FREDHOLM INTEGRAL EQUATIONS FOR FUNCTION APPROXIMATION AND THE TRAINING OF NEURAL NETWORKS

PATRICK GELSS, AIZHAN ISSAGALI, AND RALF KORNHUBER

ABSTRACT. We present a novel and mathematically transparent approach to function approximation and the training of large, high-dimensional neural networks, based on the approximate least-squares solution of associated Fredholm integral equations of the first kind by Ritz-Galerkin discretization, Tikhonov regularization and tensor-train methods. Practical application to supervised learning problems of regression and classification type confirm that the resulting algorithms are competitive with state-of-the-art neural network-based methods.

1. INTRODUCTION

Efficient and reliable methods for the training of neural networks are of fundamental importance for a multitude of applications and therefore a flourishing field of mathematical research. Stochastic gradient methods [1, 2] are known and further developed since the sixties. These methods mostly perform very well and therefore are often considered as methods of choice in the field. However, global convergence analysis is complicated and actual convergence might depend on proper parameter tuning [3]. Multilevel trust region methods [4, 5, 6] on the one hand aim at increasing efficiency by Newton-type linearization and on the other hand at increasing reliability by suitable step size control. Multilevel versions [7] are intended to additionally transfer the efficiency of multigrid methods for the minimization of energy functionals associated with elliptic partial differential equations to the minimization of loss functionals.

This paper is devoted to a novel and mathematically transparent approach to function approximation and the training of large and high-dimensional shallow neural networks with usual kernel functions ψ, training parameters $(\eta_i, u_i)_{i=1}^N$ with $\eta_i \in Q \subset \mathbb{R}^s$, $u_i \in \mathbb{R}$, $I = 1, \ldots, N$, and intentionally large s and N. The starting point of our considerations is that such neural networks can be interpreted as Monte-Carlo approximations of integrals over Q, once the parameters η_i are assumed to be equidistributed on Q (see also [8]). Resulting Fredholm integral operators, now acting on parameter functions rather than parameter vectors, give rise to so-called *Fredholm networks* providing a novel approach to function approximation. Training of such Fredholm networks amounts to least-squares solution of linear Fredholm integral equations of the first kind for a parameter function $u: Q \to \mathbb{R}$, the continuous counterpart of the discrete parameter vector $(u_i)_{i=1}^N$. Fredholm-trained parameters $(\eta_i, u(\eta_i))_{i=1}^N$ of the original discrete neural network are then obtained by suitable sampling of $\eta_i \in Q$.

A similar approach has been developed for kernel methods in supervised and semi-supervised learning [9, 10, 11] and applied to the covariance shift problem in transfer learning [12]. Recent mean-field neural networks [13] can be regarded as Fredholm networks on measures. The associated mean-field approach for the training of neural networks is based on stochastic gradient flows for the minimization of corresponding (regularized) loss functionals on the space of probability measures. It provides convergence results for well-known stochastic gradient descent [14, 15, 16] and suitable generalizations [17, 18].

In this paper, the continuous Fredholm training problem is discretized by Ritz-Galerkin methods with suitable ansatz spaces, and we apply Tikhonov regularization [19]. Instead of gradient descent, the resulting (intentionally large) linear algebraic systems are then solved approximately by tensor-train methods, such as alternating linear schemes [20, 21], in order to enhance efficiency and to allow for more general architectures. Approximately Fredholm-trained parameters $(\tilde{\eta}_i, \tilde{u}(\tilde{\eta}_i))_{i=1}^N$ of the original neural network are finally obtained from the resulting approximation $\tilde{u}$ of u and suitable sampling of $\tilde{\eta}_i \in Q$. A priori residual error estimates for the loss functional, both for trained Fredholm and neural networks, are given in the special case of piecewise constant approximation and exact linear solution. While a comprehensive numerical analysis and possible extensions to deep neural networks are devoted to forthcoming papers, a major goal of this work is to confirm the practical potential of the Fredholm approach by numerical experiments and comparison with state-of-the-art neural network-based methods.

The paper is organized as follows. After describing the architecture of the neural networks in question together with the resulting discrete training problem in the next Section 2, we introduce Fredholm networks and Fredholm training problems as their continuous counterparts and briefly discuss ill-posedness in Section 3. Ritz-Galerkin discretization and Tikhonov regularization are presented in Section 4. Utilizing suitable factorizations of the kernel functions ψ (as derived in the Appendix A), Section 5 is devoted to tensor-train methods for the iterative solution of the large linear systems associated with intentionally high-dimensional parameter domains $Q \subset \mathbb{R}^s$. In Section 6, we harvest the results on the training of Fredholm networks, briefly discuss Monte-Carlo sampling, and introduce a novel inductive importance sampling strategy. In the next Section 7, we apply our methods to three well-established test problems of regression and classification type with different computational complexity. In particular, we consider the UCI banknote authentication data set from [22], concrete compressive strength prediction [23, 24, 25, 26], and the well-known MNIST data set [27]. For all these problems, Fredholm networks and Fredholm-trained neural networks turn out to be highly competitive with state-of-the-art approaches without any problem-specific features or tuning. Note that, with some linear algebraic tricks, neural tangent kernel (NTK) methods [28] are able to achieve comparable performance even on more challenging datasets like CIFAR-10 [29]. The efficiency of Fredholm-trained neural networks however seems to rely on a careful choice of sampling strategy: While straightforward quasi-Monte-Carlo sampling appears to work perfectly well for simple problems, more advanced strategies, such as inductive importance sampling, seem to be required in more complicated cases. The final Section 8 is devoted to some concluding remarks and a brief outlook on deep Fredholm networks.

2. Large single layer neural networks

2.1. **Architecture.**
We consider a parametrized single layer neural network of the form

$$(1) \qquad X \ni x \mapsto F_N(x, \zeta) = \frac{1}{N} \sum_{i=1}^{N} \psi(x, \eta_i) u_i \in \mathbb{R}$$

defined on a compact hypercube $X \subset \mathbb{R}^d$, for simplicity, with intentionally large number of neurons $N \in \mathbb{N}$, a suitable kernel function $\psi : X \times Q \to \mathbb{R}$, and associated parameter vector

$$\zeta = (\zeta_i)_{i=1}^{N}, \quad \zeta_i = (\eta_i, u_i) \in Q \times \mathbb{R},$$

where $Q = Q_1 \times \cdots \times Q_s \subset \mathbb{R}^s$.

Ridge kernel functions

$$(2) \qquad \psi(x, \eta) = \sigma(w \cdot x + b), \quad \eta = (w, b) \in Q \subset \mathbb{R}^s, \ s = d + 1,$$

are characterized by suitable activation functions $\sigma : \mathbb{R} \to \mathbb{R}$ and give rise to so-called single hidden layer perceptron models [30]. In the simplest case $\sigma \equiv \mathrm{Id}$ and $b = 0$, this definition reduces to the Euclidean scalar product $\psi(x, \eta) = x \cdot \eta$ with $\eta = w \in \mathbb{R}^s$ and $s = d$. More relevant choices for σ include the (discontinuous) *Heaviside function*

$$(3) \qquad \sigma_\infty(z) = \begin{cases} 0 & \text{if } z < 0, \\ 1/2 & \text{if } z = 0, \\ 1 & \text{if } z > 0 \end{cases}$$

or suitable regularizations like the *logistic activation function*

$$(4) \qquad \sigma_\kappa(z) = \frac{\exp(\kappa z)}{1 + \exp(\kappa z)}$$

with $\kappa > 0$. Note that $\sigma_\kappa(z) \to \sigma_\infty(z)$ holds for $\kappa \to \infty$ and all $z \in \mathbb{R}$.

Radial kernel functions [31] are given by

$$(5) \qquad \psi(x, \eta) = \sigma(\|x - w\|), \quad \eta = w \in Q \subset \mathbb{R}^s, \ s = d,$$

with some vector norm $\| \cdot \|$ on $\mathbb{R}^d$ and suitable activation functions σ. In the case of Gaussian activation functions

$$(6) \qquad \sigma(z) = \exp\left(-\frac{z^2}{2\kappa^2}\right)$$

and the Euclidean norm $\| \cdot \| = \| \cdot \|_2$, we have

$$(7) \qquad \psi(x, \eta) = \exp\left(-\frac{\|x - \eta\|_2^2}{2\kappa^2}\right) = \prod_{k=1}^{d} \exp\left(-\frac{(x_k - \eta_k)^2}{2\kappa^2}\right).$$

2.2. **Function data, loss functional, and training.**
Throughout the following $L^2(X, \pi)$ stands for the Hilbert space of π measurable, quadratically integrable functions on X with canonical scalar product $(\cdot, \cdot)_\pi$ and induced norm $\| \cdot \|_\pi$. We write $L^2(X) = L^2(X, \pi)$ and $(\cdot, \cdot)_X = (\cdot, \cdot)_\pi$, $\| \cdot \|_X = \| \cdot \|_\pi$, if $d\pi(x) = dx$ is the Lebesgue measure. Furthermore, $L^2(Q)$ stands for the Hilbert space of Lebesgue measurable, quadratically integrable functions on Q with canonical scalar product $(\cdot, \cdot)$ and induced norm $\| \cdot \|$.

We aim at the approximation of a function $F : X \mapsto \mathbb{R}$ by parametrized ansatz functions $F_N(\cdot, \zeta)$ of the form (1). The parameters ζ are determined by minimizing the loss functional

$$(8) \qquad \mathcal{J}_N(\zeta) = \|F - F_N(\cdot, \zeta)\|_\pi^2 = \int_X (F(x) - F_N(x, \zeta))^2 \, d\pi(x)$$

where the measure π is associated with the training data available. We will consider the following two cases of given

(9) a) function data $F \in L^2(X)$ b) point data $((x_j, F(x_j)))_{j=1}^M \in (X \times \mathbb{R})^M \subset \mathbb{R}^{M(d+1)}$.

The case a) of a given function F is associated with the Lebesgue measure $d\pi(x) = dx$ while case b) of given point data corresponds to the point measure $d\pi(x) = d\mu(x)$,

$$(10) \qquad \mu = \frac{1}{M} \sum_{j=1}^M \delta_{x_j}, \quad x_1, \ldots, x_M \in X.$$

Training of the neural network $F_N(\cdot, \zeta)$ amounts to the (approximate) solution of the following discrete, non-convex, large-scale minimization problem.

Problem 2.1. *Find $\zeta^* = (\zeta_i^*)_{i=1}^N$, $\zeta_i^* = (\eta_i^*, u_i^*) \in Q \times \mathbb{R}$, such that*

$$(11) \qquad \mathcal{J}_N(\zeta^*) \leq \mathcal{J}_N(\zeta) \qquad \forall \zeta = (\zeta_i)_{i=1}^N, \ \zeta_i = (\eta_i, u_i) \in Q \times \mathbb{R}.$$

Approximation properties of neural networks of the form (1) have quite a history (see, e.g., [32] or the overview of Pinkus [30]). In particular, ridge kernels with any non-polynomial activation function σ and unconstrained parameters $\zeta \in \mathbb{R}^{N(d+2)}$ provide uniform approximation of continuous functions F on compact sets $X \subset \mathbb{R}^d$ [30, Theorem 3.1]. While similar results for constrained parameters $\zeta \in Q \times \mathbb{R}$ seem to be rare, the following proposition is a consequence of [33, Theorem 2.6].

Proposition 2.2. *For the ridge kernel (2) with logistic activation function (4), Q containing the unit ball in $\mathbb{R}^{d+1}$, and any function $F \in C(X)$, we have*

$$(12) \qquad \inf_{\zeta \in (Q \times \mathbb{R})^N} \max_{x \in X} |F(x) - F_N(x, \zeta)| \to 0 \quad for \quad N \to \infty.$$

Hence, assuming existence of a solution ζ^* for all $N \in \mathbb{N}$, we have $\mathcal{J}_N(\zeta^*) \to 0$ as $N \to \infty$ for $F \in L^2(X)$ in the case (9a) and $F \in C(X)$ in the case (9b), respectively.

3. Continuous approximations of discrete training problems

Following the ideas from [8, Section 1.2], we now introduce and investigate the asymptotic limit of $F_N(\cdot, \zeta)$ and of the corresponding Training Problem 2.1 for $N \to \infty$.

3.1. **Asymptotic integral operators.** Assume that $\psi(x, \cdot) \in L^2(Q)$ for each fixed $x \in X$. Let η_i, $i = 1, \ldots, N$, be independent, identically equidistributed random samples and let $u_i = u(\eta_i)$, $i = 1, \ldots, N$, be the values of an arbitrary function $u \in L^2(Q)$ at those sample points. Then, denoting $|Q| = \int_Q 1 \, dx$,

$$(13) \qquad \frac{1}{N} \sum_{i=1}^N \psi(x, \eta_i) u_i \ \to \ \frac{1}{|Q|} \int_Q \psi(x, \eta) u(\eta) \, d\eta \quad for \quad N \to \infty$$

holds for each fixed $x \in X$ by the strong law of large numbers.

This observation suggests to consider the Fredholm integral operator

$$(14) \qquad \mathcal{G}v = \frac{1}{|Q|} \int_Q \psi(\cdot, \eta) v(\eta) \, d\eta, \quad v \in L^2(Q),$$

Let us recall some well-known properties of $\mathcal{G}$, cf. [19, 34, 35].

Proposition 3.1. *Assume that*

$$(15) \qquad \int_X \int_Q \psi(x, \eta)^2 \, d\eta \, dx < \infty.$$

Then the integral operator

$$\mathcal{G} : L^2(Q) \to L^2(X)$$

is a compact, linear mapping.

Note that Proposition 3.1 provides compactness of $\mathcal{G}$, e.g., for the radial and ridge kernel functions mentioned in Section 2.1.

3.2. **Fredholm integral equations.**

3.2.1. *Function data.* Let us first consider the case (9a) of a given function $F \in L^2(X)$. In the light of (13), we impose the additional constraints on the unknowns $\zeta_i = (\eta_i, u_i)$, $i = 1, \ldots, N$, of the Training Problem 2.1 that

(16) η_i are samples of i.i. equidistributed random variables and $u_i = u(\eta_i)$, $i = 1, \ldots, N$,

with some function $u \in L^2(Q)$. Then, for given function F, the asymptotic coincidence (13) of the discrete neural network $F_N(\cdot, ((\eta_i, u_i)_{i=1}^N)$ and the integral operator $\mathcal{G}u$ suggests to consider the loss functional

$$(17) \qquad \mathcal{J}(v) = \int_X (F(x) - \mathcal{G}v(x))^2 \, dx.$$

and the following continuous quadratic approximation of the discrete non-convex Training Problem 2.1.

Problem 3.2. *Find $u \in L^2(Q)$ such that*

$$(18) \qquad \mathcal{J}(u) \le \mathcal{J}(v) \quad \forall v \in L^2(Q).$$

Observe that Problem 3.2 is just the least-squares formulation of the Fredholm integral equation of the first kind

$$(19) \qquad \mathcal{G}u = \frac{1}{|Q|} \int_Q \psi(\cdot, \eta) u(\eta) \, d\eta = F.$$

In particular, (18) can be equivalently rewritten as the normal equation

$$(20) \qquad \mathcal{G}^* \mathcal{G} u = \mathcal{G}^* F \quad \text{in} \quad L^2(Q)$$

utilizing the variational formulation of (18),

$$(21) \qquad (\mathcal{G}u, \mathcal{G}v)_X = (F, \mathcal{G}v)_X \qquad \forall v \in L^2(Q),$$

and the adjoint operator of $\mathcal{G}$,

$$L^2(X) \ni v \mapsto \mathcal{G}^* v = \frac{1}{|Q|} \int_X \psi(x, \cdot) v(x) \, dx \in L^2(Q).$$

This leads to the following well-known existence result (see, e.g., [19, Theorem 2.5]), where $R(\mathcal{G})$ and $N(\mathcal{G})$ stand for the range and null space of $\mathcal{G}$, respectively, $\oplus$ denotes the direct sum, and the superscript $\perp$ indicates the orthogonal complement in $L^2(X)$.

Proposition 3.3. *Assume that $F \in R(\mathcal{G}) + R(\mathcal{G})^\perp \subset L^2(X)$. Then (19) has least-squares solutions $u^\dagger \oplus N(\mathcal{G})$, with $u^\dagger = \mathcal{G}^\dagger F$ being the unique minimal-norm least squares solution (best-approximate solution), and $\mathcal{G}^\dagger \colon R(\mathcal{G}) \oplus R(\mathcal{G})^\perp \to N(\mathcal{G})^\perp \subset L^2(Q)$ denoting the generalised Moore-Penrose inverse of $\mathcal{G}$.*

If $\mathcal{G}$ is compact, then $R(\mathcal{G})$ is closed or, equivalently, $\mathcal{G}^\dagger$ is bounded, if and only if $R(\mathcal{G})$ has finite dimension (see, e.g., [19, Proposition 2.7]). As a consequence, the integral equation (19) is ill-posed, even in the least-squares sense (18) for usual kernel functions such as the radial and ridge kernels presented in Section 2.1.

Proposition 3.4. *Assume that the kernel function ψ is discriminatory in the sense that*

$$(22) \qquad N(\mathcal{G}^*) = \{0\}.$$

Then $R(\mathcal{G})$ is dense in $L^2(X)$ and therefore $\inf_{u \in L^2(Q)} \mathcal{J}(u) = 0$ holds for all $F \in L^2(X)$.

For ridge kernels with Heaviside or logistic activation function, we obtain $v = 0$, if $\mathcal{G}^* v(\eta) = 0$ holds for all $\eta \in \mathbb{R}^s$ (and not only for all $\eta \in Q$) by [36, Lemma 1]. However, similar results for the present case of bounded parameters $\eta \in Q$ seem to be unknown.

3.2.2. *Point data.* In the case (9b) of given point values of F, we replace the corresponding point values of F_N in (8) by the point values of the limiting integral operator $\mathcal{G}$ to obtain the approximation J,

$$(23) \qquad J(v) = \frac{1}{M} \sum_{j=1}^{M} (F(x_j) - \mathcal{G}v(x_j))^2,$$

of the discrete loss functional $\mathcal{J}_N$ and the following quadratic approximation of the Training Problem 2.1.

Problem 3.5. *Find $u \in L^2(Q)$ such that*

$$(24) \qquad J(u) \leq J(v) \quad \forall v \in L^2(Q).$$

Problem 3.5 is the least-squares formulation of the interpolation problem to find $u \in L^2(Q)$ such that

$$\mathcal{G}_M u = F_M$$

with $F_M = (F(x_j))_{j=1}^{M} \in \mathbb{R}^M$ and $\mathcal{G}_M \colon L^2(Q) \to \mathbb{R}^M$ defined by $\mathcal{G}_M v = (\mathcal{G}v(x_i))_{i=1}^{M} \in \mathbb{R}^M$. As a semi-discrete counterpart of (15), the condition

$$(25) \qquad \frac{1}{M} \sum_{i=1}^{M} \int_Q \psi(x_i, \eta)^2 d\eta \leq C < \infty$$

with a constant C independent of M implies that $\mathcal{G}_M$ is uniformly bounded in M. It also holds for all kernel functions mentioned in Section 2.1.

In analogy to (20), Problem 3.5 can be equivalently reformulated as a linear normal equation $\mathcal{G}_M^* \mathcal{G}_M u = \mathcal{G}_M^* F_M$ in $L^2(Q)$ utilizing the variational formulation

$$(26) \qquad \sum_{j=1}^{M}(\mathcal{G}u)(x_j)(\mathcal{G}v)(x_j) = \sum_{j=1}^{M}F(x_j)(\mathcal{G}v)(x_j) \qquad \forall v \in L^2(Q),$$

of Problem 3.5 with the adjoint operator $\mathcal{G}_M^*$ of $\mathcal{G}_M$ given by

$$\mathbb{R}^M \ni v = (v_j)_{j=1}^{M} \mapsto \mathcal{G}_M^* v = \frac{1}{|Q|}\sum_{j=1}^{M}\psi(x_j,\cdot)v_j \in L^2(Q).$$

Observe that $R(\mathcal{G}_M)$ is now finite-dimensional and thus closed. Hence, existence of a best-approximate solution $u = \mathcal{G}_M^\dagger F_M$ of Problem 3.5 for any point data $(x_i, F(x_i))$, $i = 1, \ldots, M$, follows from $R(\mathcal{G}_M) \oplus R(\mathcal{G}_M)^\perp = \mathbb{R}^M$ and the general Theorem 2.5 in [19]. As a consequence, the Moore-Penrose inverse $\mathcal{G}_M^\dagger$ is now bounded. However, its norm is expected to increase with increasing M.

Remark 3.6. *Assuming $N(\mathcal{G}_M^*) = \{0\}$, we obtain $J(u) = 0$ in analogy to Proposition 3.4.*

3.3. **Fredholm networks.** The asymptotic coincidence with discrete neural networks $F_N(\cdot, \eta)$ suggests to consider their continuous counterpart

$$(27) \qquad \mathcal{G}v = \frac{1}{|Q|}\int_Q \psi(\cdot, \eta)v(\eta)\,d\eta$$

with parameter function $v \in L^2(Q)$ directly for approximation. Observe that the continuous Training Problems 3.2 (function data) or 3.5 (point data) associated with such *Fredholm networks* are linear, in contrast to the original discrete Problem 2.1. However, linearity comes with the price of infinitely many unknown parameter function values. This suggests suitable discretizations of the continuous Training Problem 3.2 or 3.5 providing computable approximations of the corresponding Fredholm networks. For example, the integral operator $\mathcal{G}$ has been replaced by a suitable Riemann sum in [9, 10, 11].

Remark 3.7. *In order to relax the constraints (16) on the parameters ζ that provide an asymptotic limit of $F_N(\cdot, \zeta)$ of the form (13), one might assume that the parameters $\zeta_i = (\eta_i, u_i)$ are independent random variables that are identically distributed with respect to some unknown probability measure $\nu \in \mathbb{P}(Q \times \mathbb{R})$. Then, the law of large numbers implies*

$$\frac{1}{N}\sum_{i=1}^{N}\psi(x, \eta_i)u_i \;\rightarrow\; \int_{Q \times \mathbb{R}}\psi(x, \eta)u\,d\nu(\eta, u) \quad \text{for } N \longrightarrow \infty$$

in analogy to (13). Replacing $F_N(x, \zeta)$ in the loss functional $\mathcal{J}_N$ by the mean-field network [13]

$$\mathbb{P}(Q \times \mathbb{R}) \ni \nu \mapsto \int_{Q \times \mathbb{R}}\psi(\cdot, \eta)u\,d\nu(\eta, u) \in \mathbb{R},$$

the resulting training problem for the measure ν amounts to the least squares formulation of the problem to find $\nu \in \mathbb{P}(Q \times \mathbb{R})$ such that

$$(28) \qquad \int_{Q \times \mathbb{R}}\psi(\cdot, \eta)u\,d\nu(\eta, u) = F,$$

i.e., of a Fredholm integral equation on measures. Approximate solution of related problems by associated mean-field Langevin dynamics has been discussed in [13, 14, 15, 16, 17, 18]. *with entropy regularisers added in* [13, 14, 16, 18].

4. Ritz-Galerkin discretization and regularization

4.1. **Variational formulation.** Let $L \subset L^2(Q)$ denote a closed subspace. Then the corresponding Ritz-Galerkin discretization of the Fredholm Training Problems 3.2 and 3.5 reads as follows.

Problem 4.1. *Find $u_L \in L$ such that*

$$(29) \qquad (\mathcal{G}u_L, \mathcal{G}v)_\pi = (F, \mathcal{G}v)_\pi \qquad \forall v \in L,$$

with π still denoting the Lebesgue or point measure.

Again, existence of a best-approximate solution u_L follows from Theorem 2.5 in [19], and asymptotic ill-conditioning is expected to be inherited from the continuous case.

4.2. **Algebraic formulation.** Let the ansatz space $L = L_n$ have finite dimension n and let $\varphi_1, \ldots, \varphi_n$ be a basis of L_n. Then any $v \in L_n$ can be identified with the coefficient vector $V = (V_i)_{i=1}^n \in \mathbb{R}^n$ of its unique basis representation

$$(30) \qquad v = \sum_{i=1}^n V_i \varphi_i \in L_n.$$

Example 4.2. *Let $Q = \bigcup_{i=1}^n q_i$ be a decomposition of Q into disjoint hypercubes q_i with maximal diameter h. Then*

$$\varphi_i(\eta) = \begin{cases} |q_i|^{-1/2} & \text{if } \eta \in q_i \\ 0 & \text{otherwise} \end{cases}$$

is an orthonormal basis of the subspace $L_n \subset L^2(Q)$ of piecewise constant functions on q_i, and the coefficients V_i of $v \in L_n$ agree with its scaled values on q_i, $i = 1, \ldots, n$.

Introducing the semi-discrete, linear operators

$$(31) \qquad \begin{aligned} G &= (\mathcal{G}\varphi_1, \ldots, \mathcal{G}\varphi_n) &\in L^2(X, \pi) \times \mathbb{R}^n: \quad \mathbb{R}^n \to L^2(X, \pi), \\ G^* &= ((\mathcal{G}\varphi_i, \cdot)_\pi)_{i=1}^n &\in \mathbb{R}^n \times L^2(X, \pi): \quad L^2(X, \pi) \to \mathbb{R}^n \end{aligned}$$

we have $\mathcal{G}u_L = GU$ with $u_L = \sum_{i=1}^n U_i \varphi_i$ and $U = (U_i)_{i=1}^n \in \mathbb{R}^n$. Utilizing

$$(32) \qquad G^*G = ((\mathcal{G}\varphi_i, \mathcal{G}\varphi_j)_\pi)_{i,j=1}^n \in \mathbb{R}^{n,n}$$

the variational equality 29 can be equivalently rewritten as the linear system

$$(33) \qquad G^*GU = G^*F$$

for the unknown coefficient vector U of u_L. Observe that $G^*G = A^\top A$ holds with

$$(34) \qquad A = (a_{ij}) \in \mathbb{R}^{M \times n}, \ \ a_{ij} = (\mathcal{G}\varphi_j)(x_i), \ \ i = 1, \ldots M, \ j = 1, \ldots n,$$

in the case (9b) of given point data.

4.3. **Tikhonov regularization.** We only consider the case (9b) of given point data. In order to ensure uniqueness of approximate solutions, we select some $\varepsilon > 0$ and introduce the following Tikhonov regularization of the discretized Problem 4.1.

Problem 4.3. *Find a solution $U_\varepsilon \in \mathbb{R}^n$ of the regularized linear system*

$$(35) \qquad (\varepsilon I + A^\top A)U_\varepsilon = A^\top b$$

with identity matrix $I \in \mathbb{R}^{n \times n}$, $b = (F(x_i))_{i=1}^M \in \mathbb{R}^M$, and $\varepsilon > 0$.

Problem 4.3 is equivalent to minimizing the regularized discrete loss functional

$$(36) \qquad \mathbb{J}_\varepsilon(V) = \|AV - b\|_{2,\mathbb{R}^M}^2 + \varepsilon\|V\|_{2,\mathbb{R}^n}^2, \quad V \in \mathbb{R}^n.$$

with $\|\cdot\|_{2,\mathbb{R}^M}$ and $\|\cdot\|_{2,\mathbb{R}^n}$ denoting the Euclidean norm in $\mathbb{R}^M$ and $\mathbb{R}^n$, respectively. Existence and uniqueness of a solution is an immediate consequence of the Lax-Milgram lemma.

Moreover, Problem 4.3 is also equivalent with the algebraic formulation of the Ritz-Galerkin discretization of the regularized continuous problem to find $u_\varepsilon \in L^2(Q)$ such that

$$(37) \qquad J_\varepsilon(u_\varepsilon) \leq J_\varepsilon(v) \qquad \forall v \in L^2(Q)$$

with regularized loss functional

$$J_\varepsilon(v) = \varepsilon(v,v) + J(v) = \varepsilon(v,v) + \frac{1}{M}\sum_{j=1}^M (F(x_j) - \mathcal{G}v(x_j))^2$$

provided that $\varphi_1, \ldots, \varphi_n$ is an orthonormal basis of L_n (see, e.g., Example 4.2). Indeed, we have $\mathbb{J}_\varepsilon(V) = J_\varepsilon(v)$ for $v = \sum_{i=1}^n V_i\varphi_i$ in this case.

We conclude this section by a residual error estimate for the loss functional, as obtained for the ansatz space L_n of piecewise constant functions introduced in Example 4.2. Let $U_{\varepsilon,h} \in \mathbb{R}^n$ denote the solution of the corresponding Tikhonov regularized Problem 4.3, and set

$$u_{\varepsilon,h} = \sum_{i=1}^n (U_{\varepsilon,h})_i \, \varphi_i \in L_n \subset L^2(Q).$$

Proposition 4.4. *Assume that the kernel function ψ satisfies (25) and $N(\mathcal{G}_M^*) = \{0\}$. Assume further that the solution $u \in L^2(Q)$ of Problem 3.5 satisfies the smoothness condition $u \in R(\mathcal{G}_M^*\mathcal{G}_M)$ and that the solution u_ε of the regularized problem (37) satisfies the regularity condition $u_\varepsilon \in H^1(Q)$. Then the a priori residual estimate*

$$(38) \qquad 0 \leq J(u_{\varepsilon,h}) \leq C\left(\varepsilon^2 + \varepsilon^{-2}h^2\right)$$

for the continuous loss functional J introduced in (23) holds with a constant C depending only on u and u_ε.

Proof. The minimization problem (37) can be equivalently written in variational form

$$a(u_\varepsilon, v) = \ell(v) \quad \forall v \in L^2(Q)$$

with bilinear form $a(\cdot, \cdot)$ and linear functional ℓ given by

$$a(v,w) = \varepsilon(v,w) + \frac{1}{M}\sum_{i=1}^M \mathcal{G}_M v(x_i)\mathcal{G}_M w(x_i), \qquad \ell(v) = \frac{1}{M}\sum_{i=1}^M F(x_i)\mathcal{G}_M v(x_i).$$

We have $\ell \in L^2(Q)'$ and boundedness

$$|a(v,w)| \leq c\|v\|\|w\| \quad \forall v, w \in L^2(Q)$$

with a constant $c \geq \varepsilon + \|\mathcal{G}_M\|^2$ independent of M follows from uniform boundedness of $\mathcal{G}_M$ (see Subsection 3.2.2). Hence, Céa's lemma and well-known results on Poincaré inequalities on convex domains, cf., e.g. [37][Theorem 3.2] imply the discretization error bound

$$(39) \qquad \|u_\varepsilon - u_{\varepsilon,h}\| \leq c\, \varepsilon^{-1} \inf_{v \in L_n} \|u_\varepsilon - v\| \leq c\pi^{-1}\, \varepsilon^{-1} h\|u_\varepsilon\|_{H^1(Q)}$$

The regularization error bound

$$(40) \qquad\qquad\qquad\qquad\qquad \|u - u_\varepsilon\| \leq c\, \varepsilon$$

with another constant c depending only on u is a consequence of [19][Theorem 4.3]. The desired residual error estimate

$$J(u_{\varepsilon,h}) \leq (\|\mathcal{G}_M\|\|u - u_\varepsilon\| + \|\mathcal{G}_M\|\|u_\varepsilon - u_{\varepsilon,h}\|)^2 \leq c\left(\varepsilon + \varepsilon^{-1}h\right)^2 \leq 2c\left(\varepsilon^2 + \varepsilon^{-2}h^2\right)$$

with another constant c depending only on u and u_ε now follows from $J(u) = 0$, the uniform boundedness of $\mathcal{G}_M$, and the error estimates (39) and (40). □

5. Algebraic solution by tensor trains

In this section we consider the algebraic solution of the regularized linear system (35) in Problem 4.3 for high dimension d and resulting large number of unknowns n by tensor train methods.

5.1. **Tensor formats.**

Tensors are multilinear mappings that can be represented as functions $T \colon \{1, \ldots, D_1\} \times \cdots \times \{1, \ldots, D_p\} \to \mathbb{C}$ or, equivalently, as multidimensional arrays $\mathbf{T} \in \mathbb{C}^{D_1 \times \cdots \times D_p}$ with dimensions or modes $D_k \in \mathbb{N}$, $k = 1, \ldots, p$, and order p of $\mathbf{T}$. We will write $\mathbf{T} \in \mathbb{R}^{D_1 \times \cdots \times D_p}$, if all components $\mathbf{T}_{i_1,\ldots,i_p}$, $i_k = 1, \ldots, D_k$, $k = 1, \ldots, p$, are real numbers. Throughout the following, tensors are denoted by bold letters. When certain indices of a tensor are fixed, we use colon notation (cf. Python or Matlab) to indicate the remaining modes, e.g., $\mathbf{T}_{:,i_2,\ldots,i_{p-1},:} \in \mathbb{C}^{D_1 \times D_p}$.

The sum of tensors and the product with scalars is defined componentwise. Extending usual matrix multiplication, the product or index contraction $\mathbf{T} \cdot \mathbf{U} \in \mathbb{C}^{D_1 \times \cdots \times D_p \times D_1'' \times \cdots \times D_{p''}''}$ of two tensors $\mathbf{T} \in \mathbb{C}^{D_1 \times \cdots \times D_p \times D_1' \times \cdots \times D_{p'}'}$ and $\mathbf{U} \in \mathbb{C}^{D_1' \times \cdots \times D_{p'}' \times D_1'' \times \cdots \times D_{p''}''}$ is defined by

$$(\mathbf{T} \cdot \mathbf{U})_{i_1,\ldots,i_p,j_1,\ldots,j_{p''}} = \sum_{k_1=1}^{D_1'} \cdots \sum_{k_{p'}=1}^{D_{p'}'} \mathbf{T}_{i_1,\ldots,i_p,k_1,\ldots,k_{p'}} \mathbf{U}_{k_1,\ldots,k_{p'},j_1,\ldots,j_{p''}}.$$

Note that the sum of the tensors $\mathbf{T}_\ell \in \mathbb{C}^{D_1 \times \cdots \times D_p}$, $\ell = 1, \ldots, r$, can be regarded as an index contraction of $\mathbf{T} \in \mathbb{C}^{r \times D_1 \times \cdots \times D_p}$ with $\mathbf{T}_{\ell,i_1,\ldots,i_p} = (\mathbf{T}_\ell)_{i_1,\ldots,i_p}$ and $\mathbf{1} \in \mathbb{R}^r$ with $\mathbf{1}_\ell = 1$, $\ell = 1, \ldots, r$. We will also make use of the tensor product $\mathbf{T} \otimes \mathbf{U} \in \mathbb{C}^{D_1 \times \cdots \times D_p \times D_1' \times \cdots \times D_{p'}'}$ of $\mathbf{T} \in \mathbb{C}^{D_1 \times \cdots \times D_p}$ and $\mathbf{U} \in \mathbb{C}^{D_1' \times \cdots \times D_{p'}'}$ which is defined by

$$(\mathbf{T} \otimes \mathbf{U})_{i_1,\ldots,i_p,j_1,\ldots,j_{p'}} = \mathbf{T}_{i_1,\ldots,i_p} \mathbf{U}_{j_1,\ldots,j_{p'}}$$

and coincides with the dyadic product $\mathbf{T}\mathbf{U}^\top \in \mathbb{C}^{D_1 \times D_1'}$ of two vectors $\mathbf{T} \in \mathbb{C}^{D_1}$ and $\mathbf{U} \in \mathbb{C}^{D_1'}$ for $p = p' = 1$.

As the number of elements of a tensor grows exponentially with order p, the storage of high-dimensional tensors may become infeasible, if p becomes large. This motivates suitable

formats for tensor representation and approximation. Obviously, the storage of a *rank-one tensor* $\mathbf{T} \in \mathbb{C}^{D_1 \times \cdots \times D_p}$ that can be written as the tensor product

$$(41) \qquad \mathbf{T} = \mathbf{T}^{(1)} \otimes \cdots \otimes \mathbf{T}^{(p)}$$

of p vectors $\mathbf{T}^{(k)} \in \mathbb{C}^{D_k}$, $k = 1, \ldots, p$, only grows linearly with $pD_{\max}$ and $D_{\max} = \max\{D_1, \ldots, D_p\}$. As a generalization of (41), a tensor $\mathbf{T} \in \mathbb{C}^{D_1 \times \cdots \times D_p}$ is said to be in the *canonical format* [38], if it can be expressed as a finite sum of rank-one tensors according to

$$(42) \qquad \mathbf{T} = \sum_{\ell=1}^{r} \mathbf{T}^{(1)}_{\ell,:} \otimes \cdots \otimes \mathbf{T}^{(p)}_{\ell,:}.$$

Here, the tensors $\mathbf{T}^{(k)} \in \mathbb{C}^{r \times D_k}$, $k = 1, \ldots, p$, are called cores and the number r is called the canonical rank of $\mathbf{T}$. Note that the required storage grows with order $\mathcal{O}(rpD_{\max})$.

The *tensor train* (TT) format [39, 40] is a further generalization of (42). Here, the basic idea is to decompose a tensor into a chain-like network of low-dimensional tensors coupled by so-called TT ranks $r_0, \ldots r_p$ according to

$$(43) \qquad \mathbf{T} = \sum_{\ell_0=1}^{r_0} \cdots \sum_{\ell_p=1}^{r_p} \mathbf{T}^{(1)}_{\ell_0,:,\ell_1} \otimes \cdots \otimes \mathbf{T}^{(p)}_{\ell_{p-1},:,\ell_p},$$

with TT cores $\mathbf{T}^{(k)} \in \mathbb{C}^{r_{k-1} \times D_k \times r_k}$, $k = 1, \ldots, p$, and $r_0 = r_p = 1$. With maximal TT rank $r_{\max} = \max\{r_0, \ldots, r_p\}$, the storage consumption for tensors in the TT format grows like $\mathcal{O}(r_{\max}^2 p D_{\max})$ and thus is still linearly bounded in terms of the maximal dimension $D_{\max}$ and the order p of $\mathbf{T}$. The ranks $r_0, \ldots, r_p$ not only determine the required storage, but also have a strong influence on the quality of best approximation of a given tensor in full format by a tensor train. Note that such a best approximation in TT format with bounded ranks always exists [41] and is even exact for suitably chosen ranks [42, Theorem 1]. For further information about tensor formats, we refer, e.g., to [43, 44].

In order to provide a factorization of the integral operator $\mathcal{G}$ by factorization of the kernel ψ, we will make use of functional tensor networks with discrete as well as continuous modes. More precisely, a functional tensor with one discrete and one continuous mode is a function $T \colon \{1, \ldots, D\} \times Q \to \mathbb{C}$, where $Q \subseteq \mathbb{R}$ denotes the domain of the continuous mode. We write $\mathbf{T}_{i,\eta} = T(i, \eta)$, $i = 1, \ldots, D$, $\eta \in Q$, for the components of the corresponding generalized array $\mathbf{T} \in \mathbb{C}^{D \times Q}$. Replacing summation by integration, index contraction $\mathbf{T} \cdot \mathbf{U} \in \mathbb{C}^{D \times D'}$ of $\mathbf{T}$ with the functional tensor $\mathbf{U} \in \mathbb{C}^{Q \times D'}$ is defined by

$$(\mathbf{T} \cdot \mathbf{U})_{i,j} = \int_Q \mathbf{T}_{i,\eta} \mathbf{U}_{\eta,j} \, d\eta.$$

Functional tensor trains (FTTs) [45] provide a decomposition of $\mathbf{T} \in \mathbb{C}^{Q_1 \times \cdots \times Q_p}$ into cores $\mathbf{T}^{(k)} \in \mathbb{C}^{r_{k-1} \times Q_k \times r_k}$ in analogy to (43).

Figure 1 shows the diagrammatic notation of tensors in rank-one, canonical, TT, and FTT format. We depict vectors, matrices, and tensors as circles with different arms indicating the set of modes. Index contraction of two or more tensors is represented by connecting corresponding arms.
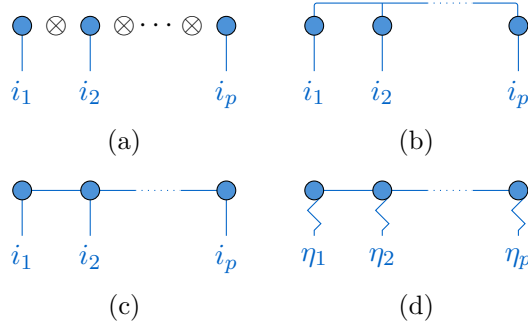
FIGURE 1. Graphical notation of tensor formats: (a) A tensor in rank-one format, given by the tensor product of p vectors. (b) A tensor in the canonical format, given by the contraction of p matrices on the common rank index. (c) A tensor in TT format, given by a network of pairwise coupled tensors. Here, the first and the last TT core are regarded as matrices, because $r_0 = r_p = 1$. (d) A tensor in FTT format, given by a TT-like network of tensors with one continuous mode. Discrete modes are represented by straight lines and continuous modes by zigzag lines.

5.2. **Tensor decomposition of basis and stiffness matrix.** As a first step towards a tensor train formulation of the regularized linear system (35), we assume that the basis functions $\varphi_1, \ldots, \varphi_n$ of the ansatz space L_n utilized in (30) are separable in the sense that

$$(44) \qquad \varphi_j(\eta) = \varphi_{j_1}^{(1)}(\eta_1) \cdot \ldots \cdot \varphi_{j_s}^{(s)}(\eta_s), \quad \eta = (\eta_1, \ldots, \eta_s) \in Q = Q_1 \times \cdots \times Q_s,$$

holds with $1 \leq j_k \leq n_k$ and some enumeration

$$(45) \qquad j = j(j_1, \ldots, j_s) = 1, \ldots, n = n_1 \cdot \ldots \cdot n_s, \qquad j_k = 1, \ldots, n_k, \quad k = 1, \ldots, s.$$

Example 5.1. *Let $Q_k = \bigcup_{j=1}^{n_k} q_j^{(k)}$ be a decomposition of $Q_k \subset \mathbb{R}^{n_k}$ into disjoint intervals $q_j^{(k)} \subset \mathbb{R}$, $j = 1, \ldots, n_k$. Then the piecewise constant basis functions $\varphi_1, \ldots, \varphi_n$ introduced in Example 4.2 can be written as a product of the form (44) with*

$$\varphi_j^{(k)}(\eta_k) = \begin{cases} |q_j^{(k)}|^{-1/2} & \text{if } \eta_k \in q_j^{(k)} \\ 0 & \text{otherwise} \end{cases}.$$

By definition of the tensor product, we have the following proposition.

Proposition 5.2. *Assume that the basis functions $\varphi_1, \ldots, \varphi_n$ of the ansatz space L_n are separable in the sense of (44). Then the functional tensor*

$$\mathbf{\Phi} \in \mathbb{R}^{n_1 \times \cdots \times n_s \times Q_1 \times \cdots \times Q_s}, \qquad \mathbf{\Phi}_{j_1, \ldots, j_s, \eta_1, \ldots, \eta_s} = \varphi_{j_1}^{(1)}(\eta_1) \cdot \ldots \cdot \varphi_{j_s}^{(s)}(\eta_s),$$

is a rank-one tensor with the decomposition

$$(46) \qquad \mathbf{\Phi} = \mathbf{\Phi}^{(1)} \otimes \cdots \otimes \mathbf{\Phi}^{(s)}$$

into $\mathbf{\Phi}^{(k)} \in \mathbb{R}^{n_k \times Q_k}$ defined by $\mathbf{\Phi}_{j_k, \eta_k}^{(k)} = \varphi_{j_k}^{(k)}(\eta_k)$, $j_k = 1, \ldots, n_k$, for $k = 1, \ldots, s$.

In the next step, we derive a decomposition of the functional tensor $\mathbf{\Psi} \in \mathbb{R}^{M \times Q_1 \times \cdots \times Q_s}$ with components

$$(47) \qquad \mathbf{\Psi}_{i, \eta_1, \ldots, \eta_s} = \psi(x_i, \eta), \qquad i = 1, \ldots, M, \quad \eta = (\eta_1, \ldots, \eta_s)^\top \in Q = Q_1 \times \cdots \times Q_s,$$

associated with the actual kernel function ψ and the given data points $x_i \in X$, $i = 1, \ldots, M$. For ease of presentation, we focus on kernels $\psi \colon X \times Q \to \mathbb{R}$ with X, $Q \subset \mathbb{R}^s$. Obviously, radial kernel functions (5) have this property with $s = d$, and in the case of ridge kernel functions (2) we can identify X with its embedding $\{(x, 1)^\top \mid x \in X\} \subset \mathbb{R}^s$, $s = d + 1$.

We further assume that the actual kernel function ψ can be expressed (or approximated) as a sum of products of univariate functions according to

$$(48) \qquad \psi(x, \eta) = \sum_{\ell=1}^{r} \psi_\ell^{(1)}(x_1, \eta_1) \cdot \ldots \cdot \psi_\ell^{(s)}(x_s, \eta_s).$$

Example 5.3. *The ridge kernel function $\psi(x, \eta) = \sigma(x \cdot \eta)$ with trivial activation function $\sigma = Id$ can be written in the form* (48) *with $r = s$ and $\psi_\ell^{(k)}(x_k, \eta_k) = x_k \eta_k$ for $k = \ell$ and $\psi_\ell^{(k)}(x_k, \eta_k) = 1$ otherwise.*

Explicit constructions of approximate factorizations of ridge kernels with more relevant activation functions can be found in Appendix A. Note that even in the case of real-valued kernel functions ψ, the factors $\psi_\ell^{(k)}$ might have complex values (see, e.g. A.3).

Example 5.4. *The radial kernel function with Gaussian activation function given in* (7) *takes the form* (48) *with $r = 1$ and*

$$\psi_1^{(k)}(x_k, \eta_k) = \exp\left(-\frac{(x_k - \eta_k)^2}{2\kappa^2}\right), \quad k = 1, \ldots, d.$$

Utilizing the given data points $x_i = (x_{i,k})_{k=1}^s \in X$, $i = 1, \ldots, M$, we now introduce the auxiliary functional tensor $\hat{\boldsymbol{\Psi}} \in \mathbb{C}^{M^s \times Q_1 \times \cdots \times Q_s}$ with canonical decomposition

$$(49) \qquad \hat{\boldsymbol{\Psi}} = \sum_{\ell=1}^{r} \hat{\boldsymbol{\Psi}}_{\ell,:,:}^{(1)} \otimes \cdots \otimes \hat{\boldsymbol{\Psi}}_{\ell,:,:}^{(s)}$$

into cores $\hat{\boldsymbol{\Psi}}^{(k)} \in \mathbb{C}^{r \times M \times Q_k}$ with components $\hat{\boldsymbol{\Psi}}_{\ell,i,\eta_k}^{(k)} = \psi_\ell^{(k)}(x_{i,k}, \eta_k)$.

Proposition 5.5. *Assume that the kernel function ψ allows for the representation* (48)*. Then the functional tensor $\boldsymbol{\Psi} \in \mathbb{C}^{M \times Q_1 \times \cdots \times Q_s}$ can be decomposed according to*

$$(50) \qquad \boldsymbol{\Psi} = \Delta \cdot \hat{\boldsymbol{\Psi}}$$

with the delta tensor $\Delta \in \mathbb{R}^{M \times M^s}$ defined by

$$(51) \qquad \Delta_{i,j_1,\ldots,j_s} = \delta_{i,j_1} \cdot \ldots \cdot \delta_{i,j_s}$$

and δ_{i,j_k} denoting the Kronecker delta.

Proof. The desired identity (50) follows from

$$
\begin{aligned}
(\Delta \cdot \hat{\boldsymbol{\Psi}})_{i,\eta_1,\ldots,\eta_s} &= \sum_{j_1=1}^{M} \cdots \sum_{j_s=1}^{M} \Delta_{i,j_1,\ldots,j_s} \hat{\boldsymbol{\Psi}}_{j_1,\ldots,j_s,\eta_1,\ldots,\eta_s} \\
&= \sum_{\ell=1}^{r} \left(\sum_{j_1=1}^{M} \delta_{i,j_1} \hat{\boldsymbol{\Psi}}_{\ell,j_1,\eta_1}^{(1)}\right) \cdots \left(\sum_{j_s=1}^{M} \delta_{i,j_s} \hat{\boldsymbol{\Psi}}_{\ell,j_s,\eta_s}^{(s)}\right) \\
&= \sum_{\ell=1}^{r} \hat{\boldsymbol{\Psi}}_{\ell,i,\eta_1}^{(1)} \cdots \hat{\boldsymbol{\Psi}}_{\ell,i,\eta_s}^{(s)} = \sum_{\ell=1}^{r} \psi_\ell^{(1)}(x_{i,1}, \eta_1) \cdots \psi_\ell^{(s)}(x_{i,s}, \eta_s) = \boldsymbol{\Psi}_{i,\eta_1,\ldots,\eta_s}. \quad \square
\end{aligned}
$$

The construction of the functional tensor $\boldsymbol{\Psi}$ is illustrated in Figure 2 (a). Note that the auxiliary tensor Δ does not explicitly appear in our numerical computations. We will also exploit the FTT format to represent the auxiliary tensor $\hat{\boldsymbol{\Psi}}$, if it allows for a low-rank decomposition. In particular, monomials of the form $(x \cdot \eta)^p$ can be expressed in a compact way using a tensor train-like coupling as we show in Appendix A. Figure 2 (b) shows the corresponding FTT network for $\boldsymbol{\Psi}$.

Now we are in the position to provide a representation of the stiffness matrix $A = (a_{ij}) \in \mathbb{R}^{M \times n}$ defined in (34) as a tensor network.
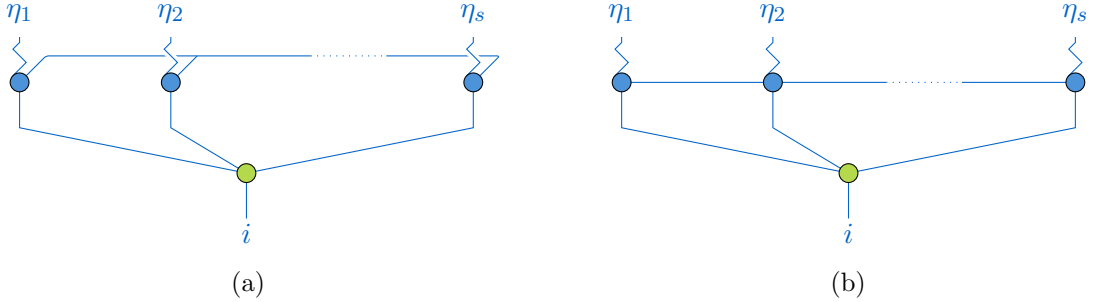

FIGURE 2. Graphical notation of the functional decomposition of $\boldsymbol{\Psi}$ corresponding to the kernel $\psi(x, \eta)$ and data points $(x_j)_{j=1}^M$: (a) Using the canonical format for $\hat{\boldsymbol{\Psi}}$ with cores coupled by a common rank. (b) Using the FTT format with a chain-like coupling. black circles represent the cores of $\hat{\boldsymbol{\Psi}}$, each having one continuous mode represented by a zigzag line. Contraction with Δ (green circle) merges all free modes of $\hat{\boldsymbol{\Psi}}$ into one.

Proposition 5.6. *Assume that the basis functions $\varphi_1, \ldots, \varphi_n$ of the ansatz space L_n are separable in the sense of (44) and that the kernel function ψ allows for the representation (48). Then the tensor*

$$\mathbf{A} \in \mathbb{R}^{M \times n_1 \times \cdots \times n_s}, \quad \mathbf{A}_{i, j_1, \ldots, j_s} = a_{i, j(j_1, \ldots, j_s)} \tag{52}$$

with enumeration $j(j_1, \ldots, j_s)$ introduced in (45) allows for the factorization

$$\mathbf{A} = \Delta \cdot \sum_{\ell=1}^r \mathbf{A}_\ell^{(1)} \otimes \cdots \otimes \mathbf{A}_\ell^{(s)}, \quad \mathbf{A}_\ell^{(k)} = (\hat{\boldsymbol{\Psi}}_{\ell,:,:}^{(k)} \cdot \boldsymbol{\Phi}^{(k)}) \in \mathbb{C}^{M \times n_k}, \quad k = 1, \ldots, s, \tag{53}$$

with $\boldsymbol{\Phi}^{(k)}$, $\hat{\boldsymbol{\Psi}}^{(k)}$, and Δ taken from (46), (49), and (51), respectively.

Proof. The desired identity follows from

$$
\begin{aligned}
\mathbf{A}_{i,j_1,\ldots,j_s} &= \sum_{i_1}^{M} \cdots \sum_{i_s}^{M} \Delta_{i,i_1,\ldots,i_s} \sum_{\ell=1}^{r} \mathbf{A}_{\ell,i_1,j_1}^{(1)} \cdots \mathbf{A}_{\ell,i_s,j_s}^{(s)} \\
&= \sum_{\ell=1}^{r} \left(\sum_{i_1}^{M} \delta_{i,i_1} \mathbf{A}_{\ell,i_1,j_1}^{(1)} \right) \cdots \left(\sum_{i_s}^{M} \delta_{i,i_s} \mathbf{A}_{\ell,i_s,j_s}^{(s)} \right) = \sum_{\ell=1}^{r} \mathbf{A}_{\ell,i,j_1}^{(1)} \cdots \mathbf{A}_{\ell,i,j_s}^{(s)} \\
&= \sum_{\ell=1}^{r} \int_{Q_1} \psi_\ell(x_{i_1},\eta_1)\varphi_{j_1}^{(1)}(\eta_1)\, d\eta_1 \cdots \int_{Q_s} \psi_\ell(x_{i_s},\eta_s)\varphi_{j_s}^{(s)}(\eta_s)\, d\eta_s \\
&= \int_Q \left(\sum_{\ell=1}^{r} \psi_\ell(x_{i_1},\eta_1)\cdots\psi_\ell(x_{i_s},\eta_s) \right) \varphi_{j_1}^{(1)}(\eta_1)\cdots\varphi_{j_s}^{(s)}(\eta_s)\, d\eta \\
&= \int_Q \psi(x_i,\eta)\varphi_{j(j_1,\ldots,j_s)}(\eta)\, d\eta = a_{i,j(j_1,\ldots,j_s)}. \qquad \square
\end{aligned}
$$

Utilizing the enumeration (45), we introduce the tensor representation

$$
\mathbf{U} \in \mathbb{R}^{n_1 \times \cdots \times n_s}, \quad \mathbf{U}_{j_1,\ldots,j_s} = U_{j(j_1,\ldots,j_s)},
$$

of the unknown solution vector $U \in \mathbb{R}^n$ in (35). Then, the regularized linear system (35) in Problem 4.3 takes the form

$$
(54) \qquad\qquad (\varepsilon\mathbf{I} + \mathbf{A}^\top\mathbf{A})\mathbf{U} = \mathbf{A}^\top\mathbf{b}
$$

with unit tensor $\mathbf{I} \in \mathbb{R}^{n_1 \times \cdots \times n_s}$ and $\mathbf{b} = b \in \mathbb{R}^M$. Solving (54) is equivalent to minimizing the tensor analogue

$$
\mathbf{J}_\varepsilon(\mathbf{V}) = \|\mathbf{A}\mathbf{V} - \mathbf{b}\|_{2,\mathbb{R}^M}^2 + \varepsilon\|\mathbf{V}\|_{2,\mathbb{R}^{n_1 \times \cdots \times n_s}}^2, \quad \mathbf{V} \in \mathbb{R}^{n_1 \times \cdots \times n_s},
$$

of the regularized discrete loss functional $\mathbb{J}_\varepsilon$ defined in (36). We now introduce the corresponding tensor-train approximation of the linear system (54).

Problem 5.7. *Find cores $\mathbf{U}^{(k)} \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$, $k = 1,\ldots,s$, with $r_0 = r_s = 1$ and prescribed ranks r_k, $k = 1,\ldots,s-1$ such that the resulting tensor train*

$$
(55) \qquad\qquad \mathbf{U}_r = \sum_{\ell_0=1}^{r_0} \cdots \sum_{\ell_s=1}^{r_s} \mathbf{U}_{\ell_0,:,\ell_1}^{(1)} \otimes \cdots \otimes \mathbf{U}_{\ell_{s-1},:,\ell_s}^{(s)} \in \mathbb{R}^{n_1 \times \cdots \times n_s},
$$

is minimizing the regularized discrete loss functional $\mathbf{J}_\varepsilon$ over the submanifold of all tensor trains $\mathbf{V}$ of the form (55).

Existence of a solution $\mathbf{U}_r$ of Problem 5.7 follows by compactness arguments. Note that we have $\mathbf{U}_r = \mathbf{U}$, if the selection $r_1,\ldots,r_{s-1}$ agrees with the minimal rank of the unique solution $\mathbf{U}$ of (54), because there exists an exact tensor train representation of $\mathbf{U}$ in this case (see [42, Theorem 1]). As a consequence of non-uniqueness of tensor train representations, uniqueness of the cores $\mathbf{U}^{(k)} \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$ can not be expected.

5.3. **Algebraic solution by alternating linear systems.** The Tensor-Train Approximation Problem 5.7 allows for a variety of different methods for direct or iterative solution, such as pseudoinversion [21, 46, 47, 48] or core-wise linear updates [20, 21, 49, 50], respectively. Here, we concentrate on iterative solution by *alternating ridge regression* (ARR), cf. [21], because this approach significantly reduces computational cost and memory consumption in comparison with explicit computation of pseudoinverses (see, e.g., [21] for details). These methods provide successive updates of the cores $\tilde{\mathbf{U}}^{(k)} \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$, $k = 1, \ldots, s$, by solving a sequence of reduced linear problems.

More precisely, for given iterate $\tilde{\mathbf{U}}$ of the form (55) and each fixed k, the cores $\tilde{\mathbf{U}}^{(1)}, \ldots, \tilde{\mathbf{U}}^{(k-1)}$ are *left-orthonormalized* and $\tilde{\mathbf{U}}^{(k+1)}, \ldots, \tilde{\mathbf{U}}^{(s)}$ are *right-orthonormalized* by successive singular value decompositions [39] such that the rows and columns, respectively, of certain reshapes of these cores form orthonormal sets. Note that, without truncation, this procedure produces a different but equivalent representation of $\tilde{\mathbf{U}}$, say $\hat{\mathbf{U}}$. The orthonormalized cores then define a retraction operator

$$\mathbf{Q}_{\hat{\mathbf{U}},k} \in \mathbb{R}^{(n_1 \times \cdots \times n_s) \times (r_{k-1} \times n_k \times r_k)},$$

which satisfies

$$\mathbf{Q}_{\hat{\mathbf{U}},k}\mathbf{V}^{(k)} = \sum_{\ell_0=1}^{r_0} \cdots \sum_{\ell_s=1}^{r_s} \bigotimes_{i=1}^{k-1} \hat{\mathbf{U}}^{(i)}_{\ell_{i-1},:,\ell_i} \otimes \mathbf{V}^{(k)}_{\ell_{k-1},:,\ell_k} \otimes \bigotimes_{i=k+1}^{s} \hat{\mathbf{U}}^{(i)}_{\ell_{i-1},:,\ell_i}$$

for all $\mathbf{V}^{(k)} \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$ and is orthonormal in the sense that

$$(56) \qquad \mathbf{Q}_{\hat{\mathbf{U}},k}^{\top} \cdot \mathbf{Q}_{\hat{\mathbf{U}},k} = \mathbf{I} \in \mathbb{R}^{(r_{k-1} \times n_k \times r_k) \times (r_{k-1} \times n_k \times r_k)}.$$

Minimization with respect to the k-th core then gives rise to the reduced problem to minimize

$$\|\mathbf{A}\mathbf{Q}_{\tilde{\mathbf{U}},k}\mathbf{V}^{(k)} - \mathbf{b}\|_{2,\mathbb{R}^M}^2 + \varepsilon\|\mathbf{Q}_{\tilde{\mathbf{U}},k}\mathbf{V}^{(k)}\|_{2,\mathbb{R}^{n_1 \times \cdots \times n_s}}^2$$

over all $\mathbf{V}^{(k)} \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$ or, equivalently, to solve the reduced linear system

$$(57) \qquad \mathbf{Q}_{\tilde{\mathbf{U}},k}^{\top}(\varepsilon\mathbf{I} + \mathbf{A}^{\top}\mathbf{A})\mathbf{Q}_{\tilde{\mathbf{U}},k}\tilde{\mathbf{U}}_{\text{new}}^{(k)} = \mathbf{Q}_{\tilde{\mathbf{U}},k}^{\top}\mathbf{A}^{\top}\mathbf{b}$$

for the new core $\tilde{\mathbf{U}}_{\text{new}}^{(k)} \in \mathbb{R}^{r_{k-1} \times n_k \times r_k}$ (see, e.g. [20, 21, 49, 50] for details). The unique solution $\tilde{\mathbf{U}}_{\text{new}}^{(k)}$ of (57) can be computed, e.g., by Cholesky decomposition of the symmetric and positive definite coefficient matrix.

The corresponding update of cores is typically applied in alternating order, i.e., from $k = 1$ to $k = s$ and then back to $k = 1$. As the same regularization parameter $\varepsilon > 0$ is used for each core k, the alternating linear systems approach [20] applied to (54) and alternating ridge regression [21] are equivalent in this case. As a consequence of the intrinsic nonlinearity of the overall iterative approach, the convergence analysis is complicated. We refer to [20, 51, 52] for further information.

6. Sampling of discrete neural network parameters

By construction, an approximate solution $\tilde{\mathbf{U}} \in \mathbb{R}^{n_1 \times \cdots \times n_s}$ of Problem 5.7 provides an approximate solution

$$\tilde{U} = (\tilde{U}_j) \in \mathbb{R}^n, \qquad \tilde{U}_{j(j_1,\ldots,j_s)} = \tilde{\mathbf{U}}_{j_1,\ldots,j_s},$$
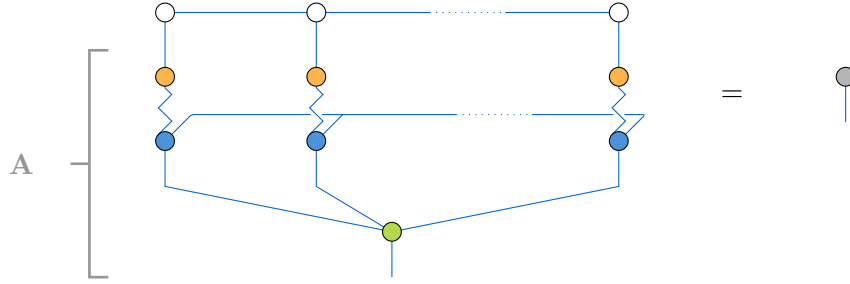
FIGURE 3. Graphical notation of tensor-based counterpart of the (underdetermined) system $AU = b$: The cores of $\boldsymbol{\Psi}$ (black circles) are contracted with the cores of $\boldsymbol{\Phi}$ (orange circles) by integrating over common modes. Together with the delta tensor (green circle), this system builds the tensor operator $\mathbf{A}$. The coefficient tensor $\mathbf{U}$ (white circles) is approximated in the TT format. Here, $\hat{\boldsymbol{\Psi}}$ is given in canonical format but could also be defined as functional tensor train, see Section 5.2.

of Problem 4.3 by utilizing the enumeration (45). In turn, $\tilde{U}$ gives rise to the approximate solution

$$\tilde{u} = \sum_{j=1}^{n} \tilde{U}_j \varphi_j \in L_n \subset L^2(Q)$$

of Problem 2.1 (function data) or Problem 3.5 (point data) and the resulting approximation

$$(58) \qquad \mathcal{G}\tilde{u} = \sum_{j=1}^{n} \tilde{U}_j \int_Q \psi(\cdot, \eta) \varphi_j(\eta) \, d\eta$$

of the trained Fredholm network $\mathcal{G}u$ introduced in Section 3.3. Note that $\mathcal{G}\tilde{u}(x)$ can be evaluated in given test or training points by contracting $\tilde{\mathbf{U}}$ with the corresponding stiffness tensor $\mathbf{A}$.

Further approximation of $\mathcal{G}\tilde{u}$ by classical Monte Carlo sampling of the integral

$$(59) \qquad \mathcal{G}\tilde{u} = \int_Q \psi(\cdot, \eta) \tilde{u}(\eta) \, d\eta$$

provides parameters $\tilde{\zeta}_i = (\tilde{\eta}_i, \tilde{u}_i)$ with i.i.d. $\tilde{\eta}_i$ and $\tilde{u}_i = \tilde{u}(\tilde{\eta}_i)$, $i = 1, \ldots, N$, and thus the trained neural network $F_N(\cdot, \tilde{\zeta})$.

Remark 6.1. *Let us consider the situation addressed in Subsection 4.3. Under the (unrealistic) assumption that the linear Problem 4.3 is solved exactly, classical Monte Carlo sampling of $\mathcal{G}u_{\varepsilon,h}$ provides the trained neural network $F_N(\cdot, \zeta_{\varepsilon,h})$. Then, in the light of Proposition 4.4, the residual error estimate*

$$\mathcal{J}_N(\zeta_{\varepsilon,h}) = \mathcal{O}(\varepsilon^2 + \varepsilon^{-2}h^2 + N^{-1/2})$$

for the discrete loss functional $\mathcal{J}_N$ introduced in (8) holds by the central limit theorem under suitable regularity assumptions.

Quasi-Monte Carlo sampling is based on quasi-random sequences in $\mathbb{R}^s$, such as, e.g., Sobol sequences, that fill the space more uniformly than pseudo-random sequences as used for classical Monte Carlo integration. In order to further enhance sampling quality, we now introduce

inductive importance sampling that is directly targeting the underlying loss functional $\mathcal{J}_N$ defined in (8). Given a set of quasi-Monte Carlo sample points $(\tilde{\eta}_i)_{i=1}^{N'}$, corresponding parameter function evaluations $(\tilde{u}(\tilde{\eta}_i))_{i=1}^{N'}$, and some $N \leq N'$, we want to inductively find parameters $(\tilde{\eta}_{i_j})_{j=1}^N$ such that $F_N(\cdot, \zeta)$ with $\zeta = (\tilde{\eta}_{i_j}, \tilde{u}(\tilde{\eta}_{i_j}))_{j=1}^N$ is reducing the loss functional $\mathcal{J}_N$ sufficiently well. For this purpose, the first sample η_{i_1} is chosen according to

$$\eta_{i_1} = \underset{\eta \in \{\eta_1, \dots, \eta_{N'}\}}{\operatorname{argmin}} \mathcal{J}_1((\eta, \tilde{u}(\eta)))$$

and for given samples $\eta_{i_1}, \dots, \eta_{i_K}$, $K < N$, we select

$$\eta_{i_{K+1}} = \underset{\eta \in \{\eta_1, \dots, \eta_{N'}\} \backslash \{\eta_{i_1}, \dots, \eta_{i_K}\}}{\operatorname{argmin}} \mathcal{J}_{K+1}\left((\eta_{i_1}, \tilde{u}(\eta_{i_1})), \dots, (\eta_{i_K}, \tilde{u}(\eta_{i_K})), (\eta, \tilde{u}(\eta))\right).$$

Note that we use exact kernel and activation functions here and not their approximate factorizations which the Fredholm network is trained on. While quasi-Monte Carlo integration performed reasonably well in simple cases, we found a considerable improvement of efficiency by inductive importance sampling in our numerical experiments to be reported below.

7. NUMERICAL EXPERIMENTS

In order to demonstrate the practical relevance of our approach, we now consider three well-established test cases of regression and classification type. In particular, we will numerically investigate different aspects of the tensor-based training of neural networks, such as adaptability, accuracy, and robustness. We emphasize that kernel, activation and even basis functions for discretization have been chosen ad hoc, and no profound parameter optimization is performed. Our implementation of ARR and other relevant tensor algorithms are available in Scikit-TT[1].

7.1. **Bank note authentication.** As a first example, we consider a classification problem to distinguish forged from genuine bank notes, utilizing the UCI banknote authentication data set[2] from [22]. It contains 1372 samples of different features that were extracted from images of genuine and forged banknote-like specimens by applying wavelet transforms. Each sample has $d = 4$ different attributes: variance, skewness, curtosis, and entropy of the wavelet transformed images. The data set is randomly divided into $M = 1098$ ($\approx 80\%$) training and $m = 274$ test samples. We first apply min-max normalization to the given training data points $\hat{x}_i = (x_{i,k})_{k=1}^d$, $i = 1, \dots, M$, to obtain

$$(60) \qquad x_{i,k} = \frac{\hat{x}_{i,k} - \min(k)}{\max(k) - \min(k)} \in [0, 1], \qquad i = 1, \dots, M, \quad k = 1, \dots, d,$$

denoting $\min(k) = \min_{i \in \{1, \dots, M\}} \hat{x}_{i,k}$ and $\max(k) = \max_{i \in \{1, \dots, M\}} \hat{x}_{i,k}$. Hence, $X = [0, 1]^d$. The test data points are normalized in the same way using the same constants.

For the neural network F_N of the form (1), we then choose a Gaussian activation function (7) with $\kappa = 1$ and the same parameter set for each dimension, i.e., $Q = [-1, 1]^s$ with $s = d = 4$.

Discretization of the continuous training Problem 3.2 is performed by piecewise constant finite elements with respect to an equidistant grid with mesh size $1/4$, see Example 4.2. The resulting ansatz space has the dimension $n = 8^d = 4096$. We select the Tikhonov regularization parameter $\varepsilon = 10^{-3}$. For the iterative solution of the resulting Problem 5.7

---

[1] https://github.com/PGelss/scikit_tt
[2] https://archive.ics.uci.edu/ml/datasets/banknote+authentication

TABLE 1. Results for the bank note authentication data set: Predictability of trained neural network $F_N(\cdot, \tilde{\zeta})$ for increasing N in terms of expectation and standard deviation of the classification rates.

N	classification rate	perfect classification
$2^6 = 64$	0.9353 ± 0.0147	0
$2^7 = 128$	0.8653 ± 0.0279	0
$2^8 = 256$	0.7827 ± 0.0136	0
$2^9 = 512$	0.9991 ± 0.0022	81
$2^{10} = 1024$	1.0000 ± 0.0000	100
$2^{11} = 2048$	0.9999 ± 0.0007	96

with ranks $r_1 = \cdots = r_{s-1} = 10$ we apply 20 sweeps of ARR (cf. Section 5.3) with all tensor core entries equal to one as initial guess. From the resulting approximation $\tilde{\mathbf{U}}$ of the exact solution $\mathbf{U}_r$, we then obtain the approximation $\tilde{u} \in L_n \subset L^2(Q)$ of the solution $u \in L^2(Q)$ of Problem 3.2, and the resulting approximately trained Fredholm network $\mathcal{G}\tilde{u}$ as explained in Section 6. From $\mathcal{G}\tilde{u}$ we extract discrete parameters $\tilde{\zeta} = (\tilde{\eta}_i, \tilde{u}(\tilde{\eta}_i))_{i=1}^N$ by quasi-Monte Carlo sampling for various numbers N of neurons to obtain the trained neural network $F_N(x, \tilde{\zeta})$, cf. Section 6. A sample $x \in X$ is then classified as 0 (genuine), if $F_N(x, \tilde{\zeta}) < 0.5$ and 1 (forged) otherwise.

In order to test the predictability of our trained neural networks for a fixed number N of neurons, we randomly select training and test sets of as described above, use quasi-Monte Carlo sampling to extract discrete parameters and evaluate the corresponding realization of the classification rate, i.e., the ratio of the number of correct predictions and the size of the selected test set. This procedure is repeated 100 times, in order to obtain expectation and standard deviation of the classification rates together with the number of perfect classifications, i.e., the number of test sets that are correctly classified without an exception. The results are shown in Table 1 for various N. As expected, the predictability is increasing with increasing N. While no test sets are perfectly classified for $N \leq 256$, we even achieve full optimality for $N = 1024$. Note that predictability deteriorates for even larger N, which might be due to numerical noise caused by round-off errors.

7.2. **Concrete compressive strength prediction.** We aim at the description of concrete compressive strength in terms of tolerated mechanical stress as a function F of the values $x = (x_k)_{k=1}^d$ of $d = 8$ parameters, i.e., the percentage of cement, fly ash, blast furnace slag, superplasticizer, fine aggregates, coarse aggregates, water, and age. To this end, we make use of a data set with 1030 samples that was collected over several years and has repeatedly been deployed for corresponding training of neural networks, see, e.g., [23, 24, 25, 26]. We randomly split the data set into $M = 824$ (80%) training and $m = 206$ test samples. As the output of each sample is distributed between 2.33 MPa and 82.60 MPa, different input variables might be of different orders of magnitude. In order to avoid large values dominating the others, min-max normalization is applied in analogy to (60), and we obtain $X = [0, 1]^d$.

We consider a neural network of the form (1) with ridge kernel and Heaviside activation function (3) together with the parameter set $Q = [-1, 1]^s \subset \mathbb{R}^s$, $s = d + 1 = 9$.

For discretization of the continuous training Problem 3.2, we choose the ansatz space L_n spanned by $n = 3^d = 6561$ tensor products of Chebyshev polynomials of the first kind up

TABLE 2. Results for the Concrete compressive strength prediction: Predictability of trained neural network $F_N(\cdot, \tilde{\zeta})$ for increasing N in terms of expectation and standard deviation of R^2 scores.

N	R_N^2 score (training data)	R_N^2 score (test data)
200	0.9023 ± 0.0051	0.8040 ± 0.0244
400	0.9406 ± 0.0027	0.8330 ± 0.0219
600	0.9563 ± 0.0023	0.8473 ± 0.0209
800	0.9648 ± 0.0019	0.8532 ± 0.0204
1000	0.9702 ± 0.0019	0.8586 ± 0.0198

to order 2. The regularization parameter is set to $\varepsilon = 10^{-8}$. Approximate factorization of the Heaviside activation function is performed as presented in the Appendix A.3, and we set $\ell_0 = 100$. Hence, the decomposition (53) of the corresponding approximate coefficient tensor $\mathbf{A} \in \mathbb{R}^{M \times n} = \mathbb{R}^{824 \times 6561}$ has the canonical rank $r = 201$. We normalize the tensor $\mathbf{A}$ such that $\|\mathbf{A}\|_{2, \mathbb{R}^{M \times n}} = 1$. The coefficient tensor corresponding to the test set is multiplied by the same normalization constant.

For the tensor-train approximation (55) in Problem 5.7 we choose the ranks $r_1 = \cdots = r_{s-1} = 10$ and apply 5 sweeps of ARR (cf. Section 5.3) with all tensor core entries equal to one as an initial guess. From the resulting approximate solution $\tilde{\mathbf{U}}$ we derive the approximate solution $\tilde{u}$ of Problem 3.2 which in turn provides the approximately trained Fredholm network $\mathcal{G}\tilde{u}$ as explained in Section 6. Discrete parameters $\tilde{\zeta} = (\tilde{\eta}_i, \tilde{u}(\tilde{\eta}_i))_{i=1}^N$ providing the corresponding neural network $F_N(\cdot, \tilde{\zeta})$ can finally be obtained by a suitable sampling strategy. In this experiment, we use both standard quasi-Monte Carlo sampling and inductive importance sampling as introduced in Section 6.

The prediction accuracy of the Fredholm network $\mathcal{G}\tilde{u}$ and of neural networks $F_N(\cdot, \tilde{\zeta})$ on given test data $(y_i, F(y_i))$, $i = 1, \ldots, m$, are measured by the usual R^2 scores

$$R_{\mathcal{G}}^2 = 1 - \frac{\sum_{i=1}^m (F(y_i) - (\mathcal{G}\tilde{u})(y_i))^2}{\sum_{i=1}^m (F(y_i) - \overline{F}(y))^2} \quad \text{and} \quad R_N^2 = 1 - \frac{\sum_{i=1}^m (F(y_i) - F_N(y_i, \tilde{\zeta}))^2}{\sum_{i=1}^m (F(y_i) - \overline{F}(y))^2},$$

respectively, where $\overline{F}(y) = \frac{1}{m} \sum_{i=1}^m F(y_i)$ stands for the algebraic mean of the test values. Note that for evaluation of $F_N(y_i, \tilde{\zeta})$ the exact Heaviside activation function σ_∞ is used and not its approximate factorization. Corresponding $R_{\mathcal{G}}^2$ and R_N^2 scores on training data are computed in the same way using the training data $(x_i, F(x_i))$, $i = 1, \ldots, M$, instead. Similar to the previous example, R^2 scores are considered as random variables, and we apply standard Monte Carlo with 100 samples each of random training and test data sets to approximately compute its expectation and standard deviation.

The resulting scores of the Fredholm network $\mathcal{G}\tilde{u}$ are

$$\text{training sets: } R_{\mathcal{G}}^2 = 0.9154 \pm 0.0029 \qquad \text{test sets: } R_{\mathcal{G}}^2 = 0.8701 \pm 0.0180.$$

with samples of the $R_{\mathcal{G}}^2$ scores on training sets ranging between 0.9093 and 0.9244 and on test sets between 0.8202 and 0.9058. These results indeed match the scores obtained by state-of-the-art ML approaches, cf. [26, 53].

However, in order to achieve similar R_N^2 scores for discrete neural networks $F_N(\cdot, \tilde{\zeta})$ as obtained by quasi-Monte Carlo sampling, it turned out in our numerical experiments that about $N = 2^{20} \approx 10^6$ sampling points were required. This motivates application of more

advanced sampling strategies such as inductive importance sampling introduced in Section 6. For this strategy, we obtained a considerable reduction of the number N of required samples, as illustrated in Table 2. As in the previous example, predictability increases with growing N, both on training and test data. Observe that the R_N^2 scores on training data are even better than for the Fredholm network $\mathcal{G}\tilde{u}$ which might be an outcome of the particular importance of training data for the applied sampling strategy. Even though the R_N^2 scores on test sets are slightly worse, they are still comparable with state-of-the-art ML approaches, cf. [26, 53].

7.3. **MNIST data set.** In our third experiment, we classify hand-written digits from the MNIST data set [27] that consists of representations of the ten digits $0, \ldots, 9$ as grayscale images of 28×28 pixels. We slightly modified the MNIST dataset by reducing the size of these images to $d = 14 \times 14 = 196$ pixels in order to lower the computational effort for classification. The data set is divided into $M = 60,000$ training images and $m = 10,000$ test images with associated labels, and we have $X = [0, 1]^d$.

For the neural network, we choose the ridge kernel ψ from (2) with logistic activation function σ_κ from (4) and $\kappa = 1$. More precisely, we use its FTT approximation based on the identification of the sample space X with its embedding $\{(x, 1)^\top \mid x \in X\} \subset \mathbb{R}^s$, $s = d + 1 = 197$ and Taylor approximation (66) at $z = 0$. See the Appendix A.2 for details. In order to provide sufficient accuracy of this approximation, we require $|z| = |x_i \cdot \eta| \leq 2$ for all data points $x_i \in X$. This is guaranteed by the choice $\eta \in Q = [-2/\rho_X, 2/\rho_X]^s$, denoting $\rho_X = \max_{i \in \{1, \ldots, M\}} \sum_{j=1}^{d+1} x_{i,j}$.

As each image $x \in X$ has to be classified as one of the ten digits, the target function F is now vector-valued. Utilizing one-hot encoding, we get $F(x) = (F_i(x))_{i=1}^{10} \in \mathbb{R}^{10}$, $x \in X$, with

$$(61) \qquad F_l(x) = \begin{cases} 1 & \text{if } x \text{ represents the digit } l \\ 0 & \text{otherwise} \end{cases}, \quad l = 0, \ldots, 9.$$

Corresponding vector-valued versions of the neural network (1) and of its continuous counterpart (3.5) are obtained by taking coefficients $u_i \in \mathbb{R}^{10}$ and parameter functions $u \in \left(L^2(Q)\right)^{10}$, respectively. Observe that each component u_l of u can be (approximately) computed separately from the scalar Problem 3.5 with $F = F_l$ and F_l defined in (61).

For the Ritz-Galerkin discretization of each of these scalar problems, we choose basis elements $\varphi_j^{(k)}$ of the form (44) with constant and linear functions $\varphi_{j_1}^{(k)}$ and $\varphi_{j_2}^{(k)}$, respectively. This leads to a subspace L_n of $L^2(Q)$ with dimension $n = 2^s = 2^{197} \approx 2 \cdot 10^{59}$. We choose the Tikhonov regularization parameter $\varepsilon = 10^{-15}$ and fix $r_1 = \cdots = r_d = 20$ in the tensor train approximation Problem 5.7. To the initial iterate of all core entries set to one (before orthonormalization) we apply 5 ARR sweeps to obtain the approximate solution $\tilde{\mathbf{U}}_l$, which provides the corresponding component $\mathcal{G}\tilde{u}_l$ of the (approximately) trained Fredholm network, cf. Section 6. We finally use the softmax function for classification, i.e. for each entry x the index $l_{\max}$ of the largest of the components $(\mathcal{G}\tilde{u}_l)(x)$, $l = 0, \ldots, 9$, determines the detected label.

For comparison with this trained Fredholm network, we consider single layer neural networks built from the well-established Keras library[3]. The input layer of the network comprises 196 nodes (one for each pixel), followed by a hidden layer with a varying number of N_K nodes and the logistic activation function. The output layer also uses one-hot encoding and the softmax function. For optimizing of the neural network parameters, 10% of the given training

---

[3] https://keras.io

set is used as a validation set. The parameters are then trained for a sufficiently large number of iterations in order to ensure convergence of the validation loss which indicates how well the model fits to unseen data.
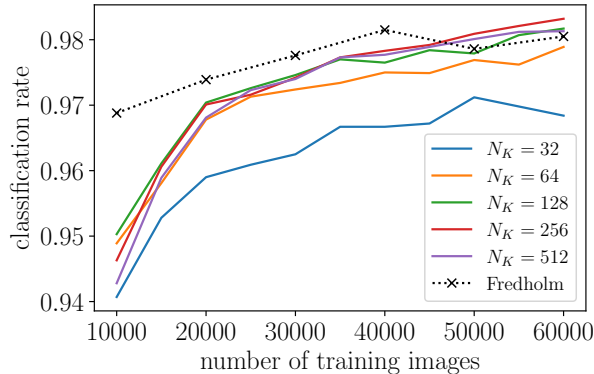

FIGURE 4. Results for the MNIST data set: Classification rates obtained from the Fredholm network over the number M of training images in comparison with Keras networks with increasing number N_K of nodes in the hidden layer.

Figure 4 shows the classification rates as obtained from our Fredholm network $\mathcal{G}\tilde{u}$ and of Keras networks of size $N_K = 34, 64, 128, 256, 512$ on the fixed MNIST test data set of $m = 10,000$ test images over the number M of utilized training images. For the largest number of $M = 60,000$ training data, the classification rates of the Keras networks are ranging from 0.9789 to 0.9832 for sufficiently large $N_K \geq 64$ and are thus comparable with the classification rate 0.9805 of the Fredholm network. However, Keras networks are clearly outperformed for smaller sets of training data.

8. CONCLUSION AND OUTLOOK

In this work, we have proposed a novel approach to function approximation and the training of large and high-dimensional shallow neural networks, based on their continuous asymptotic limit. Utilizing resulting Fredholm networks for function approximation and also for training in the original discrete case, we have thus traded highly nonlinear discrete training problems with finitely many unknown parameters for linear continuous training problems with infinitely many unknown values of a parameter function.

Fredholm training problems can be regarded as least-squares formulations of Fredholm integral equations of the first kind. Tikhonov regularization, approximation by Ritz-Galerkin methods and exact solution of the resulting linear system provide residual error estimates for the underlying loss functional. However, exact linear solution is unrealistic for most practical problems. Therefore, we particularly aimed at a significant reduction of memory consumption as well as computational costs to mitigate the curse of dimensionality as occurring for high-dimensional data. To this end, we described different tensor formats, suitable factorizations of kernel and basis functions as well as an alternating scheme for solving the tensorized linear training problems. Finally, we considered quasi Monte-Carlo sampling of discrete neural network parameters, and introduced inductive importance sampling which is directly targeting the loss functional $\mathcal{J}_N$. Note that residual error estimates extend to the trained neural network in the (unrealistic) case of exact linear solution.

The predictive quality of the resulting approximatively trained Fredholm and neural networks are illustrated by numerical experiments with three well-established test problems of regression and classification type. The results clearly confirm that, in spite of the apparently strong constraints (16), the presented Fredholm approach is highly competitive with state-of-the-art neural network-based methods concerning efficiency and reliability.

These promising results suggest future investigation of Fredholm-type counterparts of deep neural networks. For example, let σ be some continuous activation function and $\zeta^{(1)} = \left(U^{(1)}, W^{(1)}, b^{(1)}\right)$ with

$$W^{(1)} = \left(w_i^{(1)}\right)_{i=1}^{N^{(1)}} \in \mathbb{R}^{N^{(1)} \times d}, \quad b^{(1)} = \left(b_i^{(1)}\right)_{i=1}^{N^{(1)}} \in \mathbb{R}^{N^{(1)}}, \quad U^{(1)} = \frac{1}{N^{(1)}} \left(u_1^{(1)}, \ldots, u_{N^{(1)}}^{(1)}\right)$$

and $b^{(2)} \in \mathbb{R}$ be the parameters of the deep neural network

$$F_{N^{(1)}}\left(x, \zeta^{(1)}\right) = \sigma\left(U^{(1)}\sigma\left(W^{(1)}x + b^{(1)}\right) + b^{(2)}\right) = \sigma\left(\frac{1}{N^{(1)}} \sum_{i=1}^{N^{(1)}} u_i^{(1)}\sigma\left(w_i^{(1)} \cdot x + b_i^{(1)}\right) + b^{(2)}\right)$$

with input $x \in \mathbb{R}^d$, output $F_{N^{(1)}}\left(x, \zeta^{(1)}\right) \in \mathbb{R}$, one hidden layer with $N^{(1)}$ neurons, and another hidden layer with only one neuron for simplicity. Then we obtain

$$F_{N^{(1)}}\left(x, \zeta^{(1)}\right) \quad \rightarrow \quad \sigma\left(\mathcal{G}^{(1)}u^{(1)}(x) + b^{(2)}\right) \quad \text{for} \quad N^{(1)} \rightarrow \infty$$

with

$$\mathcal{G}^{(1)}v(x) = \frac{1}{|Q^{(1)}|} \int_{Q^{(1)}} v(\eta)\sigma(w \cdot x + b) \, d\eta \in \mathbb{R}, \qquad \eta = (w, b) \in Q^{(1)} \subset \mathbb{R}^{d+1},$$

by the reasoning presented in Section 3.1. A counterpart of (1) with two hidden layers is given by

$$F_{N^{(1)}, N^{(2)}}\left(x, \zeta^{(1)}, \zeta^{(2)}\right) = \frac{1}{N^{(2)}} \sum_{i=1}^{N^{(2)}} u_i^{(2)}\sigma\left(w_i^{(2)}U^{(1)}\sigma\left(W^{(1)}x + b^{(1)}\right) + b_i^{(2)}\right)$$

with additional parameters $\zeta^{(2)} = (\zeta_i^{(2)})_{i=1}^{N^{(2)}}$ and $\zeta_i^{(2)} = (u_i^{(2)}, w_i^{(2)}, b_i^{(2)}) \in \mathbb{R}^3$. Note that $W^{(2)} = (w_i^{(2)}U^{(1)})_{i=1}^{N^{(2)}} \in \mathbb{R}^{N^{(2)} \times N^{(1)}}$ is a rank-one matrix. The reasoning presented in Section 3.1 then yields

$$F_{N^{(1)}, N^{(2)}}\left(x, \zeta^{(1)}, \zeta^{(2)}\right) \quad \rightarrow \quad \mathcal{G}^{(2)}u^{(2)}\left(\mathcal{G}^{(1)}u^{(1)}(x)\right) \quad \text{for} \quad N^{(1)}, N^{(2)} \rightarrow \infty$$

with the additional Fredholm operator

$$\mathcal{G}^{(2)}v(y) = \frac{1}{|Q^{(2)}|} \int_{Q^{(2)}} v(\eta)\sigma(wy + b) \, d\eta \in \mathbb{R}, \qquad \eta = (w, b) \in Q^{(2)} \subset \mathbb{R}^2.$$

Deep Fredholm networks with more than two hidden layers are obtained in a similar way. While linearity of the minimization of corresponding loss functionals is clearly lost for deep Fredholm networks, the construction of appropriate Tensor-Gauß-Newton methods might still benefit from structural properties such as the monotonicity of activation functions or the block structure of iterated Fredholm operators.

Both shallow and deep Fredholm networks are open to promising investigations concerning analysis, numerical analysis, and algorithmic improvements, which might also shed some light on structural properties of the approximating neural networks.

Appendix A. Factorization of ridge kernel functions

According to (2), ridge kernel functions $\psi\colon X \times Q \to \mathbb{R}$, are given by

$$\psi(x,\eta) = \sigma(w \cdot x + b), \quad \eta = (w,b) \in Q \subset \mathbb{R}^s, \ s = d+1.$$

with hypercubes $X \in \mathbb{R}^d$, $Q = Q_1 \times \cdots \times Q_s \subset \mathbb{R}^s$, and ridge activation function $\sigma : \mathbb{R} \mapsto \mathbb{R}$. Identifying X with its embedding $\{(x,1)^\top \mid x \in X\} \subset \mathbb{R}^s$, ridge kernel functions take the form

$$\psi(x,\eta) = \sigma(x \cdot \eta).$$

In what follows, explicit tensor decompositons of different ridge activation functions are derived, either using the canonical format or TT format. In case of the latter, we use the core notation of the TT format [54] to represent FTT cores as two-dimensional arrays containing vector-valued functions as elements. More precisely, for a given functional tensor train $\hat{\boldsymbol{\Psi}}$ with cores $\hat{\boldsymbol{\Psi}}^{(k)} \in \mathbb{C}^{r_{k-1} \times M \times Q_k \times r_k}$ defined by

$$\hat{\boldsymbol{\Psi}}^{(k)}_{\ell_{k-1},i,\eta_k,\ell_k} = \psi^{(k)}_{\ell_{k-1},\ell_k}(x_{i,k},\eta_k)\colon \mathbb{R} \times Q_k \to \mathbb{C},$$

a single core is written as

$$(62) \qquad [\![\, \hat{\boldsymbol{\Psi}}^{(k)} \,]\!] = \left[\!\!\left[\begin{array}{ccc} \hat{\boldsymbol{\Psi}}^{(k)}_{1,:,\eta_k,1} & \cdots & \hat{\boldsymbol{\Psi}}^{(k)}_{1,:,\eta_k,r_k} \\ \vdots & \ddots & \vdots \\ \hat{\boldsymbol{\Psi}}^{(k)}_{r_{k-1},:,\eta_k,1} & \cdots & \hat{\boldsymbol{\Psi}}^{(k)}_{r_{k-1},:,\eta_k,r_k} \end{array}\right]\!\!\right],$$

where $\hat{\boldsymbol{\Psi}}^{(k)}_{\ell_{k-1},:,\eta_k,\ell_k} = (\psi^{(k)}_{\ell_{k-1},\ell_k}(x_{i,k},\eta_k))^M_{i=1}$. Thus, the FTT decomposition of the auxiliary tensor $\hat{\boldsymbol{\Psi}}$, cf. (49), takes the form $\hat{\boldsymbol{\Psi}} = [\![\, \hat{\boldsymbol{\Psi}}^{(1)} \,]\!] \otimes \cdots \otimes [\![\, \hat{\boldsymbol{\Psi}}^{(s)} \,]\!]$, denoting

$$\left([\![\, \hat{\boldsymbol{\Psi}}^{(k)} \,]\!] \otimes [\![\, \hat{\boldsymbol{\Psi}}^{(k+1)} \,]\!]\right)_{\ell_{k-1},i_k,\eta_k,i_{k+1},\eta_{k+1},\ell_{k+1}} = \sum_{\ell_k=1}^{r_k} \hat{\boldsymbol{\Psi}}^{(k)}_{\ell_{k-1},i_k,\eta_k,\ell_k} \cdot \hat{\boldsymbol{\Psi}}^{(k+1)}_{\ell_k,i_{k+1},\eta_{k+1},\ell_{k+1}}.$$

Proposition A.1. *Consider a monomial of the form $\psi(x,\eta) = (x \cdot \eta)^p$, $p \geq 0$, with $\eta \in Q$ and data points $x_i \in X$, $i = 1,\ldots,m$. The functional tensor $\boldsymbol{\Psi} \in \mathbb{R}^{M \times Q_1 \times \cdots \times Q_s}$ defined by $\boldsymbol{\Psi}_{i,\eta_1,\ldots,\eta_s} = \psi(x_i,\eta)$ can be written as $\boldsymbol{\Psi} = \Delta \cdot \hat{\boldsymbol{\Psi}}$, $\hat{\boldsymbol{\Psi}} \in \mathbb{R}^{M^s \times Q_1 \times \cdots \times Q_s}$, where*

$$(63) \quad \begin{aligned} \hat{\boldsymbol{\Psi}} = p!\, &\left[\!\!\left[\begin{array}{cccc} \tfrac{1}{p!}(x_{:,1}\eta_1)^p & \cdots & x_{:,1}\eta_1 & e \end{array}\right]\!\!\right] \otimes \left[\!\!\left[\begin{array}{cccc} e & 0 & \cdots & 0 \\ x_{:,2}\eta_2 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ \tfrac{1}{p!}(x'_{:,2}\eta_2)^p & \cdots & x_{:,2}\eta_2 & e \end{array}\right]\!\!\right] \otimes \cdots \\[2ex] &\otimes \left[\!\!\left[\begin{array}{cccc} e & 0 & \cdots & 0 \\ x_{:,s-1}\eta_{s-1} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ \tfrac{1}{p!}(x_{:,s-1}\eta_{s-1})^p & \cdots & x_{:,s-1}\eta_{s-1} & e \end{array}\right]\!\!\right] \otimes \left[\!\!\left[\begin{array}{c} e \\ x_{:,s}\eta_s \\ \vdots \\ \tfrac{1}{p!}(x_{:,s}\eta_s)^p \end{array}\right]\!\!\right] \end{aligned}$$

with $e = (1,\ldots,1)^T \in \mathbb{R}^M$.

Proof. By using the multinomial formula, we know

$$\left(\sum_{k=1}^{s} z_k\right)^p = \sum_{\substack{\ell_1+\cdots+\ell_s=p \\ \ell_1,\ldots,\ell_s\geq 0}} \frac{p!}{\ell_1!\cdots\ell_s!} z_1^{\ell_1}\cdots z_s^{\ell_s}$$

for any $i \in \{1,\ldots,M\}$. This can be written as a sequence of matrix multiplications of the form

$$\left(\sum_{k=1}^{s} z_k\right)^p = p! \begin{bmatrix} z_1^p & \cdots & z_1^0 \end{bmatrix} \begin{bmatrix} \displaystyle\sum_{\ell_2+\cdots+\ell_s=0} \frac{1}{\ell_2!\cdots\ell_s!} z_2^{\ell_2}\cdots z_s^{\ell_s} \\ \vdots \\ \displaystyle\sum_{\ell_2+\cdots+\ell_s=p} \frac{1}{\ell_2!\cdots\ell_s!} z_2^{\ell_2}\cdots z_s^{\ell_s} \end{bmatrix}$$

$$= p! \begin{bmatrix} z_2^p & \cdots & z_2^0 \end{bmatrix} \begin{bmatrix} z_2^0 & & 0 \\ \vdots & \ddots & \\ z_2^p & \cdots & z_2^0 \end{bmatrix} \begin{bmatrix} \displaystyle\sum_{\ell_3+\cdots+\ell_s=0} \frac{1}{\ell_3!\cdots\ell_s!} z_3^{\ell_3}\cdots z_s^{\ell_s} \\ \vdots \\ \displaystyle\sum_{\ell_3+\cdots+\ell_s=p} \frac{1}{\ell_3!\cdots\ell_s!} z_3^{\ell_3}\cdots z_s^{\ell_s} \end{bmatrix}$$

$$= \ldots$$

$$= p! \begin{bmatrix} z_2^p & \cdots & z_2^0 \end{bmatrix} \begin{bmatrix} z_2^0 & & 0 \\ \vdots & \ddots & \\ z_2^p & \cdots & z_2^0 \end{bmatrix} \cdots \begin{bmatrix} z_{s-1}^0 & & 0 \\ \vdots & \ddots & \\ z_{s-1}^p & \cdots & z_{s-1}^0 \end{bmatrix} \begin{bmatrix} z_s^0 \\ \vdots \\ z_s^p \end{bmatrix}.$$

Thus, for $\boldsymbol{\Psi} = \Delta\hat{\boldsymbol{\Psi}}$ with $\hat{\boldsymbol{\Psi}}$ as given in (63), we get

$$\boldsymbol{\Psi}_{i,\eta_1,\ldots,\eta_s} = \sum_{j_1=1}^{M}\cdots\sum_{j_s=1}^{M} \Delta_{i,j_1,\ldots,j_s}\,\hat{\boldsymbol{\Psi}}_{j_1,\ldots,j_s,\eta_1,\ldots,\eta_s}$$

$$= \sum_{j_1=1}^{M}\cdots\sum_{j_s=1}^{M} \delta_{i,j_1}\cdots\delta_{i,j_s}\,\hat{\boldsymbol{\Psi}}_{j_1,\ldots,j_s,\eta_1,\ldots,\eta_s}$$

$$= \hat{\boldsymbol{\Psi}}_{i,\ldots,i,\eta_1,\ldots,\eta_s} = \left(\sum_{k=1}^{s} x_{i,k}\eta_k\right)^p = (x_i\cdot\eta)^p. \qquad \square$$

A.1. **Trivial activation function.** In the simplest case $\sigma = \text{Id}$, the ridge kernel function is bilinear, i.e., $\psi(x,\eta) = x\cdot\eta$, and clearly can be written in the form (48) (see Example 5.3). Let us consider the functional tensor

$$\boldsymbol{\Psi} \in \mathbb{R}^{M\times Q_1\times\cdots\times Q_s}, \qquad \boldsymbol{\Psi}_{i,\eta_1,\ldots,\eta_s} = x_i\cdot\eta = \sum_{k=1}^{s} x_{i,k}\eta_k, \qquad s = d,$$

associated to ψ according to (47) for the given data points $x_i \in X$, $i = 1,\ldots,M$. It can be expressed as $\boldsymbol{\Psi} = \Delta\cdot\hat{\boldsymbol{\Psi}}$, see Proposition 5.5, where the auxiliary tensor $\hat{\boldsymbol{\Psi}}$ has the functional

canonical decomposition

$$(64) \quad \hat{\boldsymbol{\Psi}} = \eta_1 \begin{pmatrix} x_{1,1} \\ \vdots \\ x_{M,1} \end{pmatrix} \otimes \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \otimes \cdots \otimes \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} + \cdots + \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \otimes \cdots \otimes \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix} \otimes \eta_d \begin{pmatrix} x_{1,d} \\ \vdots \\ x_{M,d} \end{pmatrix}$$

Using the core notation (62) and Proposition A.1, we can also represent $\hat{\boldsymbol{\Psi}}$ as a functional tensor train:

$$(65) \quad \hat{\boldsymbol{\Psi}} = \llbracket\, x_{:,1}\eta_1 \quad e \,\rrbracket \otimes \begin{bmatrix} e & 0 \\ x_{:,2}\eta_2 & e \end{bmatrix} \otimes \cdots \otimes \begin{bmatrix} e & 0 \\ x_{:,d-1}\eta_{d-1} & e \end{bmatrix} \otimes \begin{bmatrix} e \\ x_{:,d}\eta_d \end{bmatrix}.$$

Note that the canonical rank of (64) is equal to d whereas the FTT ranks $r_1, \ldots, r_{d-1}$ of (65) are always equal to 2 for any state space dimension d.

A.2. **Logistic activation function.** For the logistic activation function σ_κ given in (4), the kernel function takes the form

$$(66) \quad \psi_\kappa(x, \eta) = \frac{\exp(\kappa\, x \cdot \eta)}{1 + \exp(\kappa\, x \cdot \eta))} = \frac{1}{1 + \exp(-\kappa\, x \cdot \eta)} = \frac{1}{2} + \frac{1}{2}\tanh\left(\frac{-\kappa\, x \cdot \eta}{2}\right).$$

We expand $\sigma_\kappa(z)$ into a Taylor series at $z = 0$, to obtain

$$(67) \quad \psi_\kappa(x, \eta) = \frac{1}{2} + \frac{\kappa\, x \cdot \eta}{4} - \frac{\kappa^3\, (x \cdot \eta)^3}{48} + O(x \cdot \eta)^5.$$

Assuming that the side lengths of Q are sufficiently small, we skip the quintic part in (67). The resulting approximation of ψ_κ can be written in the form of (48) with $r = 1 + s + \frac{1}{6}(s^3 + 3s^2 + 2s)$. Indeed, the zero- and the first-order term in (67) can be written as a functional canonical tensor with rank $1 + s$ summands, cf. (64), while, by the multinomial formula, the cubic term can be written as

$$\frac{\kappa^3\, (x \cdot \eta)^3}{48} = \frac{\kappa^3}{48}\left(\sum_{k=1}^s x_k\eta_k\right)^3 = \frac{\kappa^3}{48}\sum_{\ell_1 + \cdots + \ell_s = 3} \frac{6}{\ell_1! \cdots \ell_s!}(x_1\eta_1)^{\ell_1} \cdots (x_s\eta_s)^{\ell_s}$$

with $\sum_{k_3=1}^s \sum_{k_2=1}^{k_3} \sum_{k_1=1}^{k_2} 1 = \frac{1}{6}(s^3 + 3s^2 + 2s)$ summands. In contrast to that, the ranks of the corresponding FTT decomposition are always bounded by 4 for any parameter dimension s. This can be seen by using Proposition A.1 for each of the first three terms in (67) and adding the respective decompositions. The resulting FTT representation of the auxiliary tensor is then given by

$$\hat{\boldsymbol{\Psi}} = \llbracket\, \tfrac{1}{6}(\eta_1 x_{:,1})^3 \quad \tfrac{1}{2}(\eta_1 x_{:,1})^2 \quad \eta_1 x_{:,1} \quad e \,\rrbracket \otimes \begin{bmatrix} e & 0 & 0 & 0 \\ \eta_2 x_{:,2} & e & 0 & 0 \\ \tfrac{1}{2}(\eta_2 x_{:,2})^2 & \eta_2 x_{:,2} & e & 0 \\ \tfrac{1}{6}(\eta_2 x_{:,2})^3 & \tfrac{1}{2}(\eta_2 x_{:,2})^2 & \eta_2 x_{:,2} & e \end{bmatrix} \otimes \cdots$$

$$\otimes \begin{bmatrix} e & 0 & 0 & 0 \\ \eta_{d-1} x_{:,d-1} & e & 0 & 0 \\ \tfrac{1}{2}(\eta_{d-1} x_{:,d-1})^2 & \eta_{d-1} x_{:,d-1} & e & 0 \\ \tfrac{1}{6}(\eta_{d-1} x_{:,d-1})^3 & \tfrac{1}{2}(\eta_{d-1} x_{:,d-1})^2 & \eta_{d-1} x_{:,d-1} & e \end{bmatrix} \otimes \begin{bmatrix} -\tfrac{\kappa^3}{8} e \\ -\tfrac{\kappa^3}{8} \eta_d x_{:,d} \\ -\tfrac{\kappa^3}{16}(\eta_d x_{:,d})^2 + \tfrac{\kappa}{4} e \\ -\tfrac{\kappa^3}{48}(\eta_d x_{:,d})^3 + \tfrac{\kappa}{4}\eta_d x_{:,d} + \tfrac{1}{2} e \end{bmatrix}.$$

A.3. **Heaviside activation function.** For each fixed $z \in \mathbb{R}$ and $\kappa \to \infty$ the logistic activation function $\sigma_\kappa(z)$ converges to the Heaviside function $\sigma_\infty(z)$ defined in (3). However, as the interval of convergence $[-\frac{\pi}{\kappa}, \frac{\pi}{\kappa}]$ of the Taylor expansion (67) shrinks with incresing κ, we now derive a different kind of approximation of the Heaviside kernel function $\psi_\infty(x, \eta) = \sigma_\infty(x \cdot \eta)$ of the form (48). Our starting point is the Fourier epansion

$$\sigma_\infty(z) = \frac{1}{2} + \frac{2}{\pi} \sum_{\ell=1}^\infty \frac{1}{2\ell - 1} \sin\left(\frac{(2\ell - 1)\pi}{T} z\right),$$

which converges pointwise within the interval $[-T, T]$ for suitable $T > 0$. Euler's formula $\sin(z) = (\exp(\mathrm{i}z) - \exp(-\mathrm{i}z))/2\mathrm{i}$ yields

$$\sigma_\infty(z) = \frac{1}{2} + \frac{1}{\pi\mathrm{i}} \sum_{\ell=1}^\infty \frac{1}{2\ell - 1} \left(\exp\left(\frac{(2\ell - 1)\pi\mathrm{i}}{T} z\right) - \exp\left(-\frac{(2\ell - 1)\pi\mathrm{i}}{T} z\right)\right).$$

Inserting $z = x \cdot \eta$, we obtain the representation

$$\psi_\infty(x, \eta) = \frac{1}{2} + \frac{1}{\pi\mathrm{i}} \sum_{\ell=1}^\infty \frac{1}{2\ell - 1} \left(\prod_{k=1}^d \exp\left(\frac{(2\ell - 1)\pi\mathrm{i}}{T} x_k \eta_k\right) - \prod_{k=1}^d \exp\left(-\frac{(2\ell - 1)\pi\mathrm{i}}{T} x_k \eta_k\right)\right)$$

of the Heaviside kernel function for $x \cdot \eta \in [-T, T]$. Truncating this expansion at $\ell = \ell_0$, we obtain the desired factorization (48) of the resulting approximation with $r = 2\ell_0 + 1$ summands. In this case, the decomposition (50) of the associated functional tensor $\mathbf{\Psi}$ reads

$$\mathbf{\Psi} = \Delta \cdot \left[\frac{1}{2} \bigotimes_{k=1}^d e \; + \; \frac{1}{\pi\mathrm{i}} \sum_{\ell=1}^{\ell_0} \frac{1}{2\ell - 1} \bigotimes_{k=1}^d \exp\left(\frac{(2\ell - 1)\pi\mathrm{i}}{T} \eta_k x_{:,k}\right)\right.$$
$$\left. - \frac{1}{\pi\mathrm{i}} \sum_{\ell=1}^{\ell_0} \frac{1}{2\ell - 1} \bigotimes_{k=1}^d \exp\left(-\frac{(2\ell - 1)\pi\mathrm{i}}{T} \eta_k x_{:,k}\right)\right].$$

References

[1] Shun-ichi Amari. Backpropagation and stochastic gradient descent method. *Neurocomputing*, 5(4-5):185–196, 1993.
[2] Léon Bottou. Stochastic gradient descent tricks. In *Neural networks: Tricks of the trade*, pages 421–436. Springer, 2012.
[3] Patrick Cheridito, Arnulf Jentzen, and Florian Rossmannek. Non-convergence of stochastic gradient descent in the training of deep neural networks. *Journal of Complexity*, 64:101540, 2021.
[4] Qipin Chen and Wenrui Hao. Randomized Newton's method for solving differential equations based on the neural network discretization. *Journal of Scientific Computing*, 92(2):1–22, 2022.
[5] Andrew R. Conn, Nicholas I. M. Gould, and Philippe L. Toint. *Trust region methods*. SIAM, 2000.
[6] Eiji Mizutani and James W. Demmel. On structure-exploiting trust-region regularized nonlinear least squares algorithms for neural-network learning. *Neural Networks*, 16(5-6):745–753, 2003.
[7] Alena Kopanicáková and Rolf Krause. Globally convergent multilevel training of deep residual networks. *SIAM Journal on Scientific Computing*, pages S254–S280, 2022. doi:https://doi.org/10.1137/21M1434076.
[8] Grant M. Rotskoff and Eric Vanden-Eijnden. Trainability and accuracy of neural networks: An interacting particle system approach, 2019. arXiv:1805.00915.
[9] Qichao Que and Mikhail Belkin. Back to the future: Radial basis function networks revisited. In *Artificial intelligence and statistics*, pages 1375–1383. PMLR, 2016.
[10] Qichao Que, Mikhail Belkin, and Yusu Wang. Learning with Fredholm kernels. *Advances in neural information processing systems*, 27, 2014.

[11] Wei Wang, Hao Wang, Zhaoxiang Zhang, Chen Zhang, and Yang Gao. Semi-supervised domain adaptation via Fredholm integral based kernel methods. *Pattern Recognition*, 85:185–197, 2019.

[12] Qichao Que and Mikhail Belkin. Inverse density as an inverse problem: The Fredholm equation approach. *Advances in neural information processing systems*, 26, 2013.

[13] Atsushi Nitanda, Denny Wu, and Taiji Suzuki. Convex analysis of the mean field Langevin dynamics. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 9741–9757. PMLR, 2022.

[14] Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.

[15] Grant M Rotskoff and Eric Vanden-Eijnden. Neural networks as interacting particle systems: Asymptotic convexity of the loss landscape and universal scaling of the approximation error. *stat*, 1050:22, 2018.

[16] Kaitong Hu, Zhenjie Ren, David Šiška, and Łukasz Szpruch. Mean-field Langevin dynamics and energy landscape of neural networks. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 57(4):2043 – 2065, 2021. `doi:10.1214/20-AIHP1140`.

[17] Lenaic Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. *Advances in neural information processing systems*, 31, 2018.

[18] Lénaïc Chizat. Mean-field Langevin dynamics: Exponential convergence and annealing, 2022. `arXiv: 2202.01009`.

[19] Martin Hanke Heinz W. Engl and Andreas Neubauer. *Regularization of Inverse Problems*. Kluwer Academic Publishers, Dordrecht, Netherlands, 1996.

[20] Sebastian Holtz, Thorsten Rohwedder, and Reinhold Schneider. The alternating linear scheme for tensor optimization in the tensor train format. *SIAM Journal on Scientific Computing*, 34(2):A683–A713, 2012. `doi:10.1137/100818893`.

[21] Stefan Klus and Patrick Gelß. Tensor-based algorithms for image classification. *Algorithms*, 12(11), 2019. `doi:10.3390/a12110240`.

[22] Volker Lohweg, Jan Leif Hoffmann, Helene Dörksen, Roland Hildebrand, Eugen Gillich, Jürg Hofmann, and Johannes Schaede. Banknote authentication with mobile devices. In Adnan M. Alattar, Nasir D. Memon, and Chad D. Heitzenrater, editors, *Media Watermarking, Security, and Forensics 2013*, volume 8665, page 866507. International Society for Optics and Photonics, SPIE, 2013. `doi:10.1117/12.2001444`.

[23] I-Cheng Yeh. Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete Research*, 28(12):1797–1808, 1998. `doi:10.1016/S0008-8846(98)00165-3`.

[24] I-Cheng Yeh. Prediction of strength of fly ash and slag concrete by the use of artificial neural networks. *J. Chin. Inst. Civil Hydraul. Eng.*, 15:659–663, 2003.

[25] I-Cheng Yeh. Analysis of strength of concrete using design of experiments and neural networks. *Journal of Materials in Civil Engineering*, 18(4):597–604, 2006. `doi:10.1061/(ASCE)0899-1561(2006)18:4(597)`.

[26] Syyed Adnan Raheel Shah, Marc Azab, Hany M. Seif ElDin, Osama Barakat, Muhammad Kashif Anwar, and Yasir Bashir. Predicting compressive strength of blast furnace slag and fly ash based sustainable concrete using machine learning techniques: An application of advanced decision-making approaches. *Buildings*, 12(7), 2022. `doi:10.3390/buildings12070914`.

[27] Yann Lecun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. `doi:10.1109/5.726791`.

[28] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.

[29] Sanjeev Arora, Simon S Du, Zhiyuan Li, Ruslan Salakhutdinov, Ruosong Wang, and Dingli Yu. Harnessing the power of infinitely wide deep nets on small-data tasks. *arXiv preprint arXiv:1910.01663*, 2019.

[30] Allan Pinkus. Approximation theory of the mlp model in neural networks. *Acta numerica*, 8:143–195, 1999.

[31] Bekir Karlik and A Vehbi Olgac. Performance analysis of various activation functions in generalized MLP architectures of neural networks. *International Journal of Artificial Intelligence and Expert Systems*, 1:111–122, 2011.

[32] Tianping Chen and Hong Chen. Approximation capability to functions of several variables, nonlinear functionals, and operators by radial basis function neural networks. *IEEE Transactions on Neural Networks*, 6(4):904–910, 1995.

[33] Maxwell Stinchcombe and Halbert White. Approximating and learning unknown mappings using multilayer feedforward networks with bounded weights. In *1990 IJCNN International Joint Conference on Neural Networks*, pages 7–16. IEEE, 1990.

[34] Rainer Kress. *Linear Integral Equations*. Springer-Verlag, New York, 1999.

[35] Charles W. Groetsch. *The theory of Tikhonov regularization for Fredholm equations of the first kind*. Pitman Publishing Limited, London, 1984.

[36] George Cybenko. Approximation by superpositions of a sigmoidal function. *Math. Control Signal Systems*, 2:303–314, 1989. `doi:10.1007/BF02551274`.

[37] Mario Bebendorf. A note on the poincaré inequality for convex domains. *Zeitschrift für Analysis und ihre Anwendungen*, 22(4):751–756, 2003.

[38] Frank L Hitchcock. The expression of a tensor or a polyadic as a sum of products. *Journal of Mathematics and Physics*, 6(1-4):164–189, 1927.

[39] Ivan V. Oseledets and Eugene E. Tyrtyshnikov. Breaking the curse of dimensionality, or how to use SVD in many dimensions. *SIAM Journal on Scientific Computing*, 31(5):3744–3759, 2009.

[40] Ivan V. Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011. `doi:10.1137/090752286`.

[41] Antonio Falcó and Wolfgang Hackbusch. On minimal subspaces in tensor representations. *Foundations of computational mathematics*, 12(6):765–803, 2012.

[42] Sebastian Holtz, Thorsten Rohwedder, and Reinhold Schneider. On manifolds of tensors of fixed TT-rank. *Numerische Mathematik*, 120(4):701–731, 2012.

[43] Tamara G. Kolda and Brett W. Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.

[44] Wolfgang Hackbusch. *Tensor spaces and numerical tensor calculus*, volume 42. Springer, 2012.

[45] Alex Gorodetsky, Sertac Karaman, and Youssef Marzouk. A continuous analogue of the tensor-train decomposition. *Computer Methods in Applied Mechanics and Engineering*, 347:59–84, 2019. `doi:10.1016/j.cma.2018.12.015`.

[46] Patrick Gelß. *The tensor-train format and its applications: Modeling and analysis of chemical reaction networks, catalytic processes, fluid flows, and Brownian dynamics*. PhD thesis, FU Berlin, 2017.

[47] Patrick Gelß, Stefan Klus, Jens Eisert, and Christof Schütte. Multidimensional approximation of nonlinear dynamical systems. *Journal of Computational and Nonlinear Dynamics*, 14(6):061006, 2019. `doi:10.1115/1.4043148`.

[48] Kim Batselier. Enforcing symmetry in tensor network MIMO Volterra identification. *IFAC-PapersOnLine*, 54(7):469–474, 2021. `doi:10.1016/j.ifacol.2021.08.404`.

[49] Sergey V. Dolgov and Dmitry V. Savostyanov. Alternating minimal energy methods for linear systems in higher dimensions. *SIAM Journal on Scientific Computing*, 36(5):A2248–A2271, 2014. `doi:10.1137/140953289`.

[50] Lars Grasedyck and Sebastian Krämer. Stable ALS approximation in the TT-format for rank-adaptive tensor completion. *Numerische Mathematik*, 143:855–904, 2019. `doi:10.1007/s00211-019-01072-4`.

[51] Thorsten Rohwedder and André Uschmajew. On local convergence of alternating schemes for optimization of convex problems in the tensor train format. *SIAM Journal on Numerical Analysis*, 51(2):1134–1162, 2013. `doi:10.1137/110857520`.

[52] Mike Espig, Wolfgang Hackbusch, and Aram Khachatryan. On the convergence of alternating least squares optimisation in tensor format representations, 2015. `arXiv:1506.00062`.

[53] Benjamin A. Young, Alex Hall, Laurent Pilon, Puneet Gupta, and Gaurav Sant. Can the compressive strength of concrete be estimated from knowledge of the mixture proportions?: New insights from statistical analysis and machine learning methods. *Cement and Concrete Research*, 115:379–388, 2019. `doi:10.1016/j.cemconres.2018.09.006`.

[54] Patrick Gelß, Stefan Klus, Sebastian Matera, and Christof Schütte. Nearest-neighbor interaction systems in the tensor-train format. *Journal of Computational Physics*, 341:140–162, 2017.

Patrick Gelss, AI in Society, Science, and Technology, Zuse Institute Berlin, 14195 Berlin, Germany

Email address: `gelss@zib.de`

Aizhan Issagali, Institut für Mathematik, Freie Universität Berlin, 14195 Berlin, Germany
Email address: `aizhan.issagali@fu-berlin.de`

Ralf Kornhuber, Institut für Mathematik, Freie Universität Berlin, 14195 Berlin, Germany
Email address: `kornhuber@math.fu-berlin.de`