ a.e. in $L_2(D_T)$ and

$$u_\infty^* \equiv h \text{ a.e. in } L_2(D_T), \quad (164)$$

i.e., u_∞^* is a global minimizer of \mathcal{J}^* .

Proof Leveraging the adjoint PDE (163), we can derive analogously to Proposition 26 that $\frac{d}{d\tau} \mathcal{J}_\tau^* = -\alpha_\tau (\widehat{u}_\tau^*, T_{B_0} \widehat{u}_\tau^*)_{L_2} = -\alpha_\tau \mathcal{Q}_\tau^*$ for all $\tau \in [0, \infty)$. Moreover, by following the computations of Lemma 28, we can derive a uniform (in the training time τ) estimate of the form (101) for the adjoint \widehat{u}_τ^* .

Step 1: $\widehat{u}_\infty^* \equiv 0$ a.e. in $L_2(D_T)$. Since the eigenfunctions $\{e_k(t, x)\}_{k=1}^\infty$ of T_{B_0} form an orthonormal basis of $L_2(D_T)$ according to Lemma 22, \widehat{u}_∞^* has the expansion $\widehat{u}_\infty^*(t, x) = \sum_{k=1}^\infty c_k e_k(t, x)$. We now proceed by contradiction and suppose that \widehat{u}_∞^* is not 0 a.e. in $L_2(D_T)$. Then there exists at least one $\bar{k} \in \mathbb{N}$ with $c_{\bar{k}} \neq 0$. Using this, we can lower bound $\mathcal{Q}_\infty^* = (\widehat{u}_\infty^*, T_{B_0} \widehat{u}_\infty^*)_{L_2(D_T)} = \sum_{k=1}^\infty \lambda_k c_k^2 \geq \lambda_{\bar{k}} c_{\bar{k}}^2 > 0$ after recalling that $\lambda_{\bar{k}} > 0$ according to Lemma 23. We furthermore have

$$\begin{aligned} \mathcal{Q}_\tau^* &= (\widehat{u}_\tau^*, T_{B_0} \widehat{u}_\tau^*)_{L_2(D_T)} \\ &= (\widehat{u}_\tau^* - \widehat{u}_\infty^*, T_{B_0} \widehat{u}_\tau^*)_{L_2(D_T)} + (\widehat{u}_\infty^*, T_{B_0} (\widehat{u}_\tau^* - \widehat{u}_\infty^*))_{L_2(D_T)} + (\widehat{u}_\infty^*, T_{B_0} \widehat{u}_\infty^*)_{L_2(D_T)} \\ &\geq -2C_2^B C^{\widehat{u}} \|\widehat{u}_\tau^* - \widehat{u}_\infty^*\|_{L_2(D_T)} + (\widehat{u}_\infty^*, T_{B_0} \widehat{u}_\infty^*)_{L_2(D_T)}, \end{aligned} \quad (165)$$

where we used Cauchy-Schwarz inequality together with Lemma 22 in the last step. Since \widehat{u}_τ^* converges to \widehat{u}_∞^* in $L_2(D_T)$ by assumption as $\tau \rightarrow \infty$, there exists $\bar{\tau} > 0$ such that $\|\widehat{u}_\tau^* - \widehat{u}_\infty^*\|_{L_2(D_T)} \leq \lambda_{\bar{k}} c_{\bar{k}}^2 / (4C_2^B C^{\widehat{u}})$ for all $\tau > \bar{\tau}$. Thus, $\mathcal{Q}_\tau^* \geq \lambda_{\bar{k}} c_{\bar{k}}^2 / 2$ for all $\tau > \bar{\tau}$. With the fundamental theorem of calculus it then holds

$$\mathcal{J}_\tau^* = \mathcal{J}_{\bar{\tau}}^* - \int_{\bar{\tau}}^\tau \alpha_s \mathcal{Q}_s^* ds \leq \mathcal{J}_{\bar{\tau}}^* - \frac{\lambda_{\bar{k}} c_{\bar{k}}^2}{2} \int_{\bar{\tau}}^\tau \alpha_s ds \rightarrow -\infty \quad (166)$$

as $\tau \rightarrow \infty$ due to condition (6) on the learning rate α_τ . This contradicts the positivity of the loss \mathcal{J}^* . Therefore, $\widehat{u}_\infty^* \equiv 0$ a.e. in $L_2(D_T)$.

Step 2: $u_\infty^* \equiv h$ a.e. in $L_2(D_T)$. By using that \widehat{u}_∞^* is a weak solution to the adjoint PDE (163) in a sense analogous to Definition 5 with right-hand side $(u_\infty^* - h)$, we infer that the left-hand side vanishes for all test functions $\phi \in L_2(D_T)$ as in Proposition 35. Thus, $u_\infty^* \equiv h$ a.e. in $L_2(D_T)$. \blacksquare

Acknowledgments

KR would like to profusely thank Tom Hickling for many insightful discussions about practical perspectives on the topic.

This research project was supported by “DMS-EPSCRC: Asymptotic Analysis of Online Training Algorithms in Machine Learning: Recurrent, Graphical, and Deep Neural Networks”

(NSF DMS-2311500). The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work. For the purpose of Open Access, the authors have applied a CC BY public copyright license to any Author Accepted Manuscript (AAM) version arising from this submission.

Appendix A. Well-Posedness of the NN-PDE Training Dynamics

In this appendix, we show the well-posedness of the NN-PDE training dynamics in both the finite-width hidden layer regime and the infinite-width hidden layer limit. In Appendix A.1 we prove Lemma 13, which is concerned with the latter, i.e., the well-posedness of PDE system (16)–(17) coupled with the integro-differential equation (15) for g_τ^* , while Appendix A.2 is concerned with Lemma 11, i.e., the well-posedness of PDE system (1) & (8) coupled with the gradient descent update (5) for the NN parameters of the NN function $g_{\theta_\tau}^N$.

Recall that $\mathcal{S} = L_2([0, T], H^1(D)) \cap L_\infty([0, T], L_2(D))$.

A.1 Well-Posedness Proof of the NN-PDE Training Dynamics in the Infinite-Width Hidden Layer Limit

Proof [Proof of Lemma 13] Existence. The existence proof is based on a fixed point argument employing the Banach fixed point theorem. For a given training time horizon $\mathcal{T} > 0$, let us denote by $\mathcal{V}_\mathcal{T} = \mathcal{C}([0, \mathcal{T}], \mathcal{S})$ the Banach space consisting of elements with finite norm

$$\|u\|_{\mathcal{V}_\mathcal{T}} = \sup_{\tau \in [0, \mathcal{T}]} \left(\|u_\tau\|_{L_2([0, T], H^1(D))} + \|u_\tau\|_{L_\infty([0, T], L_2(D))} \right). \quad (167)$$

A solution $((u_\tau^*, \widehat{u}_\tau^*))_{\tau \in [0, \mathcal{T}]}$ to the PDE system (16)–(17) is shown in what follows to be an element of the space $\mathcal{C}([0, \mathcal{T}], \mathcal{S} \times \mathcal{S})$ (which we identify with the space $\mathcal{V}_\mathcal{T} \times \mathcal{V}_\mathcal{T}$) with additional regularity.

Step 1: Existence and regularity for given right-hand side $\tilde{g}_\tau = -\int_0^\tau \alpha_s b_s ds$. For given $\mathcal{T} > 0$, let $b : [0, \mathcal{T}] \rightarrow L_2(D_T)$ be a given function with b_τ being Lipschitz continuous on $\overline{D_T}$ for each $\tau \in [0, \mathcal{T}]$ and such that $\sup_{\tau \in [0, \mathcal{T}]} \|b_\tau\|_{L_\infty(D_T)} \leq C_b$, where C_b may depend in particular on \mathcal{T} . Consider the auxiliary PDE system

$$\begin{aligned} \partial_t \tilde{u}_\tau^* + \mathcal{L} \tilde{u}_\tau^* - q(\tilde{u}_\tau^*) &= \tilde{g}_\tau = -\int_0^\tau \alpha_s b_s ds && \text{in } D_T, \\ \tilde{u}_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\ \tilde{u}_\tau^* &= f && \text{on } \{0\} \times D \end{aligned} \quad (168)$$

and

$$\begin{aligned} -\partial_t \widehat{u}_\tau^* + \mathcal{L}^\dagger \widehat{u}_\tau^* - q_u(\tilde{u}_\tau^*) \widehat{u}_\tau^* &= (\tilde{u}_\tau^* - h) && \text{in } D_T, \\ \widehat{u}_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\ \widehat{u}_\tau^* &= 0 && \text{on } \{T\} \times D. \end{aligned} \quad (169)$$

We first prove that there exists a solution $(\tilde{u}_\tau^*, \widehat{u}_\tau^*) \in \mathcal{S} \times \mathcal{S}$ to the system (168)–(169) for all $\tau \in [0, \mathcal{T}]$ using classical existence results from (Ladyženskaja et al., 1968). Such solution,

as we show, enjoys the property that for all $\tau \in [0, \mathcal{T}]$ it holds $(\partial_t \tilde{u}_\tau^*(t, \bullet), \partial_t \widehat{u}_\tau^*(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$.

Step 1a: Existence of solution to PDE (168). For the existence of a solution to the nonlinear PDE (168), we invoke (Ladyženskaja et al., 1968, Chapter V, Theorem 6.2). To begin with, we notice that, in the notation of (Ladyženskaja et al., 1968, Chapter V, Theorem 6.2), the coefficients of the nonlinear PDE operator of the parabolic PDE (168) are $a_i(t, x, u, p) = \sum_{j=1}^d a^{ji}(t, x)p_j$ and $a(t, x, u, p) = \sum_{i=1}^d b^i(t, x)p_i + c(t, x)u - q(t, x, u) + \int_0^\tau \alpha_s b_s(t, x) ds$, and thus also $A(t, x, u, p) = \sum_{i=1}^d b^i(t, x)p_i + c(t, x)u - q(t, x, u) + \int_0^\tau \alpha_s b_s(t, x) ds - \sum_{j=1}^d \partial_{x_i} a^{ji}(t, x)p_j$. Clearly, for $(t, x) \in \overline{D_T}$ and arbitrary u it holds $\sum_{i,j=1}^d \partial_{p_j} a_i(t, x, u, p) \xi_i \xi_j \Big|_{p=0} = \sum_{i,j=1}^d a^{ji}(t, x) \xi_i \xi_j \geq \nu \|\xi\|^2 \geq 0$ by uniform parabolicity of $\partial_t + \mathcal{L}$, i.e., Assumption A3, and it holds with Young's inequality

$$\begin{aligned} A(t, x, u, 0)u &= \left(c(t, x)u - q(t, x, u) + \int_0^\tau \alpha_s b_s(t, x) ds \right) u \\ &\geq -\|c\|_{L_\infty(D_T)} u^2 - C_q(1 + |u|)|u| + \frac{1}{2} \left(\int_0^\tau \alpha_s b_s(t, x) ds \right)^2 - \frac{1}{2} u^2 \geq -b_1 u^2 - b_2 \end{aligned} \quad (170)$$

by Assumptions A4 and W1 for the first term, by Assumption W2 for the second term, and, for the last term, due to α_τ being bounded from above together with $\sup_{\tau \in [0, \mathcal{T}]} \|b_\tau\|_{L_\infty(D_T)} \leq C_b$ by assumption on b . Moreover, by Assumptions W1 and W2 the functions a_i and a are continuous w.r.t. t, x, u, p since again b_τ is continuous for every $\tau \in [0, \mathcal{T}]$. Interchanging limits in the term $\int_0^\tau \alpha_s b_s(t, x) ds$ is warranted by the dominated convergence theorem since α_τ is bounded from above and $\sup_{\tau \in [0, \mathcal{T}]} \|b_\tau\|_{L_\infty(D_T)} \leq C_b$. In addition, the functions a_i are differentiable w.r.t. x, u, p by Assumption W1. For $(t, x) \in \overline{D_T}$, $|u| \leq M$ and arbitrary p we furthermore have

$$\begin{aligned} &\sum_{i=1}^d (|a_i| + |\partial_u a_i|) (1 + \|p\|) + \sum_{i,j=1}^d |\partial_{x_j} a_i| + |a| \\ &= \sum_{i=1}^d \left| \sum_{j=1}^d a^{ji}(t, x)p_j \right| (1 + \|p\|) + \sum_{i,j=1}^d \left| \partial_{x_j} \sum_{k=1}^d a^{ki}(t, x)p_k \right| \\ &\quad + \left| \sum_{i=1}^d b^i(t, x)p_i + c(t, x)u - q(t, x, u) + \int_0^\tau \alpha_s b_s(t, x) ds \right| \leq \mu(1 + \|p\|)^2, \end{aligned} \quad (171)$$

where the last inequality holds due to Assumptions A4 and W2, and due to the last term being uniformly bounded with the same arguments as above. Furthermore, for $(t, x) \in \overline{D_T}$, $|u| \leq M$ and $\|p\| \leq \widetilde{M}$, we have the following Hölder continuity properties in (t, x, u, p) (we denote by \star the exponent if the respective function does not depend on the variable, thus being Hölder continuous with any exponent): the functions a_i are $(\gamma_1/2, \gamma_1, \star, 1)$ -Hölder continuous, the functions $\partial_{p_j} a_i$ are $(\gamma_1/2, \gamma_1, \star, \star)$ -Hölder continuous, the functions $\partial_u a_i$ are $(\star, \star, \star, \star)$ -Hölder continuous, the functions $\partial_{x_i} a_i$ are $(\gamma_1/2, \gamma_1, \star, \star)$ -Hölder continuous, and the function a is $(\min\{\gamma_1/2, 1\}, \min\{\gamma_1, 1\}, 1, 1)$ -Hölder continuous. The Hölder properties of all those functions are due to Assumption W1, except for the last function, where we further

used that firstly q is $(\gamma_1/2, \gamma_1, 1)$ -Hölder continuous in (t, x, u) by Assumption W2 for t, x and the mean-value theorem together with Assumption A5 for u , and secondly that b_τ is $(1, 1)$ -Hölder continuous for every $\tau \in [0, T]$ by assumption together with α_τ being bounded from above. Lastly, the boundary ∂D and the initial condition f and boundary condition satisfy the assumptions due to Assumptions A1 and W4, respectively. Thus, (Ladyženskaja et al., 1968, Chapter V, Theorem 6.2) ensures the existence of a solution $\tilde{u}_\tau^* \in H^{\gamma'/2, \gamma'}(\overline{D_T})$ to (168) with $\partial_{x_i} \tilde{u}_\tau^*$ being bounded in $\overline{D_T}$. Since we are on a compact domain as of Assumptions W1, where Hölder continuity implies uniform boundedness, we proved $\tilde{u}_\tau^* \in \mathcal{S}$. (Ladyženskaja et al., 1968, Chapter V, Theorem 6.2) further ensures that $\partial_t \tilde{u}_\tau^* \in H^{\gamma'/2, \gamma'}(D_T)$, and thus also $\partial_t \tilde{u}_\tau^*(t, \bullet) \in L_2(D)$ for a.e. $t \in [0, T]$ is proven.

Step 1b: Existence of solution to adjoint PDE (169). For the existence of a solution to the linear adjoint PDE (169), we invoke the classical results (Evans, 2010, Chapter 7.1, Theorem 3) and (Evans, 2010, Chapter 7.1, Theorem 4) as well as (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1) with $p = 2$. To this end, let us first reverse the adjoint parabolic backward PDE (169) in time to obtain with a time transformation for $\hat{\underline{u}}_\tau^* = \hat{\underline{u}}_\tau^*(t, x) = \hat{\underline{u}}_\tau^*(T - t, x)$ the parabolic forward PDE

$$\begin{aligned} \partial_t \hat{\underline{u}}_\tau^* + \underline{\mathcal{L}}^* \hat{\underline{u}}_\tau^* - \underline{q}_u(\tilde{u}_\tau^*(T - \bullet, \bullet)) \hat{\underline{u}}_\tau^* &= (\tilde{u}_\tau^*(T - \bullet, \bullet) - h(T - \bullet, \bullet)) && \text{in } D_T, \\ \hat{\underline{u}}_\tau^* &= 0 && \text{on } [0, T] \times \partial D, \\ \hat{\underline{u}}_\tau^* &= 0 && \text{on } \{0\} \times D, \end{aligned} \quad (172)$$

where $\underline{\mathcal{L}}^* = \underline{\mathcal{L}}^*(t, x) = \underline{\mathcal{L}}^\dagger(T - t, x)$ (analogously for the individual coefficients of the operator $\underline{\mathcal{L}}^*$) and $\underline{q} = \underline{q}(t, x, \hat{u}) = \underline{q}(T - t, x, \hat{u})$. Since the parabolic PDE (172) is linear, existence and uniqueness of a weak solution of (172) in the sense of Definition 5 follow from classical results, see, e.g., (Evans, 2010, Chapter 7.1, Theorem 3) and (Evans, 2010, Chapter 7.1, Theorem 4) for existence and uniqueness, respectively. To apply those results, note that the term $\underline{q}_u(\tilde{u}_\tau^*(T - \bullet, \bullet)) \hat{\underline{u}}_\tau^*$ can be absorbed into a PDE operator $\tilde{\underline{\mathcal{L}}}^*$ with $\tilde{\underline{c}} = \underline{c} - \sum_{i=1}^d \partial_{x_i} b^i - \underline{q}_u(\tilde{u}_\tau^*(T - \bullet, \bullet)) \in L_\infty(D_T)$ due to Assumptions A4 and A5. Moreover, since $\tilde{u}_\tau^* \in L_2(\overline{D_T})$ by the former statement and since $h \in L_2(D_T)$ by assumption, the right-hand side $(\tilde{u}_\tau^*(T - \bullet, \bullet) - h(T - \bullet, \bullet)) \in L_2(D_T)$. With this, we proved $\hat{\underline{u}}_\tau^* \in L_2([0, T], H^1(D)) \cap L_\infty([0, T], L_2(D))$. To prove additional regularity, we invoke (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1) with $p = 2$. We now notice that, in the notation of (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1), the coefficients $a_{ij} = \underline{a}^{ij}$ of the linear PDE operator of the parabolic PDE (172) are bounded continuous functions in D_T for all $i, j = 1, \dots, d$ due to Assumptions A4 and W1, while the coefficients $a_i = \underline{b}^i - \sum_{j=1}^d \partial_{x_j} \underline{a}^{ji}$ and $a = \underline{c} - \sum_{i=1}^d \partial_{x_i} b^i - \underline{q}_u(\tilde{u}_\tau^*(T - \bullet, \bullet))$ have finite norms $\|a_i\|_{L_r(D_T)}$ and $\|a\|_{L_s(D_T)}$ for any $r, s > 0$. This is due to the uniform boundedness of the coefficients per Assumptions A4 and A5 combined with the boundedness of the domain per Assumption A2, see the subsequent computations with $T' = 0$ and $\Delta T' = T$. Moreover, since it hold $\|a_i\|_{L_r(D_{T', T'+\Delta T'})} \leq (\|b^i\|_{L_\infty(D_T)} + \sum_{j=1}^d \|\partial_{x_j} \underline{a}^{ji}\|_{L_\infty(D_T)}) (\Delta T' \text{vol}(D))^{1/r}$ for all $i = 1, \dots, d$ and $\|a\|_{L_s(D_{T', T'+\Delta T'})} \leq (\|c\|_{L_\infty(D_T)} + \sum_{i=1}^d \|\partial_{x_i} b^i\|_{L_\infty(D_T)} + c_q) (\Delta T' \text{vol}(D))^{1/s}$, $\|a_i\|_{L_r(D_{T', T'+\Delta T'})}$ and $\|a\|_{L_s(D_{T', T'+\Delta T'})}$ tend to zero as $\Delta T' \rightarrow 0$. Furthermore, ∂D is sufficiently smooth as of Assumption A1. The right-hand side $f = (\tilde{u}_\tau^*(T - \bullet, \bullet) - h(T - \bullet, \bullet)) \in$

$L_2(D_T)$ as argued before. Lastly, the initial and boundary conditions $\phi = 0 \in W_2^1(D)$ and $\Phi = 0 \in W_2^{3/4, 3/2}(\partial D_T)$ satisfy the compatibility condition $\phi|_{\partial D} = \Phi|_{t=0}$. Thus, (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1) ensures the existence of a unique solution $\widehat{u}_\tau^* \in W_2^{1,2}(D_T)$ to (172) and thus also a unique solution $\widehat{u}_\tau^* \in W_2^{1,2}(D_T)$ to the parabolic backward PDE (169). We moreover have the bound

$$\|\widehat{u}_\tau^*\|_{W_2^{1,2}(D_T)} = \|\widehat{u}_\tau^*\|_{W_2^{1,2}(D_T)} \lesssim \|\widetilde{u}_\tau^*\|_{L_2(D_T)} + \|h\|_{L_2(D_T)}. \quad (173)$$

In particular, since $\widehat{u}_\tau^* \in W_2^{1,2}(D_T)$, also $\partial_t \widehat{u}_\tau^*(t, \bullet) \in L_2(D)$ for a.e. $t \in [0, T]$ is proven.

Step 1c: Explicit norm bound for the solution to PDE system (168)–(169). In this step, we compute explicit bounds on the norms $\|\widetilde{u}_\tau^*\|_{L_2([0, T], H^1(D))} + \|\widetilde{u}_\tau^*\|_{L_\infty([0, T], L_2(D))}$ as well as $\|\widehat{u}_\tau^*\|_{L_2([0, T], H^1(D))} + \|\widehat{u}_\tau^*\|_{L_\infty([0, T], L_2(D))}$, respectively.

Step 1c(i): Energy estimate for solution to (168). For the norm of a solution to the nonlinear PDE (168) we conduct the following computations. We obtain by chain rule and by using that \widetilde{u}_τ^* is a weak solution to (168) in the sense of Definition 4 that

$$\begin{aligned} \partial_t \|\widetilde{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 &= 2(\widetilde{u}_\tau^*(t, \bullet), \partial_t \widetilde{u}_\tau^*(t, \bullet))_{L_2(D)} = 2\langle \partial_t \widetilde{u}_\tau^*(t, \bullet), \widetilde{u}_\tau^*(t, \bullet) \rangle_{H^{-1}(D), H_0^1(D)} \\ &= -2\mathcal{B}[\widetilde{u}_\tau^*(t, \bullet), \widetilde{u}_\tau^*(t, \bullet); t] + 2(q(t, \bullet, \widetilde{u}_\tau^*(t, \bullet)), \widetilde{u}_\tau^*(t, \bullet))_{L_2(D)} \\ &\quad + 2(\widetilde{g}_\tau(t, \bullet), \widetilde{u}_\tau^*(t, \bullet))_{L_2(D)}, \end{aligned} \quad (174)$$

where the second step is due to the dual pairing between $H^{-1}(D)$ and $H_0^1(D)$ coinciding with the $L_2(D)$ scalar product (Evans, 2010, Chapter 5.9, Theorem 1(iii)) since $\widetilde{u}_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$ and since $\partial_t \widetilde{u}_\tau^*(t, \bullet) \in L_2(D)$ for a.e. $t \in [0, T]$. For the weak solution property in the third step of (174) to hold, we note that \widetilde{u}_τ^* can be used as a test function in the weak formulation of (168), see Definition 4, since $\widetilde{u}_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$. To estimate the right-hand side of (174) from above, we consider each of the three terms separately. For the first term of (174), by using the definition of the bilinear form \mathcal{B} as well as that by Assumption A3 the PDE operator is uniformly parabolic and that by Assumption A4 the coefficients are in L_∞ , we can estimate with Cauchy-Schwarz and Young's inequality analogously to (104)

$$-\mathcal{B}[\widetilde{u}_\tau^*(t, \bullet), \widetilde{u}_\tau^*(t, \bullet); t] \leq -\frac{\nu}{2} |\widetilde{u}_\tau^*(t, \bullet)|_{H^1(D)}^2 + \left(\frac{1}{2\nu} \sum_{i=1}^d \|b^i\|_{L_\infty(D_T)} + \|c\|_{L_\infty(D_T)} \right) \|\widetilde{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2. \quad (175)$$

For the second term of (174) we can estimate with Assumption W2 that

$$\begin{aligned} (q(t, \bullet, \widetilde{u}_\tau^*(t, \bullet)), \widetilde{u}_\tau^*(t, \bullet))_{L_2(D)} &\leq \|q(t, \bullet, \widetilde{u}_\tau^*(t, \bullet))\|_{L_2(D)} \|\widetilde{u}_\tau^*(t, \bullet)\|_{L_2(D)} \\ &\leq C_q (1 + \|\widetilde{u}_\tau^*(t, \bullet)\|_{L_2(D)}) \|\widetilde{u}_\tau^*(t, \bullet)\|_{L_2(D)} \\ &= C_q \left(\frac{1}{2} + \frac{3}{2} \|\widetilde{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 \right). \end{aligned} \quad (176)$$

For the third and last term of (174), by Cauchy-Schwarz and Young's inequality we upper bound

$$\begin{aligned} (\tilde{g}_\tau(t, \bullet), \tilde{u}_\tau^*(t, \bullet))_{L_2(D)} &\leq \|\tilde{g}_\tau(t, \bullet)\|_{L_2(D)} \|\tilde{u}_\tau^*(t, \bullet)\|_{L_2(D)} \\ &\leq \frac{1}{2} \left(\|\tilde{g}_\tau(t, \bullet)\|_{L_2(D)}^2 + \|\tilde{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 \right). \end{aligned} \quad (177)$$

Combining the bounds established in (175)–(177) and inserting them into (174), we arrive after reordering at

$$\partial_t \|\tilde{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 + \frac{\nu}{2} |\tilde{u}_\tau^*(t, \bullet)|_{H^1(D)}^2 \leq C \|\tilde{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 + \|\tilde{g}_\tau(t, \bullet)\|_{L_2(D)}^2 + C, \quad (178)$$

for a constant $C = C(\mathcal{L}, q)$. An application of Grönwall's inequality shows

$$\|\tilde{u}_\tau^*\|_{L_2([0, T], H^1(D))} + \|\tilde{u}_\tau^*\|_{L_\infty([0, T], L_2(D))} \leq C \left(\|f\|_{L_2(D)} + \|\tilde{g}_\tau\|_{L_2(D_T)} + 1 \right) \quad (179)$$

for some other, potentially larger, constant $C = C(T, \mathcal{L}, q)$. Thus, in particular,

$$\|\tilde{u}^*\|_{\mathcal{V}_T} \leq C \left(\|f\|_{L_2(D)} + \sup_{\tau \in [0, T]} \|\tilde{g}_\tau\|_{L_2(D_T)} + 1 \right). \quad (180)$$

Step 1c(ii): Energy estimate for solution to (169). For the norm of a solution to the PDE (169) we proceed as follows using the time-reversed formulation (172). We obtain again by chain rule and by using that $\widehat{\underline{u}}_\tau^*$ is a weak solution to (168) in the sense of Definition 5 that

$$\begin{aligned} \partial_t \|\widehat{\underline{u}}_\tau^*(t, \bullet)\|_{L_2(D)}^2 &= 2(\widehat{\underline{u}}_\tau^*(t, \bullet), \partial_t \widehat{\underline{u}}_\tau^*(t, \bullet))_{L_2(D)} = 2\langle \partial_t \widehat{\underline{u}}_\tau^*(t, \bullet), \widehat{\underline{u}}_\tau^*(t, \bullet) \rangle_{H^{-1}(D), H_0^1(D)} \\ &= -2\mathcal{B}^*[\widehat{\underline{u}}_\tau^*(t, \bullet), \widehat{\underline{u}}_\tau^*(t, \bullet); t] + 2(\underline{q}_u(t, \bullet, \tilde{u}_\tau^*(T-t, \bullet)) \widehat{\underline{u}}_\tau^*(t, \bullet), \widehat{\underline{u}}_\tau^*(t, \bullet))_{L_2(D)} \\ &\quad + 2(\tilde{u}_\tau^*(T-t, \bullet) - h(T-t, \bullet), \widehat{\underline{u}}_\tau^*(t, \bullet))_{L_2(D)}, \end{aligned} \quad (181)$$

where the individual steps hold as before since $\widehat{\underline{u}}_\tau^*(t, \bullet) \in H_0^1(D)$ for a.e. $t \in [0, T]$ and since $\partial_t \widehat{\underline{u}}_\tau^*(t, \bullet) \in L_2(D)$ for a.e. $t \in [0, T]$. To estimate the right-hand side of (181) from above, we consider each of the three terms separately. For the first term of (181), by using the definition of the bilinear form \mathcal{B}^\dagger as well as that by Assumption A3 the PDE operator is uniformly parabolic and that by Assumption A4 the coefficients are in L_∞ , we can estimate as in (175) that

$$\begin{aligned} -\mathcal{B}^*[\widehat{\underline{u}}_\tau^*(t, \bullet), \widehat{\underline{u}}_\tau^*(t, \bullet); t] &= -\mathcal{B}^\dagger[\widehat{\underline{u}}_\tau^*(t, \bullet), \widehat{\underline{u}}_\tau^*(t, \bullet); T-t] = -\mathcal{B}[\widehat{\underline{u}}_\tau^*(t, \bullet), \widehat{\underline{u}}_\tau^*(t, \bullet); T-t] \\ &\leq -\frac{\nu}{2} |\widehat{\underline{u}}_\tau^*(t, \bullet)|_{H^1(D)}^2 + \left(\frac{1}{2\nu} \sum_{i=1}^d \|b^i\|_{L_\infty(D_T)} + \|c\|_{L_\infty(D_T)} \right) \|\widehat{\underline{u}}_\tau^*(t, \bullet)\|_{L_2(D)}^2. \end{aligned} \quad (182)$$

For the second term of (181) we can estimate with Assumption A5 that

$$(\underline{q}_u(t, \bullet, \tilde{u}_\tau^*(T-t, \bullet)) \widehat{\underline{u}}_\tau^*(t, \bullet), \widehat{\underline{u}}_\tau^*(t, \bullet))_{L_2(D)} \leq c_q \|\widehat{\underline{u}}_\tau^*(t, \bullet)\|_{L_2(D)}^2. \quad (183)$$

For the third and last term of (181), by Cauchy-Schwarz and Young's inequality we upper bound

$$\begin{aligned} (\widehat{u}_\tau^*(T-t, \bullet) - h(T-t, \bullet), \widehat{u}_\tau^*(t, \bullet))_{L_2(D)} &\leq \|\widehat{u}_\tau^*(T-t, \bullet) - h(T-t, \bullet)\|_{L_2(D)} \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)} \\ &\leq \frac{1}{2} \left(\|\widehat{u}_\tau^*(T-t, \bullet) - h(T-t, \bullet)\|_{L_2(D)}^2 + \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 \right). \end{aligned} \quad (184)$$

Combining the bounds established in (182)–(184) and inserting them into (181), we arrive after reordering at

$$\partial_t \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 + \frac{\nu}{2} |\widehat{u}_\tau^*(t, \bullet)|_{H^1(D)}^2 \leq C \|\widehat{u}_\tau^*(t, \bullet)\|_{L_2(D)}^2 + \|\widehat{u}_\tau^*(T-t, \bullet) - h(T-t, \bullet)\|_{L_2(D)}^2, \quad (185)$$

for a constant $C = C(\mathcal{L}, q)$. Recalling that $\|\widehat{u}_\tau^*(0, \bullet)\|_{L_2(D)}^2 = 0$, an application of Grönwall's inequality shows

$$\begin{aligned} \|\widehat{u}_\tau^*\|_{L_2([0, T], H^1(D))} + \|\widehat{u}_\tau^*\|_{L_\infty([0, T], L_2(D))} &= \|\widehat{u}_\tau^*\|_{L_2([0, T], H^1(D))} + \|\widehat{u}_\tau^*\|_{L_\infty([0, T], L_2(D))} \\ &\leq C \|\widehat{u}_\tau^* - h\|_{L_2(D_T)} \leq C \left(\|\widehat{u}_\tau^*\|_{L_2(D_T)} + \|h\|_{L_2(D_T)} \right) \end{aligned} \quad (186)$$

for some other, potentially larger, constant $C = C(T, \mathcal{L}, q)$. Thus, in particular,

$$\|\widehat{u}^*\|_{\mathcal{V}_T} \leq C \left(\|\widehat{u}^*\|_{\mathcal{V}_T} + \|h\|_{L_2(D_T)} \right). \quad (187)$$

Step 1d: Existence of solution to PDE system (168)–(169). Summarizing the former results from *Steps 1a, 1b* and *1c* we thus proved that for each $\tau \in [0, T]$ there exists a solution $(\widehat{u}_\tau^*, \widehat{u}_\tau^*) \in \mathcal{S} \times \mathcal{S}$ to (168)–(169). As \widetilde{g}_τ is Lipschitz continuous in τ by the dominated convergence theorem, which can be seen since

$$\|\widetilde{g}_{\tau_2} - \widetilde{g}_{\tau_1}\|_{L_2(D_T)} = \left\| \int_{\tau_1}^{\tau_2} \alpha_s b_s ds \right\|_{L_2(D_T)} \leq C \int_{\tau_1}^{\tau_2} \|b_s\|_{L_\infty(D_T)} ds \leq C |\tau_2 - \tau_1| \quad (188)$$

for a constant $C = C(\alpha, T, D, C_b)$, the solution $(\widehat{u}^*, \widehat{u}^*)$ is in particular continuous in the training time τ , i.e., $(\widehat{u}^*, \widehat{u}^*) \in \mathcal{V}_T \times \mathcal{V}_T$. As we further showed, for each $\tau \in [0, T]$ such solution satisfies $(\partial_t \widehat{u}_\tau^*(t, \bullet), \partial_t \widehat{u}_\tau^*(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$.

Step 2: Existence for specific right-hand side $\widetilde{g}_\tau = g_\tau^* = -\int_0^\tau \alpha_s T_{B_0} \widehat{u}_s^* ds$. We now make a specific choice for the functions b_τ .

Step 2a: Choice of NN update function $b_\tau = T_{B_0} \widehat{u}_\tau^*$. For an arbitrarily given $\widehat{u}_\tau^* \in \mathcal{S}$, $\tau \in [0, T]$, with $\sup_{\tau \in [0, T]} \|\widehat{u}_\tau^*\|_{L_2(D_T)} \leq M$ (M may depend on \mathcal{T}), we set

$$b_\tau = T_{B_0} \widehat{u}_\tau^* \quad (189)$$

for all $\tau \in [0, T]$. It holds with Lemma 22 that

$$\|b_\tau\|_{L_2(D_T)} = \|T_{B_0} \widehat{u}_\tau^*\|_{L_2(D_T)} \leq C_2^B \|\widehat{u}_\tau^*\|_{L_2(D_T)} \leq C_2^B M, \quad (190)$$

which is a uniform bound in τ . In fact, a more careful estimate employing Lemma 24 shows

$$\begin{aligned} \|b_\tau\|_{L_\infty(D_T)} &= \|T_{B_0}\widehat{u}_\tau^*\|_{L_\infty(D_T)} \\ &= \sup_{(t,x)\in D_T} |[T_{B_0}\widehat{u}_\tau^*](t,x)| \leq C_\infty^{T_B} \|\widehat{u}_\tau^*\|_{L_2(D_T)} \leq C_\infty^{T_B} M. \end{aligned} \quad (191)$$

Since the right-hand side is uniform in τ , $\sup_{\tau\in[0,\mathcal{T}]} \|b_\tau\|_{L_\infty(D_T)} \leq C$, where C may depend on \mathcal{T} . Furthermore, it is immediate to see, that using the definition of T_{B_0} in (12) and that the kernel B as given in (13) is continuous in t, x on $\overline{D_T}$, the function $b_\tau = T_{B_0}\widehat{u}_\tau^*$ is continuous on $\overline{D_T}$ for each $\tau \in [0, \mathcal{T}]$ by the dominated convergence theorem. In fact, the function $b_\tau = T_{B_0}\widehat{u}_\tau^*$ is Lipschitz continuous on $\overline{D_T}$ since it holds by Lemma 25 that

$$|b_\tau(t^1, x^1) - b_\tau(t^2, x^2)| = |[T_{B_0}\widehat{u}_\tau^*](t^1, x^1) - [T_{B_0}\widehat{u}_\tau^*](t^2, x^2)| \leq L_{T_B} (|t^1 - t^2| + \|x^1 - x^2\|) \quad (192)$$

for all $(t^1, x^1), (t^2, x^2) \in \overline{D_T}$.

Step 2b: Definition of fixed point mapping. Let us consider the fixed point map

$$F : \mathcal{V}_\mathcal{T} \times \mathcal{V}_\mathcal{T} \rightarrow \mathcal{V}_\mathcal{T} \times \mathcal{V}_\mathcal{T}, \quad (u^*, \widehat{u}^*) \mapsto (\widetilde{u}^*, \widehat{\widetilde{u}}^*) \quad (193)$$

and define for given $M < \infty$ and $\mathcal{T} < \infty$ the function space $\mathcal{V}_\mathcal{T}(M) = \{u \in \mathcal{V}_\mathcal{T} : \|u\|_{\mathcal{V}_\mathcal{T}} \leq M\}$.

We will first show in *Step 2d* existence locally in the training time by proving that there exist $M_0 > 0$ and $\mathcal{T}_0 > 0$ such that F is a fixed point mapping on $\mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$, which allows to apply the Banach fixed point theorem. In *Step 2e* we will then extend the proof by a bootstrapping argument to any given (arbitrarily large) time horizon \mathcal{T} .

Step 2c: Preliminary computations. Let us start by conducting some preliminary computations on a generic space $\mathcal{V}_{\widetilde{\mathcal{T}}}(\widetilde{M}) \times \mathcal{V}_{\widetilde{\mathcal{T}}}(\widetilde{M})$.

Step 2c(i): Preliminary computations for self-mapping property of F . Consider $(\widetilde{u}^*, \widehat{\widetilde{u}}^*)$ together with its corresponding $(u^*, \widehat{u}^*) \in \mathcal{V}_{\widetilde{\mathcal{T}}}(\widetilde{M}) \times \mathcal{V}_{\widetilde{\mathcal{T}}}(\widetilde{M})$. Using (180) and (187) in the first inequality and Lemma 22 in the last step, we establish

$$\begin{aligned} &\|\widetilde{u}^*\|_{\mathcal{V}_{\widetilde{\mathcal{T}}}} + \|\widehat{\widetilde{u}}^*\|_{\mathcal{V}_{\widetilde{\mathcal{T}}}} \\ &\leq C \sup_{\tau \in [0, \widetilde{\mathcal{T}}]} \|\widetilde{g}_\tau\|_{L_2(D_T)} + C \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + 1 \right) \\ &\leq C \sup_{\tau \in [0, \widetilde{\mathcal{T}}]} \left\| \int_0^\tau \alpha_s T_{B_0} \widehat{u}_s^* ds \right\|_{L_2(D_T)} + C \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + 1 \right) \\ &\leq C \int_0^{\widetilde{\mathcal{T}}} \|\alpha_s T_{B_0} \widehat{u}_s^*\|_{L_2(D_T)} ds + C \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + 1 \right) \\ &\leq C_1 \int_0^{\widetilde{\mathcal{T}}} \|\widehat{u}_s^*\|_{L_2(D_T)} ds + C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + 1 \right) \end{aligned} \quad (194)$$

for a constant $C_1 = C_1(\alpha, \mathcal{L}, q, C_2^B)$ (to be precise, $C_1 = C \max\{\alpha_0 C_2^B, 1\}$).

Step 2c(ii): Preliminary computations for contractivity of F . Consider two pairs $(\widetilde{u}^{*,1}, \widehat{\widetilde{u}}^{*,1}), (\widetilde{u}^{*,2}, \widehat{\widetilde{u}}^{*,2})$ with their corresponding $(u^{*,1}, \widehat{u}^{*,1}), (u^{*,2}, \widehat{u}^{*,2}) \in \mathcal{V}_{\widetilde{\mathcal{T}}}(\widetilde{M}) \times \mathcal{V}_{\widetilde{\mathcal{T}}}(\widetilde{M})$.

A bound for $\|\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}\|_{L_2([0,T],H^1(D))} + \|\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}\|_{L_\infty([0,T],L_2(D))}$. Since both $\tilde{u}_\tau^{*,1}$ and $\tilde{u}_\tau^{*,2}$ weakly satisfy (168) in the sense of Definition 4, it weakly holds

$$\partial_t (\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}) + \mathcal{L} (\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}) - (q(\tilde{u}_\tau^{*,1}) - q(\tilde{u}_\tau^{*,2})) = \tilde{g}_\tau^1 - \tilde{g}_\tau^2 \quad (195)$$

with zero initial and zero boundary conditions. We obtain by chain rule and by using that $\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}$ is a weak solution to (195) in the sense of Definition 4 that

$$\begin{aligned} & \partial_t \left\| \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet) \right\|_{L_2(D)}^2 \\ &= 2 \left(\tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet), \partial_t (\tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet)) \right)_{L_2(D)} \\ &= 2 \left\langle \partial_t (\tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet)), \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} \\ &= -2\mathcal{B}[\tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet), \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet); t] \\ &\quad + 2 \left(q(t, \bullet, \tilde{u}_\tau^{*,1}(t, \bullet)) - q(t, \bullet, \tilde{u}_\tau^{*,2}(t, \bullet)), \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet) \right)_{L_2(D)} \\ &\quad + 2 \left(\tilde{g}_\tau^1(t, \bullet) - \tilde{g}_\tau^2(t, \bullet), \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet) \right)_{L_2(D)}, \end{aligned} \quad (196)$$

where the individual steps hold as previously described. To estimate the right-hand side of (196) from above, we consider each of the three terms separately. For the first term, by using the definition of the bilinear form \mathcal{B} as well as that by Assumption A3 the PDE operator is uniformly parabolic and that by Assumption A4 the coefficients are in L_∞ , we can estimate as in (175) that

$$\begin{aligned} & -\mathcal{B}[\tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet), \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet); t] \\ & \leq -\frac{\nu}{2} \left| \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet) \right|_{H^1(D)}^2 \\ & \quad + \left(\frac{1}{2\nu} \sum_{i=1}^d \|b^i\|_{L_\infty(D_T)} + \|c\|_{L_\infty(D_T)} \right) \left\| \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet) \right\|_{L_2(D)}^2. \end{aligned} \quad (197)$$

For the second term, we first note that by the mean-value theorem, for any $(t, x) \in D_T$ there exists a $\xi(t, x)$ such that

$$q(t, x, \tilde{u}_\tau^{*,1}(t, x)) - q(t, x, \tilde{u}_\tau^{*,2}(t, x)) = q_u(t, x, \xi(t, x)) (\tilde{u}_\tau^{*,1}(t, x) - \tilde{u}_\tau^{*,2}(t, x)). \quad (198)$$

Leveraging this while using that by Assumption A5 q_u is bounded, we can estimate

$$\begin{aligned} & \left(q(t, \bullet, \tilde{u}_\tau^{*,1}(t, \bullet)) - q(t, \bullet, \tilde{u}_\tau^{*,2}(t, \bullet)), \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet) \right)_{L_2(D)} \\ &= \int_D (q_u(t, x, \xi(t, x)) (\tilde{u}_\tau^{*,1}(t, x) - \tilde{u}_\tau^{*,2}(t, x))) (\tilde{u}_\tau^{*,1}(t, x) - \tilde{u}_\tau^{*,2}(t, x)) dx \\ & \leq c_q \int_D (\tilde{u}_\tau^{*,1}(t, x) - \tilde{u}_\tau^{*,2}(t, x))^2 dx = c_q \left\| \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet) \right\|_{L_2(D)}^2. \end{aligned} \quad (199)$$

For the third and last term, by Cauchy-Schwarz and Young's inequality we upper bound

$$\begin{aligned} & \left(\tilde{g}_\tau^1(t, \bullet) - \tilde{g}_\tau^2(t, \bullet), \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet) \right)_{L_2(D)} \\ & \leq \left\| \tilde{g}_\tau^1(t, \bullet) - \tilde{g}_\tau^2(t, \bullet) \right\|_{L_2(D)} \left\| \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet) \right\|_{L_2(D)} \\ & \leq \frac{1}{2} \left(\left\| \tilde{g}_\tau^1(t, \bullet) - \tilde{g}_\tau^2(t, \bullet) \right\|_{L_2(D)}^2 + \left\| \tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet) \right\|_{L_2(D)}^2 \right). \end{aligned} \quad (200)$$

Combining the bounds established in (197)–(200) and inserting them into (196), we arrive after reordering at

$$\begin{aligned} & \partial_t \|\tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet)\|_{L_2(D)}^2 + \frac{\nu}{2} \|\tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet)\|_{H^1(D)}^2 \\ & \leq C \|\tilde{u}_\tau^{*,1}(t, \bullet) - \tilde{u}_\tau^{*,2}(t, \bullet)\|_{L_2(D)}^2 + \|\tilde{g}_\tau^1(t, \bullet) - \tilde{g}_\tau^2(t, \bullet)\|_{L_2(D)}^2, \end{aligned} \quad (201)$$

for a constant $C = C(\mathcal{L}, q)$. Recalling that $\tilde{u}_\tau^{*,1}(0, \bullet) = \tilde{u}_\tau^{*,2}(0, \bullet)$, an application of Grönwall's inequality shows

$$\|\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}\|_{L_2([0,T], H^1(D))} + \|\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}\|_{L_\infty([0,T], L_2(D))} \leq C \|\tilde{g}_\tau^1 - \tilde{g}_\tau^2\|_{L_2(D_T)} \quad (202)$$

for some other, potentially larger, constant $C = C(T, \mathcal{L}, q)$. Thus, in particular,

$$\|\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}\|_{\mathcal{V}_T} \leq C \sup_{\tau \in [0, T]} \|\tilde{g}_\tau^1 - \tilde{g}_\tau^2\|_{L_2(D_T)}. \quad (203)$$

A bound for $\|\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}\|_{L_\infty(D_T)}$. For later use, let us further provide an L_∞ bound for the weak solution to (195) by employing Morrey's inequality after leveraging (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1) for any $p > 1$. Therefore notice, that by the mean-value theorem, for any $(t, x) \in D_T$ there exists a $\xi(t, x)$ such that, in place of (195), it weakly holds

$$\partial_t (\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}) + \mathcal{L} (\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}) - q_u(\xi) (\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}) = \tilde{g}_\tau^1 - \tilde{g}_\tau^2 \quad (204)$$

with zero initial and zero boundary conditions. We now notice that, in the notation of (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1), the coefficients $a_{ij}(t, x) = a^{ij}(t, x)$ of the linear PDE operator of the parabolic PDE (204) are bounded continuous functions in D_T for all $i, j = 1, \dots, d$, while the coefficients $a_i(t, x) = b^i(t, x) - \sum_{j=1}^d \partial_{x_j} a^{ji}(t, x)$ and $a(t, x) = c(t, x) - \sum_{i=1}^d \partial_{x_i} b^i(t, x) - q_u(t, x, \xi(t, x))$ have finite norms $\|a_i\|_{L_r(D_T)}$ and $\|a\|_{L_s(D_T)}$ for any $r, s > 0$. This is due to the uniform boundedness of the coefficients per Assumptions A4 and A5 combined with the boundedness of the domain per Assumption A2, see the subsequent computations with $T' = 0$ and $\Delta T' = T$. Moreover, since it hold $\|a_i\|_{L_r(D_{T', T'+\Delta T'})} \leq (\|b^i\|_{L_\infty(D_T)} + \sum_{j=1}^d \|\partial_{x_j} a^{ji}\|_{L_\infty(D_T)}) (\Delta T' \text{vol}(D))^{1/r}$ for all $i = 1, \dots, d$ and $\|a\|_{L_s(D_{T', T'+\Delta T'})} \leq (\|c\|_{L_\infty(D_T)} + \sum_{i=1}^d \|\partial_{x_i} b^i\|_{L_\infty(D_T)} + c_q) (\Delta T' \text{vol}(D))^{1/s}$, $\|a_i\|_{L_r(D_{T', T'+\Delta T'})}$ and $\|a\|_{L_s(D_{T', T'+\Delta T'})}$ tend to zero as $\Delta T' \rightarrow 0$. Furthermore, ∂D is sufficiently smooth as of Assumption A1. The right-hand side $f = \tilde{g}_\tau^1 - \tilde{g}_\tau^2 \in L_p(D_T)$ for any

$p \geq 2$ since

$$\begin{aligned}
 \|\tilde{g}_\tau^1 - \tilde{g}_\tau^2\|_{L_p(D_T)} &= \left\| \int_0^\tau \alpha_s T_{B_0}(\hat{u}_s^{*,1} - \hat{u}_s^{*,2}) ds \right\|_{L_p(D_T)} \\
 &\leq \alpha_0 \int_0^\tau \|T_{B_0}(\hat{u}_s^{*,1} - \hat{u}_s^{*,2})\|_{L_p(D_T)} ds = \alpha_0 \int_0^\tau \left(\int_0^T \int_D |[T_{B_0}(\hat{u}_s^{*,1} - \hat{u}_s^{*,2})](t, x)|^p dx dt \right)^{1/p} ds \\
 &\leq \alpha_0 C_\infty^{TB} \int_0^\tau \left(\int_0^T \int_D \|\hat{u}_s^{*,1} - \hat{u}_s^{*,2}\|_{L_2(D_T)}^p dx dt \right)^{1/p} ds \\
 &= \alpha_0 C_\infty^{TB} (T \operatorname{vol}(D))^{1/p} \int_0^\tau \|\hat{u}_s^{*,1} - \hat{u}_s^{*,2}\|_{L_2(D_T)} ds \\
 &\leq \alpha_0 \mathcal{T} C_\infty^{TB} (T \operatorname{vol}(D))^{1/p} \sup_{\tau \in [0, T]} \|\hat{u}_\tau^{*,1} - \hat{u}_\tau^{*,2}\|_{L_2(D_T)} ds,
 \end{aligned} \tag{205}$$

where we used Lemma 24 in the third line. Moreover, the initial and boundary conditions $\phi = 0 \in W_p^{2-2/p}(D)$ and $\Phi = 0 \in W_p^{1-1/(2p), 2-1/p}(\partial D_T)$ satisfy the compatibility condition $\phi|_{\partial D} = \Phi|_{t=0}$. Thus, (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1) ensures that the unique solution $(\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}) \in W_p^{1,2}(D_T)$ to (204) and thus (195) obeys the bound

$$\|\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}\|_{W_p^{1,2}(D_T)} \leq C \|\tilde{g}_\tau^1 - \tilde{g}_\tau^2\|_{L_p(D_T)} \tag{206}$$

for a constant $C = C(T, \mathcal{L}, q)$.

With the conditions of (Ladyženskaja et al., 1968, Chapter IV, Theorem 9.1) being fulfilled for any $p \geq 2$ as we verified before, they are in particular fulfilled for $p > d + 1$. Since we have for such p the continuous embedding $W_p^{1,2}(D_T) \hookrightarrow W_p^{1,1}(D_T) \hookrightarrow L_\infty(\bar{D}_T)$ by Morrey's inequality (Brezis, 2011, Theorem 9.12), we have the first inequality in

$$\|\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}\|_{L_\infty(D_T)} \leq c(d, p) \|\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}\|_{W_p^{1,2}(D_T)} \leq c(d, p) C \|\tilde{g}_\tau^1 - \tilde{g}_\tau^2\|_{L_p(D_T)}, \tag{207}$$

with the second one being due to (206). Hence, using the last two lines of (205), we have

$$\begin{aligned}
 \|\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}\|_{L_\infty(D_T)} &\leq C \int_0^\tau \|\hat{u}_s^{*,1} - \hat{u}_s^{*,2}\|_{L_2(D_T)} ds \\
 &\leq C \mathcal{T} \sup_{\tau \in [0, T]} \|\hat{u}_\tau^{*,1} - \hat{u}_\tau^{*,2}\|_{L_2(D_T)} ds,
 \end{aligned} \tag{208}$$

for a constant $C = C(\alpha, T, D, \mathcal{L}, q, C_\infty^{TB})$.

A bound for $\|\hat{u}_\tau^{*,1} - \hat{u}_\tau^{*,2}\|_{L_2([0, T], H^1(D))} + \|\hat{u}_\tau^{*,1} - \hat{u}_\tau^{*,2}\|_{L_\infty([0, T], L_2(D))}$. Since both $\hat{u}_\tau^{*,1}$ and $\hat{u}_\tau^{*,2}$ weakly satisfy (169) in the sense of Definition 5, it weakly holds

$$-\partial_t(\hat{u}_\tau^{*,1} - \hat{u}_\tau^{*,2}) + \mathcal{L}^\dagger(\hat{u}_\tau^{*,1} - \hat{u}_\tau^{*,2}) - (q_u(\hat{u}_\tau^{*,1})\hat{u}_\tau^{*,1} - q_u(\hat{u}_\tau^{*,2})\hat{u}_\tau^{*,2}) = \tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2} \tag{209}$$

with zero terminal and zero boundary conditions, or equivalently

$$\begin{aligned}
 -\partial_t(\hat{u}_\tau^{*,1} - \hat{u}_\tau^{*,2}) + \mathcal{L}^\dagger(\hat{u}_\tau^{*,1} - \hat{u}_\tau^{*,2}) - (q_u(\hat{u}_\tau^{*,1})(\hat{u}_\tau^{*,1} - \hat{u}_\tau^{*,2})) \\
 = (\tilde{u}_\tau^{*,1} - \tilde{u}_\tau^{*,2}) + (q_u(\hat{u}_\tau^{*,1}) - q_u(\hat{u}_\tau^{*,2}))\hat{u}_\tau^{*,2}.
 \end{aligned} \tag{210}$$

Let us now first reverse (210) in time to obtain with a time transformation for $\widehat{\underline{u}}_\tau^{*,k} = \widehat{\underline{u}}_\tau^{*,k}(t, x) = \widehat{\underline{u}}_\tau^{*,k}(T - t, x)$ for $k = 1, 2$ the parabolic forward PDE

$$\begin{aligned} \partial_t(\widehat{\underline{u}}_\tau^{*,1} - \widehat{\underline{u}}_\tau^{*,2}) + \mathcal{L}^*(\widehat{\underline{u}}_\tau^{*,1} - \widehat{\underline{u}}_\tau^{*,2}) - (q_u(\widetilde{u}_\tau^{*,1}(T - \cdot, \cdot))(\widehat{\underline{u}}_\tau^{*,1} - \widehat{\underline{u}}_\tau^{*,2})) \\ = (\widetilde{u}_\tau^{*,1}(T - \cdot, \cdot) - \widetilde{u}_\tau^{*,2}(T - \cdot, \cdot)) + (q_u(\widetilde{u}_\tau^{*,1}(T - \cdot, \cdot)) - q_u(\widetilde{u}_\tau^{*,2}(T - \cdot, \cdot)))\widehat{\underline{u}}_\tau^{*,2} \end{aligned} \quad (211)$$

with zero initial and zero boundary conditions. We obtain by chain rule and by using that $\widehat{\underline{u}}_\tau^{*,1} - \widehat{\underline{u}}_\tau^{*,2}$ is a weak solution to (211) in the sense of Definition 5 that

$$\begin{aligned} \partial_t \|\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot)\|_{L_2(D)}^2 \\ = 2(\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot), \partial_t(\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot)))_{L_2(D)} \\ = 2\langle \partial_t(\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot)), \widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot) \rangle_{H^{-1}(D), H_0^1(D)} \\ = -2\mathcal{B}^*[\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot), \widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot); t] \\ + 2(q_u(t, \cdot, \widetilde{u}_\tau^{*,1}(T - t, \cdot))(\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot)), \widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot))_{L_2(D)} \\ + 2(\widetilde{u}_\tau^{*,1}(T - t, \cdot) - \widetilde{u}_\tau^{*,2}(T - t, \cdot), \widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot))_{L_2(D)} \\ + 2((q_u(t, \cdot, \widetilde{u}_\tau^{*,1}(T - t, \cdot)) - q_u(t, \cdot, \widetilde{u}_\tau^{*,2}(T - t, \cdot)))\widehat{\underline{u}}_\tau^{*,2}(t, \cdot), \widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot))_{L_2(D)}, \end{aligned} \quad (212)$$

where the individual steps hold as previously described. To estimate the right-hand side of (212) from above, we consider each of the four terms separately. For the first term, by using the definition of the bilinear form \mathcal{B} as well as that by Assumption A3 the PDE operator is uniformly parabolic and that by Assumption A4 the coefficients are in L_∞ , we can estimate as in (182) that

$$\begin{aligned} -\mathcal{B}^*[\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot), \widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot); t] \\ \leq -\frac{\nu}{2} \|\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot)\|_{H^1(D)}^2 \\ + \left(\frac{1}{2\nu} \sum_{i=1}^d \|b^i\|_{L_\infty(D_T)} + \|c\|_{L_\infty(D_T)} \right) \|\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot)\|_{L_2(D)}^2. \end{aligned} \quad (213)$$

For the second term, using that by Assumption A5 q_u is bounded, we can estimate directly

$$\begin{aligned} (q_u(t, \cdot, \widetilde{u}_\tau^{*,1}(T - t, \cdot))(\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot)), \widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot))_{L_2(D)} \\ \leq c_q \|\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot)\|_{L_2(D)}^2. \end{aligned} \quad (214)$$

For the third term, by Cauchy-Schwarz and Young's inequality we upper bound

$$\begin{aligned} (\widetilde{u}_\tau^{*,1}(T - t, \cdot) - \widetilde{u}_\tau^{*,2}(T - t, \cdot), \widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot))_{L_2(D)} \\ \leq \|\widetilde{u}_\tau^{*,1}(T - t, \cdot) - \widetilde{u}_\tau^{*,2}(T - t, \cdot)\|_{L_2(D)} \|\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot)\|_{L_2(D)} \\ \leq \frac{1}{2} \left(\|\widetilde{u}_\tau^{*,1}(T - t, \cdot) - \widetilde{u}_\tau^{*,2}(T - t, \cdot)\|_{L_2(D)}^2 + \|\widehat{\underline{u}}_\tau^{*,1}(t, \cdot) - \widehat{\underline{u}}_\tau^{*,2}(t, \cdot)\|_{L_2(D)}^2 \right). \end{aligned} \quad (215)$$

For the fourth and final term, we first note that by the mean-value theorem, for any $(t, x) \in D_T$ there exists a $\xi(t, x)$ such that

$$\underline{q}_u(t, x, \widetilde{u}_\tau^{*,1}(T-t, x)) - \underline{q}_u(t, x, \widetilde{u}_\tau^{*,2}(T-t, x)) = \underline{q}_{uu}(t, x, \xi(t, x)) (\widetilde{u}_\tau^{*,1}(T-t, x) - \widetilde{u}_\tau^{*,2}(T-t, x)). \quad (216)$$

Leveraging this while using that by Assumption A6 q_{uu} is bounded, we can estimate by Cauchy-Schwarz and Young's inequality

$$\begin{aligned} & ((\underline{q}_u(t, \bullet, \widetilde{u}_\tau^{*,1}(T-t, \bullet)) - \underline{q}_u(t, \bullet, \widetilde{u}_\tau^{*,2}(T-t, \bullet))) \widehat{\underline{u}}_\tau^{*,2}(t, \bullet), \widehat{\underline{u}}_\tau^{*,1}(t, \bullet) - \widehat{\underline{u}}_\tau^{*,2}(t, \bullet))_{L_2(D)} \\ & \leq \|(\underline{q}_u(t, \bullet, \widetilde{u}_\tau^{*,1}(T-t, \bullet)) - \underline{q}_u(t, \bullet, \widetilde{u}_\tau^{*,2}(T-t, \bullet))) \widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)} \\ & \quad \cdot \|\widehat{\underline{u}}_\tau^{*,1}(t, \bullet) - \widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)} \\ & \leq \|\underline{q}_u(t, \bullet, \widetilde{u}_\tau^{*,1}(T-t, \bullet)) - \underline{q}_u(t, \bullet, \widetilde{u}_\tau^{*,2}(T-t, \bullet))\|_{L_\infty(D)} \|\widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)} \\ & \quad \cdot \|\widehat{\underline{u}}_\tau^{*,1}(t, \bullet) - \widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)} \\ & \leq c'_q \|\widetilde{u}_\tau^{*,1}(T-t, \bullet) - \widetilde{u}_\tau^{*,2}(T-t, \bullet)\|_{L_\infty(D)} \|\widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)} \|\widehat{\underline{u}}_\tau^{*,1}(t, \bullet) - \widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)} \\ & \leq \frac{1}{2} \left(\|\widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)}^2 \|\widetilde{u}_\tau^{*,1}(T-t, \bullet) - \widetilde{u}_\tau^{*,2}(T-t, \bullet)\|_{L_\infty(D)}^2 \right. \\ & \quad \left. + (c'_q)^2 \|\widehat{\underline{u}}_\tau^{*,1}(t, \bullet) - \widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)}^2 \right). \end{aligned} \quad (217)$$

Combining the bounds established in (213)–(217) and inserting them into (212), we arrive after reordering at

$$\begin{aligned} & \partial_t \|\widehat{\underline{u}}_\tau^{*,1}(t, \bullet) - \widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)}^2 + \frac{\nu}{2} \|\widehat{\underline{u}}_\tau^{*,1}(t, \bullet) - \widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{H^1(D)}^2 \\ & \leq C \|\widehat{\underline{u}}_\tau^{*,1}(t, \bullet) - \widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)}^2 + \|\widetilde{u}_\tau^{*,1}(T-t, \bullet) - \widetilde{u}_\tau^{*,2}(T-t, \bullet)\|_{L_2(D)}^2 \\ & \quad + \|\widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)}^2 \|\widetilde{u}_\tau^{*,1}(T-t, \bullet) - \widetilde{u}_\tau^{*,2}(T-t, \bullet)\|_{L_\infty(D)}^2 \\ & \leq C \|\widehat{\underline{u}}_\tau^{*,1}(t, \bullet) - \widehat{\underline{u}}_\tau^{*,2}(t, \bullet)\|_{L_2(D)}^2 + \|\widetilde{u}_\tau^{*,1}(T-t, \bullet) - \widetilde{u}_\tau^{*,2}(T-t, \bullet)\|_{L_2(D)}^2 \\ & \quad + \|\widehat{\underline{u}}_\tau^{*,2}\|_{L_\infty([0, T], L_2(D))}^2 \|\widetilde{u}_\tau^{*,1}(T-t, \bullet) - \widetilde{u}_\tau^{*,2}(T-t, \bullet)\|_{L_\infty(D)}^2 \end{aligned} \quad (218)$$

for a constant $C = C(\mathcal{L}, q)$. Recalling that $\widehat{\underline{u}}_\tau^{*,1}(0, \bullet) = \widehat{\underline{u}}_\tau^{*,2}(0, \bullet)$, an application of Grönwall's inequality shows

$$\begin{aligned} & \|\widehat{\underline{u}}_\tau^{*,1} - \widehat{\underline{u}}_\tau^{*,2}\|_{L_2([0, T], H^1(D))} + \|\widehat{\underline{u}}_\tau^{*,1} - \widehat{\underline{u}}_\tau^{*,2}\|_{L_\infty([0, T], L_2(D))} \\ & \leq C \left(\|\widetilde{u}_\tau^{*,1} - \widetilde{u}_\tau^{*,2}\|_{L_2(D_T)} + \|\widehat{\underline{u}}_\tau^{*,2}\|_{L_\infty([0, T], L_2(D))} \|\widetilde{u}_\tau^{*,1} - \widetilde{u}_\tau^{*,2}\|_{L_\infty(D_T)} \right) \end{aligned} \quad (219)$$

for some other, potentially larger, constant $C = C(T, \mathcal{L}, q)$. Thus, in particular,

$$\begin{aligned} & \|\widehat{u}^{*,1} - \widehat{u}^{*,2}\|_{\mathcal{V}_T} \\ & \leq C \left(\sup_{\tau \in [0, T]} \|\widetilde{u}_\tau^{*,1} - \widetilde{u}_\tau^{*,2}\|_{L_2(D_T)} + \sup_{\tau \in [0, T]} \|\widehat{u}_\tau^{*,2}\|_{L_\infty([0, T], L_2(D))} \sup_{\tau \in [0, T]} \|\widetilde{u}_\tau^{*,1} - \widetilde{u}_\tau^{*,2}\|_{L_\infty(D_T)} \right). \end{aligned} \quad (220)$$

Combination of the bounds. Using (203) and (220) in the first inequality together with the first step of (208) in the second inequality and Lemma 22 in the next-to-last step, we have

$$\begin{aligned} & \|\widetilde{u}^{*,1} - \widetilde{u}^{*,2}\|_{\mathcal{V}_{\widetilde{T}}} + \|\widehat{u}^{*,1} - \widehat{u}^{*,2}\|_{\mathcal{V}_{\widetilde{T}}} \\ & \leq C \sup_{\tau \in [0, \widetilde{T}]} \|\widetilde{g}_\tau^1 - \widetilde{g}_\tau^2\|_{L_2(D_T)} \\ & \quad + C \sup_{\tau \in [0, \widetilde{T}]} \|\widehat{u}_\tau^{*,2}\|_{L_\infty([0, T], L_2(D))} \sup_{\tau \in [0, \widetilde{T}]} \|\widetilde{u}_\tau^{*,1} - \widetilde{u}_\tau^{*,2}\|_{L_\infty(D_T)} \\ & \leq C \sup_{\tau \in [0, \widetilde{T}]} \left\| \int_0^\tau \alpha_s T_{B_0} [\widehat{u}_s^{*,1} - \widehat{u}_s^{*,2}] ds \right\|_{L_2(D_T)} \\ & \quad + C \sup_{\tau \in [0, \widetilde{T}]} \|\widehat{u}_\tau^{*,2}\|_{L_\infty([0, T], L_2(D))} \sup_{\tau \in [0, \widetilde{T}]} \int_0^\tau \|\widehat{u}_s^{*,1} - \widehat{u}_s^{*,2}\|_{L_2(D_T)} ds \quad (221) \\ & \leq C \int_0^{\widetilde{T}} \|\alpha_s T_{B_0} [\widehat{u}_s^{*,1} - \widehat{u}_s^{*,2}]\|_{L_2(D_T)} ds \\ & \quad + C \sup_{\tau \in [0, \widetilde{T}]} \|\widehat{u}_\tau^{*,2}\|_{L_\infty([0, T], L_2(D))} \int_0^{\widetilde{T}} \|\widehat{u}_s^{*,1} - \widehat{u}_s^{*,2}\|_{L_2(D_T)} ds \\ & \leq C_2 \left(1 + \sup_{\tau \in [0, \widetilde{T}]} \|\widehat{u}_\tau^{*,2}\|_{L_\infty([0, T], L_2(D))} \right) \int_0^{\widetilde{T}} \|\widehat{u}_s^{*,1} - \widehat{u}_s^{*,2}\|_{L_2(D_T)} ds \end{aligned}$$

for a constant $C_2 = C_2(\alpha, \mathcal{L}, q, C_2^B)$ (to be precise, $C_2 = C \max\{\alpha_0 C_2^B, 1\}$).

Step 2d: Existence locally in training time. Let us choose

$$M_0 = 2C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + 1 \right) \quad \text{and} \quad \mathcal{T}_0 = \min \left\{ \frac{1}{4C_1}, \frac{1}{2C_2(1 + M_0)} \right\} \quad (222)$$

where the constants C_1 and C_2 are as given implicitly in (194) and (221), respectively. We show in what follows that there exists a unique solution $(u^*, \widehat{u}^*) \in \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$.

Step 2d(i): Self-mapping property of F . Consider $(\widetilde{u}^*, \widehat{u}^*)$ together with its corresponding $(u^*, \widehat{u}^*) \in \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$. Using the definitions of M_0 and \mathcal{T}_0 in (222), respectively, we can derive from (194) that

$$\begin{aligned} \|\widetilde{u}^*\|_{\mathcal{V}_{\mathcal{T}_0}} + \|\widehat{u}^*\|_{\mathcal{V}_{\mathcal{T}_0}} & \leq C_1 \int_0^{\mathcal{T}_0} \|\widehat{u}_s^*\|_{L_2(D_T)} ds + C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + 1 \right) \\ & \leq C_1 \mathcal{T}_0 \|\widehat{u}^*\|_{\mathcal{V}_{\mathcal{T}_0}} + \frac{M_0}{2} \leq C_1 \mathcal{T}_0 M_0 + \frac{M_0}{2} \leq \frac{M_0}{4} + \frac{M_0}{2} \leq M_0. \end{aligned} \quad (223)$$

Thus, $(\tilde{u}^*, \hat{u}^*) \in \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$.

Step 2d(ii): Contractivity of F . Consider two pairs $(\tilde{u}^{*,1}, \hat{u}^{*,1})$, $(\tilde{u}^{*,2}, \hat{u}^{*,2})$ together with their corresponding $(u^{*,1}, \hat{u}^{*,1})$, $(u^{*,2}, \hat{u}^{*,2}) \in \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$. According to *Step 2d(i)*, $(\tilde{u}^{*,1}, \hat{u}^{*,1})$, $(\tilde{u}^{*,2}, \hat{u}^{*,2}) \in \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$. Using the definitions of M_0 and \mathcal{T}_0 in (222), we can derive from (221) that

$$\begin{aligned}
 & \|\tilde{u}^{*,1} - \tilde{u}^{*,2}\|_{\mathcal{V}_{\mathcal{T}_0}} + \|\hat{u}^{*,1} - \hat{u}^{*,2}\|_{\mathcal{V}_{\mathcal{T}_0}} \\
 & \leq C_2 \left(1 + \sup_{\tau \in [0, \mathcal{T}_0]} \|\hat{u}_\tau^{*,2}\|_{L_\infty([0, T], L_2(D))} \right) \int_0^{\mathcal{T}_0} \|\hat{u}_s^{*,1} - \hat{u}_s^{*,2}\|_{L_2(D_T)} ds \\
 & \leq C_2 (1 + M_0) \int_0^{\mathcal{T}_0} \|\hat{u}_s^{*,1} - \hat{u}_s^{*,2}\|_{L_2(D_T)} ds \leq C_2 (1 + M_0) \mathcal{T}_0 \|\hat{u}^{*,1} - \hat{u}^{*,2}\|_{\mathcal{V}_{\mathcal{T}_0}} \\
 & \leq \frac{1}{2} \|\hat{u}^{*,1} - \hat{u}^{*,2}\|_{\mathcal{V}_{\mathcal{T}_0}} \leq \frac{1}{2} \left(\|u^{*,1} - u^{*,2}\|_{\mathcal{V}_{\mathcal{T}_0}} + \|\hat{u}^{*,1} - \hat{u}^{*,2}\|_{\mathcal{V}_{\mathcal{T}_0}} \right),
 \end{aligned} \tag{224}$$

showing that the map F is a contraction.

Step 2d(iii): Banach fixed point theorem. Hence, the Banach fixed point theorem guarantees that there exists a unique solution $(u^*, \hat{u}^*) \in \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$, which satisfies $(u^*, \hat{u}^*) = F(u^*, \hat{u}^*)$. We have thus established the existence of a unique local-in-training-time solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) for $\tilde{g}_\tau = g_\tau^*$ on the training time domain $[0, \mathcal{T}_0]$. In particular, $(u_\tau^*, \hat{u}_\tau^*) \in \mathcal{S} \times \mathcal{S}$ for every $\tau \in [0, \mathcal{T}_0]$. Reapplying the classical existence and regularity results from *Steps 1a, 1b* and *1c* further shows that for each $\tau \in [0, \mathcal{T}]$ such solution satisfies $(\partial_t \tilde{u}_\tau^*(t, \bullet), \partial_t \hat{u}_\tau^*(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$.

Step 2e: Existence globally in training time. Leveraging a bootstrapping argument, we now extend this argument to obtain a solution on a training time domain $[0, \mathcal{T}]$ for an arbitrary $\mathcal{T} < \infty$. To do so, we proceed inductively. Suppose we have a solution $(u^*, \hat{u}^*) \in \mathcal{V}_{\mathcal{T}_{k-1}}(M_{k-1}) \times \mathcal{V}_{\mathcal{T}_{k-1}}(M_{k-1})$ which is such that for each $\tau \in [0, \mathcal{T}_{k-1}]$ it satisfies $(\partial_t u_\tau^*(t, \bullet), \partial_t \hat{u}_\tau^*(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$. (We showed in *Step 2d* before that this is the case for the induction start $k = 1$.)

On the training time interval $I = [0, \mathcal{T}_{k-1}]$ we can now employ Proposition 26 which ensures that $\frac{d}{d\tau} \mathcal{J}_\tau^* \leq 0$ for all $\tau \in I = [0, \mathcal{T}_{k-1}]$. Thanks to this, Lemma 28 (applied in the setting $I = [0, \mathcal{T}_{k-1}]$) provides a uniform (in the training time τ and on the training time interval $[0, \mathcal{T}_{k-1}]$) bound $\sup_{\tau \in [0, \mathcal{T}_{k-1}]} \|\hat{u}_\tau^*\|_{L_2(D_T)} \leq C^{\hat{u}}$, where $C^{\hat{u}}$ does not depend on \mathcal{T}_{k-1} but only on \mathcal{J}_0^* . Let us now choose

$$M_k = kC^{\hat{u}} + 2C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + 1 \right) \quad \text{and} \quad \mathcal{T}_k = \mathcal{T}_{k-1} + \min \left\{ \frac{1}{4C_1}, \frac{1}{2C_2(1 + M_k)} \right\}. \tag{225}$$

We show in what follows that there exists a unique solution $(u^*, \hat{u}^*) \in \mathcal{V}_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k)$.

Step 2e(i): Self-mapping property of F . Consider (\tilde{u}^*, \hat{u}^*) together with its corresponding $(u^*, \hat{u}^*) \in \mathcal{V}_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k)$. Using the definitions of M_k and \mathcal{T}_k in (225), we can derive from (194) that

$$\|\tilde{u}^*\|_{\mathcal{V}_{\mathcal{T}_k}} + \|\hat{u}^*\|_{\mathcal{V}_{\mathcal{T}_k}} \leq C_1 \int_0^{\mathcal{T}_k} \|\hat{u}_s^*\|_{L_2(D_T)} ds + C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + 1 \right)$$

$$\begin{aligned}
 &\leq C_1 \int_0^{\mathcal{T}_{k-1}} \|\widehat{u}_s^*\|_{L_2(D_T)} ds + C_1 \int_{\mathcal{T}_{k-1}}^{\mathcal{T}_k} \|\widehat{u}_s^*\|_{L_2(D_T)} ds + C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + 1 \right) \\
 &\leq C_1 \mathcal{T}_{k-1} C^{\widehat{u}} + C_1 (\mathcal{T}_k - \mathcal{T}_{k-1}) \|\widehat{u}^*\|_{\mathcal{V}_{\mathcal{T}_k}} + \frac{M_k}{2} \\
 &\leq C_1 \mathcal{T}_{k-1} C^{\widehat{u}} + C_1 (\mathcal{T}_k - \mathcal{T}_{k-1}) M_k + \frac{M_k}{2} \leq \frac{M_k}{4} + \frac{M_k}{4} + \frac{M_k}{2} \leq M_k,
 \end{aligned}$$

where we used for the first step in the last line that (with $\mathcal{T}_{-1} := 0$) according to the definition of \mathcal{T}_k in (225) it holds $C_1 \mathcal{T}_{k-1} C^{\widehat{u}} = C_1 \sum_{\ell=0}^{k-1} (\mathcal{T}_\ell - \mathcal{T}_{\ell-1}) C^{\widehat{u}} \leq \frac{1}{4} k C^{\widehat{u}} \leq \frac{M_k}{4}$. Thus, $(\widetilde{u}^*, \widehat{u}^*) \in \mathcal{V}_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k)$.

Step 2e(ii): Contractivity of F . Consider two pairs $(\widetilde{u}^{*,1}, \widehat{u}^{*,1})$, $(\widetilde{u}^{*,2}, \widehat{u}^{*,2})$ together with their corresponding $(u^{*,1}, \widehat{u}^{*,1})$, $(u^{*,2}, \widehat{u}^{*,2}) \in \mathcal{V}_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k)$. Using the definitions of M_k and \mathcal{T}_k in (225), we can derive from (221) that

$$\begin{aligned}
 &\|\widetilde{u}^{*,1} - \widetilde{u}^{*,2}\|_{\mathcal{V}_{\mathcal{T}_k}} + \|\widehat{u}^{*,1} - \widehat{u}^{*,2}\|_{\mathcal{V}_{\mathcal{T}_k}} \\
 &\leq C_2 \left(1 + \sup_{\tau \in [0, \mathcal{T}_k]} \|\widehat{u}_\tau^{*,2}\|_{L_\infty([0, T], L_2(D))} \right) \int_0^{\mathcal{T}_k} \|\widehat{u}_s^{*,1} - \widehat{u}_s^{*,2}\|_{L_2(D_T)} ds \\
 &\leq C_2 (1 + M_k) \int_0^{\mathcal{T}_k} \|\widehat{u}_s^{*,1} - \widehat{u}_s^{*,2}\|_{L_2(D_T)} ds \\
 &= C_2 (1 + M_k) \int_{\mathcal{T}_{k-1}}^{\mathcal{T}_k} \|\widehat{u}_s^{*,1} - \widehat{u}_s^{*,2}\|_{L_2(D_T)} ds \leq C_2 (1 + M_k) (\mathcal{T}_k - \mathcal{T}_{k-1}) \|\widehat{u}^{*,1} - \widehat{u}^{*,2}\|_{\mathcal{V}_{\mathcal{T}_k}} \\
 &\leq \frac{1}{2} \|\widehat{u}^{*,1} - \widehat{u}^{*,2}\|_{\mathcal{V}_{\mathcal{T}_k}} \leq \frac{1}{2} \left(\|u^{*,1} - u^{*,2}\|_{\mathcal{V}_{\mathcal{T}_k}} + \|\widehat{u}^{*,1} - \widehat{u}^{*,2}\|_{\mathcal{V}_{\mathcal{T}_k}} \right),
 \end{aligned} \tag{226}$$

where the third step is due to the uniqueness of the solution on the training time interval $[0, \mathcal{T}_{k-1}]$. Thus, the map F is a contraction.

Step 2e(iii): Banach fixed point theorem. Hence, the Banach fixed point theorem guarantees that there exists a unique solution $(u^*, \widehat{u}^*) \in \mathcal{V}_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k)$, which satisfies $(u^*, \widehat{u}^*) = F(u^*, \widehat{u}^*)$. We have thus established the existence of a unique solution to the PDE system (16)–(17) coupled with the integro-differential equation (15) for $\widetilde{g}_\tau = g_\tau^*$ on the training time domain $[0, \mathcal{T}_k]$. In particular, $(u_\tau^*, \widehat{u}_\tau^*) \in \mathcal{S} \times \mathcal{S}$ for every $\tau \in [0, \mathcal{T}_k]$. Reapplying the classical existence and regularity results from *Steps 1a, 1b* and *1c* further shows that for each $\tau \in [0, \mathcal{T}_k]$ such solution satisfies $(\partial_t \widetilde{u}_\tau^*(t, \bullet), \partial_t \widehat{u}_\tau^*(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$.

Step 2e(iv): Globality of the construction in training time. It remains to notice that, due to the definition of the times \mathcal{T}_k in (225), the telescopic sum

$$\begin{aligned}
 \sum_{k=1}^{\infty} (\mathcal{T}_k - \mathcal{T}_{k-1}) &= \sum_{k=1}^{\infty} \min \left\{ \frac{1}{4C_1}, \frac{1}{2C_2(1 + M_k)} \right\} \geq \sum_{k=K}^{\infty} \frac{1}{2C_2(1 + M_k)} \\
 &= \sum_{k=K}^{\infty} \frac{1}{2C_2(1 + kC^{\widehat{u}} + 2C_1(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + 1))}
 \end{aligned} \tag{227}$$

diverges. To see this, simply note that the definition of the bounds M_k in (225) grows linearly in k , thus the last term in (227) being a harmonic series for some sufficiently large integer K . This ensures that the above construction in *Step 2d* and *2e* gives a solution in $\mathcal{V}_{\mathcal{T}}(M) \times \mathcal{V}_{\mathcal{T}}(M)$ for any given $\mathcal{T} < \infty$ and suitable associated $M > 0$.

Uniqueness. It remains to prove the uniqueness of a solution $(u_{\tau}^*, \widehat{u}_{\tau}^*)$ to the PDE system (16)–(17) coupled with the integro-differential equation (15) for g_{τ}^* . For this purpose, suppose that there are two weak solutions $(u^{*,1}, \widehat{u}^{*,1}), (u^{*,2}, \widehat{u}^{*,2}) \in \mathcal{V}_{\mathcal{T}}(M) \times \mathcal{V}_{\mathcal{T}}(M)$. This means we have $(u_{\tau}^{*,1}, \widehat{u}_{\tau}^{*,1}), (u_{\tau}^{*,2}, \widehat{u}_{\tau}^{*,2}) \in \mathcal{S} \times \mathcal{S}$ with satisfy $(\partial_t u_{\tau}^{*,1}, \partial_t \widehat{u}_{\tau}^{*,1}), (\partial_t u_{\tau}^{*,2}, \partial_t \widehat{u}_{\tau}^{*,2}) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$ and where $g_{\tau}^{*,1}, g_{\tau}^{*,2} \in L_2(D_T)$ denote the corresponding integral terms (15) for each $\tau \in [0, \mathcal{T}]$. By repeating the computations of the existence proof in *Step 2c(ii)* we obtain analogously to (202) and (207) that

$$\|u_{\tau}^{*,1} - u_{\tau}^{*,2}\|_{L_{\infty}([0,T],H^1(D))} + \|u_{\tau}^{*,1} - u_{\tau}^{*,2}\|_{L_2([0,T],L_2(D))} \leq C \|g_{\tau}^{*,1} - g_{\tau}^{*,2}\|_{L_2(D_T)} \quad (228)$$

and for some $p > d + 1$ that

$$\|u_{\tau}^{*,1} - u_{\tau}^{*,2}\|_{L_{\infty}(D_T)} \leq C \|g_{\tau}^{*,1} - g_{\tau}^{*,2}\|_{L_p(D_T)}, \quad (229)$$

as well as analogously to (219) that

$$\begin{aligned} & \|\widehat{u}_{\tau}^{*,1} - \widehat{u}_{\tau}^{*,2}\|_{L_2([0,T],H^1(D))} + \|\widehat{u}_{\tau}^{*,1} - \widehat{u}_{\tau}^{*,2}\|_{L_{\infty}([0,T],L_2(D))} \\ & \leq C \left(\|u_{\tau}^{*,1} - u_{\tau}^{*,2}\|_{L_2(D_T)} + M \|u_{\tau}^{*,1} - u_{\tau}^{*,2}\|_{L_{\infty}(D_T)} \right), \end{aligned} \quad (230)$$

where we used directly that $\widehat{u}_{\tau}^{*,2} \in \mathcal{V}_{\mathcal{T}}(M)$. Since it holds

$$\|g_{\tau}^{*,1} - g_{\tau}^{*,2}\|_{L_p(D_T)} \leq C \int_0^{\tau} \|\widehat{u}_s^{*,1} - \widehat{u}_s^{*,2}\|_{L_2(D_T)} ds, \quad (231)$$

according to the next-to-last step in (205), we get

$$\begin{aligned} & \|g_{\tau}^{*,1} - g_{\tau}^{*,2}\|_{L_2(D_T)} + \|g_s^{*,1} - g_s^{*,2}\|_{L_p(D_T)} \leq C \int_0^{\tau} \|\widehat{u}_s^{*,1} - \widehat{u}_s^{*,2}\|_{L_2(D_T)} ds \\ & \leq C \int_0^{\tau} \|u_s^{*,1} - u_s^{*,2}\|_{L_2(D_T)} + \|u_s^{*,1} - u_s^{*,2}\|_{L_{\infty}(D_T)} ds \\ & \leq C \int_0^{\tau} \|g_s^{*,1} - g_s^{*,2}\|_{L_2(D_T)} + \|g_s^{*,1} - g_s^{*,2}\|_{L_p(D_T)} ds, \end{aligned} \quad (232)$$

where C may depend in particular on M .

Recalling that $g_0^1 = g_0^2 = 0$, we can now employ Grönwall's inequality in its integral form to obtain

$$\|g_{\tau}^{*,1} - g_{\tau}^{*,2}\|_{L_2(D_T)} + \|g_{\tau}^{*,1} - g_{\tau}^{*,2}\|_{L_p(D_T)} = 0 \quad (233)$$

for every $\tau \in [0, \mathcal{T}]$. With (228) and (230) we hence conclude that for every $\tau \in [0, \mathcal{T}]$ it hold

$$\|u_{\tau}^{*,1} - u_{\tau}^{*,2}\|_{L_2([0,T],H^1(D))} + \|u_{\tau}^{*,1} - u_{\tau}^{*,2}\|_{L_{\infty}([0,T],L_2(D))} = 0 \quad (234)$$

and

$$\|\widehat{u}_{\tau}^{*,1} - \widehat{u}_{\tau}^{*,2}\|_{L_2([0,T],H^1(D))} + \|\widehat{u}_{\tau}^{*,1} - \widehat{u}_{\tau}^{*,2}\|_{L_{\infty}([0,T],L_2(D))} = 0. \quad (235)$$

Thus $\|u^{*,1} - u^{*,2}\|_{\mathcal{V}_{\mathcal{T}}} = 0$ and $\|\widehat{u}^{*,1} - \widehat{u}^{*,2}\|_{\mathcal{V}_{\mathcal{T}}} = 0$, proving uniqueness in $\mathcal{V}_{\mathcal{T}}(M)$. \blacksquare

A.2 Well-Posedness Proof of the NN-PDE Training Dynamics in the Finite-Width Hidden Layer Regime

Proof [Proof of Lemma 11] **Existence.** As in the proof of Lemma 13, existence is shown using a fixed point argument. We denote for a given training time horizon $\mathcal{T} > 0$ by $\Theta_{\mathcal{T}} = \mathcal{C}([0, \mathcal{T}], \Theta := (\mathbb{R}^N \times (L_2(D_T))^N \times (L_2(D_T))^N)$ the Banach space consisting of elements with finite norm

$$\|(c, n, m)\|_{\Theta_{\mathcal{T}}} = \sup_{\tau \in [0, \mathcal{T}]} \sum_{i=1}^N \left(|c_{\tau}^i| + \|n_{\tau}^i\|_{L_2(D_T)} + \|m_{\tau}^i\|_{L_2(D_T)} \right), \quad (236)$$

where $n_{\tau}^i(t, x) = \sigma(w_{\tau}^{t,i}t + (w_{\tau}^i)^T x + \eta_{\tau}^i)$ and $m_{\tau}^i(t, x) = \sigma'(w_{\tau}^{t,i}t + (w_{\tau}^i)^T x + \eta_{\tau}^i)$, and by $\mathcal{V}_{\mathcal{T}} = \mathcal{C}([0, \mathcal{T}], \mathcal{S})$ again the Banach space consisting of elements with finite norm $\|\cdot\|_{\mathcal{V}_{\mathcal{T}}}$ as defined in (167).

A solution $((c_{\tau}, n_{\tau}, m_{\tau}), u_{\theta_{\tau}}^N, \hat{u}_{\theta_{\tau}}^N)_{\tau \in [0, \mathcal{T}]}$ to the PDE system (1) & (8) coupled with the gradient descent update (5) within the above definitions is shown in what follows to be an element of the space $\mathcal{C}([0, \mathcal{T}], \Theta \times \mathcal{S} \times \mathcal{S})$ (which we identify with the space $\Theta_{\mathcal{T}} \times \mathcal{V}_{\mathcal{T}} \times \mathcal{V}_{\mathcal{T}}$) with additional regularity.

Step 1: Existence and regularity for given NN parameter updates $\tilde{c}_{\tau}^i = c_0^i - \int_0^{\tau} \alpha_s b_{\theta_s}^{c^i} ds$ etc. For given $\mathcal{T} > 0$, let $b_{\theta}^{c^i}, b_{\theta}^{w^{t,i}}, b_{\theta}^{\eta^i} : [0, \mathcal{T}] \rightarrow \mathbb{R}$ and $b_{\theta}^{w^i} : [0, \mathcal{T}] \rightarrow \mathbb{R}^d$ be given functions with $b_{\theta}^{c^i}, b_{\theta}^{w^{t,i}}, b_{\theta}^{w^i}, b_{\theta}^{\eta^i}$ being such that $\sup_{\tau \in [0, \mathcal{T}]} \sum_{i=1}^N (|b_{\theta_{\tau}}^{c^i}| + |b_{\theta_{\tau}}^{w^{t,i}}| + \|b_{\theta_{\tau}}^{w^i}\| + |b_{\theta_{\tau}}^{\eta^i}|) \leq C_b$, where C_b may depend in particular on \mathcal{T} . Consider the auxiliary NN parameter update

$$\tilde{c}_{\tau}^i = c_0^i - \int_0^{\tau} \alpha_s b_{\theta_s}^{c^i} ds, \quad (237a)$$

$$\tilde{w}_{\tau}^{t,i} = w_0^{t,i} - \int_0^{\tau} \alpha_s b_{\theta_s}^{w^{t,i}} ds, \quad (237b)$$

$$\tilde{w}_{\tau}^i = w_0^i - \int_0^{\tau} \alpha_s b_{\theta_s}^{w^i} ds, \quad (237c)$$

$$\tilde{\eta}_{\tau}^i = \eta_0^i - \int_0^{\tau} \alpha_s b_{\theta_s}^{\eta^i} ds, \quad (237d)$$

as well as $\tilde{n}_{\tau}^i(t, x) = \sigma(\tilde{w}_{\tau}^{t,i}t + (\tilde{w}_{\tau}^i)^T x + \tilde{\eta}_{\tau}^i)$ and $\tilde{m}_{\tau}^i(t, x) = \sigma'(\tilde{w}_{\tau}^{t,i}t + (\tilde{w}_{\tau}^i)^T x + \tilde{\eta}_{\tau}^i)$ together with the auxiliary PDE system

$$\begin{aligned} \partial_t \tilde{u}_{\theta_{\tau}}^N + \mathcal{L} \tilde{u}_{\theta_{\tau}}^N - q(\tilde{u}_{\theta_{\tau}}^N) &= \tilde{g}_{\theta_{\tau}}^N && \text{in } D_T, \\ \tilde{u}_{\theta_{\tau}}^N &= 0 && \text{on } [0, T] \times \partial D, \\ \tilde{u}_{\theta_{\tau}}^N &= f && \text{on } \{0\} \times D, \end{aligned} \quad (238)$$

where $\tilde{g}_{\theta_{\tau}}^N(t, x) = \frac{1}{N^{\beta}} \sum_{i=1}^N \tilde{c}_{\tau}^i \tilde{n}_{\tau}^i(t, x) = \frac{1}{N^{\beta}} \sum_{i=1}^N \tilde{c}_{\tau}^i \sigma(\tilde{w}_{\tau}^{t,i}t + (\tilde{w}_{\tau}^i)^T x + \tilde{\eta}_{\tau}^i)$, and

$$\begin{aligned} -\partial_t \hat{u}_{\theta_{\tau}}^N + \mathcal{L}^{\dagger} \hat{u}_{\theta_{\tau}}^N - q_u(\hat{u}_{\theta_{\tau}}^*) \hat{u}_{\theta_{\tau}}^N &= (\hat{u}_{\theta_{\tau}}^N - h) && \text{in } D_T, \\ \hat{u}_{\theta_{\tau}}^N &= 0 && \text{on } [0, T] \times \partial D, \\ \hat{u}_{\theta_{\tau}}^N &= 0 && \text{on } \{T\} \times D. \end{aligned} \quad (239)$$

We first prove that there exists a solution $(\tilde{u}_{\theta_\tau}^N, \widehat{u}_{\theta_\tau}^N) \in \mathcal{S} \times \mathcal{S}$ to the system (238)–(239) for all $\tau \in [0, \mathcal{T}]$ using classical existence results from (Ladyženskaja et al., 1968). Such solution, as we show, enjoys the property that for all $\tau \in [0, \mathcal{T}]$ it holds $(\partial_t \tilde{u}_{\theta_\tau}^N(t, \bullet), \partial_t \widehat{u}_{\theta_\tau}^N(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$.

Computing $\frac{d}{d\tau} \tilde{g}_{\theta_\tau}^N$ by taking the training time derivative in $\tilde{g}_{\theta_\tau}^N$ as defined above and combining it with the expressions (237), we obtain by the fundamental theorem of calculus that

$$\tilde{g}_{\theta_\tau}^N = g_{\theta_0}^N - \int_0^\tau \alpha_s b_s^N ds = g_{\theta_0}^N - \int_0^\tau \frac{\alpha_s}{N^\beta} \sum_{i=1}^N b_{\theta_s^i}^{\sigma^i}(\tilde{\star}) + \tilde{c}_s^i \sigma'(\tilde{\star}) \left(b_{\theta_s^i}^{w^{t,i}} t + (b_{\theta_s^i}^{w^i})^T x + b_{\theta_s^i}^{\eta^i} \right) ds, \quad (240)$$

where we abbreviated $\tilde{\star} = \tilde{w}_s^{t,i} t + (\tilde{w}_s^i)^T x + \tilde{\eta}_s^i$ and defined b^N implicitly. We now notice that $b^N : [0, \mathcal{T}] \rightarrow L_2(D_T)$ is such that b_τ^N is Lipschitz continuous on $\overline{D_T}$ for each $\tau \in [0, \mathcal{T}]$ and such that $\sup_{\tau \in [0, \mathcal{T}]} \|b_\tau^N\|_{L_\infty(D_T)} \leq C_{b^N}$, where C_{b^N} may depend in particular on \mathcal{T} .

For the latter, uniform boundedness, we note that $b_{\theta_s^i}^{\sigma^i}, b_{\theta_s^i}^{w^{t,i}}, b_{\theta_s^i}^{w^i}, b_{\theta_s^i}^{\eta^i}$ are bounded by C_b , that σ and σ' are bounded as of Assumptions B1 and B2, that the domain D is bounded as of Assumption A2, and that by Cauchy-Schwarz inequality it holds with (237a) that $|\tilde{c}_\tau^i| \leq |c_0^i| + \left| \int_0^\tau \alpha_s b_{\theta_s^i}^{\sigma^i} ds \right| \leq |c_0^i| + \int_0^\tau \alpha_s^2 ds \int_0^\tau |b_{\theta_s^i}^{\sigma^i}|^2 ds$ is bounded as of Assumption B3(ii) and due to condition (6) on the learning rate α_τ . Since σ and σ' are further Lipschitz continuous as of Assumptions B1 and B2, and since also $\tilde{w}_\tau^{t,i}, \tilde{w}_\tau^i$ and $\tilde{\eta}_\tau^i$ are bounded with the above argument and for given initial conditions $w_0^{t,i}, w_0^i$ and η_0^i , it is straightforward to check that b_τ^N is Lipschitz continuous on $\overline{D_T}$ for each $\tau \in [0, \mathcal{T}]$. We can thus follow *Steps 1a* and *b* in the proof of Lemma 13 to show that there exists a solution $(\tilde{u}_{\theta_\tau}^N, \widehat{u}_{\theta_\tau}^N) \in \mathcal{S} \times \mathcal{S}$ for all $\tau \in [0, \mathcal{T}]$, which enjoys the property that for all $\tau \in [0, \mathcal{T}]$ it holds $(\partial_t \tilde{u}_{\theta_\tau}^N(t, \bullet), \partial_t \widehat{u}_{\theta_\tau}^N(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$. For this, one only needs to additionally notice that, by definition (3), $g_{\theta_0}^N$ is continuous w.r.t. t, x , uniformly bounded due to Assumptions B1 and B3(ii), i.e., $\|g_{\theta_0}^N\|_{L_\infty(D_T)} \leq C$ for a constant $C = C(N, \sigma, \mu_0)$, which may depend on N (since it is fixed throughout the proof), and (1, 1)-Hölder continuous due to Assumptions B1.

Repeating further the energy estimates of *Step 1c* in the proof of Lemma 13 in the above setting, we can derive the estimates

$$\|\tilde{u}_{\theta_\tau}^N\|_{L_2([0, T], H^1(D))} + \|\tilde{u}_{\theta_\tau}^N\|_{L_\infty([0, T], L_2(D))} \leq C \left(\|f\|_{L_2(D)} + \|\tilde{g}_{\theta_\tau}^N\|_{L_2(D_T)} + 1 \right) \quad (241)$$

and, in particular,

$$\|\tilde{u}_\theta^N\|_{\mathcal{V}_\tau} \leq C \left(\|f\|_{L_2(D)} + \sup_{\tau \in [0, \mathcal{T}]} \|\tilde{g}_{\theta_\tau}^N\|_{L_2(D_T)} + 1 \right) \quad (242)$$

for some constant $C = C(T, \mathcal{L}, q)$, as well as

$$\begin{aligned} \|\widehat{u}_\tau^N\|_{L_2([0, T], H^1(D))} + \|\widehat{u}_\tau^N\|_{L_\infty([0, T], L_2(D))} &\leq C \|\widehat{u}_{\theta_\tau}^N - h\|_{L_2(D_T)} \\ &\leq C \left(\|\tilde{u}_{\theta_\tau}^N\|_{L_2(D_T)} + \|h\|_{L_2(D_T)} \right) \end{aligned} \quad (243)$$

and, in particular,

$$\|\widehat{u}_\theta^N\|_{\mathcal{V}_T} \leq C \left(\|\widetilde{u}_\theta^N\|_{\mathcal{V}_T} + \|h\|_{L_2(D_T)} \right) \quad (244)$$

for some other, potentially larger, constant $C = C(T, \mathcal{L}, q)$.

Since $\widetilde{g}_{\theta_\tau}^N$ is Lipschitz continuous in τ with the argument from *Step 1d* in the proof of Lemma 13 together with the formerly established uniform boundedness of b^N in (240), we proved that there exists a continuous in the training time τ solution $(\widetilde{u}^*, \widehat{u}^*) \in \mathcal{V}_T \times \mathcal{V}_T$ to (238)–(239). As we further ensured, for each $\tau \in [0, \mathcal{T}]$ such solution satisfies $(\partial_t \widetilde{u}_\tau^*(t, \bullet), \partial_t \widehat{u}_\tau^*(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$.

Step 2: Existence for specific NN parameter updates $\widetilde{c}_\tau^i = c_0^i - \frac{1}{N^{1-\beta}} \int_0^\tau \alpha_s \int_0^T \int_D n_s^i(t, x) \widehat{u}_{\theta_s}^N(t, x) dx dt ds$ etc. We now make specific choices for the functions $b_{\theta_\tau}^{c^i}, b_{\theta_\tau}^{w^{t,i}}, b_{\theta_\tau}^{w^i}, b_{\theta_\tau}^{\eta^i}$.

Step 2a: Choice of NN parameter update functions $b_{\theta_\tau}^{c^i} = \frac{1}{N^{1-\beta}} \int_0^T \int_D n_\tau^i(t, x) \widehat{u}_{\theta_\tau}^N(t, x) dx dt$ etc. For arbitrarily given $(c_\tau, n_\tau, m_\tau) \in \Theta$ and $\widehat{u}_{\theta_\tau}^N \in \mathcal{S}$, $\tau \in [0, \mathcal{T}]$, with $\sup_{\tau \in [0, \mathcal{T}]} \sum_{i=1}^N (|c_\tau^i| + \|n_\tau^i\|_{L_2(D_T)} + \|m_\tau^i\|_{L_2(D_T)}) \leq M$ and $\sup_{\tau \in [0, \mathcal{T}]} \|\widehat{u}_{\theta_\tau}^N\|_{L_2(D_T)} \leq M$ (M may depend on \mathcal{T}), we set

$$b_{\theta_\tau}^{c^i} = \frac{1}{N^{1-\beta}} \int_0^T \int_D n_\tau^i(t, x) \widehat{u}_{\theta_\tau}^N(t, x) dx dt, \quad (245a)$$

$$b_{\theta_\tau}^{w^{t,i}} = \frac{1}{N^{1-\beta}} \int_0^T \int_D c_\tau^i m_\tau^i(t, x) t \widehat{u}_{\theta_\tau}^N(t, x) dx dt \quad (245b)$$

$$b_{\theta_\tau}^{w^i} = \frac{1}{N^{1-\beta}} \int_0^T \int_D c_\tau^i m_\tau^i(t, x) x \widehat{u}_{\theta_\tau}^N(t, x) dx dt, \quad (245c)$$

$$b_{\theta_\tau}^{\eta^i} = \frac{1}{N^{1-\beta}} \int_0^T \int_D c_\tau^i m_\tau^i(t, x) \widehat{u}_{\theta_\tau}^N(t, x) dx dt, \quad (245d)$$

for all $\tau \in [0, \mathcal{T}]$. By Cauchy-Schwarz inequality it holds

$$\begin{aligned} & \sum_{i=1}^N \left(|b_{\theta_\tau}^{c^i}| + |b_{\theta_\tau}^{w^{t,i}}| + \|b_{\theta_\tau}^{w^i}\| + |b_{\theta_\tau}^{\eta^i}| \right) \\ & \leq C \sum_{i=1}^N \left(\|n_\tau^i\|_{L_2(D_T)} \|\widehat{u}_{\theta_\tau}^N\|_{L_2(D_T)} + |c_\tau^i| \|m_\tau^i\|_{L_2(D_T)} \|\widehat{u}_{\theta_\tau}^N\|_{L_2(D_T)} \right) \\ & \leq C (M^2 + M^3) \end{aligned} \quad (246)$$

for a constant $C = C(N, T, D)$, which is a uniform bound in τ . Since the right-hand side is uniform in τ , $\sup_{\tau \in [0, \mathcal{T}]} \sum_{i=1}^N (|b_{\theta_\tau}^{c^i}| + |b_{\theta_\tau}^{w^{t,i}}| + \|b_{\theta_\tau}^{w^i}\| + |b_{\theta_\tau}^{\eta^i}|) \leq C_{b^N}$, where C_{b^N} may depend on \mathcal{T} .

Step 2b: Definition of fixed point mapping. Let us consider the fixed point map

$$F : \Theta_{\mathcal{T}} \times \mathcal{V}_{\mathcal{T}} \times \mathcal{V}_{\mathcal{T}} \rightarrow \Theta_{\mathcal{T}} \times \mathcal{V}_{\mathcal{T}} \times \mathcal{V}_{\mathcal{T}}, \quad ((c, n, m), u^N, \widehat{u}^N) \mapsto ((\widetilde{c}, \widetilde{n}, \widetilde{m}), \widetilde{u}^N, \widehat{u}^N) \quad (247)$$

and define for given $M < \infty$ and $\mathcal{T} < \infty$ the function spaces $\Theta_{\mathcal{T}}(M) = \{(c, n, m) \in \Theta_{\mathcal{T}} : \|(c, n, m)\|_{\Theta_{\mathcal{T}}} \leq M\}$ and $\mathcal{V}_{\mathcal{T}}(M) = \{u \in \mathcal{V}_{\mathcal{T}} : \|u\|_{\mathcal{V}_{\mathcal{T}}} \leq M\}$.

We will first show in *Step 2d* existence locally in the training time by proving that there exist $M_0 > 0$ and $\mathcal{T}_0 > 0$ such that F is a fixed point mapping on $\Theta_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$, which allows to apply the Banach fixed point theorem. In *Step 2e* we will then extend the proof by a bootstrapping argument to any given (arbitrarily large) time horizon \mathcal{T} .

Step 2c: Preliminary computations. Let us start by conducting some preliminary computations on a generic space $\Theta_{\tilde{\mathcal{T}}}(\tilde{M}) \times \mathcal{V}_{\tilde{\mathcal{T}}}(\tilde{M}) \times \mathcal{V}_{\tilde{\mathcal{T}}}(\tilde{M})$.

Step 2c(i): Preliminary computations for self-mapping property of F . Consider the triple $((\tilde{c}, \tilde{n}, \tilde{m}), \tilde{u}^N, \tilde{\hat{u}}^N)$ together with its corresponding $((c, n, m), u^N, \hat{u}^N) \in \Theta_{\tilde{\mathcal{T}}}(\tilde{M}) \times \mathcal{V}_{\tilde{\mathcal{T}}}(\tilde{M}) \times \mathcal{V}_{\tilde{\mathcal{T}}}(\tilde{M})$.

A bound for $|\tilde{c}_\tau^i|$. Recalling (237a), we can estimate with triangle inequality and two applications of Cauchy-Schwarz inequality, while using the boundedness of σ as of Assumptions B1 and that the domain D has bounded volume as of Assumption A2 together with the additional (and with (6) compatible) assumption $\int_0^\infty \alpha_\tau^{4/3} d\tau < \infty$ on the learning rate, that

$$\begin{aligned}
 |\tilde{c}_\tau^i| &\leq |c_0^i| + C \left| \int_0^\tau \alpha_s \int_0^T \int_D \sigma(w_s^{t,i}t + (w_s^i)^T x + \eta_s^i) \hat{u}_{\theta_s^N}^N(t, x) dx dt ds \right| \\
 &\leq |c_0^i| + C \left| \int_0^\tau \alpha_s \sqrt{\int_0^T \int_D (\sigma(w_s^{t,i}t + (w_s^i)^T x + \eta_s^i))^2 dx dt} \|\hat{u}_{\theta_s^N}^N\|_{L_2(D_T)} ds \right| \\
 &\leq |c_0^i| + C \left| \int_0^\tau \alpha_s \|\hat{u}_{\theta_s^N}^N\|_{L_2(D_T)} ds \right| \\
 &\leq |c_0^i| + C \left(\int_0^\tau \alpha_s^{4/3} ds \right)^{3/4} \left(\int_0^\tau \|\hat{u}_{\theta_s^N}^N\|_{L_2(D_T)}^4 ds \right)^{1/4} \\
 &\leq |c_0^i| + C \left(\int_0^\tau \|\hat{u}_{\theta_s^N}^N\|_{L_2(D_T)}^4 ds \right)^{1/4}
 \end{aligned} \tag{248}$$

for a constant $C = C(\alpha, N, T, D, \mathcal{L}, q, \sigma)$.

A bound for $\|\tilde{n}_\tau^i\|_{L_2(D_T)}$. Using the boundedness of σ as of Assumptions B1 and that the domain D has bounded volume as of Assumption A2, clearly,

$$\|\tilde{n}_\tau^i\|_{L_2(D_T)} \leq C \tag{249}$$

for a constant $C = C(T, D, \sigma)$.

A bound for $\|\tilde{m}_\tau^i\|_{L_2(D_T)}$. Since also σ' is bounded as of Assumptions B2,

$$\|\tilde{m}_\tau^i\|_{L_2(D_T)} \leq C. \tag{250}$$

A bound for $\|\tilde{g}_{\theta_\tau^N}^N\|_{L_2(D_T)}$. Using that σ is bounded as of Assumption B1 to obtain the first inequality, that the domain D has bounded volume as of Assumption A2 to get the

equality in the second line, and (248) in the last step, we can upper bound

$$\begin{aligned}
 \|\tilde{g}_{\tilde{\theta}_\tau}^N\|_{L_2(D_T)} &= \left\| \frac{1}{N^\beta} \sum_{i=1}^N \tilde{c}_\tau^i \sigma(\tilde{w}_\tau^{t,i} t + (\tilde{w}_\tau^i)^T x + \tilde{\eta}_\tau^i) \right\|_{L_2(D_T)} \\
 &\leq C \left\| \sum_{i=1}^N |\tilde{c}_\tau^i| \right\|_{L_2(D_T)} = C \sum_{i=1}^N |\tilde{c}_\tau^i| \\
 &\leq \|\theta_0\| + C \left(\int_0^\tau \|\hat{u}_{\theta_s}^N\|_{L_2(D_T)}^4 ds \right)^{1/4}
 \end{aligned} \tag{251}$$

for a constant $C = C(N, T, D, \sigma)$.

Combination of the bounds. Using (242) and (244) in the first inequality, together with (248), (249) and (251) in the second step, we establish (under the additional assumption $\int_0^\infty \alpha_\tau^{4/3} d\tau < \infty$ on the learning rate)

$$\begin{aligned}
 &\|(\tilde{c}, \tilde{n}, \tilde{m})\|_{\Theta_{\tilde{\tau}}} + \|\tilde{u}_\theta^N\|_{\mathcal{V}_{\tilde{\tau}}} + \|\hat{u}_\theta^N\|_{\mathcal{V}_{\tilde{\tau}}} \\
 &\leq \|(\tilde{c}, \tilde{n}, \tilde{m})\|_{\Theta_{\tilde{\tau}}} + C \sup_{\tau \in [0, \tilde{\tau}]} \|\tilde{g}_{\tilde{\theta}_\tau}^N\|_{L_2(D_T)} + C \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + 1 \right) \\
 &\leq C_1 \sup_{\tau \in [0, \tilde{\tau}]} \left(\int_0^\tau \|\hat{u}_{\theta_s}^N\|_{L_2(D_T)}^4 ds \right)^{1/4} + C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + \|\theta_0\| + 1 \right) \\
 &\leq C_1 \left(\int_0^{\tilde{\tau}} \|\hat{u}_{\theta_s}^N\|_{L_2(D_T)}^4 ds \right)^{1/4} + C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + \|\theta_0\| + 1 \right)
 \end{aligned} \tag{252}$$

for a constant $C_1 = C_1(\alpha, N, T, D, \mathcal{L}, q, \sigma)$. (Note that the bound on the right-hand side only grows like $\tilde{\mathcal{T}}^{1/2}$ instead of $\tilde{\mathcal{T}}^{1/4}$ due to the additional assumption on the learning rate.)

Step 2c(ii): Preliminary computations for contractivity of F . Consider two pairs of triples $((\tilde{c}^1, \tilde{n}^1, \tilde{m}^1), \tilde{u}_{\tilde{\theta}_1}^N, \hat{u}_{\tilde{\theta}_1}^N)$, $((\tilde{c}^2, \tilde{n}^2, \tilde{m}^2), \tilde{u}_{\tilde{\theta}_2}^N, \hat{u}_{\tilde{\theta}_2}^N)$ with their corresponding two triples $((c^1, n^1, m^1), u_{\tilde{\theta}_1}^N, \hat{u}_{\tilde{\theta}_1}^N)$, $((c^2, n^2, m^2), u_{\tilde{\theta}_2}^N, \hat{u}_{\tilde{\theta}_2}^N) \in \Theta_{\tilde{\mathcal{T}}}(\tilde{M}) \times \mathcal{V}_{\tilde{\mathcal{T}}}(\tilde{M}) \times \mathcal{V}_{\tilde{\mathcal{T}}}(\tilde{M})$.

A bound for $|\tilde{c}_\tau^{i,1} - \tilde{c}_\tau^{i,2}|$. Recalling (237a), we can estimate after inserting mixed terms with two applications of Cauchy-Schwarz inequality, while using the boundedness and Lipschitz continuity of σ as of Assumption B1 and that the domain D has bounded volume

as of Assumption A2 together with the fact that the learning rate is decreasing, that

$$\begin{aligned}
 |\tilde{c}_\tau^{i,1} - \tilde{c}_\tau^{i,2}| &= C \left| \int_0^\tau \alpha_s \int_0^T \int_D \sigma(w_s^{t,i,1}t + (w_s^{i,1})^T x + \eta_s^{i,1}) \widehat{u}_{\theta_s^1}^N(t, x) \right. \\
 &\quad \left. - \sigma(w_s^{t,i,2}t + (w_s^{i,2})^T x + \eta_s^{i,2}) \widehat{u}_{\theta_s^2}^N(t, x) dx dt ds \right| \\
 &= C \left| \int_0^\tau \alpha_s \int_0^T \int_D \left(\sigma(w_s^{t,i,1}t + (w_s^{i,1})^T x + \eta_s^{i,1}) - \sigma(w_s^{t,i,2}t + (w_s^{i,2})^T x + \eta_s^{i,2}) \right) \widehat{u}_{\theta_s^1}^N(t, x) \right. \\
 &\quad \left. + \sigma(w_s^{t,i,2}t + (w_s^{i,2})^T x + \eta_s^{i,2}) \left(\widehat{u}_{\theta_s^1}^N(t, x) - \widehat{u}_{\theta_s^2}^N(t, x) \right) dx dt ds \right| \\
 &\leq C \int_0^\tau \alpha_s \left(\sqrt{\int_0^T \int_D \left(\sigma(w_s^{t,i,1}t + (w_s^{i,1})^T x + \eta_s^{i,1}) - \sigma(w_s^{t,i,2}t + (w_s^{i,2})^T x + \eta_s^{i,2}) \right)^2 dx dt} \right. \\
 &\quad \cdot \left. \|\widehat{u}_{\theta_s^1}^N\|_{L_2(D_T)} \right. \\
 &\quad \left. + \sqrt{\int_0^T \int_D \left(\sigma(w_s^{t,i,2}t + (w_s^{i,2})^T x + \eta_s^{i,2}) \right)^2 dx dt} \|\widehat{u}_{\theta_s^1}^N - \widehat{u}_{\theta_s^2}^N\|_{L_2(D_T)} \right) ds \\
 &\leq C \int_0^\tau \alpha_s \left(\|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} \|\widehat{u}_{\theta_s^1}^N\|_{L_2(D_T)} + \|\widehat{u}_{\theta_s^1}^N - \widehat{u}_{\theta_s^2}^N\|_{L_2(D_T)} \right) ds \\
 &\leq C \int_0^\tau \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} \|\widehat{u}_{\theta_s^1}^N\|_{L_2(D_T)} + \|\widehat{u}_{\theta_s^1}^N - \widehat{u}_{\theta_s^2}^N\|_{L_2(D_T)} ds.
 \end{aligned} \tag{253}$$

Proceeding analogously for (237b)–(237d) while using that also σ' is bounded and Lipschitz continuous as of Assumption B2, we obtain

$$\begin{aligned}
 &|\tilde{w}_\tau^{t,i,1} - \tilde{w}_\tau^{t,i,2}| + \|\tilde{w}_\tau^{i,1} - \tilde{w}_\tau^{i,2}\| + |\tilde{\eta}_\tau^{i,1} - \tilde{\eta}_\tau^{i,2}| \\
 &\leq C \int_0^\tau |c_s^{i,1} - c_s^{i,2}| \|\widehat{u}_{\theta_s^1}^N\|_{L_2(D_T)} + |c_s^{i,2}| \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} \|\widehat{u}_{\theta_s^1}^N\|_{L_2(D_T)} \\
 &\quad + |c_s^{i,2}| \|\widehat{u}_{\theta_s^1}^N - \widehat{u}_{\theta_s^2}^N\|_{L_2(D_T)} ds
 \end{aligned} \tag{254}$$

A bound for $\|\tilde{n}_\tau^{i,1} - \tilde{n}_\tau^{i,2}\|_{L_2(D_T)}$. Using that σ is Lipschitz as of Assumption B1 and that the domain D is bounded and has bounded volume as of Assumption A2, we can estimate

$$\begin{aligned}
 \|\tilde{n}_\tau^{i,1} - \tilde{n}_\tau^{i,2}\|_{L_2(D_T)} &= \|\sigma(\tilde{w}_\tau^{t,i,1}t + (\tilde{w}_\tau^{i,1})^T x + \tilde{\eta}_\tau^{i,1}) - \sigma(\tilde{w}_\tau^{t,i,2}t + (\tilde{w}_\tau^{i,2})^T x + \tilde{\eta}_\tau^{i,2})\|_{L_2(D_T)} \\
 &\leq C \left(|\tilde{w}_\tau^{t,i,1} - \tilde{w}_\tau^{t,i,2}| + \|\tilde{w}_\tau^{i,1} - \tilde{w}_\tau^{i,2}\| + |\tilde{\eta}_\tau^{i,1} - \tilde{\eta}_\tau^{i,2}| \right),
 \end{aligned} \tag{255}$$

i.e., resulting in the same bound as in (254).

A bound for $\|\tilde{m}_\tau^{i,1} - \tilde{m}_\tau^{i,2}\|_{L_2(D_T)}$. Since also σ' is Lipschitz as of Assumption B2,

$$\|\tilde{m}_\tau^{i,1} - \tilde{m}_\tau^{i,2}\|_{L_2(D_T)} \leq C \left(|\tilde{w}_\tau^{t,i,1} - \tilde{w}_\tau^{t,i,2}| + \|\tilde{w}_\tau^{i,1} - \tilde{w}_\tau^{i,2}\| + |\tilde{\eta}_\tau^{i,1} - \tilde{\eta}_\tau^{i,2}| \right), \tag{256}$$

i.e., resulting in the same bound as in (254).

A bound for $\|\tilde{g}_{\theta_1^1}^N - \tilde{g}_{\theta_2^2}^N\|_{L_p(D_T)}$. Using that σ is bounded and Lipschitz continuous as of Assumption B1 and that the domain D has bounded volume as of Assumption A2 to obtain the second inequality, the latter again to get the equality thereafter, (248) in the subsequent step, and (253)–(254) in the next-to-last step, we can upper bound for $p \geq 2$ with triangle inequality to obtain the first inequality that

$$\begin{aligned}
 & \|\tilde{g}_{\theta_1^1}^N - \tilde{g}_{\theta_2^2}^N\|_{L_p(D_T)} \\
 &= \left\| \frac{1}{N\beta} \sum_{i=1}^N \tilde{c}_\tau^{i,1} \sigma(\tilde{w}_\tau^{t,i,1} t + (\tilde{w}_\tau^{i,1})^T x + \tilde{\eta}_\tau^{i,1}) - \tilde{c}_\tau^{i,2} \sigma(\tilde{w}_\tau^{t,i,2} t + (\tilde{w}_\tau^{i,2})^T x + \tilde{\eta}_\tau^{i,2}) \right\|_{L_p(D_T)} \\
 &\leq \left\| \frac{1}{N\beta} \sum_{i=1}^N (\tilde{c}_\tau^{i,1} - \tilde{c}_\tau^{i,2}) \sigma(\tilde{w}_\tau^{t,i,1} t + (\tilde{w}_\tau^{i,1})^T x + \tilde{\eta}_\tau^{i,1}) \right\|_{L_p(D_T)} \\
 &\quad + \left\| \frac{1}{N\beta} \sum_{i=1}^N \tilde{c}_\tau^{i,2} (\sigma(\tilde{w}_\tau^{t,i,1} t + (\tilde{w}_\tau^{i,1})^T x + \tilde{\eta}_\tau^{i,1}) - \sigma(\tilde{w}_\tau^{t,i,2} t + (\tilde{w}_\tau^{i,2})^T x + \tilde{\eta}_\tau^{i,2})) \right\|_{L_p(D_T)} \\
 &\leq C \left\| \sum_{i=1}^N |\tilde{c}_\tau^{i,1} - \tilde{c}_\tau^{i,2}| \right\|_{L_p(D_T)} \\
 &\quad + C \left\| \sum_{i=1}^N |\tilde{c}_\tau^{i,2}| (|\tilde{w}_s^{t,i,1} - \tilde{w}_s^{t,i,2}| + \|\tilde{w}_s^{i,1} - \tilde{w}_s^{i,2}\| + |\tilde{\eta}_s^{i,1} - \tilde{\eta}_s^{i,2}|) \right\|_{L_p(D_T)} \\
 &= C \sum_{i=1}^N |\tilde{c}_\tau^{i,1} - \tilde{c}_\tau^{i,2}| + |\tilde{c}_\tau^{i,2}| (|\tilde{w}_s^{t,i,1} - \tilde{w}_s^{t,i,2}| + \|\tilde{w}_s^{i,1} - \tilde{w}_s^{i,2}\| + |\tilde{\eta}_s^{i,1} - \tilde{\eta}_s^{i,2}|) \\
 &\leq C \sum_{i=1}^N |\tilde{c}_\tau^{i,1} - \tilde{c}_\tau^{i,2}| + (|\tilde{w}_s^{t,i,1} - \tilde{w}_s^{t,i,2}| + \|\tilde{w}_s^{i,1} - \tilde{w}_s^{i,2}\| + |\tilde{\eta}_s^{i,1} - \tilde{\eta}_s^{i,2}|) \\
 &\quad \cdot \left(\|\theta_0\| + \left(\int_0^\tau \|\hat{u}_{\theta_s^2}^N\|_{L_2(D_T)}^4 ds \right)^{1/4} \right) \\
 &\leq C \int_0^\tau \sum_{i=1}^N \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} \|\hat{u}_{\theta_s^1}^N\|_{L_2(D_T)} + \|\hat{u}_{\theta_s^1}^N - \hat{u}_{\theta_s^2}^N\|_{L_2(D_T)} \\
 &\quad + \left(|c_s^{i,1} - c_s^{i,2}| \|\hat{u}_{\theta_s^1}^N\|_{L_2(D_T)} + |c_s^{i,2}| \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} \|\hat{u}_{\theta_s^1}^N\|_{L_2(D_T)} \right. \\
 &\quad \left. + |c_s^{i,2}| \|\hat{u}_{\theta_s^1}^N - \hat{u}_{\theta_s^2}^N\|_{L_2(D_T)} \right) \left(\|\theta_0\| + \left(\int_0^\tau \|\hat{u}_{\theta_s^2}^N\|_{L_2(D_T)}^4 ds \right)^{1/4} \right) ds \\
 &\leq C \int_0^\tau \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} + \|\hat{u}_{\theta_s^1}^N - \hat{u}_{\theta_s^2}^N\|_{L_2(D_T)} \right) \\
 &\quad \cdot \left(1 + |c_s^{i,2}| + \|\hat{u}_{\theta_s^1}^N\|_{L_2(D_T)} + |c_s^{i,2}| \|\hat{u}_{\theta_s^1}^N\|_{L_2(D_T)} \right) \\
 &\quad \cdot \left(1 + \|\theta_0\| + \left(\int_0^\tau \|\hat{u}_{\theta_s^2}^N\|_{L_2(D_T)}^4 ds \right)^{1/4} \right) ds \tag{257}
 \end{aligned}$$

for a constant $C = C(p, N, T, D, \sigma)$. The last inequality is just a rough upper bound.

A bound for $\|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{L_2([0,T],H^1(D))} + \|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{L_\infty([0,T],L_2(D))}$. By following the computations of *Step 2c(ii)* in the proof of Lemma 13 in Appendix A that lead to (202), it holds

$$\|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{L_2([0,T],H^1(D))} + \|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{L_\infty([0,T],L_2(D))} \leq C \|\tilde{g}_{\theta_1^N}^N - \tilde{g}_{\theta_2^N}^N\|_{L_2(D_T)} \quad (258)$$

for a constant $C = C(T, \mathcal{L}, q)$. Thus, in particular,

$$\|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{\mathcal{V}_T} \leq C \sup_{\tau \in [0, T]} \|\tilde{g}_{\theta_1^N}^N - \tilde{g}_{\theta_2^N}^N\|_{L_2(D_T)}. \quad (259)$$

A bound for $\|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{L_\infty(D_T)}$. By following the computations of *Step 2c(ii)* in the proof of Lemma 13 in Appendix A that lead to (208), it holds for $p > d + 1$ that

$$\|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{L_\infty(D_T)} \leq c(d, p) C \|\tilde{g}_{\theta_1^N}^N - \tilde{g}_{\theta_2^N}^N\|_{L_p(D_T)} \quad (260)$$

for a constant $C = C(d, p, N, T, D, \mathcal{L}, q, \sigma)$.

A bound for $\|\widehat{u}_{\theta_1^N}^N - \widehat{u}_{\theta_2^N}^N\|_{L_2([0,T],H^1(D))} + \|\widehat{u}_{\theta_1^N}^N - \widehat{u}_{\theta_2^N}^N\|_{L_\infty([0,T],L_2(D))}$. By following the computations of *Step 2c(ii)* in the proof of Lemma 13 in Appendix A that lead to (219), it holds

$$\begin{aligned} & \|\widehat{u}_{\theta_1^N}^N - \widehat{u}_{\theta_2^N}^N\|_{L_2([0,T],H^1(D))} + \|\widehat{u}_{\theta_1^N}^N - \widehat{u}_{\theta_2^N}^N\|_{L_\infty([0,T],L_2(D))} \\ & \leq C \left(\|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{L_2(D_T)} + \|\widehat{u}_{\theta_2^N}^N\|_{L_\infty([0,T],L_2(D))} \|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{L_\infty(D_T)} \right) \end{aligned} \quad (261)$$

for a constant $C = C(T, \mathcal{L}, q)$. Thus, in particular,

$$\begin{aligned} & \|\widehat{u}_{\theta_1^N}^N - \widehat{u}_{\theta_2^N}^N\|_{\mathcal{V}_T} \\ & \leq C \left(\sup_{\tau \in [0, T]} \|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{L_2(D_T)} + \sup_{\tau \in [0, T]} \|\widehat{u}_{\theta_2^N}^N\|_{L_\infty([0,T],L_2(D))} \sup_{\tau \in [0, T]} \|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{L_\infty(D_T)} \right). \end{aligned} \quad (262)$$

Combination of the bounds. Using (259) and (262) in the first inequality together with (260) in the second inequality as well as (253)–(257) in the third, we have

$$\begin{aligned} & \|(\tilde{c}^1, \tilde{n}^1, \tilde{m}^1) - (\tilde{c}^2, \tilde{n}^2, \tilde{m}^2)\|_{\Theta_{\tilde{\tau}}} + \|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{\mathcal{V}_{\tilde{\tau}}} + \|\widehat{u}_{\theta_1^N}^N - \widehat{u}_{\theta_2^N}^N\|_{\mathcal{V}_{\tilde{\tau}}} \\ & \leq \|(\tilde{c}^1, \tilde{n}^1, \tilde{m}^1) - (\tilde{c}^2, \tilde{n}^2, \tilde{m}^2)\|_{\Theta_{\tilde{\tau}}} + C \sup_{\tau \in [0, \tilde{T}]} \|\tilde{g}_{\theta_1^N}^N - \tilde{g}_{\theta_2^N}^N\|_{L_2(D_T)} \\ & \quad + C \sup_{\tau \in [0, \tilde{T}]} \|\widehat{u}_{\theta_2^N}^N\|_{L_\infty([0,T],L_2(D))} \sup_{\tau \in [0, \tilde{T}]} \|\tilde{u}_{\theta_1^N}^N - \tilde{u}_{\theta_2^N}^N\|_{L_\infty(D_T)} \\ & \leq \|(\tilde{c}^1, \tilde{n}^1, \tilde{m}^1) - (\tilde{c}^2, \tilde{n}^2, \tilde{m}^2)\|_{\Theta_{\tilde{\tau}}} + \\ & \quad + C \left(1 + \sup_{\tau \in [0, \tilde{T}]} \|\widehat{u}_{\theta_2^N}^N\|_{L_\infty([0,T],L_2(D))} \right) \sup_{\tau \in [0, \tilde{T}]} \|\tilde{g}_{\theta_1^N}^N - \tilde{g}_{\theta_2^N}^N\|_{L_2(D_T)} \end{aligned}$$

$$\begin{aligned}
 &\leq C_2 \left(1 + \sup_{\tau \in [0, \tilde{\mathcal{T}}]} \|\widehat{u}_{\theta^2}^N\|_{L_\infty([0, T], L_2(D))} \right) \left(1 + \|\theta_0\| + \left(\int_0^{\tilde{\mathcal{T}}} \|\widehat{u}_{\theta^2}^N\|_{L_2(D_T)}^4 ds \right)^{1/4} \right) \\
 &\quad \cdot \int_0^{\tilde{\mathcal{T}}} \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} + \|\widehat{u}_{\theta^1}^N - \widehat{u}_{\theta^2}^N\|_{L_2(D_T)} \right) \\
 &\quad \cdot \left(1 + |c_s^{i,2}| + \|\widehat{u}_{\theta^1}^N\|_{L_2(D_T)} + |c_s^{i,2}| \|\widehat{u}_{\theta^1}^N\|_{L_2(D_T)} \right) ds
 \end{aligned} \tag{263}$$

for a constant $C_2 = C_2(\alpha, N, T, D, \mathcal{L}, q, \sigma)$. (Note that multiple higher-order (up to fourth-order) product terms appear on the right-hand side due to the NTK varying during training.)

Step 2d: Existence locally in training time. Let us choose

$$M_0 = 2C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + \|\theta_0\| + 1 \right) \tag{264}$$

and

$$\mathcal{T}_0 = \min \left\{ \frac{1}{2^5 C_1^4}, \frac{1}{2C_2 c_2 N (1 + M_0)^4} \right\} \tag{265}$$

where the constants C_1 and C_2 are as given implicitly in (252) and (263), respectively, and where $c_2 = \max\{1 + \|\theta_0\|, 1/(2C_1)\}$. We show in what follows that there exists a unique solution $((c, n, m), u_\theta^N, \widehat{u}_\theta^N) \in \Theta_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$.

Step 2d(i): Self-mapping property of F . Consider $((\tilde{c}, \tilde{n}, \tilde{m}), \tilde{u}_\theta^N, \widehat{\tilde{u}}_\theta^N)$ together with its corresponding $((c, n, m), u_\theta^N, \widehat{u}_\theta^N) \in \Theta_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$. Using the definitions of M_0 and \mathcal{T}_0 in (264) and (265), respectively, we can derive from (252) that

$$\begin{aligned}
 &\|(\tilde{c}, \tilde{n}, \tilde{m})\|_{\Theta_{\mathcal{T}_0}} + \|\tilde{u}_\theta^N\|_{\mathcal{V}_{\mathcal{T}_0}} + \|\widehat{\tilde{u}}_\theta^N\|_{\mathcal{V}_{\mathcal{T}_0}} \\
 &\leq C_1 \left(\int_0^{\mathcal{T}_0} \|\widehat{u}_{\theta^2}^N\|_{L_2(D_T)}^4 ds \right)^{1/4} + C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + \|\theta_0\| + 1 \right) \\
 &\leq C_1 \mathcal{T}_0^{1/4} \|\widehat{u}_\theta^N\|_{\mathcal{V}_{\mathcal{T}_0}} + \frac{M_0}{2} \leq C_1 \mathcal{T}_0^{1/4} M_0 + \frac{M_0}{2} \leq \frac{M_0}{2} + \frac{M_0}{2} \leq M_0.
 \end{aligned} \tag{266}$$

Thus, $((\tilde{c}, \tilde{n}, \tilde{m}), \tilde{u}_\theta^N, \widehat{\tilde{u}}_\theta^N) \in \Theta_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$.

Step 2d(ii): Contractivity of F . Consider two pairs of triples $((\tilde{c}^1, \tilde{n}^1, \tilde{m}^1), \tilde{u}_{\theta^1}^N, \widehat{\tilde{u}}_{\theta^1}^N)$, $((\tilde{c}^2, \tilde{n}^2, \tilde{m}^2), \tilde{u}_{\theta^2}^N, \widehat{\tilde{u}}_{\theta^2}^N)$ with their corresponding $((c^1, n^1, m^1), u_{\theta^1}^N, \widehat{u}_{\theta^1}^N)$, $((c^2, n^2, m^2), u_{\theta^2}^N, \widehat{u}_{\theta^2}^N) \in \Theta_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$. According to *Step 2d(i)*, we have that $((\tilde{c}^1, \tilde{n}^1, \tilde{m}^1), \tilde{u}_{\theta^1}^N, \widehat{\tilde{u}}_{\theta^1}^N)$, $((\tilde{c}^2, \tilde{n}^2, \tilde{m}^2), \tilde{u}_{\theta^2}^N, \widehat{\tilde{u}}_{\theta^2}^N) \in \Theta_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$. Using the definitions of M_0 and \mathcal{T}_0

in (264) and (265), we can derive from (263) that

$$\begin{aligned}
 & \|(\tilde{c}^1, \tilde{n}^1, \tilde{m}^1) - (\tilde{c}^2, \tilde{n}^2, \tilde{m}^2)\|_{\Theta_{\mathcal{T}_0}} + \|\tilde{u}_{\theta^1}^N - \tilde{u}_{\theta^2}^N\|_{\mathcal{V}_{\mathcal{T}_0}} + \|\widehat{u}_{\theta^1}^N - \widehat{u}_{\theta^2}^N\|_{\mathcal{V}_{\mathcal{T}_0}} \\
 & \leq C_2 \left(1 + \sup_{\tau \in [0, \mathcal{T}_0]} \|\widehat{u}_{\theta_\tau^N}\|_{L_\infty([0, T], L_2(D))} \right) \left(1 + \|\theta_0\| + \left(\int_0^{\mathcal{T}_0} \|\widehat{u}_{\theta_s^N}\|_{L_2(D_T)}^4 ds \right)^{1/4} \right) \\
 & \quad \cdot \int_0^{\mathcal{T}_0} \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} + \|\widehat{u}_{\theta_s^1}^N - \widehat{u}_{\theta_s^2}^N\|_{L_2(D_T)} \right) \\
 & \quad \cdot \left(1 + |c_s^{i,2}| + \|\widehat{u}_{\theta_s^1}^N\|_{L_2(D_T)} + |c_s^{i,2}| \|\widehat{u}_{\theta_s^1}^N\|_{L_2(D_T)} \right) ds \\
 & \leq C_2 (1 + M_0) \left(1 + \|\theta_0\| + \mathcal{T}_0^{1/4} M_0 \right) \\
 & \quad \cdot \int_0^{\mathcal{T}_0} \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} + \|\widehat{u}_{\theta_s^1}^N - \widehat{u}_{\theta_s^2}^N\|_{L_2(D_T)} \right) \\
 & \quad \cdot (1 + 2M_0 + M_0^2) ds \\
 & \leq C_2 c_2 (1 + M_0)^4 \\
 & \quad \cdot \int_0^{\mathcal{T}_0} \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} \right. \\
 & \quad \left. + \|\widehat{u}_{\theta_s^1}^N - \widehat{u}_{\theta_s^2}^N\|_{L_2(D_T)} \right) ds \\
 & \leq C_2 c_2 (1 + M_0)^4 \mathcal{T}_0 \left(\|(c^1, n^1, m^1) - (c^2, n^2, m^2)\|_{\Theta_{\mathcal{T}_0}} + N \|\widehat{u}_{\theta^1}^N - \widehat{u}_{\theta^2}^N\|_{\mathcal{V}_{\mathcal{T}_0}} \right) ds \\
 & \leq \frac{1}{2} \left(\|(c^1, n^1, m^1) - (c^2, n^2, m^2)\|_{\Theta_{\mathcal{T}_0}} + \|u_{\theta^1}^N - u_{\theta^2}^N\|_{\mathcal{V}_{\mathcal{T}_0}} + \|\widehat{u}_{\theta^1}^N - \widehat{u}_{\theta^2}^N\|_{\mathcal{V}_{\mathcal{T}_0}} \right), \tag{267}
 \end{aligned}$$

showing that the map F is a contraction.

Step 2d(iii): Banach fixed point theorem. Hence, the Banach fixed point theorem guarantees that there exists a unique solution $((c, n, m), u_\theta^N, \widehat{u}_\theta^N) \in \Theta_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0) \times \mathcal{V}_{\mathcal{T}_0}(M_0)$, which satisfies $((c, n, m), u_\theta^N, \widehat{u}_\theta^N) = F((c, n, m), u_\theta^N, \widehat{u}_\theta^N)$. We have thus established the existence of a unique local-in-training-time solution to the PDE system (1) & (8) coupled with the gradient descent update (5) on the training time domain $[0, \mathcal{T}_0]$. In particular, $(u_{\theta_\tau}^N, \widehat{u}_{\theta_\tau}^N) \in \mathcal{S} \times \mathcal{S}$ for every $\tau \in [0, \mathcal{T}_0]$. Reapplying the classical existence and regularity results from *Steps 1a, 1b* and *1c* further shows that for each $\tau \in [0, \mathcal{T}_0]$ such solution satisfies $(\partial_t u_{\theta_\tau}^N(t, \bullet), \partial_t \widehat{u}_{\theta_\tau}^N(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$.

Step 2e: Existence globally in training time. Leveraging a bootstrapping argument, we now extend this argument to obtain a solution on a training time domain $[0, \mathcal{T}]$ for an arbitrary $\mathcal{T} < \infty$. To do so, we proceed inductively. Suppose we have a solution $(u^N, \widehat{u}^N) \in \mathcal{V}_{\mathcal{T}_{k-1}}(M_{k-1}) \times \mathcal{V}_{\mathcal{T}_{k-1}}(M_{k-1})$ which is such that for each $\tau \in [0, \mathcal{T}_{k-1}]$ it satisfies $(\partial_t u_\tau^N(t, \bullet), \partial_t \widehat{u}_\tau^N(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$. (We showed in *Step 2d* before that this is the case for the induction start $k = 1$.)

On the training time interval $I = [0, \mathcal{T}_{k-1}]$ we can now employ Lemma 38 which ensures that $\frac{d}{d\tau} \mathcal{J}_{\theta_\tau}^N \leq 0$ for all $\tau \in I = [0, \mathcal{T}_{k-1}]$. Thanks to this, Lemma 40 (applied in the setting

$I = [0, \mathcal{T}_{k-1}]$) provides a uniform (in the training time τ and on the training time interval $[0, \mathcal{T}_{k-1}]$) bound $\sup_{\tau \in [0, \mathcal{T}_{k-1}]} \|\widehat{u}_{\theta_\tau}^N\|_{L_2(D_T)} \leq C^{\widehat{u}^N}$, where $C^{\widehat{u}^N}$ does not depend on \mathcal{T}_{k-1} but only on $\mathcal{J}_{\theta_0}^N$. Let us now choose

$$M_k = k^{1/4} C^{\widehat{u}^N} + 2C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + \|\theta_0\| + 1 \right) \quad (268)$$

(note that here we can leverage the slower growth of the worst-case bound and impose it in the proposed bound M_k) and

$$\mathcal{T}_k = \mathcal{T}_{k-1} + \min \left\{ \frac{1}{2^5 C_1^4}, \frac{1}{2C_2 c_2 N(1 + M_k)^4} \right\}. \quad (269)$$

We show in what follows that there exists a unique solution $((c, n, m), u_\theta^N, \widehat{u}_\theta^N) \in \Theta_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k)$.

Step 2e(i): Self-mapping property of F . Consider $((\widetilde{c}, \widetilde{n}, \widetilde{m}), \widetilde{u}_\theta^N, \widehat{u}_\theta^N)$ together with its corresponding $((c, n, m), u_\theta^N, \widehat{u}_\theta^N) \in \Theta_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k)$. Using the definitions of M_k and \mathcal{T}_k in (268) and (269), we can derive from (252) that

$$\begin{aligned} & \|(\widetilde{c}, \widetilde{n}, \widetilde{m})\|_{\Theta_{\mathcal{T}_k}} + \|\widetilde{u}_\theta^N\|_{\mathcal{V}_{\mathcal{T}_k}} + \|\widehat{u}_\theta^N\|_{\mathcal{V}_{\mathcal{T}_k}} \\ & \leq C_1 \left(\int_0^{\mathcal{T}_k} \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^4 ds \right)^{1/4} + C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + \|\theta_0\| + 1 \right) \\ & \leq C_1 \left(\int_0^{\mathcal{T}_{k-1}} \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^4 ds + \int_{\mathcal{T}_{k-1}}^{\mathcal{T}_k} \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^4 ds \right)^{1/4} \\ & \quad + C_1 \left(\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + \|\theta_0\| + 1 \right) \\ & \leq C_1 \left(\mathcal{T}_{k-1} (C^{\widehat{u}^N})^4 + (\mathcal{T}_k - \mathcal{T}_{k-1}) M_k^4 \right)^{1/4} + \frac{M_k}{2} \\ & \leq \left(\frac{1}{2^5} k (C^{\widehat{u}^N})^4 + \frac{1}{2^5} M_k^4 \right)^{1/4} + \frac{M_k}{2} \leq \left(\frac{1}{2^5} M_k^4 + \frac{1}{2^5} M_k^4 \right)^{1/4} + \frac{M_k}{2} \leq \frac{M_k}{2} + \frac{M_k}{2} \leq M_k, \end{aligned} \quad (270)$$

where we used for the first step in the last line that (with $\mathcal{T}_{-1} := 0$) according to the definition of \mathcal{T}_k in (269) it holds $\mathcal{T}_{k-1} (C^{\widehat{u}^N})^4 = \sum_{\ell=0}^{k-1} (\mathcal{T}_\ell - \mathcal{T}_{\ell-1}) (C^{\widehat{u}^N})^4 \leq \frac{1}{2^5 C_1^4} k (C^{\widehat{u}^N})^4$.

Thus, $((\widetilde{c}, \widetilde{n}, \widetilde{m}), \widetilde{u}_\theta^N, \widehat{u}_\theta^N) \in \Theta_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k)$.

Step 2e(ii): Contractivity of F . Consider two pairs of triples $((\widetilde{c}^1, \widetilde{n}^1, \widetilde{m}^1), \widetilde{u}_{\theta^1}^N, \widehat{u}_{\theta^1}^N)$, $((\widetilde{c}^2, \widetilde{n}^2, \widetilde{m}^2), \widetilde{u}_{\theta^2}^N, \widehat{u}_{\theta^2}^N)$ with their corresponding $((c^1, n^1, m^1), u_{\theta^1}^N, \widehat{u}_{\theta^1}^N)$, $((c^2, n^2, m^2), u_{\theta^2}^N, \widehat{u}_{\theta^2}^N) \in \Theta_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k)$. Using the definitions of M_k and \mathcal{T}_k in (268) and (269), we can derive from (263) that

$$\begin{aligned} & \|(\widetilde{c}^1, \widetilde{n}^1, \widetilde{m}^1) - (\widetilde{c}^2, \widetilde{n}^2, \widetilde{m}^2)\|_{\Theta_{\mathcal{T}_k}} + \|\widetilde{u}_{\theta^1}^N - \widetilde{u}_{\theta^2}^N\|_{\mathcal{V}_{\mathcal{T}_k}} + \|\widehat{u}_{\theta^1}^N - \widehat{u}_{\theta^2}^N\|_{\mathcal{V}_{\mathcal{T}_k}} \\ & \leq C_2 \left(1 + \sup_{\tau \in [0, \mathcal{T}_k]} \|\widehat{u}_{\theta_\tau}^N\|_{L_\infty([0, T], L_2(D))} \right) \left(1 + \|\theta_0\| + \left(\int_0^{\mathcal{T}_k} \|\widehat{u}_{\theta_s}^N\|_{L_2(D_T)}^4 ds \right)^{1/4} \right) \end{aligned}$$

$$\begin{aligned}
 & \cdot \int_0^{\mathcal{T}_k} \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} + \|\widehat{u}_{\theta_1^N} - \widehat{u}_{\theta_2^N}\|_{L_2(D_T)} \right) \\
 & \quad \cdot \left(1 + |c_s^{i,2}| + \|\widehat{u}_{\theta_1^N}\|_{L_2(D_T)} + |c_s^{i,2}| \|\widehat{u}_{\theta_1^N}\|_{L_2(D_T)} \right) ds \\
 & \leq C_2 c_2 (1 + M_k)^4 \\
 & \quad \cdot \int_0^{\mathcal{T}_k} \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} \right. \\
 & \quad \left. \|\widehat{u}_{\theta_1^N} - \widehat{u}_{\theta_2^N}\|_{L_2(D_T)} \right) ds \\
 & = C_2 c_2 (1 + M_k)^4 \\
 & \quad \cdot \int_{\mathcal{T}_{k-1}}^{\mathcal{T}_k} \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} \right. \\
 & \quad \left. + \|\widehat{u}_{\theta_1^N} - \widehat{u}_{\theta_2^N}\|_{L_2(D_T)} \right) ds \\
 & \leq C_2 c_2 (1 + M_k)^4 (\mathcal{T}_k - \mathcal{T}_{k-1}) \left(\|(c^1, n^1, m^1) - (c^2, n^2, m^2)\|_{\Theta_{\mathcal{T}_k}} + N \|\widehat{u}_{\theta_1^N} - \widehat{u}_{\theta_2^N}\|_{\mathcal{V}_{\mathcal{T}_k}} \right) ds \\
 & \leq \frac{1}{2} \left(\|(c^1, n^1, m^1) - (c^2, n^2, m^2)\|_{\Theta_{\mathcal{T}_k}} + \|u_{\theta_1^N} - u_{\theta_2^N}\|_{\mathcal{V}_{\mathcal{T}_k}} + \|\widehat{u}_{\theta_1^N} - \widehat{u}_{\theta_2^N}\|_{\mathcal{V}_{\mathcal{T}_k}} \right), \quad (271)
 \end{aligned}$$

where the second step reuses from (270) that $(\int_0^{\mathcal{T}_k} \|\widehat{u}_{\theta_s^N}\|_{L_2(D_T)}^4 ds)^{1/4} \leq M_k/(2C_1)$, while the third step is due to the uniqueness of the solution on the training time interval $[0, \mathcal{T}_{k-1}]$. Thus, the map F is a contraction.

Step 2e(iii): Banach fixed point theorem. Hence, the Banach fixed point theorem guarantees that there exists a unique solution $((c, n, m), u_{\theta}^N, \widehat{u}_{\theta}^N) \in \Theta_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k) \times \mathcal{V}_{\mathcal{T}_k}(M_k)$, which satisfies $((c, n, m), u_{\theta}^N, \widehat{u}_{\theta}^N) = F((c, n, m), u_{\theta}^N, \widehat{u}_{\theta}^N)$. We have thus established the existence of a unique local-in-training-time solution to the PDE system (1) & (8) coupled with the gradient descent update (5) on the training time domain $[0, \mathcal{T}_k]$. In particular, $(u_{\theta}^N, \widehat{u}_{\theta}^N) \in \mathcal{S} \times \mathcal{S}$ for every $\tau \in [0, \mathcal{T}_k]$. Reapplying the classical existence and regularity results from *Steps 1a, 1b* and *1c* further shows that for each $\tau \in [0, \mathcal{T}_k]$ such solution satisfies $(\partial_t u_{\theta}^N(t, \bullet), \partial_t \widehat{u}_{\theta}^N(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$.

Step 2e(iv): Globality of the construction in training time. It remains to notice that, due to the definition of the times \mathcal{T}_k in (269), the telescopic sum

$$\begin{aligned}
 \sum_{k=1}^{\infty} (\mathcal{T}_k - \mathcal{T}_{k-1}) &= \sum_{k=1}^{\infty} \min \left\{ \frac{1}{2^5 C_1^4}, \frac{1}{2C_2 c_2 N (1 + M_k)^4} \right\} \geq \sum_{k=K}^{\infty} \frac{1}{2C_2 c_2 N (1 + M_k)^4} \\
 &= \sum_{k=K}^{\infty} \frac{1}{2C_2 c_2 N (1 + k^{1/4} C \widehat{u}^N + 2C_1 (\|h\|_{L_2(D_T)} + \|f\|_{L_2(D)} + \|\theta_0\| + 1))^4} \quad (272)
 \end{aligned}$$

diverges. To see this, simply note that the last term in (272) is a harmonic series for some sufficiently large integer K . This ensures that the above construction in *Step 2d* and *2e* gives a solution in $\mathcal{V}_{\mathcal{T}}(M) \times \mathcal{V}_{\mathcal{T}}(M)$ for any given $\mathcal{T} < \infty$ and suitable associated $M > 0$. (Note that here we exploit that we correctly balanced the appearing higher-order product

terms with the slower worst-case growth, enabling to get a diverging series, and thus a global existence result. With only (6) and without the additional assumption $\int_0^\infty \alpha_\tau^{4/3} d\tau < \infty$ on the learning rate, one would have only been able to get $M_k \propto k^{1/2}$, which would have lead to a geometric series of the form $\sum_{k=K}^\infty \frac{c}{1+k^2}$, which does not diverge, thus leading to no global existence result.)

Uniqueness. It remains to prove the uniqueness of a solution $(u_{\theta_\tau}^N, \widehat{u}_{\theta_\tau}^N)$ to the PDE system (1) & (8) coupled with the gradient descent update (5). For this purpose, suppose that there are two weak solutions $((c^1, n^1, m^1), u_{\theta_1}^N, \widehat{u}_{\theta_1}^N), ((c^2, n^2, m^2), u_{\theta_2}^N, \widehat{u}_{\theta_2}^N) \in \Theta_{\mathcal{T}}(M) \times \mathcal{V}_{\mathcal{T}}(M) \times \mathcal{V}_{\mathcal{T}}(M)$. This means we have $(u_{\theta_1}^N, \widehat{u}_{\theta_1}^N), (u_{\theta_2}^N, \widehat{u}_{\theta_2}^N) \in \mathcal{S} \times \mathcal{S}$ which satisfy $(\partial_t u_{\theta_\tau}^N(t, \bullet), \partial_t \widehat{u}_{\theta_\tau}^N(t, \bullet)), (\partial_t u_{\theta_2}^N(t, \bullet), \partial_t \widehat{u}_{\theta_2}^N(t, \bullet)) \in L_2(D) \times L_2(D)$ for a.e. $t \in [0, T]$ and where $g_{\theta_\tau}^N, g_{\theta_2}^N \in L_2(D_T)$ denote the corresponding NN functions (3) for each $\tau \in [0, \mathcal{T}]$. By repeating the computations of the existence proof in *Step 2c(ii)* we obtain analogously to (258) and (260) that

$$\|u_{\theta_1}^N - u_{\theta_2}^N\|_{L_2([0,T], H^1(D))} + \|u_{\theta_1}^N - u_{\theta_2}^N\|_{L_\infty([0,T], L_2(D))} \leq C \|g_{\theta_1}^N - g_{\theta_2}^N\|_{L_2(D_T)} \quad (273)$$

and for some $p > d + 1$ that

$$\|u_{\theta_1}^N - u_{\theta_2}^N\|_{L_\infty(D_T)} \leq c(d, p) C \|g_{\theta_1}^N - g_{\theta_2}^N\|_{L_p(D_T)}, \quad (274)$$

as well as analogously to (261) that

$$\begin{aligned} & \|\widehat{u}_{\theta_1}^N - \widehat{u}_{\theta_2}^N\|_{L_2([0,T], H^1(D))} + \|\widehat{u}_{\theta_1}^N - \widehat{u}_{\theta_2}^N\|_{L_\infty([0,T], L_2(D))} \\ & \leq C \left(\|u_{\theta_1}^N - u_{\theta_2}^N\|_{L_2(D_T)} + M \|u_{\theta_1}^N - u_{\theta_2}^N\|_{L_\infty(D_T)} \right), \end{aligned} \quad (275)$$

where we used directly that $\widehat{u}_{\theta_2}^N \in \mathcal{V}_{\mathcal{T}}(M)$. Since it holds

$$\begin{aligned} & \|g_{\theta_1}^N - g_{\theta_2}^N\|_{L_p(D_T)} \\ & \leq C \int_0^\tau \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} + \|\widehat{u}_{\theta_1}^N - \widehat{u}_{\theta_2}^N\|_{L_2(D_T)} \right) \\ & \quad \cdot (1 + M)^2 (1 + \|\theta_0\| + \mathcal{T}^{1/4} M) ds \end{aligned} \quad (276)$$

by repeating the computations (257) of the existence proof in *Step 2c(ii)* and using directly that $(c^i, n^i, m^i) \in \Theta_{\mathcal{T}}(M)$ as well as $\widehat{u}_{\theta_i}^N \in \mathcal{V}_{\mathcal{T}}(M)$, and since it furthermore holds

$$\begin{aligned} & \sum_{i=1}^N \left(|c_\tau^{i,1} - c_\tau^{i,2}| + \|n_\tau^{i,1} - n_\tau^{i,2}\|_{L_2(D_T)} + \|m_\tau^{i,1} - m_\tau^{i,2}\|_{L_2(D_T)} \right) \\ & \leq C \int_0^\tau \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} + \|\widehat{u}_{\theta_1}^N - \widehat{u}_{\theta_2}^N\|_{L_2(D_T)} \right) \\ & \quad \cdot (1 + M)^2 ds \end{aligned} \quad (277)$$

by repeating and combining the computations (253)–(256) of the existence proof in *Step 2c(ii)* and using directly that $(c^i, n^i, m^i) \in \Theta_{\mathcal{T}}(M)$ as well as $\widehat{u}_{\theta_\tau^i}^N \in \mathcal{V}_{\mathcal{T}}(M)$, we get

$$\begin{aligned}
 & \|g_{\theta_\tau^1}^N - g_{\theta_\tau^2}^N\|_{L_2(D_T)} + \|g_{\theta_\tau^1}^N - g_{\theta_\tau^2}^N\|_{L_p(D_T)} + \sum_{i=1}^N \left(|c_\tau^{i,1} - c_\tau^{i,2}| + \|n_\tau^{i,1} - n_\tau^{i,2}\|_{L_2(D_T)} + \|m_\tau^{i,1} - m_\tau^{i,2}\|_{L_2(D_T)} \right) \\
 & \leq C \int_0^\tau \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} + \|\widehat{u}_{\theta_s^1}^N - \widehat{u}_{\theta_s^2}^N\|_{L_2(D_T)} \right) ds \\
 & \leq C \int_0^\tau \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} \right. \\
 & \quad \left. + \|u_{\theta_s^1}^N - u_{\theta_s^2}^N\|_{L_2(D_T)} + \|u_{\theta_s^1}^N - u_{\theta_s^2}^N\|_{L_\infty(D_T)} \right) ds \\
 & \leq C \int_0^\tau \sum_{i=1}^N \left(|c_s^{i,1} - c_s^{i,2}| + \|n_s^{i,1} - n_s^{i,2}\|_{L_2(D_T)} + \|m_s^{i,1} - m_s^{i,2}\|_{L_2(D_T)} \right. \\
 & \quad \left. + \|g_{\theta_s^1}^N - g_{\theta_s^2}^N\|_{L_2(D_T)} + \|g_{\theta_s^1}^N - g_{\theta_s^2}^N\|_{L_p(D_T)} \right) ds,
 \end{aligned} \tag{278}$$

where C may depend in particular on M and \mathcal{T} .

Recalling that $\theta_0^1 = \theta_0^2$ and consequently $c_0^{i,1} = c_0^{i,2}$, $n_0^{i,1} = n_0^{i,2}$, $m_0^{i,1} = m_0^{i,2}$ as well as $g_{\theta_0^1}^N = g_{\theta_0^2}^N = 0$, we can now employ Grönwall's inequality in its integral form to obtain

$$\|g_{\theta_\tau^1}^N - g_{\theta_\tau^2}^N\|_{L_2(D_T)} + \|g_{\theta_\tau^1}^N - g_{\theta_\tau^2}^N\|_{L_p(D_T)} = 0 \tag{279}$$

for every $\tau \in [0, \mathcal{T}]$. With (273) and (275) we hence conclude that for every $\tau \in [0, \mathcal{T}]$ it hold

$$\|u_{\theta_\tau^1}^N - u_{\theta_\tau^2}^N\|_{L_2([0, \mathcal{T}], H^1(D))} + \|u_{\theta_\tau^1}^N - u_{\theta_\tau^2}^N\|_{L_\infty([0, \mathcal{T}], L_2(D))} = 0 \tag{280}$$

and

$$\|\widehat{u}_{\theta_\tau^1}^N - \widehat{u}_{\theta_\tau^2}^N\|_{L_2([0, \mathcal{T}], H^1(D))} + \|\widehat{u}_{\theta_\tau^1}^N - \widehat{u}_{\theta_\tau^2}^N\|_{L_\infty([0, \mathcal{T}], L_2(D))} = 0. \tag{281}$$

Thus $\|u_{\theta_\tau^1}^N - u_{\theta_\tau^2}^N\|_{\mathcal{V}_{\mathcal{T}}} = 0$ and $\|\widehat{u}_{\theta_\tau^1}^N - \widehat{u}_{\theta_\tau^2}^N\|_{\mathcal{V}_{\mathcal{T}}} = 0$, proving uniqueness in $\mathcal{V}_{\mathcal{T}}(M)$. \blacksquare

We now provide auxiliary results that were used in the proof of Lemma 11. First, we establish that the loss $\mathcal{J}_{\theta_\tau}^N$ defined in (4) is monotonically non-increasing in the training time τ .

Lemma 38 (Decay of the loss $\mathcal{J}_{\theta_\tau}^N$) *Assume that the learning rate satisfies additionally $\int_0^\infty \alpha_\tau^{4/3} d\tau < \infty$. Let $((u_{\theta_\tau}^N, \widehat{u}_{\theta_\tau}^N))_{\tau \in I} \in \mathcal{C}(I, \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (1) & (8) coupled with the gradient descent update (5) in the sense of Lemma 11 on the training time interval I . Define the loss $\mathcal{J}_{\theta_\tau}^N$ as in (4). Then, for the training time derivative $\frac{d}{d\tau} \mathcal{J}_{\theta_\tau}^N$ it holds*

$$\begin{aligned}
 \frac{d}{d\tau} \mathcal{J}_{\theta_\tau}^N &= -\alpha_\tau (\widehat{u}_{\theta_\tau}^N, T_{B(\mu_\tau^N)} \widehat{u}_{\theta_\tau}^N)_{L_2(D_T)} \\
 &= -\alpha_\tau \int_0^T \int_D \widehat{u}_{\theta_\tau}^N(t, x) \int_0^T \int_D B(t, x, t', x'; \mu_\tau^N) \widehat{u}_{\theta_\tau}^N(t', x') dx' dt' dx dt
 \end{aligned} \tag{282}$$

for all $\tau \in I$ with the operator $T_{B(\mu_\tau^N)}$ defined in (12) and where the kernel $B(\mu_\tau^N)$ is as in (13). In particular, we have $\frac{d}{d\tau} \mathcal{J}_{\theta_\tau}^N \leq 0$ for all $\tau \in I$.

Proof Taking the training time derivative of our loss $\mathcal{J}_{\theta_\tau}^N$, i.e., the derivative w.r.t. the training time τ , we obtain by chain rule and by using that $\widehat{u}_{\theta_\tau}^N$ is a weak solution to the adjoint PDE (8) in the sense of Definition 5 with right-hand side $(u_{\theta_\tau}^N - h)$ that

$$\begin{aligned}
 \frac{d}{d\tau} \mathcal{J}_{\theta_\tau}^N &= \frac{d}{d\tau} \frac{1}{2} \int_0^T \int_D (u_{\theta_\tau}^N(t, x) - h(t, x))^2 dx dt \\
 &= \int_0^T \int_D (u_{\theta_\tau}^N(t, x) - h(t, x)) \frac{d}{d\tau} u_{\theta_\tau}^N(t, \bullet) dx dt = \int_0^T \left(u_{\theta_\tau}^N(t, \bullet) - h(t, \bullet), \frac{d}{d\tau} u_{\theta_\tau}^N(t, \bullet) \right)_{L_2(D)} dt \\
 &= \int_0^T \left\langle -\partial_t \widehat{u}_{\theta_\tau}^N(t, \bullet), \frac{d}{d\tau} u_{\theta_\tau}^N(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B}^\dagger \left[\widehat{u}_{\theta_\tau}^N(t, \bullet), \frac{d}{d\tau} u_{\theta_\tau}^N(t, \bullet); t \right] \\
 &\quad - \left(q_u(t, \bullet, u_{\theta_\tau}^N(t, \bullet)) \widehat{u}_{\theta_\tau}^N(t, \bullet), \frac{d}{d\tau} u_{\theta_\tau}^N(t, \bullet) \right)_{L_2(D)} dt \\
 &= \int_0^T \left\langle \partial_t \frac{d}{d\tau} u_{\theta_\tau}^N(t, \bullet), \widehat{u}_{\theta_\tau}^N(t, \bullet) \right\rangle_{H^{-1}(D), H_0^1(D)} + \mathcal{B} \left[\frac{d}{d\tau} u_{\theta_\tau}^N(t, \bullet), \widehat{u}_{\theta_\tau}^N(t, \bullet); t \right] \\
 &\quad - \left(q_u(t, \bullet, u_{\theta_\tau}^N(t, \bullet)) \frac{d}{d\tau} u_{\theta_\tau}^N(t, \bullet), \widehat{u}_{\theta_\tau}^N(t, \bullet) \right)_{L_2(D)} dt \\
 &= \int_0^T \left(\frac{d}{d\tau} g_{\theta_\tau}^N(t, \bullet), \widehat{u}_{\theta_\tau}^N(t, \bullet) \right)_{L_2(D)} dt = \int_0^T \int_D \left(\frac{d}{d\tau} g_{\theta_\tau}^N(t, x) \right) \widehat{u}_{\theta_\tau}^N(t, x) dx dt,
 \end{aligned} \tag{283}$$

where the individual steps are analogous to the ones taken in (92) in the proof of Proposition 26.

Now, recalling the representation of the right-hand side $g_{\theta_\tau}^N$ from (11) and taking its training time derivative to obtain $\frac{d}{d\tau} g_{\theta_\tau}^N = -\alpha_\tau T_{B(\mu_\tau^N)} \widehat{u}_{\theta_\tau}^N$, as well as recalling the definition of the operator $T_{B(\mu_\tau^N)}$ from (12), we can continue (283) to obtain

$$\begin{aligned}
 \frac{d}{d\tau} \mathcal{J}_{\theta_\tau}^N &= -\alpha_\tau \int_0^T \int_D \widehat{u}_{\theta_\tau}^N(t, x) \int_0^T \int_D B(t, x, t', x'; \mu_\tau^N) \widehat{u}_{\theta_\tau}^N(t', x') dx' dt' dx dt \\
 &= -\alpha_\tau (\widehat{u}_{\theta_\tau}^N, T_{B(\mu_\tau^N)} \widehat{u}_{\theta_\tau}^N)_{L_2(D_T)},
 \end{aligned} \tag{284}$$

which concludes the first part of the proof.

The second part now follows immediately thanks to the operator $T_{B(\mu_\tau^N)}$ being positive semi-definite for every $\tau \in I$, as can be seen by noting that with computations analogous to (83) it holds

$$\begin{aligned}
 (\widehat{u}, T_{B(\mu_\tau^N)} \widehat{u})_{L_2(D_T)} &= \int_0^T \int_D \widehat{u}(t, x) \int_0^T \int_D B(t, x, t', x'; \mu_\tau^N) \widehat{u}(t', x') dx' dt' dx dt \\
 &\geq \int \left(\int_0^T \int_D \sigma(w^t t + w^T x + \eta) \widehat{u}(t, x) dx dt \right)^2 d\mu_{\tau, (w^t, w, \eta)}^N(w^t, w, \eta) \\
 &\geq 0,
 \end{aligned} \tag{285}$$

which concludes the proof. ■

An immediate consequence of the loss $\mathcal{J}_{\theta_\tau}^N$ being monotonically non-increasing, are uniform (in the training time τ) bounds on the L_2 norm of the PDE solution $u_{\theta_\tau}^N$ and the $L_2([0, T], H^1(D))$ and $L_\infty([0, T], L_2(D))$ norm of the adjoint $\widehat{u}_{\theta_\tau}^N$.

Lemma 39 *Assume that the learning rate satisfies additionally $\int_0^\infty \alpha_\tau^{4/3} d\tau < \infty$. Let $((u_{\theta_\tau}^N, \widehat{u}_{\theta_\tau}^N))_{\tau \in I} \in \mathcal{C}(I, \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (1) & (8) coupled with the gradient descent update (5) in the sense of Lemma 11 on the training time interval I . Then the solution $u_{\theta_\tau}^N$ is uniformly (in the training time τ) bounded in $L_2(D_T)$ on that interval I , i.e., it holds*

$$\sup_{\tau \in I} \|u_{\theta_\tau}^N\|_{L_2(D_T)} \leq C^{u^N} \tag{286}$$

for the constant $C^{u^N} = 4\mathcal{J}_{\theta_0}^N + 2\|h\|_{L_2(D_T)}^2$.

Proof Using Lemma 38, the proof follows the one of Lemma 27. ■

Lemma 40 *Assume that the learning rate satisfies additionally $\int_0^\infty \alpha_\tau^{4/3} d\tau < \infty$. Let $((u_{\theta_\tau}^N, \widehat{u}_{\theta_\tau}^N))_{\tau \in I} \in \mathcal{C}(I, \mathcal{S} \times \mathcal{S})$ denote the unique weak solution to the PDE system (1) & (8) coupled with the gradient descent update (5) in the sense of Lemma 11 on the training time interval I . Then the adjoint $\widehat{u}_{\theta_\tau}^N$ in (8) is uniformly (in the training time τ) bounded in $L_2([0, T], H^1(D))$ and $L_\infty([0, T], L_2(D))$ on that interval I , i.e., it holds*

$$\sup_{\tau \in I} \left(\|\widehat{u}_{\theta_\tau}^N\|_{L_2([0, T], H^1(D))} + \|\widehat{u}_{\theta_\tau}^N\|_{L_\infty([0, T], L_2(D))} \right) \leq C^{\widehat{u}^N} \tag{287}$$

for a constant $C^{\widehat{u}^N} = C^{\widehat{u}^N}(T, \mathcal{L}, \mathcal{J}_0^*)$.

Proof Using Lemma 38, the proof follows the one of Lemma 28. ■

References

- C. Aarset, M. Holler, and T. T. N. Nguyen. Learning-informed parameter identification in nonlinear time-dependent PDEs. *Appl. Math. Optim.*, 88(3):Paper No. 76, 53, 2023.
- D. P. Bertsekas and J. N. Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM J. Optim.*, 10(3):627–642, 2000.
- T. Bosse, N. R. Gauger, A. Griewank, S. Günther, and V. Schulz. One-shot approaches to design optimization. In *Trends in PDE constrained optimization*, volume 165 of *Internat. Ser. Numer. Math.*, pages 43–66. Birkhäuser/Springer, Cham, 2014.

- N. Boullé and A. Townsend. Chapter 3 – A mathematical guide to operator learning. In S. Mishra and A. Townsend, editors, *Numerical Analysis Meets Machine Learning*, volume 25 of *Handbook of Numerical Analysis*, pages 83–125. Elsevier, 2024.
- C. Brandenburg, F. Lindemann, M. Ulbrich, and S. Ulbrich. A continuous adjoint approach to shape optimization for Navier Stokes flow. In *Optimal control of coupled systems of partial differential equations*, volume 158 of *Internat. Ser. Numer. Math.*, pages 35–56. Birkhäuser Verlag, Basel, 2009.
- H. Brezis. *Functional analysis, Sobolev spaces and partial differential equations*. Universitext. Springer, New York, 2011.
- S. L. Brunton and J. N. Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2019.
- S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences*, 113(15):3932–3937, 2016.
- S. L. Brunton, B. R. Noack, and P. Koumoutsakos. Machine learning for fluid mechanics. *Annual review of fluid mechanics*, 52(1):477–508, 2020.
- A. Bueno-Orovio, C. Castro, F. Palacios, and E. Zuazua. Continuous adjoint approach for the spalart-allmaras model in aerodynamic optimization. *AIAA journal*, 50(3):631–646, 2012.
- F. Cagnetti, D. Gomes, and H. V. Tran. Adjoint methods for obstacle problems and weakly coupled systems of PDE. *ESAIM Control Optim. Calc. Var.*, 19(3):754–779, 2013.
- K. Champion, B. Lusch, J. N. Kutz, and S. L. Brunton. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116(45):22445–22451, 2019.
- L. Chizat, E. Oyallon, and F. R. Bach. On lazy training in differentiable programming. In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 2933–2943, 2019.
- J. Cohen, S. Kaur, Y. Li, J. Z. Kolter, and A. Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021.
- J. M. Cohen, B. Ghorbani, S. Krishnan, N. Agarwal, S. Medapati, M. Badura, D. Suo, D. Cardoze, Z. Nado, G. E. Dahl, and J. Gilmer. Adaptive gradient methods at the edge of stability. *arXiv preprint arXiv:2207.14484*, 2022.
- S. N. Cohen, D. Jiang, and J. Sirignano. Neural Q-learning for solving PDEs. *J. Mach. Learn. Res.*, 24:Paper No. [236], 49, 2023a.

- S. N. Cohen, C. Reisinger, and S. Wang. Arbitrage-free neural-SDE market models. *Applied Mathematical Finance*, 30(1):1–46, 2023b.
- S. Cuomo, V. Schiano Di Cola, F. Giampaolo, G. Rozza, M. Raissi, and F. Piccialli. Scientific machine learning through physics-informed neural networks: where we are and what's next. *J. Sci. Comput.*, 92(3):Paper No. 88, 62, 2022.
- G. Cybenko. Approximation by superpositions of a sigmoidal function. *Math. Control. Signals Syst.*, 2(4):303–314, 1989.
- A. Damian, E. Nichani, and J. D. Lee. Self-stabilization: The implicit bias of gradient descent at the edge of stability. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*, 2023.
- G. Dong, M. Hintermüller, and K. Papafitsoros. Optimization with learning-informed differential equation constraints and its applications. *ESAIM Control Optim. Calc. Var.*, 28:Paper No. 3, 44, 2022.
- G. Dong, M. Hintermüller, and K. Papafitsoros. A descent algorithm for the optimal control of relu neural network informed PDEs based on approximate directional derivatives. *SIAM Journal on Optimization*, 34(3):2314–2349, 2024.
- K. Duraisamy. Perspectives on machine learning-augmented Reynolds-averaged and large eddy simulation models of turbulence. *Physical Review Fluids*, 6(5):050504, 2021.
- K. Duraisamy, G. Iaccarino, and H. Xiao. Turbulence modeling in the age of data. *Annual review of fluid mechanics*, 51(1):357–377, 2019.
- M. Duta, M. Giles, and M. Campobasso. The harmonic adjoint approach to unsteady turbomachinery design. *International Journal for Numerical Methods in Fluids*, 40(3-4): 323–332, 2002.
- W. E and B. Yu. The deep Ritz method: a deep learning-based numerical algorithm for solving variational problems. *Commun. Math. Stat.*, 6(1):1–12, 2018.
- L. C. Evans. *Partial differential equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, second edition, 2010.
- L. Fan and J. Sirignano. Machine learning methods for pricing financial derivatives. *arXiv preprint arXiv:2406.00459*, 2024.
- N. Gauger, A. Griewank, A. Hamdi, C. Kratzenstein, E. Özkaya, and T. Slawig. Automated extension of fixed point PDE solvers for optimal design with bounded retardation. In *Constrained optimization and optimal control for partial differential equations*, volume 160 of *Internat. Ser. Numer. Math.*, pages 99–122. Birkhäuser/Springer Basel AG, Basel, 2012.
- P. Gierjatowicz, M. Sabate-Vidales, D. Šiška, L. Szpruch, and Ž. Žurič. Robust pricing and hedging via neural SDEs. *arXiv preprint arXiv:2007.04154*, 2020.

- M. Giles and S. Ulbrich. Convergence of linearized and adjoint approximations for discontinuous solutions of conservation laws. Part 1: Linearized approximations and linearized output functionals. *SIAM J. Numer. Anal.*, 48(3):882–904, 2010a.
- M. Giles and S. Ulbrich. Convergence of linearized and adjoint approximations for discontinuous solutions of conservation laws. Part 2: Adjoint approximations and extensions. *SIAM J. Numer. Anal.*, 48(3):905–921, 2010b.
- M. B. Giles and N. A. Pierce. An introduction to the adjoint approach to design. *Flow, turbulence and combustion*, 65:393–415, 2000.
- A. Goswami, S. Rajani, and A. Tanksale. Data-driven option pricing using single and multi-asset supervised learning. *International Journal of Financial Engineering*, 8(02):2141001, 2021.
- S. B. Hazra. Direct treatment of state constraints in aerodynamic shape optimization using simultaneous pseudo-time-stepping. *AIAA journal*, 45(8):1988–1997, 2007.
- S. B. Hazra and V. Schulz. Simultaneous pseudo-timestepping for PDE-model based optimization problems. *Bit Numerical Mathematics*, 44:457–472, 2004.
- T. Hickling, J. Sirignano, and J. F. MacArt. Large eddy simulation of airfoil flows using adjoint-trained deep learning closure models. In *AIAA SCITECH 2024 Forum*, page 0296, 2024.
- G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, and B. Kingsbury. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Process. Mag.*, 29(6):82–97, 2012.
- M. Hinze, R. Pinnau, M. Ulbrich, and S. Ulbrich. *Optimization with PDE constraints*, volume 23. Springer Science & Business Media, 2008.
- J. R. Holland, J. D. Baeder, and K. Duraisamy. Towards integrated field inversion and machine learning with embedded neural networks for rans modeling. In *AIAA Scitech 2019 forum*, page 1884, 2019.
- M. Holler and E. Morina. On uniqueness in structured model learning. *arXiv preprint arXiv:2410.22009*, 2024.
- K. Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257, 1991.
- A. Jacot, C. Hongler, and F. Gabriel. Neural tangent kernel: Convergence and generalization in neural networks. In S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 8580–8589, 2018.

- A. Jameson. Aerodynamic shape optimization using the adjoint method. *Lectures at the Von Karman Institute, Brussels*, 6, 2003.
- A. Jameson and S. Kim. Reduction of the adjoint gradient formula in the continuous limit. In *41st Aerospace Sciences Meeting and Exhibit*, page 40, 2003.
- A. Jameson, L. Martinelli, and N. A. Pierce. Optimum aerodynamic design using the Navier–Stokes equations. *Theoretical and computational fluid dynamics*, 10(1):213–237, 1998.
- D. Jiang, J. Sirignano, and S. N. Cohen. Global convergence of deep galerkin and pinns methods for solving partial differential equations. *arXiv preprint arXiv:2305.06000*, 2023.
- J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589, 2021.
- P. Kakka and J. F. MacArt. Neural network-augmented eddy viscosity closures for turbulent premixed jet flames. *Combustion and Flame*, 278:114241, 2025.
- L. Kaland, J. C. De Los Reyes, and N. R. Gauger. One-shot methods in function space for PDE-constrained optimal control problems. *Optim. Methods Softw.*, 29(2):376–405, 2014.
- G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- D. A. Knopoff, D. R. Fernández, G. A. Torres, and C. V. Turner. Adjoint method for a tumor growth PDE-constrained optimization problem. *Computers & Mathematics with Applications*, 66(6):1104–1119, 2013.
- D. Kochkov, J. A. Smith, A. Alieva, Q. Wang, M. P. Brenner, and S. Hoyer. Machine learning-accelerated computational fluid dynamics. *Proc. Natl. Acad. Sci. USA*, 118(21): Paper No. e2101784118, 8, 2021.
- N. B. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. M. Stuart, and A. Anandkumar. Neural operator: Learning maps between function spaces with applications to PDEs. *J. Mach. Learn. Res.*, 24:89:1–89:97, 2023.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In P. L. Bartlett, F. C. N. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States*, pages 1106–1114, 2012.
- A. Kumar, L. Owen, N. R. Chowdhury, and F. Güra. Zclip: Adaptive spike mitigation for llm pre-training. *arXiv preprint arXiv:2504.02507*, 2025.
- O. A. Ladyženskaja, V. A. Solonnikov, and N. N. Ural’ceva. *Linear and quasilinear equations of parabolic type*, volume Vol. 23 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 1968. Translated from the Russian by S. Smith.

- Y. LeCun, Y. Bengio, and G. E. Hinton. Deep learning. *Nat.*, 521(7553):436–444, 2015.
- Z. Li, N. B. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. M. Stuart, and A. Anandkumar. Fourier neural operator for parametric partial differential equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021.
- L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nat. Mach. Intell.*, 3(3):218–229, 2021a.
- L. Lu, X. Meng, Z. Mao, and G. E. Karniadakis. Deepxde: A deep learning library for solving differential equations. *SIAM Review*, 63(1):208–228, 2021b.
- J. F. MacArt, J. Sirignano, and J. B. Freund. Embedded training of neural-network subgrid-scale turbulence models. *Physical Review Fluids*, 6(5):050502, 2021.
- S. Nadarajah and A. Jameson. A comparison of the continuous and discrete adjoint approach to automatic aerodynamic optimization. In *38th Aerospace sciences meeting and exhibit*, page 667, 2000.
- S. Nadarajah and A. Jameson. Studies of the continuous and discrete adjoint approaches to viscous automatic aerodynamic shape optimization. In *15th AIAA computational fluid dynamics conference*, page 2530, 2001.
- A. S. Nair, J. Sirignano, M. Panesi, and J. F. MacArt. Deep learning closure of the Navier–Stokes equations for transition-continuum flows. *AIAA journal*, 61(12):5484–5497, 2023.
- A. S. Nair, D. Waidmann, J. Sirignano, N. Singh, M. Panesi, and J. F. MacArt. Adjoint-trained deep-learning closures of the Navier–Stokes equations for 2d nonequilibrium flows. In *AIAA SCITECH 2024 Forum*, page 2860, 2024.
- N. A. Pierce and M. B. Giles. Adjoint recovery of superconvergent functionals from PDE approximations. *SIAM Review*, 42(2):247–264, 2000.
- M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- J. Reuther, A. Jameson, J. Farmer, L. Martinelli, and D. Saunders. Aerodynamic shape optimization of complex aircraft configurations via an adjoint formulation. In *34th aerospace sciences meeting and exhibit*, page 94, 1996.
- H. Robbins and S. Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
- H. Schaeffer. Learning partial differential equations via data discovery and sparse optimization. *Proc. R. Soc. A.*, 473(2197):20160446, 20, 2017.

- J. Sirignano and J. F. MacArt. Deep learning closure models for large-eddy simulation of flows around bluff bodies. *J. Fluid Mech.*, 966:Paper No. A26, 23, 2023.
- J. Sirignano and K. Spiliopoulos. DGM: a deep learning algorithm for solving partial differential equations. *J. Comput. Phys.*, 375:1339–1364, 2018.
- J. Sirignano and K. Spiliopoulos. Online adjoint methods for optimization of PDEs. *Appl. Math. Optim.*, 85(2):Paper No. 18, 29, 2022.
- J. Sirignano, J. F. MacArt, and J. B. Freund. DPM: a deep learning PDE augmentation method with application to large-eddy simulation. *J. Comput. Phys.*, 423:109811, 21, 2020.
- J. Sirignano, J. MacArt, and K. Spiliopoulos. PDE-constrained models with neural network terms: optimization and global convergence. *J. Comput. Phys.*, 481:Paper No. 112016, 35, 2023.
- V. Srivastava and K. Duraisamy. Generalizable physics-constrained modeling using learning and inference assisted by feature-space engineering. *Physical Review Fluids*, 6(12):124602, 2021.
- Y. Sun, L. Zhang, and H. Schaeffer. NeuPDE: Neural network based ordinary and partial differential equations for modeling time-dependent data. In J. Lu and R. A. Ward, editors, *Proceedings of Mathematical and Scientific Machine Learning, MSML 2020, 20-24 July 2020, Virtual Conference / Princeton, NJ, USA*, volume 107 of *Proceedings of Machine Learning Research*, pages 352–372. PMLR, 2020.
- B. D. Tracey, K. Duraisamy, and J. J. Alonso. A machine learning strategy to assist turbulence model development. In *53rd AIAA aerospace sciences meeting*, page 1287, 2015.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017.
- S. Wang, S. Sankaran, H. Wang, and P. Perdikaris. An expert’s guide to training physics-informed neural networks. *arXiv preprint arXiv:2308.08468*, 2023.
- W. Zhao and T. Luo. Convergence guarantees for gradient-based training of neural PDE solvers: From linear to nonlinear PDEs. *arXiv preprint arXiv:2505.14002*, 2025.