

Statistical guarantees for denoising reflected diffusion models

Asbjørn Holk

*Department of Mathematics
Aarhus University
Aarhus, Denmark*

A.HOLK@MATH.AU.DK

Claudia Strauch

*Institute for Mathematics
Heidelberg University
Heidelberg, Germany*

STRAUCH@MATH.UNI-HEIDELBERG.DE

Lukas Trottner*

*Department of Mathematics
University of Stuttgart
Stuttgart, Germany*

LUKAS.TROTTNER@ISA.UNI-STUTTART.DE

Editor: Qiang Liu

Abstract

In recent years, denoising diffusion models have become a crucial area of research due to their abundance in the rapidly expanding field of generative AI. While recent statistical advances have delivered explanations for the generation ability of idealised denoising diffusion models for high-dimensional target data, implementations introduce thresholding procedures for the generating process to overcome issues arising from the unbounded state space of such models. This mismatch between theoretical design and implementation of diffusion models has been addressed empirically by using a *reflected* diffusion process as the driver of noise instead. In this paper, we study statistical guarantees of these denoising reflected diffusion models. In particular, under Sobolev smoothness assumptions, we establish rates of convergence in total variation which, up to a polylogarithmic factor, match the minimax lower bound. Our main contributions include the statistical analysis of this novel class of denoising reflected diffusion models and a refined score approximation method in both time and space, leveraging spectral decomposition and rigorous neural network analysis.

Keywords: diffusion models, score-based generative models, reflected diffusions, minimax optimality, deep neural networks

1. Introduction

Deep generative models (DGMs) are a broad class of models that train deep neural networks to generate synthetic samples from a target distribution representing a given training data set. Examples that have shown tremendous empirical success in the last decade include the classes of Generative Adversarial Networks (GANs) (Goodfellow et al., 2020; Nowozin et al., 2016; Arjovsky et al., 2017; Zhao et al., 2017), Variational Autoencoders (Kingma and Welling, 2014) and normalising flows (Papamakarios et al., 2021; Rezende and Mohamed,

*. Corresponding author

2015; Tabak and Turner, 2013). Most recently, dynamic generative methods, which work by learning the reverse dynamics of a forward noising process that gradually evolves the original data distribution into a simple and easy-to-sample-from distribution, have gained much traction in the machine learning community. The most prominent class of such models are Denoising Diffusion Models (DDMs) in discrete and continuous time (Song and Ermon, 2019; Ho et al., 2020; Song et al., 2021; Sohl-Dickstein et al., 2015), which add Gaussian noise to the data and for which the backward dynamics is determined by the gradient of the log likelihood of the forward marginals, also known as the score. The score is intractable as it depends on the unknown initial data distribution and therefore needs to be learned, which is why these models are also referred to as score-based generative models.

These models have been adapted and improved for enhanced efficiency and performance in numerous applications. However, the statistical theory for standard DGMs is still under active development with many basic questions left unanswered. The statistical theory for GANs is perhaps the most mature among DGMs, with significant efforts in the recent past (Puchkin et al., 2024; Tang and Yang, 2023; Schreuder et al., 2021; Chen et al., 2022; Chae, 2022; Liang, 2021; Stéphanovitch et al., 2024b,a) to build a minimax theory that, among other things, draws on deeper connections to optimal transport theory (Villani, 2009; Chewi et al., 2025).

Related work Starting only recently with the seminal paper Oko et al. (2023), the study of statistical convergence guarantees for standard denoising diffusion models has experienced rapid growth in the past two years. For denoising diffusion models in \mathbb{R}^d that transform data into noise via a forward Ornstein–Uhlenbeck process, Oko et al. (2023) prove total variation and Wasserstein-1 rates that match the minimax lower bound over Besov classes up to log factors for the total variation and arbitrarily small polynomial loss for the Wasserstein-1 distance. Their assumptions on the initial distribution with density p_0 can be summarised by three key components:

- (i) p_0 is compactly supported on the d -dimensional hypercube $[-1, 1]^d$;
- (ii) p_0 is bounded away from zero on its support;
- (iii) p_0 has Besov smoothness of order s away from the support boundary (where s is allowed to be sufficiently small to not necessarily imply continuity of p_0) and is infinitely differentiable close to the boundary.

These assumptions provide a natural starting point from a statistical perspective, since they allow one to easily control certain quantities in the approximation analysis that would otherwise obscure the central mathematical ideas. However, these assumptions have practical limitations since, in many empirical applications, data distributions tend to be multimodal and extremely high-dimensional. Consequently, the resulting optimal rates, derived in terms of the ambient dimension d , often do not align with the empirical success of diffusion models in generating high-quality samples in such complex scenarios. This gap has motivated recent studies to extend these results to cases where data distributions are constrained to lower-dimensional structures, aligning with the widely held *manifold hypothesis*, which suggests that popular high-dimensional training data sets are supported on lower-dimensional manifolds. Chen et al. (2023), Tang and Yang (2024) and Azangulov et al. (2024) have

expanded the analysis to such structured settings, refining the rates in Wasserstein-1 distance and providing theoretical foundations for understanding the adaptivity of DDMs to low-dimensional structures.

Our contribution Our work offers a different type of extension: we provide statistical convergence guarantees for *denoising reflected diffusion models* (DRDMs). This class of generative models was introduced and empirically implemented in Lou and Ermon (2023), motivated by the observation that, in practice, the implementation of the generation process of standard DDMs relies significantly on the incorporation of thresholding procedures to prevent the backward process, which has unbounded state space, from exiting the permitted range of target samples (say, e.g., tensors containing bounded RGB values of pixels in an image). Since such thresholding procedures have no theoretical justification in the design of unconstrained diffusion models, Lou and Ermon (2023) suggest using reflected diffusion processes on bounded domains D as the forward noising process and demonstrate competitive performance to state-of-the-art models with comparable training and inference times. Specifically, Lou and Ermon (2023) implement a time-changed reflected Brownian motion as forward model, which yields a uniform ergodic distribution on the bounded domain \bar{D} that is used as initialisation of the backward generative process with learned score obtained by denoising score matching, see below. We also refer to Fishman et al. (2023a) and Fishman et al. (2023b) for related work on constrained diffusion models. To help the reader quickly identify the main technical novelties and to place our results in the context of the existing literature, let us briefly summarise core contributions of this paper:

- We derive an explicit upper bound on the expected total variation distance between the data distribution p_0 and the distribution induced by the DRDM with estimated score, with a resulting convergence rate which matches the well-known minimax lower bounds over Sobolev classes up to logarithmic factors.
- We derive a spectral characterisation of the time-dependent score function for reflected diffusions generated by self-adjoint operators with Neumann boundary conditions, yielding an explicit representation amenable to statistical analysis.
- We construct a calibrated neural network approximation scheme for these spectral score representations, combining truncation of eigenexpansions with space-time interpolation to balance approximation accuracy and model complexity.
- We analyse how the semi-explicit nature of transition densities in the reflected setting impacts score estimation and show how the resulting approximation and estimation errors can be controlled despite the absence of closed-form Gaussian expressions.

Having outlined these contributions, we now describe the construction and analysis of the reflected denoising dynamics in more detail. Time-reversal of the forward noising process motivates the formulation of the backward generative model, which under suitable regularity assumptions is again given by a reflected diffusion whose drift is determined by the score function

$$\mathfrak{s}^\circ(x, t) := \nabla \log p_t(x), \quad p_t(x) := \int_D q_t(y, x) \mathbb{P}(X_0 \in dy),$$

where $q_t(x, y)$ are the transition densities of the reflected forward diffusion and $\mathbb{P}(X_0 \in \cdot)$ is the data distribution. Since the latter is unknown, the score must be learned from a given i.i.d. data sample $(X_{0,i})_{i=1, \dots, n}$, which, similarly to the unconstrained diffusion case, is implemented by minimising an empirical version of the denoising score matching loss

$$\int_{\underline{T}}^{\overline{T}} \int_{D^2} |\mathfrak{s}(y, t) - \nabla_y \log q_t(x, y)|^2 q_t(x, y) \mathbb{P}(X_0 \in dx) dy dt,$$

in a class of feedforward ReLU neural networks $\mathcal{S} \ni \mathfrak{s}$, where $\overline{T} - \underline{T} \in (0, \overline{T}]$ is the stopping time of the backward reflected diffusion. By the equivalence of denoising and explicit score matching (Vincent, 2011), it therefore becomes crucial to precisely calibrate the neural network class in terms of the number of observations n , the dimension d and the smoothness s of the data to balance the explicit score approximation error

$$\min_{\mathfrak{s} \in \mathcal{S}} \int_{\underline{T}}^{\overline{T}} \int_D |\mathfrak{s}(x, t) - \nabla \log p_t(x)|^2 p_t(x) dx dt$$

and the complexity of the class \mathcal{S} , measured in terms of its covering number w.r.t. the supremum norm. While the approximation analysis for standard diffusion models with Ornstein–Uhlenbeck forward noising process (Oko et al., 2023; Chen et al., 2023; Tang and Yang, 2024; Azangulov et al., 2024) can rely on the explicitly known Gaussian transition densities $q_t(x, y)$, this situation dramatically changes in the reflected model, where the transition densities are at best semi-explicit and need to be numerically approximated. We choose a forward reflected diffusion with self-adjoint weighted Laplacian $\nabla \cdot f \nabla$ subject to a Neumann boundary condition as a generator, which provides us with a space-time spectral decomposition of the transition densities of the form

$$q_t(x, y) = \sum_{j=0}^{\infty} e^{-t\lambda_j} e_j(x) e_j(y), \quad t > 0, x, y \in D, \quad (1)$$

for eigenpairs $(\lambda_j, e_j)_{j \in \mathbb{N}_0}$ of $-\nabla \cdot f \nabla$ with known growth behaviour. The statistical estimation of the conductivity f based on such reflected diffusion data has been recently studied for low- and high-frequency observations in Nickl (2024) and Hoffmann and Ray (2024), respectively. In the low-frequency statistical setting, the computational challenges arising from the semi-explicit nature of (1) have been subsequently analysed in Giordano and Wang (2025). Let us note that, for the specific choice of $f \equiv 1/2$, the generator is given by $\frac{1}{2}\Delta$ and thus yields a reflected Brownian motion as a forward model. Up to a time change, we therefore cover the reflected diffusion models implemented in Lou and Ermon (2023). Independently of our choice for the heat conductivity $f: \mathbb{R}^d \rightarrow [f_{\min}, \infty) \subset (0, \infty)$, the forward noising process has a uniform ergodic distribution from which we sample to initialise the backward generating process. The conductivity $f(x)$ controls the speed at which the forward process diffuses around x , with larger values pushing the process faster towards equilibrium, but also increasing the oscillation of the eigenfunctions e_j around x . This implies a tradeoff between the convergence speed of the forward model and the approximation difficulty of the backward model. For some potential $V: \mathbb{R}^d \rightarrow \mathbb{R}$ we could also consider the generator $\mathcal{A}_V := e^V \nabla \cdot e^{-V} f \nabla$ with Neumann boundary conditions instead, resulting

in a normally reflected forward diffusion with stationary distribution $\mu_V \propto e^{-V}$ and a spectral decomposition as in (1) with respect to the dominating measure $\mu_V(y) dy$, see Nickl (2024). For the particular choice of $f = 1/2$, this yields a reflected (overdamped) Langevin diffusion process. Our general approximation approach, which is described below, could be adapted to this scenario. However, in the context of generative modelling, we are primarily interested in a stationary forward distribution that is easy to sample from. As long as the domain D is sufficiently nice, this is clearly satisfied for the uniform ergodic distribution in the model considered here, but it is generally difficult for space-dependent choices of V , which require sampling from $\mu_V \propto e^{-V}$ via MCMC methods, which are expensive in the dimension d . From a practical perspective, it is therefore quite natural to consider reflected diffusion models with uniform stationary distribution.

To get a grasp on the neural network approximation of the score obtained from (1), we impose assumptions on the data distribution $\mathbb{P}(X_0 \in \cdot)$, which in the technical setting of Nickl (2024) provide a natural analogue to the three essential conditions assumed for the data distribution in the statistical analysis of unconstrained diffusion models alluded to above. More precisely, let $H_c^s(D)$ be the class of compactly supported Sobolev functions of order s on D , where we extend any function $\varphi \in H_c^s(D)$ to \bar{D} by setting $\varphi|_{\partial D} = 0$. We assume that

(H0) there exist a nonnegative function $\tilde{p}_0 \in H_c^s(D)$, for $s \in \mathbb{N} \cap (d/2, \infty)$, and $\alpha > 0$ such that $\mathbb{P}(X_0 \in \cdot)$ has a density $p_0: \bar{D} \rightarrow [\alpha, \infty)$ given by $p_0 = \tilde{p}_0 + \alpha$.

This simplifying assumption provides a strictly positive lower bound $\alpha > 0$ on the data distribution, while at the same time avoiding boundary issues associated with reflection thanks to $p_0 \in H_c^s(D)/\mathbb{R}$ and guaranteeing Hölder continuity of p_0 by the Sobolev smoothness condition $s > d/2$. These conditions are directly comparable to the assumptions made on the data density p_0 in Oko et al. (2023):

- in their unconstrained model, they allow the noising and denoising processes to leave the support $[-1, 1]^d$ and assume a lower bound $p_0|_{[-1, 1]^d} \geq \alpha > 0$. We enforce the support constraint by restricting the generative model to $\bar{D} = \text{supp } p_0$ and therefore assume $p_0 = p_0|_{\bar{D}} \geq \alpha > 0$.
- Oko et al. (2023) assumes Besov smoothness $p_0 \in B_{p,q}^s$ for $s > (1/p - 1/2) \vee 0$, with the additional restriction that p_0 is arbitrarily smooth near the boundary. Since $H^s = B_{2,2}^s$, our Sobolev assumption on p_0 is more restrictive, and our requirement that $s > d/2$ implies continuity of p_0 . For the particular case $p = q = 2$, however, Oko et al. (2023) impose no restrictions on the smoothness parameter $s > 0$. We enforce arbitrary smoothness close to the boundary in a specific way by modelling p_0 as the shift of a compactly supported function on the open domain D .

Given the preceding discussion, our focus in this paper is therefore *not* to optimise our initial data assumptions with regard to what is observed in practice for popular data sets. We rather opt for a set of conditions that shares common principles with previous studies for standard DDMs, but allows us to highlight the central mathematical ideas for score approximation in the reflected setting and to develop a fairly compact approximation theory that we regard as a versatile starting point for future statistical investigations.

Given $(\mathcal{H}0)$, the score can be written as

$$\mathfrak{s}(x, t) = \nabla \log p_t(x) = \frac{\sum_{j=0}^{\infty} e^{-t\lambda_j} \langle p_0, e_j \rangle_{L^2} \nabla e_j(x)}{\sum_{j=0}^{\infty} e^{-t\lambda_j} \langle p_0, e_j \rangle_{L^2} e_j(x)}, \quad x \in D, t > 0.$$

Our neural network approximation strategy is then broken down into the following steps:

1. truncate the series representation of $p_t(x)$ at $j = N$ to obtain an approximation $h_N(x, t)$ and corresponding truncated score $\nabla \log h_N(x, t)$, where N needs to be appropriately chosen depending on n, s and d ;
2. for an appropriately chosen discrete set of time points $\{t_i\}$, use the spatial smoothness of $h_N(x, t_i)$ induced by the Sobolev smoothness of p_0 to obtain an efficient neural network approximation of $h_N(\cdot, t_i)$, based on general approximation results from Suzuki (2019);
3. approximate the space-time function $h_N(x, t)$ by constructing a neural network approximation of the time interpolation of the neural networks from Step 2., where the interpolation degree is adapted to the parameters N, s and d .

All steps require a careful calibration of the approximation parameters and accordingly the neural network sizes, which is made possible by a precise analysis of the different levels of numerical errors induced by our stepwise approach. These techniques may be of independent interest since they are applicable to generic spectral decompositions of semigroups associated to self-adjoint Markov generators. In particular, this is potentially relevant for extending the statistical analysis to the unifying class of *denoising Markov models* that has been recently introduced in Benton et al. (2024), see also Ren et al. (2025): following our previous discussion of generalised reflected forward models with self-adjoint generator $\mathcal{A}_V = e^V \nabla \cdot f e^{-V} \nabla$ and associated spectral decomposition of the transition density $q_t(x, y)$ relative to the invariant density $\mu_V \propto e^{-V}$, for a general self-adjoint Markov generator \mathcal{A} with invariant distribution μ and eigendecomposition $(\lambda_j, e_j)_{j \in \mathbb{N}_0}$ in $L^2(\mu)$, the unknown forward marginals that are needed to express the backward generator in denoising Markov models are decomposed as

$$p_t(x) = \sum_{j=0}^{\infty} e^{-t\lambda_j} \langle p_0, e_j \rangle_{L^2(\mu)} e_j(x).$$

Thus, under appropriate assumptions on p_0 and controls on the eigenvalues and eigenfunctions, as well as with detailed knowledge of the invariant distribution μ , a neural-network approximation of the space-time function $p_t(x)$ and functionals thereof appears plausible, using the general strategy introduced in this paper.

This strategy allows us to determine an explicit calibration of the neural network class \mathcal{S} , the forward terminal time \bar{T} and the backward early stopping time $\bar{T} - \underline{T}$, such that for the empirical denoising score loss minimiser $\widehat{\mathfrak{s}}_n$ in this class of neural networks, we obtain the convergence rate

$$\mathbb{E}[\text{TV}(p_0, \widehat{p}_{\bar{T}-\underline{T}}^{\widehat{\mathfrak{s}}_n})] \lesssim n^{-\frac{s}{2s+d}} (\log n)^3 (\log \log n)^{1/2}.$$

Here, $\overleftarrow{p}_t^{\widehat{\mathfrak{s}}_n}$ denotes the density at time t of the backward generating reflected diffusion that has drift term determined by $\widehat{\mathfrak{s}}_n$ and is started in the uniform distribution on D . This establishes an expected total variation convergence rate (up to small log-factors) for denoising reflected diffusion models that matches the well-known minimax lower bound over Sobolev classes (Yang and Barron, 1999).

Organisation of the paper The paper is structured as follows. In Section 2, we provide the necessary technical background on denoising diffusion models. Here, we briefly outline the fundamentals underlying standard unconstrained DDMs, before providing a precise mathematical framework for denoising reflected diffusion models. Section 3 introduces the exact specification of the score estimator via denoising score matching and the associated generative model. We then present our main result, Theorem 1, with the remainder of the paper dedicated to its proof. The proof preparation proceeds along three sections. In Section 3.1, we present the basic error decomposition of the expected total variation risk and establish bounds on the first two sources of error, arising from early stopping and the initiation of the backwards generating process in the invariant uniform distribution on \overline{D} . The score matching error as the last component of the error decomposition is discussed in Section 3.2, before in Section 3.3, we construct the neural network approximation of the score that allows us to optimally bound the score matching error. Section 3.4 is then dedicated to proving Theorem 1, based on the results of the previous subsections. In the concluding Section 4, we discuss our findings and highlight promising directions for future research, motivated by the insights and limitations of this paper.

Notation For an open set D and $s \in \mathbb{N}$, we denote by $H^s(D)$ the Sobolev space of functions having weak partial derivatives up to order s in $L^2(D)$. We denote the inner product and corresponding norm on $L^2(D)$ by $\langle \cdot, \cdot \rangle_{L^2(D)}$ and $\|\cdot\|_{L^2(D)}$, or simply $\langle \cdot, \cdot \rangle_{L^2}$ and $\|\cdot\|_{L^2}$ if the domain D is fixed and there is no room for confusion. $|\cdot|$ denotes the Euclidean norm on \mathbb{R}^d for any $d \in \mathbb{N}$, and for a function $f: \mathcal{X} \rightarrow \mathbb{R}^d$ for some space \mathcal{X} , $\|f\|_{\mathcal{X}} = \sup_{x \in \mathcal{X}} |f(x)|$ is its supremum norm. $\mathcal{C}^s(D)$ denotes the space of functions with continuous partial derivatives of order s in \overline{D} , and we let $\mathcal{C}^\infty(D) = \bigcap_{s \in \mathbb{N}} \mathcal{C}^s(D)$. For $\beta \in (0, 1]$, $\mathcal{C}^{0,\beta}(D)$ is the space of β -Hölder continuous functions on D . For $1 \leq p, q \leq \infty$ and $\alpha > 0$, $\mathcal{B}_{p,q}^\alpha(D)$ denote the usual Besov spaces on D , cf. Triebel (2010) for details.

2. Theoretical background

Before introducing our framework for DRDMS, we first recall the basic ideas of unconstrained diffusion models. Based on observing a finite number of samples corresponding to an *unknown* distribution p_0 on \mathbb{R}^d , DDMs provide an iterative generative algorithm to create new samples that approximately match the target distribution p_0 . The general idea is to find a stochastic process that perturbs p_0 to a new distribution p_T in such a way that 1) p_T or a good approximation thereof is easy to sample from, and 2) the perturbation is reversible in the sense that we know how to simulate the time-reversed process. In an idealised setting where we have access to the exact specifications of the backward dynamics, samples from p_0 can be generated exactly by first sampling from p_T and then running the backward process. However, these dynamics must naturally adapt to the information

contained in the unknown p_0 , so the true backward dynamics must be estimated from the data.

In the framework of DDMS, this perturbation is done via an SDE, i.e., for some fixed time $T > 0$ and suitable drift $b: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and diffusion coefficient $\sigma: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, we consider the forward model

$$dX_t = b(t, X_t) dt + \sigma(t, X_t) dW_t, \quad t \in [0, T], X_0 \sim p_0,$$

where $W = (W_t)_{t \in [0, T]}$ is a standard d -dimensional Brownian motion. Under sufficient regularity conditions (Anderson, 1982; Haussmann and Pardoux, 1986), the forward model has a solution $X = (X_t)_{t \in [0, T]}$ with marginal densities $(p_t)_{t \in [0, T]}$ such that the time-reversed process $\tilde{X}_t = X_{T-t}$, $t \in [0, T]$, solves

$$d\tilde{X}_t = -\bar{b}(T-t, \tilde{X}_t) dt + \sigma(T-t, \tilde{X}_t) d\bar{W}_t, \quad t \in [0, T], \tilde{X}_0 \sim p_T, \quad (2)$$

for some Brownian motion $(\bar{W}_t)_{t \in [0, T]}$ and drift $\bar{b}: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ given by

$$\bar{b}_i(t, x) = b_i(t, x) - \frac{1}{p_t(x)} \sum_{j,k=1}^d \frac{\partial}{\partial x_j} [p_t(x) \sigma_{ik}(t, x) \sigma_{jk}(t, x)], \quad i = 1, \dots, d.$$

Thus, the time-reversed process solves a time-inhomogeneous SDE, with drift $-\bar{b}(T - \cdot, \cdot)$ and diffusion coefficient $\sigma(T - \cdot, \cdot)$. In many practical implementations as well as in the statistical studies Oko et al. (2023); Chen et al. (2023); Tang and Yang (2024); Azangulov et al. (2024), the diffusion coefficient is set to $\sigma(t, x) = \gamma(t) \mathbb{I}_d$ for some scalar function γ , which implies that the forward model is given by a (possibly time-inhomogeneous) Ornstein–Uhlenbeck process with explicit transition densities and the backward drift becomes

$$\bar{b}(t, x) = b(t, x) - \gamma^2(t) \nabla \log p_t(x),$$

where $\nabla \log p_t$ is referred to as the *score* of the forward model. Substituting this into (2), we obtain the dynamics

$$d\tilde{X}_t = (-b(T-t, \tilde{X}_t) + \gamma^2(T-t) \nabla \log p_{T-t}(\tilde{X}_t)) dt + \gamma(T-t) d\bar{W}_t \quad t \in [0, T], \tilde{X}_0 \sim p_T,$$

of the backward process. Then, when $t \rightarrow T$, the density of \tilde{X}_t approaches p_0 , so that simulating the reverse process generates new data samples corresponding to the target p_0 . While we are free to choose the coefficients of our forward process (i.e., b and σ), the score function $\nabla \log p_t$ depends on p_0 and hence needs to be estimated from the data, which is referred to as *score matching*.

Reflected generative diffusion models RGDMs follow the same generative principle, but constrain both forward and backward dynamics to a bounded domain. To avoid some technicalities, let us assume that $D \subset \mathbb{R}^d$ is an open, connected and bounded set with C^∞ boundary ∂D . We consider the reflected time-homogeneous forward model

$$dX_t = b(X_t) dt + \sigma(X_t) dW_t + \nu(X_t) d\ell_t, \quad X_0 \in \bar{D},$$

with smooth and bounded coefficients $b: \bar{D} \rightarrow \mathbb{R}^d$, $\sigma: \bar{D} \rightarrow \mathbb{R}^{d \times d}$ and conormal reflection determined by

$$\nu(x) := \frac{1}{2}a(x)n(x) = \frac{1}{2} \sum_{i=1}^d \langle \sigma_{\cdot,i}(x), n(x) \rangle \sigma_{\cdot,i}(x) > 0, \quad x \in \partial D,$$

where $a := \sigma \sigma^\top$ is assumed to be uniformly elliptic. Here, n is the inward unit normal vector at the boundary ∂D , and the process $(\ell_t)_{t \geq 0}$ is the local time at ∂D , which is a nondecreasing continuous process of bounded variation that increases only when the solution X hits the boundary, i.e., $\ell_t = \int_0^t \mathbf{1}_{\partial D}(X_s) d\ell_s$ almost surely. Because of the conormal reflection, the boundary local time can equivalently be characterised by

$$\ell_t = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \int_0^t \mathbf{1}_{(\partial D)_\varepsilon}(X_s) ds, \quad (3)$$

where $(\partial D)_\varepsilon := \{x \in \mathbb{R}^d : \text{dist}(x, \partial D) \leq \varepsilon\}$, and the limit holds both in L^2 and almost surely, uniformly on $[0, T]$, cf. Cattiaux (1988, Proposition 1.3). The boundary reflection process $L_t := \int_0^t \varphi(X_s) d\ell_s$ reflects X in a conormal direction whenever it hits the boundary ∂D , thus constraining the state space of the diffusion to the compact set \bar{D} . The process X is a time-homogeneous Markov process with transition semigroup $(Q_t)_{t \geq 0}$ determined by transition densities $(q_t)_{t \geq 0}$, given as the fundamental solutions of the PDE with conormal Neumann boundary condition,

$$\begin{cases} \frac{\partial}{\partial t} u(x, t) = \mathcal{A}u(x, t), & (x, t) \in D \times (0, T], \\ \frac{\partial}{\partial \nu} u(x, t) = 0, & (x, t) \in \partial D \times (0, T], \end{cases}$$

where \mathcal{A} is the second-order differential operator given by

$$\mathcal{A} = \sum_{i=1}^d b_i(x) \partial_{x_i} + \frac{1}{2} \sum_{i,j=1}^d a_{i,j}(x) \partial_{x_i} \partial_{x_j}.$$

We denote the density of the forward process at time t by p_t , i.e.,

$$p_t(x) dx = \mathbb{P}(X_t \in dx \mid X_0 \sim p_0) = \int_{\bar{D}} p_0(y) Q_t(y, dx) dy = \int_{\bar{D}} p_0(y) q_t(y, x) dy dx, \quad x \in \bar{D}.$$

Similarly to the unconstrained model, Cattiaux (1988, Theorem 2.5) shows that time-reversion of the reflected diffusion process yields a time-inhomogeneous reflected diffusion process, whose drift is reminiscent of the unconstrained model. More precisely, letting $\tilde{X}^T = (\tilde{X}_t)_{t \in [0, T]}$, $\tilde{X}_t = X_{T-t}$ and $\bar{\ell}^T = (\bar{\ell}_t)_{t \in [0, T]}$, $\bar{\ell}_t = \ell_T - \ell_{T-t}$, $t \in [0, T]$, there exists a Brownian motion $\bar{W}^T = (\bar{W}_t)_{t \in [0, T]}$ w.r.t. an enlargement of the natural filtration generated by \tilde{X}^T such that \tilde{X}^T solves

$$d\tilde{X}_t = -\bar{b}(\tilde{X}_t) dt + \sigma(\tilde{X}_t) d\bar{W}_t + \nu(\tilde{X}_t) d\bar{\ell}_t, \quad \tilde{X}_0 \sim p_T, \quad (4)$$

on $[0, T]$, where $\bar{b}: [0, T] \times \bar{D} \rightarrow \mathbb{R}^d$ is given by

$$\bar{b}_i(t, x) = b_i(t, x) - \frac{1}{p_t(x)} \sum_{j,k=1}^d \frac{\partial}{\partial x_j} [p_t(x) \sigma_{ik}(t, x) \sigma_{jk}(t, x)], \quad i = 1, \dots, d.$$

Note that, by definition, $\bar{\ell}^T$ is nondecreasing, has bounded variation and satisfies $\bar{\ell}_t = \int_0^t \mathbf{1}_{\partial D}(\bar{X}_s) d\bar{\ell}_s$ for $t \in [0, T]$, i.e., $\bar{\ell}^T$ is the local time at the boundary of the backward process \bar{X}^T .

A fundamental requirement for diffusion generative modeling is a precise understanding of the limiting behaviour of the forward process to evaluate the required run time of the forward process for the backward initialisation to be a sufficiently good approximation of the true terminal forward distribution p_T . To this end, we choose $b = \nabla f$ and $\sigma = \sqrt{2f} \mathbb{I}_{d \times d}$ for some potential $f: \mathbb{R}^d \rightarrow [f_{\min}, \infty) \subset (0, \infty)$. To avoid technicalities, we assume that $f \in C^\infty(\bar{D})$ and that the smooth and bounded domain D is also convex. This enables us to remain within the technical framework of Nickl (2024), which provides useful technical results required for our analysis. With our choice of coefficients, the time-homogeneous forward dynamics are described by the divergence form L^2 -generator

$$\mathcal{A} = \nabla \cdot f \nabla = \langle \nabla f, \nabla \cdot \rangle + f \Delta \quad (5)$$

with core

$$H_\nu^1(D) := \{\varphi \in H^1(D) : \partial\varphi/\partial\nu = 0 \text{ on } \partial D\}$$

for $\nu = fn$, corresponding to the constrained SDE

$$dX_t = \nabla f(X_t) dt + \sqrt{2f(X_t)} dW_t + \nu(X_t) d\ell_t. \quad (6)$$

Note that the ellipticity condition $f \geq f_{\min} > 0$ implies that the conormal Neumann boundary condition $\langle \nu(x), \nabla\varphi(x) \rangle = \frac{\partial\varphi}{\partial\nu}(x) = 0$ for $x \in \partial D$ is equivalent to a normal Neumann boundary condition, i.e., $H_\nu^1(D) = H_n^1(D)$. Thus, both the reflected forward and backward SDEs exhibit normal reflection at the boundary, and (6) induces a space-dependent scaling of local time that ensures that the occupation limit (3) holds true. Moreover, the specific choice $f \equiv 1/2$ yields a normally reflected Brownian motion.

By the divergence theorem, it follows that the invariant distribution of the forward Markov process X is the easy-to-sample-from uniform distribution on \bar{D} , i.e., $\mu = \frac{\text{Leb}|_{\bar{D}}}{\text{Leb}(\bar{D})}$. Furthermore, there exist orthonormal eigenpairs $(\lambda_j, e_j)_{j \geq 0}$ of the nonnegative operator $-\nabla \cdot f \nabla$ satisfying $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots$ and obeying the Weyl asymptotics $\lambda_j \asymp j^{2/d}$ and $e_0 = \frac{1}{\text{Leb}(\bar{D})^{1/2}} \mathbf{1}$ and $(e_j)_{j \geq 1} \subset H_\nu^1(D) \cap L_0^2(D)$ such that

$$q_t(x, y) = \sum_{j \geq 0} e^{-t\lambda_j} e_j(x) e_j(y), \quad x, y \in D.$$

See Nickl (2024, Section 3) for a detailed discussion of these properties. Because $f \in C^\infty(\bar{D})$, Nickl (2024, Corollary 1) yields the bounds

$$\begin{aligned} \|e_j\|_{H^k} &\lesssim \lambda_j^{k/2} \asymp j^{k/d}, \quad j \geq 1, \\ \|e_j\|_\infty &\lesssim j^\tau, \quad \text{for any } \tau > 1/2. \end{aligned}$$

This implies the smoothing property that, for any bounded initial density p_0 ,

$$\|p_t\|_{H^k} \lesssim \|p_0\|_\infty \sum_{j \geq 0} e^{-t\lambda_j} \|e_j\|_\infty \|e_j\|_{H^k} \lesssim \|p_0\|_\infty e^{-tj^{2/d}} j^{\tau+k/d} < \infty, \quad t > 0,$$

for arbitrary $\tau > 1/2$. By the Sobolev imbedding theorem, we therefore have $p_t \in \mathcal{C}^\infty(D)$, regardless of the smoothness properties of p_0 , and we can identify the weak derivatives of p_t with its classical derivatives. The SDE (4) governing the backward dynamics becomes

$$d\tilde{X}_t = (\nabla f(\tilde{X}_t) + 2f(\tilde{X}_t)\nabla \log p_{T-t}(\tilde{X}_t)) dt + \sqrt{2f(\tilde{X}_t)} d\bar{W}_t + \nu(\tilde{X}_t) d\bar{\ell}_t,$$

with initialisation $\tilde{X}_0 \sim p_T$. Thus, even though the diffusion coefficient is state-dependent, the particular interplay between drift and diffusion coefficient ensures that the backward drift is fully determined by the forward drift and the score $(x, t) \mapsto \nabla \log p_t(x)$, as in the unconstrained Ornstein–Uhlenbeck forward model with state-independent diffusion coefficient. By the spectral decomposition of the transition densities, the score is explicitly given by

$$\nabla \log p_t(x) = \frac{\sum_{j \geq 0} e^{-t\lambda_j} \langle p_0, e_j \rangle_{L^2} \nabla e_j(x)}{\sum_{j \geq 0} e^{-t\lambda_j} \langle p_0, e_j \rangle_{L^2} e_j(x)}, \quad x \in D, t > 0, \quad (7)$$

which will be instrumental in analysing the score approximation properties of neural networks underlying the algorithm described in the next section.

Neural network classes We will construct an estimator of the score via minimising the denoising score matching error in an appropriate class of neural networks. For doing so, we introduce parameterised classes of neural networks with ReLU (Rectified Linear Unit) activation function. In particular, for any $b, x \in \mathbb{R}^m$, let

$$\sigma_b(x) = \begin{bmatrix} \sigma(x_1 - b_1) \\ \sigma(x_2 - b_2) \\ \vdots \\ \sigma(x_m - b_m) \end{bmatrix}, \quad \sigma(y) = y \vee 0,$$

and denote for $L \in \mathbb{N}$, $W \in \mathbb{N}^{L+2}$, $S \in \mathbb{N}$ and $B > 0$ by $\Phi(L, W, S, B)$ the class of neural networks with depth (i.e., number of hidden layers) L , layer widths (including input and output layers) W , sparsity constraint S , and norm constraint B . We thus consider functions of the form

$$\varphi(x) = A_L \sigma_{b_L} A_{L-1} \sigma_{b_{L-1}} \cdots A_1 \sigma_{b_1} A_0 x,$$

where $A_i \in \mathbb{R}^{W_{i+1} \times W_i}$, $b_i \in \mathbb{R}^{W_{i+1}}$ for $i = 0, \dots, L$ (to ease notation, we always set $b_0 = 0$), and where there are at most a total of S non-zero entries of the A_i 's and b_i 's and all entries are numerically at most B . In an abuse of notation, we denote σ_0 simply by σ . This can be written succinctly as

$$\Phi(L, W, S, B) := \left\{ A_L \sigma_{b_L} A_{L-1} \sigma_{b_{L-1}} \cdots A_1 \sigma_{b_1} A_0 \mid A_i \in \mathbb{R}^{W_{i+1} \times W_i}, b_i \in \mathbb{R}^{W_{i+1}}, \right. \\ \left. \sum_{i=0}^L (\|A_i\|_0 + \|b_i\|_0) \leq S, \max_{i \in \{0, \dots, L\}} (\|A_i\|_\infty \vee \|b_i\|_\infty) \leq B \right\}.$$

3. Generative modelling with reflected diffusions

Denote the true score by $\mathfrak{s}^\circ(x, t) := \nabla \log p_t(x)$, and assume we are given data samples $(X_{0,i})_{i \in [n]} \stackrel{\text{i.i.d.}}{\sim} p_0$. For a hypothesis class \mathcal{S} of neural networks with ReLU activation func-

tion, to be exactly calibrated later, and $\mathfrak{s} \in \mathcal{S} \cup \{\mathfrak{s}^\circ\}$, we define

$$\begin{aligned} L_{\mathfrak{s}}(x) &:= \int_{\underline{T}}^{\overline{T}} \int_{\overline{D}} |\mathfrak{s}(y, t) - \nabla_y \log q_t(x, y)|^2 q_t(x, y) \, dy \, dt \\ &= \mathbb{E} \left[\int_{\underline{T}}^{\overline{T}} |\mathfrak{s}(X_t, t) - \nabla_y \log q_t(x, X_t)|^2 \mid X_0 = x \right], \end{aligned} \quad (8)$$

where \overline{T} is the terminal runtime of the reflected forward process and $\underline{T} \in (0, \overline{T})$ is such that we run the reflected backward process, which is initialised with distribution $\mathcal{U}(\overline{D})$, until $\overline{T} - \underline{T}$. We then denote the empirical score matching loss associated to \mathfrak{s} by

$$\widehat{L}_{\mathfrak{s}, n} := \frac{1}{n} \sum_{i=1}^n L_{\mathfrak{s}}(X_{0,i}),$$

and we define the empirical score minimiser by

$$\widehat{\mathfrak{s}}_n := \arg \min_{\mathfrak{s} \in \mathcal{S}} \widehat{L}_{\mathfrak{s}, n}. \quad (9)$$

We let $\overline{X}^{\mathfrak{s}}$ be a solution of the reflected SDE

$$\begin{aligned} d\overline{X}_t^{\mathfrak{s}} &= (\nabla f(\overline{X}_t^{\mathfrak{s}}) + 2f(\overline{X}_t^{\mathfrak{s}})s(\overline{X}^{\mathfrak{s}}, t)) \, dt + \sqrt{2f(\overline{X}_t^{\mathfrak{s}})} \, d\overline{W}_t + \nu(\overline{X}_t^{\mathfrak{s}}) \, d\bar{\ell}_t, \quad t \in [0, \overline{T} - \underline{T}], \\ \overline{X}_0^{\mathfrak{s}} &\sim \mathcal{U}(\overline{D}), \end{aligned} \quad (10)$$

for some Brownian motion $(\overline{W}_t)_{t \in [0, \overline{T} - \underline{T}]}$ and local time $(\bar{\ell}_t)_{t \in [0, \overline{T} - \underline{T}]}$ at the boundary ∂D , and we denote its density at time t by $\tilde{p}_t^{\mathfrak{s}}$. Here, the initialisation $X_0^{\widehat{\mathfrak{s}}_n} \sim \mathcal{U}(\overline{D})$ and the Brownian motion \overline{W} are chosen independently of the data $(X_{0,i})_{i=1, \dots, n}$. Then, $(\tilde{p}_t^{\widehat{\mathfrak{s}}_n})_{t \in [0, \overline{T}]}$ are the densities of the backward process driven by the score estimate $\widehat{\mathfrak{s}}_n$. In particular, $\tilde{p}_{\overline{T} - \underline{T}}^{\widehat{\mathfrak{s}}_n}$ is the density of the generated data obtained from stopping the backward process early at time $\overline{T} - \underline{T}$. Assessing the quality of the generated samples therefore boils down to analysing the distance between the distribution induced by p_0 and the (random) distribution induced by $\tilde{p}_{\overline{T} - \underline{T}}^{\widehat{\mathfrak{s}}_n}$.

In this paper, we use the total variation distance as the divergence measure. For two probability measures \mathbf{P}, \mathbf{Q} on \overline{D} with Lebesgue densities p, q , the total variation distance is denoted by

$$\text{TV}(p, q) \equiv \text{TV}(\mathbf{P}, \mathbf{Q}) := \sup_{A \in \mathcal{B}(\overline{D})} |\mathbf{P}(A) - \mathbf{Q}(A)|.$$

Our main result is the following.

Theorem 1 *Assume that $p_0 = \tilde{p}_0 + \alpha$ for some nonnegative $\tilde{p}_0 \in H_c^s(D)$ and $\alpha > 0$, where $s \in \mathbb{N} \cap (d/2, \infty)$. Let*

$$\underline{T} \asymp n^{-\frac{2s}{\beta(2s+d)}} \quad \text{and} \quad \overline{T} = \frac{s}{\lambda_1(2s+d)} \log n,$$

where $\beta = 1$ if $s > d/2 + 1$, $\beta = s - d/2$ if $s \in (d/2, d/2 + 1)$ and for $s = d/2 + 1$, β can be arbitrarily chosen in $(0, 1)$. Then, there exists a class of neural networks

$$\mathcal{S} = \{\mathfrak{s} \in \Phi(L(n), W(n), S(n), B(n)) : \|\mathfrak{s}(\cdot, t)\|_D \leq C(t^{-1/2} \vee 1) \forall t \in [\underline{T}, \overline{T}]\},$$

for some global constant $C > 0$, with network sizes

$$\begin{aligned} L(n) &\lesssim \log n \log \log n, \\ \|W(n)\|_\infty &\lesssim n^{\frac{d}{2s+d}} (\log n)^2, \\ S(n) &\lesssim n^{\frac{d}{2s+d}} (\log n)^3, \quad \text{and} \\ B(n) &\lesssim n^{\frac{2s}{\beta(2s+d)}}, \end{aligned}$$

such that, for $\widehat{\mathfrak{s}}_n$ given by (9), it holds for n large enough that

$$\mathbb{E}[\text{TV}(p_0, \widetilde{p}_{\underline{T}-\underline{T}}^{\widehat{\mathfrak{s}}_n})] \lesssim n^{-\frac{s}{2s+d}} (\log n)^3 (\log \log n)^{1/2}. \quad (11)$$

The proof of this theorem is prepared in the following sections. There, we will always work under the assumption on p_0 from the theorem, that is, we assume

($\mathcal{H}0$) there exist a nonnegative function $\widetilde{p}_0 \in H_c^s(D)$, for $s \in \mathbb{N} \cap (d/2, \infty)$, and $\alpha > 0$ such that $p_0 = \widetilde{p}_0 + \alpha$,

without further comment. Note that since $s > d/2$, the Sobolev imbedding theorem yields the continuous imbedding $H_c^s(D) \hookrightarrow \mathcal{C}^{0,\beta}(D)$, where β is defined depending on s and d as stated in Theorem 1. Thus, we may consider p_0 as a continuous function such that there exist some Hölder constants $\beta \in (0, 1]$, $c_\beta \in (0, \infty)$, fixed throughout the rest of the paper, such that

$$|p_0(x) - p_0(y)| \leq c_\beta |x - y|^\beta, \quad x, y \in D. \quad (12)$$

3.1 Error decomposition

Let \overline{p}_t^s be the density at time t of the solution to the SDE (10) with initial condition $\overline{X}_0^s \sim p_{\overline{T}}$ (independent of the data $(X_{0,i})_{i=1,\dots,n}$) replacing $\overline{X}_0^s \sim \mathcal{U}(\overline{D})$. In particular, we then notice that $\overline{p}_t^{\circ} = p_{\overline{T}-t}$ and therefore $\overline{p}_{\overline{T}-\underline{T}}^{\circ} = p_{\underline{T}}$. As for unconstrained diffusion models, cf. Oko et al. (2023), the triangle inequality for the total variation distance then implies the error decomposition bound

$$\begin{aligned} \mathbb{E}[\text{TV}(p_0, \widetilde{p}_{\underline{T}-\underline{T}}^{\widehat{\mathfrak{s}}_n})] &\leq \text{TV}(p_0, p_{\underline{T}}) + \text{TV}(\mathbb{P}(X_{\overline{T}} \in \cdot \mid X_0 \sim p_0), \mathcal{U}(\overline{D})) \\ &\quad + \mathbb{E}[\text{TV}(\overline{p}_{\overline{T}-\underline{T}}^{\circ}, \widetilde{p}_{\overline{T}-\underline{T}}^{\widehat{\mathfrak{s}}_n})]. \end{aligned} \quad (13)$$

The generalisation error therefore splits into three separate error contributions:

1. the first term represents the error induced by stopping early the backward process initialised by the true forward terminal density $p_{\overline{T}}$ at time $\overline{T} - \underline{T}$;
2. the second term is the error associated to starting the backward process in its stationary distribution instead of $p_{\overline{T}}$;

3. the third term quantifies the error coming from running the backward process with the drift determined by the estimated score $\widehat{\mathfrak{s}}_n$ instead of the true score \mathfrak{s}° .

We start with controlling the first two terms before treating the most challenging score approximation error. The early stopping contribution to the error decomposition is controlled via small time heat kernel bounds for the transition densities in the following lemma.

Lemma 2 *There exists a constant C depending only on f, d, D, β and c_β such that*

$$\text{TV}(p_0, p_{\underline{T}}) \leq C\underline{T}^{\beta/2}, \quad \underline{T} \leq 1.$$

Proof Fix $0 < t \leq 1$. We need to show that $\frac{1}{2}\|p_t - p_0\|_{L^1} \leq Ct^{\beta/2}$. To do so, note that the reversibility of X implies the symmetry of its transition densities. Hence,

$$p_t(x) = \int p_0(y)q_t(y, x) dy = \int p_0(y)q_t(x, y) dy = \mathbb{E}[p_0(X_t) | X_0 = x], \quad x \in D.$$

Using the Hölder continuity stated in (12), we therefore obtain

$$|p_t(x) - p_0(x)| = |\mathbb{E}[p_0(X_t) - p_0(X_0) | X_0 = x]| \leq c_\beta \mathbb{E}[|X_t - X_0|^\beta | X_0 = x], \quad x \in D.$$

Furthermore, by Davies (1989, Corollary 3.2.9), we have the small time Gaussian heat kernel bound

$$q_t(x, y) \leq C_0 \frac{1}{t^{d/2}} \exp\left(-C_1 \frac{|x-y|^2}{t}\right), \quad x, y \in D,$$

where $C_0, C_1 \geq 0$ are constants depending only on f, d and D . Thus, for $x \in D$,

$$\mathbb{E}[|X_t - X_0|^\beta | X_0 = x] = \int_D q_t(x, y)|x-y|^\beta dy \leq C_0 \frac{1}{t^{d/2}} \int_D |x-y|^\beta \exp\left(-C_1 \frac{|x-y|^2}{t}\right) dy.$$

Next, note that

$$\begin{aligned} \int_D |x-y|^\beta \exp\left(-C_1 \frac{|x-y|^2}{t}\right) dy &\leq \int_{\mathbb{R}^d} |y|^\beta \exp\left(-C_1 \frac{|y|^2}{t}\right) dy \\ &= \kappa_d \int_0^\infty r^{\beta+d-1} \exp\left(-C_1 \frac{r^2}{t}\right) dr \\ &= \frac{\kappa_d}{2} t^{\frac{\beta+d}{2}} C_1^{-\frac{\beta+d}{2}} \int_0^\infty u^{\frac{\beta+d}{2}-1} e^{-u} du \\ &= C_2 t^{\frac{\beta+d}{2}}, \end{aligned}$$

where κ_d is the surface area of S^{d-1} and $C_2 = \frac{\kappa_d}{2} C_1^{-\frac{\beta+d}{2}} \Gamma(\frac{\beta+d}{2})$. Finally, since D is bounded, the result follows by setting $C = c_\beta C_0 C_2 \text{Leb}(D)/2$. \blacksquare

The error contribution from starting the backward process uniformly on D is controlled in terms of the spectral gap λ_1 of \mathcal{A} defined in (5), which can be lower bounded by $\lambda_1 \geq f_{\min}/C_P(D)$, where $C_P(D)$ is the Poincaré constant of the domain D .

Lemma 3 *It holds that*

$$\text{TV}(\mathbb{P}(X_{\bar{T}} \in \cdot \mid X_0 \sim p_0), \mathcal{U}(\bar{D})) \leq \frac{\sqrt{\text{Leb}(\bar{D})}}{2} \|p_0\|_{L^2} e^{-\lambda_1 \bar{T}}, \quad \bar{T} > 0.$$

Proof Let $t > 0$. The Cauchy–Schwarz inequality implies that

$$\begin{aligned} \text{TV}(\mathbb{P}(X_t \in \cdot \mid X_0 \sim p_0), \mathcal{U}(\bar{D})) &= \frac{1}{2} \int_{\bar{D}} \left| p_t(x) - \frac{1}{\text{Leb}(\bar{D})} \right| dx \\ &\leq \frac{\sqrt{\text{Leb}(\bar{D})}}{2} \left(\int_{\bar{D}} \left| p_t(x) - \frac{1}{\text{Leb}(\bar{D})} \right|^2 dx \right)^{1/2}. \end{aligned}$$

By the spectral decomposition of $q_t(y, x)$, it follows for $a_k := \langle p_0, e_k \rangle_{L^2}$ that

$$\begin{aligned} \int_{\bar{D}} \left| p_t(x) - \frac{1}{\text{Leb}(\bar{D})} \right|^2 dx &= \int_{\bar{D}} \left| \sum_{k \geq 1} e^{-\lambda_k t} e_k(x) a_k \right|^2 dx \\ &= \sum_{k \geq 1} a_k^2 e^{-2\lambda_k t} \\ &\leq \|p_0\|_{L^2}^2 e^{-2\lambda_1 t}, \end{aligned}$$

where we used $\lambda_0 = 0$, $e_0 \equiv \text{Leb}(\bar{D})^{-1/2}$ and $\text{Leb}(\bar{D})^{1/2} a_0 = \langle p_0, 1 \rangle_{L^2} = 1$ for the first line and orthonormality of $(e_k)_{k \geq 0}$ for the second one. This yields the claim. \blacksquare

We now move on to the treatment of score approximation error in the next section.

3.2 Score matching error

In this section, \mathcal{S} denotes a generic neural network class that shall be exactly calibrated at the end of the section for our score approximation purposes. Recall that the score estimator $\widehat{\mathfrak{s}} = \widehat{\mathfrak{s}}_n$ is defined according to (9).

By Girsanov’s theorem, cf. Theorem 14, and Pinsker’s inequality, the third term in the decomposition (13) is controlled by

$$\begin{aligned} &\left(\frac{1}{2} \mathbb{E} \left[\int_{\underline{T}}^{\bar{T}} \int_{\bar{D}} f(x) |\widehat{\mathfrak{s}}(x, t) - \nabla \log p_t(x)|^2 p_t(x) dx dt \right] \right)^{1/2} \\ &\asymp \left(\mathbb{E} \left[\int_{\underline{T}}^{\bar{T}} \int_{\bar{D}} |\widehat{\mathfrak{s}}(x, t) - \nabla \log p_t(x)|^2 p_t(x) dx dt \right] \right)^{1/2}, \end{aligned} \tag{14}$$

where we used that $0 < f_{\min} \leq f(x) \leq \|f\|_{\bar{D}} < \infty$ for all $x \in \bar{D}$. The key to bounding this term is the equivalence between explicit and denoising score matching, i.e.,

$$\begin{aligned} &\int_{\underline{T}}^{\bar{T}} \int_{\bar{D}} |\mathfrak{s}(y, t) - \nabla \log p_t(y)|^2 p_t(y) dy dt \\ &= \int_{\underline{T}}^{\bar{T}} \int_{\bar{D}^2} |\mathfrak{s}(y, t) - \nabla_y \log q_t(x, y)|^2 q_t(x, y) p_0(x) dx dy dt + C \\ &= \mathbb{E}[L_{\mathfrak{s}}(X_0)] + C, \end{aligned} \tag{15}$$

where

$$C = \int_{\underline{T}}^{\overline{T}} \int_D |\nabla \log p_t(y)|^2 p_t(y) dy dt - \int_{\underline{T}}^{\overline{T}} \int_{D^2} |\nabla_y \log q_t(x, y)|^2 q_t(x, y) p_0(x) dx dy dt \leq 0$$

is a constant that is independent of \mathfrak{s} . Note that (15) is valid in our reflected diffusion model by the same arguments as in Vincent (2011), see also the proof of Lemma C.3 in Oko et al. (2023).

Using (15), the generalisation loss (14) can be bounded in terms of the minimal score approximation error over the class \mathcal{S} and the complexity of the induced function class $\mathcal{L} := \{L_{\mathfrak{s}} : \mathfrak{s} \in \mathcal{S}\}$ for a desired precision level δ .

Theorem 4 *Suppose that $\sup_{\mathfrak{s} \in \mathcal{S} \cup \{\mathfrak{s}^\circ\}} \|L_{\mathfrak{s}}\|_D \leq C(\mathcal{L}) < \infty$. Then, for any $\delta > 0$ such that $\mathcal{N}(\mathcal{L}, \|\cdot\|_D, \delta) \geq 3$, it holds that*

$$\begin{aligned} & \mathbb{E} \left[\int_{\underline{T}}^{\overline{T}} \int_D |\widehat{\mathfrak{s}}(x, t) - \nabla \log p_t(x)|^2 p_t(x) dx dt \right] \\ & \leq 2 \inf_{\mathfrak{s} \in \mathcal{S}} \int_{\underline{T}}^{\overline{T}} \int_D |\mathfrak{s}(x, t) - \nabla \log p_t(x)|^2 p_t(x) dx dt + 2 \frac{C(\mathcal{L})}{n} \left(\frac{145}{9} \log \mathcal{N}(\mathcal{L}, \|\cdot\|_D, \delta) + 160 \right) \\ & \quad + 5\delta. \end{aligned}$$

The proof of this theorem is in principle the same as that of Theorem C.4 in Oko et al. (2023), but let us emphasize that larger numeric constants appear in our statement. Apart from a few simple typos in Oko et al. (2023), the main reason for this is a small gap in the proof of Oko et al. (2023, Theorem C.4), which has recently been pointed out in Yakovlev and Puchkin (2025), and which requires fixing. More precisely, Oko et al. (2023) claim that the excess loss satisfies

$$\mathbb{E}[(L_{\mathfrak{s}}(X_0) - L_{\mathfrak{s}^\circ}(X_0))^2] \leq C(\mathcal{L}) \mathbb{E}[L_{\mathfrak{s}}(X_0) - L_{\mathfrak{s}^\circ}(X_0)],$$

which is unjustified because $L_{\mathfrak{s}}(x) - L_{\mathfrak{s}^\circ}(x)$ is not necessarily non-negative pointwise. This is not dramatic however, since we can show that instead

$$\mathbb{E}[(L_{\mathfrak{s}}(X_0) - L_{\mathfrak{s}^\circ}(X_0))^2] \leq 4C(\mathcal{L}) \mathbb{E}[L_{\mathfrak{s}}(X_0) - L_{\mathfrak{s}^\circ}(X_0)], \quad (16)$$

holds, i.e., the bound from Oko et al. (2023) is true up to a universal multiplicative constant that does not matter in an essential way for the remainder of the proof. Using the terminology of Bartlett and Mendelson (2006), (16) shows that the denoising score matching excess loss satisfies a *Bernstein condition*, which for a variety of problems in the empirical risk minimisation literature has been identified as a crucial ingredient to obtain minimax optimal convergence rates for empirical risk minimisers. For reasons of reproducibility in other modelling contexts, we prove (16) in Appendix B in a general Markovian framework.

To deal with the stochastic error in the upper bound, it is essential to control both the uniform loss upper bound $C(\mathcal{L})$ and the covering number $\mathcal{N}(\mathcal{L}, \|\cdot\|_D, \delta)$. The size of the networks from the generic class \mathcal{S} which is required to have a sufficient bound on the approximation error

$$\inf_{\mathfrak{s} \in \mathcal{S}} \int_{\underline{T}}^{\overline{T}} \int_D |\mathfrak{s}(x, t) - \nabla \log p_t(x)|^2 p_t(x) dx dt$$

translates directly into covering number bounds on $\mathcal{N}(\mathcal{L}, \|\cdot\|_D, \delta)$, as the following result shows.

Lemma 5 *If for any $t > 0$, $\sup_{\mathfrak{s} \in \mathcal{S}} \|\mathfrak{s}(\cdot, t)\|_D \leq C(\mathcal{S})(t^{-1/2} \vee 1)$, for some finite constant $C(\mathcal{S})$, then there exists a constant c , depending only on d, f and D , such that, for any $\delta > 0$,*

$$\mathcal{N}(\mathcal{L}, \|\cdot\|_D, \delta) \leq \mathcal{N}(\mathcal{S}, \|\cdot\|_{D \times [\underline{T}, \bar{T}]}, \frac{\delta}{cC(\mathcal{S})\bar{T}}).$$

Proof For $\tilde{\delta} > 0$, let $\mathfrak{s}_1, \dots, \mathfrak{s}_N: \mathbb{R}^d \rightarrow \mathbb{R}$ be a $\tilde{\delta}$ -net for \mathcal{S} w.r.t. $\|\cdot\|_{D \times [\underline{T}, \bar{T}]}$. Denote $h_{\mathfrak{s}}(x; y, t) := \mathfrak{s}(y, t) - \nabla_y \log q_t(x, y)$ for $\mathfrak{s} \in \mathcal{S}$. Let $\mathfrak{s} \in \mathcal{S}$, and choose $\mathfrak{s}' \in \{\mathfrak{s}_1, \dots, \mathfrak{s}_N\}$ such that $\|\mathfrak{s} - \mathfrak{s}'\|_{D \times [\underline{T}, \bar{T}]} \leq \tilde{\delta}$. Using definition (8) and the inequality $\|h_{\mathfrak{s}} - h_{\mathfrak{s}'}\|(x; y, t) \leq |\mathfrak{s} - \mathfrak{s}'|(y, t)$, we obtain

$$\begin{aligned} |L_{\mathfrak{s}} - L_{\mathfrak{s}'}|(x) &\leq \int_{\underline{T}}^{\bar{T}} \int_D \left| |h_{\mathfrak{s}}|^2 - |h_{\mathfrak{s}'}|^2 \right|(x; y, t) q_t(x, y) \, dy \, dt \\ &= \int_{\underline{T}}^{\bar{T}} \int_D \left| |h_{\mathfrak{s}}| - |h_{\mathfrak{s}'}| \right|(x; y, t) (|h_{\mathfrak{s}}| + |h_{\mathfrak{s}'}|)(x; y, t) q_t(x, y) \, dy \, dt \\ &\lesssim \|\mathfrak{s} - \mathfrak{s}'\|_{D \times [\underline{T}, \bar{T}]} \int_{\underline{T}}^{\bar{T}} \int_D (|h_{\mathfrak{s}}| + |h_{\mathfrak{s}'}|)(x; y, t) q_t(x, y) \, dy \, dt \tag{17} \\ &\lesssim \tilde{\delta} \left(\int_{\underline{T}}^{\bar{T}} \sup_{\mathfrak{s} \in \mathcal{S}} \sup_{z \in D} |\mathfrak{s}(z, t)| \, dt + \int_{\underline{T}}^{\bar{T}} \int_D |\nabla_y \log q_t(x, y)| q_t(x, y) \, dy \, dt \right) \\ &\lesssim \tilde{\delta} \left(C(\mathcal{S})\bar{T} + \int_{\underline{T}}^{\bar{T}} \int_D |\nabla_y \log q_t(x, y)| q_t(x, y) \, dy \, dt \right). \end{aligned}$$

Using Ouhabaz (2005, Theorem 6.19) and symmetry of q_t , we obtain for some constants $C, \gamma > 0$ only depending on d, f and D ,

$$\begin{aligned} \int_D |\nabla_y \log q_t(x, y)| q_t(x, y) \, dy &= \int_D |\nabla_y q_t(x, y)| \, dy \\ &= \int_D |\nabla_y q_t(y, x)| \, dy \\ &\leq Ct^{-1/2} e^{\gamma t}. \end{aligned}$$

This shows that

$$\int_{\underline{T}}^{\bar{T}} \int_D |\nabla_y \log q_t(x, y)| q_t(x, y) \, dy \, dt \lesssim 1. \tag{18}$$

Furthermore, as in the proof of Nickl (2024, Proposition 3), we have

$$\sup_{(x, y) \in D^2} |\nabla_y q_t(x, y)| \lesssim \sum_{j \geq 1} j^{\tau+1/d} e^{-ctj^{2/d}}, \quad t > 0,$$

for some constants $c > 0, \tau > 1/2$, showing also that

$$\int_1^{\bar{T}} \int_D |\nabla \log q_t(x, y)| q_t(x, y) \, dy \, dt \lesssim 1. \tag{19}$$

Plugging (18) and (19) into (17), it follows that $\{L_{\mathfrak{s}_i} : i \in [N]\}$ is an $\widetilde{\delta}cC(\mathcal{S})\overline{T}$ -covering of \mathcal{L} w.r.t. $\|\cdot\|_D$, where c is some constant depending only on f, d and D . This implies the claimed result. \blacksquare

For the uniform loss upper bound $C(\mathcal{L})$ we again need to deal with the challenge of not having access to a simple analytic expression for the transition densities $q_t(x, y)$. While upper and lower heat kernel bounds are available for q_t and its gradient under specific assumptions on the domain, these are not sufficient to yield appropriate pointwise estimates for the log-gradient $\nabla_y q_t(x, y)$. Instead, we exploit that for our purposes it suffices to have (time)-integrated bounds. The basic idea is best illustrated for the particular case of a constant diffusivity $f \equiv 1$. Then, the generator of the forward process is just the Neumann Laplacian Δ on \overline{D} which implies that for any $t > 0$ and $x, y \in D$

$$\Delta_y q_t(x, y) = \partial_t q_t(x, y),$$

because $q_t(x, y)$ is a symmetric fundamental solution to the Neumann heat equation. Furthermore,

$$\Delta_y \log q_t(x, y) = \frac{\Delta_y q_t(x, y)}{q_t(x, y)} - |\nabla_y \log q_t(x, y)|^2,$$

which together with the above establishes the fundamental relation

$$|\nabla_y \log q_t(x, y)|^2 - \partial_t \log q_t(x, y) = -\Delta_y \log q_t(x, y),$$

between spatial and temporal log-gradients. Based on this observation, the famous Li–Yau estimate (Li and Yau, 1986) establishes that for $f \equiv 1$ it holds that

$$|\nabla_y \log q_t(x, y)|^2 - \partial_t \log q_t(x, y) \leq \frac{d}{2t}.$$

This result could be directly used for $f \equiv 1$ to establish the bound in the following lemma, but if f is not constant the situation becomes a bit more tricky. While we may always interpret $\Delta_f = \nabla \cdot f \nabla$ as a weighted Neumann Laplacian on the manifold \overline{D} equipped with a Riemannian metric induced by the Riemannian tensor associated to f , corresponding Li–Yau type estimates from the literature (Qian, 1995; Bakry and Qian, 1999) require the validation of certain curvature conditions on $\nabla \cdot f \nabla$, which may be hard to check for specific choices of f and D . To circumvent this problem, we follow a more elementary approach, which is however still based on the type of reasoning outlined above.

Lemma 6 *Assume that $\underline{T} \leq 1$ and $\overline{T} \geq 1$. If for any $t > 0$, $\sup_{\mathfrak{s} \in \mathcal{S}} \|\mathfrak{s}(\cdot, t)\|_D \leq C(\mathcal{S})(t^{-1/2} \vee 1)$, then*

$$\sup_{\mathfrak{s} \in \mathcal{S} \cup \{\mathfrak{s}^\circ\}} \|L_{\mathfrak{s}}\|_D \lesssim (C(\mathcal{S})^2 \vee 1)(|\log \underline{T}| + \overline{T}).$$

Proof Let $\mathfrak{s} \in \mathcal{S} \cup \{\mathfrak{s}^\circ\}$ and $x \in D$ be arbitrarily chosen. By the assumption on the uniform temporal growth of \mathfrak{s} , and using that $\int_D q_t(x, y) dy = 1$

$$L_{\mathfrak{s}}(x) \leq 2 \left(\int_{\underline{T}}^{\overline{T}} \int_D (|\mathfrak{s}(y, t)|^2 + |\nabla_y \log q_t(x, y)|^2) q_t(x, y) dy \right)$$

$$\begin{aligned}
 &\lesssim C(\mathcal{S})^2 \int_{\underline{T}}^{\bar{T}} t^{-1} \vee 1 \, dt + \int_{\underline{T}}^{\bar{T}} \int_D |\nabla_y \log q_t(x, y)|^2 q_t(x, y) \, dy \\
 &\leq C(\mathcal{S})^2 (|\log \underline{T}| + \bar{T}) + \int_{\underline{T}}^{\bar{T}} \int_D |\nabla_y \log q_t(x, y)|^2 q_t(x, y) \, dy.
 \end{aligned}$$

To bound the remaining integral involving the log-gradient of the transition density, we first observe that $\nabla_y \cdot f \nabla_y \log q_t(x, y)$ satisfies

$$\begin{aligned}
 \nabla_y \cdot f \nabla_y \log q_t(x, y) &= \nabla_y \cdot f(y) \frac{\nabla_y q_t(x, y)}{q_t(x, y)} \\
 &= \frac{\nabla_y \cdot f \nabla_y q_t(x, y)}{q_t(x, y)} - f(y) \frac{|\nabla_y q_t(x, y)|^2}{q_t(x, y)^2}.
 \end{aligned}$$

Since $q_t(x, y) = q_t(y, x)$ is a fundamental solution to the elliptic PDE $(\nabla \cdot f \nabla - \partial_t)u(y, t) = 0$ on D with Neumann boundary conditions, we can write

$$\begin{aligned}
 \int_D \int_{\underline{T}}^{\bar{T}} \frac{\nabla_y \cdot f \nabla_y q_t(x, y)}{q_t(x, y)} q_t(x, y) \, dt \, dy &= - \int_D \int_{\underline{T}}^{\bar{T}} \underbrace{\frac{\partial_t q_t(x, y)}{q_t(x, y)}}_{=\partial_t \log q_t(x, y)} q_t(x, y) \, dy \, dt \\
 &= - \int_D \int_{\underline{T}}^{\bar{T}} \partial_t q_t(x, y) \, dt \, dy = 0.
 \end{aligned}$$

Combining these two observations with $0 < f_{\min} \leq f(y) \leq \|f\|_{\bar{D}} < \infty$ for all $y \in \bar{D}$, and denoting the generator by $\Delta_f = \nabla \cdot f \nabla$, we obtain

$$\begin{aligned}
 \int_{\underline{T}}^{\bar{T}} \int_D \frac{|\nabla_y q_t(x, y)|^2}{q_t(x, y)^2} q_t(x, y) \, dy \, dt &\asymp - \int_{\underline{T}}^{\bar{T}} \int_D (\nabla_y \cdot f(y) \nabla_y \log q_t(x, y)) q_t(x, y) \, dy \, dt \\
 &= - \int_{\underline{T}}^{\bar{T}} \int_D ((\Delta_f + \partial_t) \log q_t(x, y)) q_t(x, y) \, dy \, dt \\
 &= -\mathbb{E}^x \left[\int_{\underline{T}}^{\bar{T}} (\Delta_f + \partial_t) \log q_t(x, X_t) \, dt \right] \\
 &= \mathbb{E}^x [\log q_{\underline{T}}(x, X_{\underline{T}})] - \mathbb{E}^x [\log q_{\bar{T}}(x, X_{\bar{T}})].
 \end{aligned}$$

Noting that $q_t(x, \cdot)$ satisfies the Neumann boundary condition at ∂D , the last equality is a consequence of Itô's formula, by which

$$\begin{aligned}
 &\log q_{\bar{T}}(x, X_{\bar{T}}) - \log q_{\underline{T}}(x, X_{\underline{T}}) \\
 &= \int_{\underline{T}}^{\bar{T}} (\partial_t + \Delta_f) \log q_t(x, X_t) \, dt + \int_{\underline{T}}^{\bar{T}} \sqrt{2f(X_t)} \langle \nabla \log q_t(x, X_t), dW_t \rangle \\
 &\quad + \int_{\underline{T}}^{\bar{T}} \underbrace{\langle \nabla_y \log q_t(x, X_t), \nu(X_t) \rangle}_{=0} \, d\ell_t,
 \end{aligned}$$

where the expectation of the stochastic integral is zero because $(y, t) \mapsto |\sqrt{f(y)}\nabla_y \log q_t(x, y)|$ is bounded on $\bar{D} \times [\underline{T}, \bar{T}]$, which follows from the lower bound

$$\inf_{t \geq t_0, x, y \in \bar{D}} q_t(x, y) > 0, \tag{20}$$

for any $t_0 > 0$, see, e.g., Itô (1992, p.166), and the space-time smoothness of $(t, x, y) \mapsto q_t(x, y)$ on the compact set $\bar{D} \times [\underline{T}, \bar{T}]$. We finish the proof by noting that for some constant C , the upper heat kernel bound

$$q_t(x, y) \lesssim (t^{-d/2} \vee 1) \exp\left(-C \frac{|x-y|^2}{t}\right), \quad t > 0, x, y \in \bar{D},$$

from Davies (1989, Theorem 3.29) yields

$$\sup_{x, y \in \bar{D}} \log q_{\underline{T}}(x, y) \lesssim \frac{d}{2} \log \underline{T}^{-1},$$

and, moreover, (20) implies $\inf_{x, y \in \bar{D}} q_{\bar{T}}(x, y) \geq c$, for some constant c not depending on \underline{T}, \bar{T} since $\bar{T} \geq 1$. Putting these bounds together, we conclude that

$$\mathbb{E}^x[\log q_{\underline{T}}(x, X_{\underline{T}})] - \mathbb{E}^x[\log q_{\bar{T}}(x, X_{\bar{T}})] \lesssim 1 + \log \underline{T}^{-1},$$

and therefore also

$$\int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} \frac{|\nabla_y q_t(x, y)|^2}{q_t(x, y)^2} q_t(x, y) dy dt \lesssim 1 + \log \underline{T}^{-1}.$$

■

3.3 Bounding the approximation error

Before going into details, we first outline here our general strategy for bounding the approximation error:

1. approximate the true score by truncation of the spectral decomposition h_N , i.e., approximate $\nabla_x \log p_t(x)$ by $\nabla_x \log h_N(x, t)$;
2. approximate h_N and $\nabla_x h_N$ by neural networks on $[\underline{T}, \bar{T}]$ via
 - (a) dividing $[\underline{T}, \bar{T}]$ into sub-intervals of increasing length, totalling a number of intervals on the order of $\log N$;
 - (b) on each sub-interval, fix a number of time-points $\{t_i\}$, also on the order of $\log N$, and at each of these, make an approximation of $h_N(t_i)$ and $\nabla_x h_N(t_i)$ using existing results;
 - (c) extend these discrete approximations to the entire sub-interval using polynomial interpolation;

- (d) combine approximations on each sub-interval into one final approximation via a partition of unity;
- 3. combine the first two steps, along with general results on neural networks, to achieve a neural network approximation of $\nabla_x \log p_t$.

At this point, we comment on how and why our approximation strategy differs from those in, for example, Oko et al. (2023). The authors there assume a Gaussian transition kernel with a density that is comparatively easy to approximate using neural networks. The difficulty then lies in approximating the initial density p_0 and the convolution of the two. By contrast, in our spectral composition, the influence of p_0 enters the function only as the weights $\langle p_0, e_j \rangle_{L^2}$. The growth of these weights encodes the smoothness of p_0 and thereby influences the truncation parameter N but they need no approximation by neural networks. Thus, the difficulty lies instead in approximating the eigenfunctions e_j themselves. One might think that one could use existing results, e.g. from Suzuki (2019); Schmidt-Hieber (2020), to approximate each e_j and the individual time components $t \mapsto e^{-t\lambda_j}$ separately, and then sum them together. However, using these existing results, the sparsity constraint of each summand would be of order at least N , and thus the sparsity constraint of the network as a whole would be at least N^2 . This is problematic, since, according to Oko et al. (2023, Lemma C.2), the sparsity constraint enters exponentially in the covering number of the associated class and, consequently, by Theorem 4, linearly in the generalisation error. This is why we employ a different strategy: we approximate the entire sum at fixed time points and use polynomials to interpolate between them over time. This ultimately gives us a sparsity constraint of order $N \text{Poly}(\log N)$. Following this general strategy, we can prove the following score approximation result.

Theorem 7 *Let $0 < \underline{T} < \bar{T}$ and $n \in \mathbb{N}$ sufficiently large be given with $\underline{T} \in \text{Poly}(n^{-1})$. Then, there exists a neural network $\varphi_{\mathfrak{s}} \in \Phi(L(n), W(n), S(n), B(n))$ satisfying*

$$\int_{\underline{T}}^{\bar{T}} \int_D |\varphi_{\mathfrak{s}}(x, t) - \nabla_x \log p_t(x)|^2 p_t(x) dx dt \lesssim n^{-\frac{2s}{2s+d}} (\log n)^2 (\bar{T} + \log(\underline{T}^{-1})). \quad (21)$$

The size of the network is evaluated as

$$\begin{aligned} L(n) &\lesssim \log n \log \log n, \\ \|W(n)\|_{\infty} &\lesssim M n^{\frac{d}{2s+d}} \log n, \\ S(n) &\lesssim M n^{\frac{d}{2s+d}} (\log n)^2, \quad \text{and} \\ B(n) &\lesssim \frac{\sqrt{n}}{\log n} \vee \frac{1}{\underline{T}}, \end{aligned}$$

where $M \in O(|\log \frac{\bar{T}}{\underline{T}}|)$. Furthermore, the network can be chosen such that there exists a constant $C < \infty$ depending only on p_0 and D such that $|\varphi_{\mathfrak{s}}(x, t)| \leq \frac{C}{\sqrt{t}}$ for all $t \in [\underline{T}, \bar{T}]$ and $x \in D$.

As alluded to above, the idea is to break up the score approximation error by using the spectral score representation (7) and to reduce this to the problem of approximating

$\nabla \log h_N(x, t) = \nabla h_N(x, t)/h_N(x, t)$, where, for $N \in \mathbb{N}$ and $t \in [0, T]$, the truncated series $h_N(t) = h_N(\cdot, t)$ is given by

$$h_N(t) := \sum_{j=0}^N e^{-t\lambda_j} \langle p_0, e_j \rangle_{L^2} e_j. \quad (22)$$

It will then be crucial to choose the cutoff value N of the right order in terms of n to balance the tradeoff between the error incurred through the truncation procedure and the increased approximation quality of $\nabla \log h_N$ in terms of neural networks for smaller N . Our analysis will demonstrate that for the desired approximation accuracy $n^{-s/(2s+d)}$, the choice $N \asymp n^{d/(2s+d)}$ is appropriate.

Step 1: Bounding the truncation loss We start by considering the approximation properties of h_N and $\nabla_x h_N$. To this end, let us introduce the homogeneous Sobolev space of order s that is induced by the eigendecomposition of $-\nabla \cdot f \nabla$ via

$$\bar{H}^s(D) := \left\{ \phi \in L_0^2(D) : \|\phi\|_{\bar{H}^s}^2 := \sum_{j \geq 1} \lambda_j^s \langle \phi, e_j \rangle_{L^2}^2 < \infty \right\}.$$

Lemma 8 *It holds that*

$$\int_{\underline{T}}^{\bar{T}} \|p_t - h_N(t)\|_{H^1}^2 dt \lesssim \|p_0 - \frac{1}{\text{Leb}(D)}\|_{\bar{H}^s}^2 N^{-2s/d}.$$

Proof Arguing as in Nickl (2024, Proposition 3), we see that $p_t, h_N(t) \in H^k(D)$ for any $k \in \mathbb{N}$. Since $p_t - h_N(t) \in L_0^2$, Nickl (2024, Proposition 2) shows that $p_t - h_N(t) \in \bar{H}^1(D)$ and $\|p_t - h_N(t)\|_{\bar{H}^1} \asymp \|p_t - h_N(t)\|_{H^1}$. Thus,

$$\begin{aligned} \int_{\underline{T}}^{\bar{T}} \|p_t - h_N(t)\|_{H^1}^2 dt &\asymp \int_{\underline{T}}^{\bar{T}} \|p_t - h_N(t)\|_{\bar{H}^1}^2 dt \\ &= \sum_{j \geq N+1} \int_{\underline{T}}^{\bar{T}} \lambda_j e^{-2\lambda_j t} dt \langle p_0, e_j \rangle_{L^2}^2 \\ &\leq \sum_{j \geq N+1} \langle p_0, e_j \rangle_{L^2}^2. \end{aligned}$$

As $p_0 \in H_c^s(D)/\mathbb{R}$, it holds $p_0 - \frac{1}{\text{Leb}(D)} \in H_c^s(D)/\mathbb{R} \cap L_0^2(D)$. Furthermore, by Nickl (2024, Proposition 2), we have $H_c^s(D)/\mathbb{R} \cap L_0^2 \subset \bar{H}^s(D)$ and $\|\phi\|_{H^s} \asymp \|\phi\|_{\bar{H}^s}$ for $\phi \in \bar{H}^s(D)$, implying that $\|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s} \asymp \|p_0 - \frac{1}{\text{Leb}(D)}\|_{\bar{H}^s}$. Using this and $\lambda_j \asymp j^{2/d}$, it follows

$$\begin{aligned} \sum_{j \geq N+1} \langle p_0, e_j \rangle_{L^2}^2 &\leq \frac{\sum_{j=1}^{\infty} \langle p_0, e_j \rangle_{L^2}^2 j^{2s/d}}{(N+1)^{2s/d}} = \frac{\sum_{j=1}^{\infty} \langle p_0 - \frac{1}{\text{Leb}(D)}, e_j \rangle_{L^2}^2 j^{2s/d}}{(N+1)^{2s/d}} \\ &\lesssim \|p_0 - \frac{1}{\text{Leb}(D)}\|_{\bar{H}^s}^2 N^{-2s/d}. \end{aligned} \quad (23)$$

■

This result allows us to study the approximation quality of (a truncated version of) $\nabla \log h_N$.

Proposition 9 *It holds that*

$$\int_{\underline{T}}^{\bar{T}} \int_D \left| \nabla_x \log p_t(x) - \frac{\nabla_x h_N(x, t)}{h_N(x, t) \vee \alpha} \right|^2 p_t(x) dx dt \lesssim C(p_0)(1 + \lambda_1^{-1})\alpha^{-4} \log(\underline{T}^{-1})N^{-2s/d},$$

where

$$C(p_0) = \|p_0\|_\infty \|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s}^2 (1 + \|p_0\|_{H^s}^2).$$

Proof Symmetry of the transition densities implies that, for any $t \geq 0, x \in D$,

$$p_t(x) = \int_D p_0(y) q_t(y, x) dy = \int_D p_0(y) q_t(x, y) dy \leq \|p_0\|_\infty \int_D q_t(x, y) dy = \|p_0\|_\infty,$$

and, similarly,

$$p_t(x) \geq \alpha.$$

It follows that

$$\begin{aligned} & \int_{\underline{T}}^{\bar{T}} \int_D \left| \nabla_x \log p_t(x) - \frac{\nabla_x h_N(x, t)}{h_N(x, t) \vee \alpha} \right|^2 p_t(x) dx dt \\ & \leq \|p_0\|_\infty \int_{\underline{T}}^{\bar{T}} \left\| \nabla_x \log p_t - \frac{\nabla_x h_N(t)}{h_N(t) \vee \alpha} \right\|_{L^2}^2 dt \\ & \leq 2\|p_0\|_\infty \left(\int_{\underline{T}}^{\bar{T}} \left\| \frac{\nabla_x(p_t - h_N(t))}{h_N(t) \vee \alpha} \right\|_{L^2}^2 dt + \int_{\underline{T}}^{\bar{T}} \left\| \nabla_x p_t \frac{p_t - (h_N(t) \vee \alpha)}{p_t(h_N(t) \vee \alpha)} \right\|_{L^2}^2 dt \right) \\ & \leq 2\|p_0\|_\infty \left(\frac{1}{\alpha^2} \int_{\underline{T}}^{\bar{T}} \left\| \nabla_x(p_t - h_N(t)) \right\|_{L^2}^2 dt + \frac{1}{\alpha^4} \int_{\underline{T}}^{\bar{T}} \left\| \nabla_x p_t(p_t - (h_N(t) \vee \alpha)) \right\|_{L^2}^2 dt \right). \end{aligned} \tag{24}$$

By Lemma 8, we obtain for the first term

$$\begin{aligned} \frac{1}{\alpha^2} \int_{\underline{T}}^{\bar{T}} \left\| \nabla_x(p_t - h_N(t)) \right\|_{L^2}^2 dt & \leq \alpha^{-2} \int_{\underline{T}}^{\bar{T}} \|p_t - h_N(t)\|_{H^1}^2 dt \\ & \lesssim \|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s}^2 \alpha^{-2} N^{-2s/d}. \end{aligned} \tag{25}$$

To bound the second term, note first that

$$\begin{aligned} \|p_t - \frac{1}{\text{Leb}(D)}\|_{H^{s+1}}^2 & = \sum_{j=1}^{\infty} e^{-2\lambda_j t} \lambda_j^{s+1} \langle p_0, e_j \rangle_{L^2}^2 \leq \frac{1}{2t} \sum_{j \geq 1} \lambda_j^s \langle p_0, e_j \rangle_{L^2}^2 = \frac{2}{t} \|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s}^2 \\ & \asymp \frac{2}{t} \|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s}^2 < \infty, \quad t > 0. \end{aligned}$$

Therefore, for any $i \in [d]$,

$$\begin{aligned} \|\partial_{x_i} p_t\|_{H^s} & = \|\partial_{x_i}(p_t - \frac{1}{\text{Leb}(D)})\|_{H^s} \leq \|p_t - \frac{1}{\text{Leb}(D)}\|_{H^{s+1}} \asymp \|p_t - \frac{1}{\text{Leb}(D)}\|_{H^{s+1}} \\ & \lesssim \frac{1}{\sqrt{t}} \|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s}. \end{aligned}$$

Since $s > d/2$, the Sobolev imbedding theorem yields

$$\sup_{i \in [d]} \|\partial_{x_i} p_t\|_\infty \lesssim t^{-1/2} \|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s}.$$

Consequently, for N as above,

$$\begin{aligned} \int_{\underline{T}}^{\overline{T}} \|\|\nabla_x p_t(p_t - (h_N(t) \vee \alpha))\|\|_{L^2}^2 dt &\lesssim \|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s}^2 \int_{\underline{T}}^{\overline{T}} \frac{1}{t} \|p_t - (h_N(t) \vee \alpha)\|_{L^2}^2 dt \\ &\leq \|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s}^2 \int_{\underline{T}}^{\overline{T}} \frac{1}{t} \|p_t - h_N(t)\|_{L^2}^2 dt \\ &= \|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s}^2 \int_{\underline{T}}^{\overline{T}} \frac{1}{t} e^{-2\lambda_1 t} dt \sum_{j \geq N+1} \langle p_0, e_j \rangle_{L^2}^2 \\ &\lesssim \|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s}^4 N^{-2s/d} \left(\int_{\underline{T}}^1 \frac{1}{t} dt + \int_1^{\overline{T}} e^{-\lambda_1 t} dt \right) \\ &\leq \|p_0 - \frac{1}{\text{Leb}(D)}\|_{H^s}^4 (1 + 1/\lambda_1) \log(\overline{T}^{-1}) N^{-2s/d}, \end{aligned} \tag{26}$$

where we used (23) and that $|p_t(x) - (h_N(x, t) \vee \alpha)| \leq |p_t(x) - h_N(x, t)|$ since $p_t(x) \geq \alpha$. Combining (24), (25) and (26) yields the assertion. \blacksquare

Step 2: Approximation of truncated score by neural networks Given the previous result, it remains to show that (the truncated version of) $\nabla_x \log h_N$ with h_N as defined in (22) can be well approximated by a neural network whose size is quantified in terms of N . To this end, we will make repeated use of the following fundamental observations about compositions and parallelisations of the type of ReLU neural networks introduced in Section 2 that we are dealing with.

First, we observe that the concatenation of such neural networks is itself simply another neural network. In particular, if $\varphi_1 \in \Phi(L_1, W_1, S_1, B_1)$ and $\varphi_2 \in \Phi(L_2, W_2, S_2, B_2)$ are such that $W_{2, L_2+2} = W_{1,1}$, then $\varphi_1 \circ \varphi_2 \in \Phi(L_1 + L_2 + 1, W_{1,2}^{\text{cat}}, S_1 + S_2, B_1 \vee B_2)$, where

$$W_{1,2}^{\text{cat}} = [W_{2,1} \quad W_{2,2} \quad \cdots \quad W_{2, L_2+1} \quad W_{1,1} \quad W_{1,2} \quad \cdots \quad W_{1, L_1+2}]^\top.$$

In general, if for some $k \in \mathbb{N}$, $\varphi_i \in \Phi(L_i, W_i, S_i, B_i)$ for $i = 1, \dots, k$, then

$$\varphi_1 \circ \varphi_2 \circ \cdots \circ \varphi_k \in \Phi\left(\sum_{i=1}^k L_i + k, W_{[k]}^{\text{cat}}, \sum_{i=1}^k S_i, \max_{i \in \{1, \dots, k\}} B_i\right),$$

where $W_{[k]}^{\text{cat}}$ is defined recursively as above. Similarly, if φ_1, φ_2 are as before but with $L_1 = L_2 = L$, we can parallelise the two into one network $\varphi_{1,2}^{\text{par}} \in \Phi(L, W_{1,2}^{\text{par}}, S_1 + S_2, B_1 \vee B_2)$ such that $\varphi_{1,2}^{\text{par}}(x, y) = [\varphi_1(x) \quad \varphi_2(y)]^\top$ for $x \in \mathbb{R}^{W_{1,1}}$ and $y \in \mathbb{R}^{W_{2,1}}$. The simplest way to construct this is using block matrices, i.e., by

$$\varphi_{1,2}^{\text{par}} = \begin{bmatrix} A_{1,L} & 0 \\ 0 & A_{2,L} \end{bmatrix} \sigma \begin{bmatrix} b_{1,L} \\ b_{2,L} \end{bmatrix} \begin{bmatrix} A_{1, L-1} & 0 \\ 0 & A_{2, L-1} \end{bmatrix} \sigma \begin{bmatrix} b_{1, L-1} \\ b_{2, L-1} \end{bmatrix} \cdots \begin{bmatrix} A_{1,1} & 0 \\ 0 & A_{2,1} \end{bmatrix} \sigma \begin{bmatrix} b_{1,1} \\ b_{2,1} \end{bmatrix} \begin{bmatrix} A_{1,0} & 0 \\ 0 & A_{2,0} \end{bmatrix},$$

which would mean $W_{1,2}^{\text{par}} = W_1 + W_2$. However, if φ_1 and φ_2 share some inputs (i.e., if the first, say, $m \in \mathbb{N}$ entries of x, y are the same), then the rightmost matrix in the above may be altered to

$$\begin{bmatrix} A_{1,0} & 0 \\ 0 & A_{2,0} \end{bmatrix} \begin{bmatrix} I_m & 0 & 0 \\ 0 & I_{W_{1,1}-m} & 0 \\ I_m & 0 & 0 \\ 0 & 0 & I_{W_{2,1}-m} \end{bmatrix},$$

whereby $(W_{1,2}^{\text{par}})_1 = W_{1,1} + W_{2,1} - m$ instead. Again, this can of course naturally be generalised to k networks of equal depth, where we then have

$$\varphi_{[k]}^{\text{par}} \in \Phi\left(L, W_{[k]}^{\text{par}}, \sum_{i=1}^k S_k, \max_{i \in \{1, \dots, k\}} B_i\right), \quad (W_{[k]}^{\text{par}})_j = \sum_{i=1}^k W_{i,j} \text{ for } j > 1.$$

Finally, note that multiplying the network $\varphi_{[k]}^{\text{par}}$ with the vector $[1 \cdots 1]$ from the left sums the entries of $\varphi_{[k]}^{\text{par}}$ without changing the size of the network substantially, whence

$$\sum_{i=1}^k \varphi_i \in \Phi\left(L, W_{[k]}^{\text{sum}}, k + \sum_{i=1}^k S_k, 1 \vee \max_{i \in \{1, \dots, k\}} B_i\right), \quad (W_{[k]}^{\text{sum}})_j = \sum_{i=1}^k W_{i,j} \text{ for } 1 < j < k.$$

For larger and more complicated neural networks, their exact sizes are often unavailable, and we only have access to their asymptotic sizes. Due to this, we also introduce the following class of neural networks that eases network size analysis in the proofs that follow,

$$\tilde{\Phi}(\tilde{L}, \tilde{W}, \tilde{S}, \tilde{B}) := \left\{ \varphi \in \Phi(L, W, S, B) : L \lesssim \tilde{L}, \|W\|_\infty \lesssim \tilde{W}, S \lesssim \tilde{S} \text{ and } B \lesssim \tilde{B} \right\}.$$

With this notation, we have for arbitrary networks $\varphi_i \in \tilde{\Phi}(L_i, W_i, S_i, B_i)$, $i = 1, 2$, that $\varphi_1 \circ \varphi_2 \in \tilde{\Phi}(L_1 + L_2, W_1 \vee W_2, S_1 + S_2, B_1 \vee B_2)$ and $\varphi_{1,2}^{\text{par}} \in \tilde{\Phi}(L_1 \vee L_2, W_1 + W_2, S_1 + S_2, B_1 \vee B_2)$.

With this class established, we can begin to approximate $\nabla_x \log h_N$ with a suitable network. We do this by approaching the problem in smaller pieces, first noting that, for $t \geq 0$ and $x \in D$,

$$\nabla_x \log h_N(x, t) = \frac{\sum_{j=1}^N e^{-t\lambda_j} \langle p_0, e_j \rangle_{L^2} \nabla_x e_j(x)}{\frac{1}{\text{Leb}(D)} + \sum_{j=1}^N e^{-t\lambda_j} \langle p_0, e_j \rangle_{L^2} e_j(x)}.$$

Thus, in order to approximate $\nabla_x \log h_N$, we need to be able to approximate products and quotients of functions. Such approximation results already exist in the literature, see, e.g., Oko et al. (2023); Schmidt-Hieber (2020); Telgarsky (2017); Yarotsky (2017), but here we give slightly stronger versions with optimised neural network sizes, which lead to improved convergence rates (in terms of log factors). More detailed statements and their proofs are given in Appendix C, but for our purposes the following will suffice.

Lemma 10 *For $m \in \mathbb{N}$ and $C \geq 1$, there exist neural networks $\varphi_m^{\text{mult}} \in \tilde{\Phi}(m, 1, m, C)$ and $\varphi_m^{\text{mult}, d} \in \tilde{\Phi}(m, d, dm, C)$ satisfying*

$$|\varphi_m^{\text{mult}}(x, y) - xy| \leq C2^{-m}, \quad x \in [0, 1], y \in [-C, C],$$

and

$$|\varphi_m^{\text{mult},d}(x, y) - xy| \leq \sqrt{d}C2^{-m}, \quad x \in [0, 1], y \in [-C, C]^d.$$

These also satisfy $\varphi_m^{\text{mult}}(x, 0) = \varphi_m^{\text{mult}}(0, y) = 0$.

Lemma 11 For $m, \underline{k}, \bar{k} \in \mathbb{N}$, there exists a neural network $\varphi_m^{\text{rec}} \in \tilde{\Phi}((k+m) \log(k+m), k, (k+m) \log(k+m), 2^k)$, where $k = \underline{k} + \bar{k}$, satisfying

$$|\varphi_m^{\text{rec}}(x) - x^{-1}| \leq 2^{-m}, \quad x \in [2^{-\bar{k}}, 2^{\underline{k}}].$$

Lemma 12 For each $m \in \mathbb{N}$, there exists a neural network $\varphi^{\text{cap}} \in \tilde{\Phi}(m \log m, m, m \log m, 2^{m/2})$ satisfying $\varphi^{\text{cap}}(t) \asymp \frac{1}{\sqrt{t}}$ for all $t \in [2^{-m}, 1]$.

With these, we are ready to prove the following key approximation result on h_N and $\nabla_x h_N$.

Lemma 13 Let $0 < \underline{T} < \bar{T}$ with $\underline{T} \in \text{Poly}(N^{-1})$ be given, and let f denote either h_N or $\partial_{x_i} h_N$, for some $i \in \{1, \dots, d\}$ and $N \in \mathbb{N}$ sufficiently large. Then, there exists a neural network $\varphi_f \in \tilde{\Phi}(L(N), W(N), S(N), B(N))$ satisfying

$$\forall t \in [\underline{T}, \bar{T}] : \|\varphi_f(\cdot, t) - f(\cdot, t)\|_{L^2}^2 \lesssim \begin{cases} N^{-\frac{2s}{d}} (\log N)^2, & \text{if } f = h_N \\ \varepsilon(t) N^{-\frac{2s}{d}} (\log N)^2, & \text{if } f = \partial_{x_i} h_N, \end{cases}$$

where

$$\int_{\underline{T}}^{\bar{T}} \varepsilon(t) dt \lesssim \bar{T} + \log(\underline{T}^{-1})$$

and whose network size is evaluated as

$$\begin{aligned} L(N) &= \log N \log \log N, \\ W(N) &= MN \log N, \\ S(N) &= MN (\log N)^2, \quad \text{and} \\ B(N) &= \frac{N^{\frac{2s+d}{2d}}}{\log N} \vee \frac{1}{\underline{T}}, \end{aligned}$$

where $M \in O(\log \frac{\bar{T}}{\underline{T}})$. Furthermore, there exists a constant $C < \infty$ depending only on p_0 and D such that $\sup_{x \in D} |\varphi_{\partial_{x_i} h_N}(x, t)| \leq C(1 \vee \frac{1}{\sqrt{t}})$ for all $t \in [\underline{T}, \bar{T}]$.

Proof We first construct a network φ_f with the desired error rate and specify its size at the end. To this end, suppose that there exist neural networks $\varphi_f^{(1)}, \dots, \varphi_f^{(M)}$, where $M = \lfloor \log_2 \frac{\bar{T}}{\underline{T}} \rfloor$ such that

$$\|\varphi_f^{(m)}(\cdot, t) - f(\cdot, t)\|_{L^2}^2 \lesssim \begin{cases} N^{-\frac{2s}{d}} (\log N)^2, & \text{if } f = h_N, \\ \left(\frac{1}{2^{m-1}\underline{T}} \vee 1\right) N^{-\frac{2s}{d}} (\log N)^2, & \text{if } f = \partial_{x_i} h_N, \end{cases}$$

for $m = 1, \dots, M$ and $t \in [2^{m-1}\underline{T}, 2^{m+1}\underline{T}]$. Then, consider the partition of unity $\{\pi_m\}_{m=1}^M$ given by

$$\pi_m(t) = 0 \vee \left(\frac{t - 2^{m-1}\underline{T}}{2^{m-1}\underline{T}} \wedge \frac{2^{m+1}\underline{T} - t}{2^m\underline{T}} \right) = \begin{cases} \frac{t}{2^{m-1}\underline{T}} - 1, & \text{if } t \in [2^{m-1}\underline{T}, 2^m\underline{T}] \\ 2 - \frac{t}{2^m\underline{T}}, & \text{if } t \in [2^m\underline{T}, 2^{m+1}\underline{T}], \\ 0, & \text{otherwise} \end{cases}$$

for $m = 2, \dots, M-1$, while $\pi_1(t) = 0 \vee \left(1 \wedge \frac{4\underline{T} - t}{2\underline{T}}\right)$ and $\pi_M(t) = 0 \vee \left(1 \wedge \frac{t - 2^{M-1}\underline{T}}{2^{M-1}\underline{T}}\right)$. Since for $a, b \in \mathbb{R}$, $a \vee b = a + \sigma(b - a)$ and $a \wedge b = a - \sigma(a - b)$, each π_m is representable as a neural network. We then claim that φ_f , defined as

$$\varphi_f(x, t) = \sum_{m=1}^M \varphi_{\ell_1}^{\text{mult}}(\pi_m(t), \varphi_f^{(m)}(x, t)), \quad \ell_1 = \left\lceil \frac{s}{d} \log_2 N \right\rceil,$$

yields the desired network. Indeed, we first note that since at most two of the π_m 's are non-zero for any $t \in [\underline{T}, \bar{T}]$ and $\varphi_{\ell}^{\text{mult}}(0, y) = 0$ for all $y \in \mathbb{R}$, we have for $m = 2, \dots, M$ and $t \in [2^{m-1}\underline{T}, 2^m\underline{T}]$ that

$$\begin{aligned} \|\varphi_f(\cdot, t) - f(\cdot, t)\|_{L^2} &= \|\varphi_{\ell_1}^{\text{mult}}(\pi_{m-1}(t), \varphi_f^{(m-1)}(\cdot, t)) + \varphi_{\ell_1}^{\text{mult}}(\pi_m(t), \varphi_f^{(m)}(\cdot, t)) - f(\cdot, t)\|_{L^2} \\ &\lesssim 2^{-\ell_1} + \|\pi_{m-1}(t)\varphi_f^{(m-1)}(\cdot, t) + \pi_m(t)\varphi_f^{(m)}(\cdot, t) - f(\cdot, t)\|_{L^2} \\ &\lesssim \begin{cases} 2^{-\ell_1} + N^{-\frac{s}{d}} \log N, & \text{if } f = h_N \\ 2^{-\ell_1} + \left(\frac{1}{\sqrt{2^{m-2}\underline{T}}} \vee 1\right) N^{-\frac{s}{d}} \log N, & \text{if } f = \partial_{x_i} h_N, \end{cases} \end{aligned}$$

where in the last inequality we used

$$\begin{aligned} &\|\pi_{m-1}(t)\varphi_f^{(m-1)}(\cdot, t) + \pi_m(t)\varphi_f^{(m)}(\cdot, t) - f(\cdot, t)\|_{L^2} \\ &= \|\pi_{m-1}(t)(\varphi_f^{(m-1)}(\cdot, t) - f(\cdot, t)) + \pi_m(t)(\varphi_f^{(m)}(\cdot, t) - f(\cdot, t))\|_{L^2} \\ &\leq \pi_{m-1}(t)\|\varphi_f^{(m-1)}(\cdot, t) - f(\cdot, t)\|_{L^2} + \pi_m(t)\|\varphi_f^{(m)}(\cdot, t) - f(\cdot, t)\|_{L^2} \\ &\lesssim \begin{cases} N^{-\frac{s}{d}} \log N, & \text{if } f = h_N \\ \left(\frac{1}{\sqrt{2^{m-2}\underline{T}}} \vee 1\right) N^{-\frac{s}{d}} \log N, & \text{if } f = \partial_{x_i} h_N, \end{cases} \end{aligned}$$

Setting

$$\varepsilon(t) = \sum_{m=1}^{M+1} \left(\frac{1}{2^{m-2}\underline{T}} \vee 1 \right) \mathbf{1}_{[2^{m-1}\underline{T}, 2^m\underline{T}]}(t),$$

and by choice of ℓ_1 , squaring both sides of the inequality yields the desired error rate. A similar but simpler analysis shows that this also holds for $t \in [\underline{T}, 2\underline{T}]$ and $t \in [2^M\underline{T}, 2^{M+1}\underline{T}]$, whence the inequality holds for all $t \in [\underline{T}, 2^{M+1}\underline{T}] \supseteq [\underline{T}, \bar{T}]$. Furthermore, we have

$$\int_{\underline{T}}^{\bar{T}} \varepsilon(t) dt \leq \sum_{m=1}^{M+1} \left(\frac{1}{2^{m-2}\underline{T}} \vee 1 \right) (2^m\underline{T} - 2^{m-1}\underline{T})$$

$$\begin{aligned}
 &= \sum_{m=1}^{\lfloor \log_2(\underline{T}^{-1}) \rfloor + 2} \frac{2^{m-1} \underline{T}}{2^{m-2} \underline{T}} + \sum_{m=\lfloor \log_2(\underline{T}^{-1}) \rfloor + 3}^{M+1} 2^{m-1} \underline{T} \\
 &= 2(\lfloor \log_2(\underline{T}^{-1}) \rfloor + 2) + \underline{T} \left(\sum_{m=0}^M 2^m - \sum_{m=0}^{\lfloor \log_2(\underline{T}^{-1}) \rfloor + 2} 2^m \right) \\
 &= 2(\lfloor \log_2(\underline{T}^{-1}) \rfloor + 2) + 2\underline{T} \left(2^M - 2^{\lfloor \log_2(\underline{T}^{-1}) \rfloor} \right) \\
 &\lesssim \bar{T} + \log(\underline{T}^{-1})
 \end{aligned}$$

as claimed. For that reason, we only need to construct the networks $\varphi_f^{(m)}$ for all $m \in \{1, \dots, M\}$, so let some such m be fixed. Then, let $a_m = 3 \cdot 2^{m-2} \underline{T}$ and $b_m = 5 \cdot 2^{m-2} \underline{T}$, and set $f_m(x, t) = f(x, a_m t + b_m)$ such that $f_m(x, [-1, 1]) = f(x, [2^{m-1} \underline{T}, 2^{m+1} \underline{T}])$ for all $x \in D$. As in the proof of Lemma 8, we see that for each fixed $t \in [\underline{T}, \bar{T}]$, $h_N(\cdot, t) \in H^{s+1}(D)$, whence $f(\cdot, t) \in H^s(D) = B_{2,2}^s(D)$. Furthermore, $\|f(\cdot, t)/(1 \vee \|f(\cdot, t)\|_{B_{2,2}^s(D)})\|_{B_{2,2}^s(D)} \leq 1$. Thus, since D is bounded, a slight modification of Suzuki (2019, Proposition 1) using the Sobolev extension theorem yields the existence of a neural network $\tilde{\varphi}_{f,t} \in \tilde{\Phi}(\log N, N, N \log N, N^{1/d})$ satisfying $\|\tilde{\varphi}_{f,t} - f(\cdot, t)/(1 \vee \|f(\cdot, t)\|_{B_{2,2}^s(D)})\|_{L^2(D)} \lesssim N^{-\frac{s}{d}}$. Then, setting $\varphi_{f,t} = (1 \vee \|f(\cdot, t)\|_{B_{2,2}^s(D)}) \tilde{\varphi}_{f,t}$, we have

$$\begin{aligned}
 \|\varphi_{f,t} - f(\cdot, t)\|_{L^2(D)} &= (1 \vee \|f(\cdot, t)\|_{B_{2,2}^s(D)}) \left\| \tilde{\varphi}_{f,t} - \frac{f(\cdot, t)}{(1 \vee \|f(\cdot, t)\|_{B_{2,2}^s(D)})} \right\|_{L^2(D)} \\
 &\lesssim (1 \vee \|f(\cdot, t)\|_{B_{2,2}^s(D)}) N^{-\frac{s}{d}}.
 \end{aligned}$$

Noting that

$$\|f(\cdot, t)\|_{B_{2,2}^s(D)} \asymp \|f(\cdot, t)\|_{H^s(D)} \lesssim \begin{cases} \|p_0\|_{H^s}, & \text{if } f = h_N, \\ \frac{1}{\sqrt{t}} \|p_0\|_{H^s}, & \text{if } f = \partial_{x_i} h_N, \end{cases}$$

we thus have that $\varphi_{f,t} \in \tilde{\Phi}(\log N, N, N \log N, N^{1/d} \vee \frac{1}{\sqrt{\underline{T}}})$, and

$$\|\varphi_{f,t} - f(\cdot, t)\|_{L^2(D)} \lesssim \begin{cases} N^{-\frac{s}{d}}, & \text{if } f = h_N, \\ \left(\frac{1}{\sqrt{t}} \vee 1\right) N^{-\frac{s}{d}}, & \text{if } f = \partial_{x_i} h_N. \end{cases}$$

Then, for each $t \in [-1, 1]$, let $\varphi_{f_m,t} = \varphi_{f, a_m t + b_m}$ such that $\varphi_{f_m,t}$ is an approximation of $f_m(t, \cdot)$. We will then approximate f_m by polynomial interpolation in time and by $\varphi_{f_m,t}$ in space. There are then three main sources of error: the error in approximating a polynomial with a neural network, the error from polynomial interpolation, and finally the error from approximating $f_m(\cdot, t)$ by $\varphi_{f_m,t}$. To separate these sources of error, we now let $\{t_i\}_{i=0}^k$ be the first $k+1$ Chebyshev nodes on $[-1, 1]$ for some k to be determined later, i.e. $t_i = \cos \frac{i\pi}{k}$. Then, for $i = 0, \dots, k$, let $p_i(t) = \prod_{j \neq i} (t - t_j)$ and set $c_i = \frac{1}{p_i(t_i)}$. Furthermore, set

$$\varphi_{f_m}(x, t) = \sum_{i=0}^k c_i \varphi_{\ell_2}^{\text{mult}}(\varphi_{p_i}(t), \varphi_{f_m, t_i}(x)),$$

$$\begin{aligned}\psi_m(x, t) &= \sum_{i=0}^k c_i p_i(t) \varphi_{f_m, t_i}(x), \\ P_m(x, t) &= \sum_{i=0}^k c_i p_i(t) f_m(x, t_i),\end{aligned}$$

where φ_{p_i} is a neural network approximation of p_i satisfying $|\varphi_{p_i}(t) - p_i(t)| \lesssim k2^{-\ell_3}$ to be constructed later. We then have that

$$\begin{aligned}\|\varphi_{f_m}(\cdot, t) - f_m(\cdot, t)\|_{L^2} &\leq \|\varphi_{f_m}(\cdot, t) - \psi_m(\cdot, t)\|_{L^2} + \|\psi_m(\cdot, t) - P_m(\cdot, t)\|_{L^2} \\ &\quad + \|P_m(\cdot, t) - f_m(\cdot, t)\|_{L^2}.\end{aligned}\tag{27}$$

By Trefethen (2020, Theorem 5.2), it holds that $|c_i| \leq \frac{2^{k-1}}{k}$, and so the first term of (27) is upper bounded by

$$\begin{aligned}\|\varphi_{f_m}(\cdot, t) - \psi_m(\cdot, t)\|_{L^2} &\leq \sum_{i=0}^k |c_i| \|\varphi_{\ell_2}^{\text{mult}}(\varphi_{p_i}(t), \varphi_{f_m, t_i}) - p_i(t) \varphi_{f_m, t_i}\|_{L^2} \\ &\lesssim \sum_{i=0}^k |c_i| (2^{-\ell_2} + k2^{-\ell_3} \|\varphi_{f_m, t_i}\|_{L^2}) \\ &\lesssim \sum_{i=0}^k |c_i| (2^{-\ell_2} + k2^{-\ell_3} (\underline{T}^{-\frac{1}{2}} N^{-\frac{s}{d}} + \|f_m(\cdot, t_i)\|_{L^2})) \\ &\leq 2^{k-1} (2^{-\ell_2} + k2^{-\ell_3} (\underline{T}^{-\frac{1}{2}} N^{-\frac{s}{d}} + \|p_0\|_{H^1})),\end{aligned}$$

and choosing $\ell_2 = \lceil \frac{s}{d} \log_2 N + k \rceil$ and $\ell_3 = \ell_2 + \lceil \log_2(k + \underline{T}^{-\frac{1}{2}}) \rceil$ bounds this term by $N^{-s/d}$. For the second term of (27), it is a well-known property of Chebyshev nodes that $|p_i(t)c_i| \leq 2$, whence

$$\begin{aligned}\|\psi_m(\cdot, t) - P_m(\cdot, t)\|_{L^2} &\leq \sum_{i=0}^k |c_i p_i(t)| \|\varphi_{f_m, t_i} - f_m(\cdot, t_i)\|_{L^2} \\ &\lesssim \begin{cases} kN^{-\frac{s}{d}}, & \text{if } f = h_N \\ k \left(\frac{1}{\sqrt{2^{m-1}\underline{T}}} \vee 1 \right) N^{-\frac{s}{d}}, & \text{if } f = \partial_{x_i} h_N, \end{cases}\end{aligned}$$

where the last inequality in the case of $f = \partial_{x_i} h_N$ follows from the fact that for all $t \in [-1, 1]$

$$\begin{aligned}\|\varphi_{f_m, t} - f_m(\cdot, t)\|_{L^2} &= \|\varphi_{f, a_m t + b_m} - f(\cdot, a_m t + b_m)\|_{L^2} \\ &\lesssim \left(\frac{1}{\sqrt{a_m t + b_m}} \vee 1 \right) N^{-\frac{s}{d}} \\ &\leq \left(\frac{1}{\sqrt{2^{m-1}\underline{T}}} \vee 1 \right) N^{-\frac{s}{d}}.\end{aligned}$$

Finally, for the third term of (27), we note that for each $x \in D$, the function $t \mapsto f_m(x, t)$ is entire on \mathbb{C} as an affine combination of exponentials. It then follows from Trefethen (2020,

Theorem 8.2) that, for $\rho > 1$,

$$|f_m(x, t) - P_m(x, t)| \leq \frac{4M_{m,\rho}(x)\rho^{-k}}{\rho - 1}, \quad t \in [-1, 1],$$

where

$$M_{m,\rho}(x) = \max_{z \in \partial E_\rho} |f_m(x, z)|, \quad \text{and} \quad \partial E_\rho = \left\{ \frac{z + z^{-1}}{2} \mid |z| = \rho \right\}.$$

For $z \in \partial E_\rho$, we have, letting \tilde{e}_n denote either e_n or $\partial_{x_i} e_n$, depending on f ,

$$f_m(x, z) = f(x, a_m z + b_m) = \sum_{n=0}^N e^{-\lambda_n(a_m \operatorname{Re} z + b_m)} \langle p_0, e_n \rangle \tilde{e}_n(x) e^{-i\lambda_n a_m \operatorname{Im} z} = \langle \mathbf{r}(x, \operatorname{Re} z), \boldsymbol{\theta}(\operatorname{Im} z) \rangle,$$

where $(\mathbf{r}(x, y))_n = e^{-\lambda_n(a_m y + b_m)} \langle p_0, e_n \rangle \tilde{e}_n(x)$ and $\boldsymbol{\theta}(y)_n = e^{-i\lambda_n a_m y}$, so $|\boldsymbol{\theta}(y)| = N + 1$. Thus, by Cauchy-Schwarz inequality,

$$\begin{aligned} M_{m,\rho}(x) &\leq \left((N + 1) \max_{z \in \partial E_\rho} \sum_{n=0}^N e^{-2\lambda_n(a_m \operatorname{Re} z + b_m)} \langle p_0, e_n \rangle^2 \tilde{e}_n(x)^2 \right)^{1/2} \\ &= \left((N + 1) \sum_{n=0}^N e^{-2\lambda_n(a_m(\frac{-\rho-\rho^{-1}}{2}) + b_m)} \langle p_0, e_n \rangle^2 \tilde{e}_n(x)^2 \right)^{1/2}. \end{aligned}$$

Consequently,

$$\begin{aligned} \|M_{m,\rho}\|_{L^2}^2 &\leq (N + 1) \sum_{n=0}^N e^{-2\lambda_n(a_m(\frac{-\rho-\rho^{-1}}{2}) + b_m)} \langle p_0, e_n \rangle^2 \|\tilde{e}_n\|_{L^2}^2 \\ &\leq (N + 1) \sum_{n=0}^N e^{-2\lambda_n(a_m(\frac{-\rho-\rho^{-1}}{2}) + b_m)} \langle p_0, e_n \rangle^2 \|e_n\|_{H^1}^2 \\ &\lesssim (N + 1) \sum_{n=0}^N e^{-2\lambda_n(a_m(\frac{-\rho-\rho^{-1}}{2}) + b_m)} \lambda_n \langle p_0, e_n \rangle^2, \end{aligned}$$

and if $a_m(\frac{-\rho-\rho^{-1}}{2}) + b_m = 0$, the right hand side is bounded by $(N + 1)\|p_0\|_{H^1}^2$. This is exactly the case when $\rho = 3$, in which case we have

$$\|f_m(\cdot, t) - P_m(\cdot, t)\|_{L^2}^2 \lesssim N3^{-2k},$$

and setting $k = \lceil (\frac{s}{d} + \frac{1}{2}) \log N \rceil$ yields an approximation φ_{f_m} of f_m with the correct error rate. Setting $\varphi_f^{(m)}(x, t) = \varphi_{f_m}(\frac{t-b_m}{a_m}, x)$ for $t \in [2^{m-2}\underline{T}, 2^{m+2}\underline{T}]$ then gives the desired network. We thus only need to construct the network φ_{p_i} as detailed above. To this end, note that the neural network

$$\varphi_{p_i}^{(0)} : t \mapsto \begin{bmatrix} I_k & -I_k & 0 \\ 0 & 0 & I_{2^{\lceil \log_2 k \rceil - k}} \end{bmatrix} \sigma \begin{bmatrix} \mathbf{t}_i \\ -\mathbf{t}_i \\ -1 \end{bmatrix} \begin{bmatrix} I_k \\ -I_k \\ 0 \end{bmatrix} t, \quad \mathbf{t}_i = [t_0 \quad \cdots \quad t_{i-1} \quad t_{i+1} \quad \cdots \quad t_k]^\top,$$

maps t to the vector

$$\left[t - t_0 \quad \cdots \quad t - t_{i-1} \quad t - t_{i+1} \quad \cdots \quad t - t_k \quad 1 \quad \cdots \quad 1 \right]^\top \in [-2, 2]^{2^{\lceil \log_2 k \rceil}},$$

regardless of the signs of its entries. Without altering the size of $\varphi_{p_i}^{(0)}$, we can swap the entries of $\varphi_{p_i}^{(0)}(t)$ such that adjacent entries correspond to opposing t_j 's, e.g., such that the first two entries are $t - t_0$ and $t - t_k$ rather than $t - t_0$ and $t - t_1$ and so on. This ensures that the products of adjacent entries stay uniformly bounded rather than some products growing and some shrinking, better bounding the size of the network. A slight modification of (the proof of) Lemma 17 then yields a network $\bar{\varphi}_{\ell_3}^{\text{mult}}$ of the same asymptotic size as in Lemma 17 such that $|\varphi_{\ell_3}^{\text{mult}}(x, y) - xy| \leq 2^{-\ell_3}$ for all $x, y \in [-2, 2]$. Let $\varphi_{p_i}^{(j)}$ be a parallelisation of $2^{\lceil \log_2 k \rceil - j}$ copies of this network, and set

$$\varphi_{p_i} = \varphi_{p_i}^{(\lceil \log_2 k \rceil)} \circ \cdots \circ \varphi_{p_i}^{(1)} \circ \varphi_{p_i}^{(0)}.$$

To ensure that φ_{p_i} satisfies $|\varphi_{p_i} - p_i| \lesssim k2^{-\ell_3}$, fix $t \in [-1, 1]$, and for notation, set $\xi_n^{(j)} = (\varphi_{p_i}^{(j)} \circ \cdots \circ \varphi_{p_i}^{(1)} \circ \varphi_{p_i}^{(0)}(t))_n$ and $y_l = (\varphi_{p_i}^{(0)}(t))_l$. We then claim that

$$\left| \xi_n^{(j)} - \prod_{l=(n-1)2^j+1}^{n2^j} y_l \right| \leq (2^j - 1)2^{-\ell_3}.$$

This is true by construction for $j = 1$, so assume that this holds for some $j \geq 1$. We then have

$$\begin{aligned} \left| \xi_n^{(j+1)} - \prod_{l=(n-1)2^{j+1}+1}^{n2^{j+1}} y_l \right| &= \left| \varphi_{\ell_3}^{\text{mult}}(\xi_{2n-1}^{(j)}, \xi_{2n}^{(j)}) - \prod_{l=(n-1)2^{j+1}+1}^{n2^{j+1}} y_l \right| \\ &\leq 2^{-\ell_3} + \left| \xi_{2n-1}^{(j)} \xi_{2n}^{(j)} - \left(\prod_{l=(2n-2)2^j+1}^{(2n-1)2^j} y_l \right) \left(\prod_{l=(2n-1)2^j+1}^{2n2^j} y_l \right) \right| \\ &\leq 2^{-\ell_3} \left(1 + \left(|\xi_{2n-1}^{(j)}| + \prod_{l=(2n-1)2^j+1}^{2n2^j} |y_l| \right) (2^j - 1) \right). \end{aligned}$$

By our previous rearranging of entries in $\varphi_{p_i}^{(0)}(t)$, we have that $|\xi_{2n-1}^{(j)}| + \prod_{l=(2n-1)2^j+1}^{2n2^j} |y_l| \leq 2$ for all $j \geq 1$, and the claim follows. This then implies that $|\varphi_{p_i}(t) - p_i(t)| \lesssim k2^{-\ell_3}$ for all $t \in [-1, 1]$ as desired.

We now shift from analysing the error of the network to its size instead. First, it is apparent from their construction that $\varphi_{p_i}^{(0)} \in \tilde{\Phi}(1, k, k, 1)$, while it also holds that $\varphi_{p_i}^{(j)} \in \tilde{\Phi}(\ell_3, 2^{\lceil \log_2 k \rceil - (j-1)}, 2^{\lceil \log_2 k \rceil - j} \ell_3, 1)$. Hence, since $\sum_{j=1}^{\lceil \log_2 k \rceil} 2^{\lceil \log_2 k \rceil - j} = \sum_{j=0}^{\lceil \log_2 k \rceil - 1} 2^j = 2^{\lceil \log_2 k \rceil} - 1$,

$$\begin{aligned} \varphi_{p_i} &\in \tilde{\Phi}\left(\lceil \log_2 k \rceil \ell_3, k, (2^{\lceil \log_2 k \rceil} - 1) \ell_3 + k, 1\right) \\ &= \tilde{\Phi}(\log N \log \log N, \log N, (\log N)^2, 1). \end{aligned}$$

Parallellising this with φ_{f_m, t_i} , we thus get a network in $\tilde{\Phi}(\log N \log \log N, N, N \log N, N^{\frac{1}{d}} \vee \frac{1}{\sqrt{T}})$. Since this dominates the size of $\varphi_{\ell_2}^{\text{mult}}$ and since $|c_i| \leq \frac{2^{k-1}}{k} \leq \frac{N^{\frac{2s+d}{2d}}}{\frac{2s+d}{d} \log N}$, it follows that each term of φ_{f_m} is in $\tilde{\Phi}(L_M, W_M, S_M, B_M)$, where

$$\begin{aligned} L_M &= \log N \log \log N, \\ W_M &= N \log N, \\ S_M &= N(\log N)^2, \quad \text{and} \\ B_M &= \frac{N^{\frac{2s+d}{2d}}}{\log N} \vee \frac{1}{\sqrt{T}}. \end{aligned}$$

Next, by construction of the π_m 's, $\pi_m \in \tilde{\Phi}(1, 1, 1, \frac{1}{T})$, where parallellising $\varphi_f^{(m)}$ with π_m does not change the asymptotic size of the network. Since the size of $\varphi_{\ell_1}^{\text{mult}}$ is also dominated by that of φ_{f_m} , it follows that each term of φ_f is included in $\tilde{\Phi}(L_M, W_M, S_M, B_M)$ as well. Parallellising all of these and summing them yields

$$\begin{aligned} L(N) &= \log N \log \log N, \\ W(N) &= MN \log N, \\ S(N) &= MN(\log N)^2, \quad \text{and} \\ B(N) &= \frac{N^{\frac{2s+d}{2d}}}{\log N} \vee \frac{1}{T}. \end{aligned}$$

Finally, since $s > \frac{d}{2}$, we have by similar calculations as those in the proof of Lemma 9 that for some $\tilde{C} < \infty$ depending only on D and p_0 , it holds that $\sup_{x \in D} |\partial_{x_i} h_N(x, t)| \leq \tilde{C} \frac{1}{\sqrt{t}}$. Letting φ^{cap} be as in Lemma 12 (with $m = \log_2(\underline{T}^{-1}) \asymp \log N$), it follows that also $\sup_{x \in D} |\partial_{x_i} h_N(x, t)| \leq C \varphi^{\text{cap}}(t)$ for all $t \in [\underline{T}, \overline{T}]$, whence we get a no worse approximation by replacing $\varphi_{\partial_{x_i} h_N}$ with $\varphi_{\partial_{x_i} h_N} \wedge (C \varphi^{\text{cap}})$ and this network has the desired bound. Furthermore, since $\varphi^{\text{cap}} \in \tilde{\Phi}(\log N \log \log N, \log N, \log N \log \log N, 1/\sqrt{T})$, and this is dominated by the network size of $\varphi_{\partial_{x_i} h_N}$, taking this minimum does not alter the size of the network. This finishes the proof. \blacksquare

Step 3: Putting things together With the essential preparations from Step 1 and 2, we can now finally prove Theorem 7.

Proof of Theorem 7 Let $\varphi_{h_N}, \varphi_{\partial_{x_i} h_N}$, $i \in \{1, \dots, d\}$ and $N \in \mathbb{N}$, be as in Lemma 13. Parallellising the latter of these yields a network $\varphi_{\nabla_x h_N} \in \tilde{\Phi}(L(N), W(N), S(N), B(N))$ approximating $\nabla_x h_N$. We then claim that φ_s defined as

$$\varphi_s(x, t) = \varphi_\ell^{\text{mult}, d}(\varphi_\ell^{\text{rec}}(\varphi_{h_N}(x, t) \vee \alpha), \varphi_{\nabla_x h_N}(x, t)), \quad \ell = \left\lceil \frac{s}{d} \log_2 N \right\rceil,$$

has the desired properties. Indeed, for the size of the network, notice that $\varphi_{h_N} \vee \alpha$ is bounded above (by $\|p_0\|_\infty$) and below, and hence $\varphi_\ell^{\text{rec}} \circ (\varphi_{h_N} \vee \alpha)$ is bounded above by $2\alpha^{-1}$

for N large enough, while the entries of $\varphi_{\nabla_x h_N}$ are all bounded numerically by $\frac{C}{\sqrt{T}}$ some $C < \infty$ since $s > \frac{d}{2}$. Thus, $\varphi_\ell^{\text{rec}} \in \tilde{\Phi}(\log N \log \log N, \log N \log \log N, \log N \log \log N, 1)$ and $\varphi_\ell^{\text{mult}, d} \in \tilde{\Phi}(\log N, \log N, \log N, \frac{1}{\sqrt{T}})$, whereby φ_s has the same asymptotic size as $\varphi_{\nabla_x h_N}$. Note also that $|\varphi_s(x, t)| \leq \frac{4C\alpha^{-1}}{\sqrt{t}} \vee 1$ for all $x \in D$, $t \in [\underline{T}, \bar{T}]$ and N large enough. So all that remains is to show that φ_s , as defined above, satisfies (21). By Proposition 9 and the triangle inequality, this is equivalent to verifying

$$\int_{\underline{T}}^{\bar{T}} \int_D \left| \varphi_s(x, t) - \frac{\nabla_x h_N(x, t)}{h_N(x, t) \vee \alpha} \right|^2 p_t(x) dx dt \lesssim N^{-\frac{2s}{d}} (\log N)^2 (\bar{T} + \log(\underline{T}^{-1})).$$

which follows if we can show that $\|\varphi_s(\cdot, t) - \frac{\nabla_x h_N(\cdot, t)}{h_N(\cdot, t) \vee \alpha}\|_{L^2}^2 \lesssim \varepsilon(t) N^{-\frac{2s}{d}} (\log N)^2$ for each $t \in [\underline{T}, \bar{T}]$, where $\varepsilon(t)$ is as in Lemma 13. For doing so, first note that

$$\begin{aligned} \left\| \varphi_s(\cdot, t) - \frac{\nabla_x h_N(\cdot, t)}{h_N(\cdot, t) \vee \alpha} \right\|_{L^2} &\leq \left\| \varphi_s(\cdot, t) - \varphi_\ell^{\text{rec}}(\varphi_{h_N}(\cdot, t) \vee \alpha) \varphi_{\nabla_x h_N}(\cdot, t) \right\|_{L^2} \\ &\quad + \left\| \left(\varphi_\ell^{\text{rec}}(\varphi_{h_N}(\cdot, t) \vee \alpha) - \frac{1}{\varphi_{h_N}(\cdot, t) \vee \alpha} \right) \varphi_{\nabla_x h_N}(\cdot, t) \right\|_{L^2} \\ &\quad + \left\| \frac{\varphi_{\nabla_x h_N}(\cdot, t)}{\varphi_{h_N}(\cdot, t) \vee \alpha} - \frac{\nabla_x h_N(\cdot, t)}{h_N(\cdot, t) \vee \alpha} \right\|_{L^2}. \end{aligned}$$

The first term is simply evaluated as

$$\left\| \varphi_s(\cdot, t) - \varphi_\ell^{\text{rec}}(\varphi_{h_N}(\cdot, t) \vee \alpha) \varphi_{\nabla_x h_N}(\cdot, t) \right\|_{L^2} \leq 2^{-\ell} \left(4\sqrt{d}\alpha^{-1} \left\| p_0 - \frac{1}{\text{Leb}(D)} \right\|_{H^s} \text{Leb}(D) \right) \lesssim N^{-\frac{s}{d}},$$

while for the second term, we have

$$\begin{aligned} \left\| \left(\varphi_\ell^{\text{rec}}(\varphi_{h_N}(\cdot, t) \vee \alpha) - \frac{1}{\varphi_{h_N}(\cdot, t) \vee \alpha} \right) \varphi_{\nabla_x h_N}(\cdot, t) \right\|_{L^2} &\leq 2^{-\ell} \|\varphi_{\nabla_x h_N}(\cdot, t)\|_{L^2} \\ &\lesssim 2^{-\ell} \left(\sqrt{d\varepsilon(t)} N^{-\frac{s}{d}} \log N + \|p_0\|_{H^1} \right) \\ &\lesssim \sqrt{\varepsilon(t)} N^{-\frac{s}{d}}. \end{aligned}$$

For the final term, one obtains, similarly to the proof of Proposition 9,

$$\begin{aligned} \left\| \frac{\varphi_{\nabla_x h_N}(\cdot, t)}{\varphi_{h_N}(\cdot, t) \vee \alpha} - \frac{\nabla_x h_N(\cdot, t)}{h_N(\cdot, t) \vee \alpha} \right\|_{L^2} &\leq \alpha^{-1} \|\varphi_{\nabla_x h_N}(\cdot, t) - \nabla_x h_N(\cdot, t)\|_{L^2} \\ &\quad + \alpha^{-2} \|\nabla_x h_N(\cdot, t)(h_N(\cdot, t) - \varphi_{h_N}(\cdot, t))\|_{L^2} \\ &\lesssim \alpha^{-1} \sqrt{d\varepsilon(t)} N^{-\frac{s}{d}} \log N + \alpha^{-2} \|p_0\|_{H^1} \sqrt{\varepsilon(t)} N^{-\frac{s}{d}} \log N \\ &\lesssim \sqrt{\varepsilon(t)} N^{-\frac{s}{d}} \log N. \end{aligned}$$

Finally, setting $N = n^{\frac{d}{2s+d}}$ yields both the desired network size and error rate, which finishes the proof. \blacksquare

3.4 Proof of the main result

With the previous preparations, we can now prove our main result on the generative error, Theorem 1. To this end, we use the general error decomposition (13) in combination with Lemma 2 and Lemma 3 to control the early stopping and ergodic error contributions, as well as with the approximation result Theorem 7 that allows us to obtain an optimised upper bound on the empirical score loss via Theorem 4.

Proof of Theorem 1 Choose $\delta = n^{-2s/(2s+d)}$ and $N = n^{d/(2s+d)}$. By the choices for \underline{T}, \bar{T} and N , Theorem 7 implies that there exists a family of neural networks \mathcal{S} with the specified size constraints, such that for some $\mathfrak{s} \in \mathcal{S}$ we have

$$\int_{\underline{T}}^{\bar{T}} \int_D |\mathfrak{s}(x, t) - \nabla_x \log p_t(x)|^2 p_t(x) dx dt \lesssim n^{-\frac{2s}{2s+d}} (\log n)^3.$$

With these network size constraints, we get for $C(\mathcal{S}) := C$ and c from Lemma 5, by using a straightforward modification of Oko et al. (2023, Lemma C.2), that

$$\begin{aligned} \log \mathcal{N}(\mathcal{S}, \|\cdot\|_{D \times [\underline{T}, \bar{T}]}, \delta / (cC\bar{T})) &\lesssim S(n)L(n) \log(\delta^{-1}\bar{T}^2 L(n) \|W(n)\|_\infty B(n)) \\ &\lesssim n^{\frac{d}{2s+d}} (\log n)^5 \log \log n. \end{aligned}$$

Lemma 5 therefore implies that

$$\log \mathcal{N}(\mathcal{L}, \|\cdot\|_D, \delta) \lesssim n^{\frac{d}{2s+d}} (\log n)^5 \log \log n$$

as well. By Lemma 6 and the choices of \bar{T}, \underline{T} , we can choose $C(\mathcal{L}) \lesssim \log n$ so that

$$\frac{C(\mathcal{L})}{n} \log \mathcal{N}(\mathcal{L}, \|\cdot\|_D, \delta) \lesssim n^{-\frac{2s}{2s+d}} (\log n)^6 \log \log n.$$

Using Theorem 7, it follows from the above and Theorem 4 by our choice of N that

$$\begin{aligned} &\mathbb{E} \left[\int_{\underline{T}}^{\bar{T}} \int_D |\widehat{\mathfrak{s}}_n(x, t) - \nabla \log p_t(x)|^2 p_t(x) dx dt \right] \\ &\leq 2 \inf_{\mathfrak{s} \in \mathcal{S}} \int_{\underline{T}}^{\bar{T}} \int_D |\mathfrak{s}(x, t) - \nabla \log p_t(x)|^2 p_t(x) dx dt + 2 \frac{C(\mathcal{L})}{n} \left(\frac{37}{9} \log \mathcal{N}(\mathcal{L}, \|\cdot\|_D, \delta) + 32 \right) + 3\delta \\ &\lesssim n^{-\frac{2s}{2s+d}} (\log n)^3 + n^{-\frac{2s}{2s+d}} (\log n)^6 \log \log n + n^{-\frac{2s}{2s+d}} \\ &\lesssim n^{-\frac{2s}{2s+d}} (\log n)^6 \log \log n. \end{aligned}$$

Thus,

$$\begin{aligned} \mathbb{E} \left[\text{TV}(\bar{p}_{\bar{T}-\underline{T}}^{\mathfrak{s}^\circ}, \bar{p}_{\bar{T}-\underline{T}}^{\widehat{\mathfrak{s}}_n}) \right] &\leq \sqrt{\frac{1}{2} \mathbb{E} \left[\text{KL}(\bar{p}_{\bar{T}-\underline{T}}^{\mathfrak{s}^\circ} \| \bar{p}_{\bar{T}-\underline{T}}^{\widehat{\mathfrak{s}}_n}) \right]} \\ &= \sqrt{\mathbb{E} \left[\int_{\underline{T}}^{\bar{T}} \int_D f(x) |\widehat{\mathfrak{s}}_n(x, t) - \nabla \log p_t(x)|^2 p_t(x) dx dt \right]} \end{aligned}$$

$$\begin{aligned} &\leq \sqrt{\|f\|_{\overline{D}} \mathbb{E} \left[\int_{\underline{T}}^{\overline{T}} \int_{\overline{D}} |\widehat{\mathfrak{s}}_n(x, t) - \nabla \log p_t(x)|^2 p_t(x) \, dx \, dt \right]} \\ &\lesssim n^{-\frac{s}{2s+d}} (\log n)^3 (\log \log n)^{1/2}, \end{aligned}$$

where we used Pinsker's inequality and Jensen's inequality together with concavity of $x \mapsto \sqrt{x}$ for the first line, while the second line follows from Theorem 14 (Girsanov), together with independence of the driving Brownian motion \overline{W} and the initialisation $\overline{X}_0 \sim p_T$ from the data $(X_{0,i})_{i=1,\dots,n}$. The proof is concluded by applying the error decomposition (13) and using that, by Lemma 2 and Lemma 3, we have

$$\text{TV}(p_0, p_{\underline{T}}) + \text{TV}(\mathbb{P}(X_{\overline{T}} \in \cdot \mid X_0 \sim p_0), \mathcal{U}(\overline{D})) \lesssim \underline{T}^{\beta/2} + e^{-\lambda_1 \overline{T}} \lesssim n^{-\frac{s}{2s+d}},$$

by our choices of $\underline{T}, \overline{T}$. ■

4. Discussion

This paper presents a rigorous investigation of the non-standard class of denoising reflected diffusion models (DRDMs), focusing on statistical convergence guarantees and approximation within constrained settings, as it is relevant for scenarios involving bounded state spaces. This is a first step in extending the statistical analysis of standard denoising diffusion models to the generalised class of *denoising Markov models* proposed in Benton et al. (2024).

A key result of our analysis is the derivation of convergence rates in total variation that match the minimax lower bound on Sobolev classes up to a logarithmic factor. More precisely, from Yang and Barron (1999, Theorem 4), see also Oko et al. (2023, Proposition 5.2) for the verification of their assumptions for s -smooth Sobolev functions on $[0, 1]^d$, we have that

$$\inf_{\widehat{p}_n} \sup_{p_0 \in H^s(D)} \mathbb{E}_{p_0} [\text{TV}(p_0, \widehat{p}_n)] \asymp n^{-\frac{s}{2s+d}},$$

where the infimum ranges over all estimators \widehat{p}_n based on n i.i.d. data points having density p_0 under \mathbb{P}_{p_0} . Our upper bound stated in (11) therefore establishes that DRDMs attain the minimax optimal rate of convergence up to log-factors for specific Sobolev densities. Note that our logarithmic loss is comparatively small relative to that of unconstrained DDMs in Oko et al. (2023), where it is of order $(\log n)^8$.

However, convergence rates (even optimal ones) expressed in terms of the ambient dimension d fall short of capturing the empirical success of DDMs. This gap is related to the *manifold hypothesis*, a prominent idea that real-world high-dimensional data often reside on lower-dimensional manifolds, to which well-trained generative models are believed to adapt. Developing a theoretical underpinning for this hypothesis has therefore been one of the central goals in statistical theory for generative models. In the pioneering paper Oko et al. (2023), the authors also take a first step towards investigating statistical convergence guarantees of DDMs for data distributed on such lower-dimensional structures by extending their analysis to initial distributions supported on a lower-dimensional hyperplane, where they obtain the almost minimax optimal rate $n^{-\frac{s+1}{2s+d} + \varepsilon}$ in the Wasserstein-1 distance in

terms of the sample size n and the subspace dimension \tilde{d} . Related statistical results under linear subspace assumptions are given in Chen et al. (2023). In the recent work Tang and Yang (2024), Tang and Yang significantly extend this result by establishing (up to log factors) the minimax convergence rate $C(d)n^{-\frac{s+1}{2s+\tilde{d}}}$ in Wasserstein-1 distance for distributions p_0 such that

- (i) p_0 is supported on a compact and β -smooth \tilde{d} -dimensional submanifold \mathcal{M} with positive reach, where $\beta \geq 2$;
- (ii) p_0 is bounded away from zero on \mathcal{M} ;
- (iii) p_0 has smoothness of order $s \in [0, \beta - 1]$ w.r.t. the volume measure on \mathcal{M} .

Note that these three conditions mirror the three assumptions from Oko et al. (2023) mentioned in the introductory Section 1 in the manifold setting. The multiplicative factor $C(d)$ in Tang and Yang (2024)'s convergence rate is of order $d^{s+\tilde{d}/2}$ and thus potentially very large for high ambient dimension d . Most recently, Azangulov et al. (2024) show that this multiplicative factor can be significantly reduced to the order \sqrt{d} by appropriately choosing the neural network class for score approximation. Extending the DRDM framework to support data on submanifolds (and thus improving the convergence rate) presents additional mathematical challenges. For example, enforcing a reflecting boundary for data supported on lower-dimensional submanifolds would require non-trivial modifications to the spectral score representation and a revised analysis of the associated Sobolev bounds. Such adjustments are beyond the scope of this study, yet our current work may provide a foundational approach for future research in the reflected diffusion context.

The assumptions on p_0 in our model (see $(\mathcal{H}0)$) play a key role in controlling the approximation error in our DRDM framework under bounded domain constraints. Such assumptions, although somewhat limiting in a practical context, are comparable to those made in Oko et al. (2023) for the total variation convergence analysis of unconstrained models while maintaining compatibility with the spectral methods that we employ for our statistical analysis. In this context, our more stringent smoothness assumption $s > d/2$ implies Hölder continuity of the data density p_0 , but allows us to avoid some technical difficulties arising due to the less explicit analytical nature of DRDMs compared to standard DDMs.

How to remove the lower bound assumption $p_0|_{\text{supp } p_0} \geq \alpha$ that is present in all the works discussed above is a highly relevant and conceptually challenging question. Recent works by Zhang et al. (2024); Stéphanovitch et al. (2025), and Yakovlev and Puchkin (2025) prove minimax optimal rates for particular unconstrained diffusion models without lower bound assumptions for sub-Gaussian densities p_0 and push-forward distributions on the ambient space \mathbb{R}^D of the form $g_{\#}\mathcal{U}[0, 1]^d$ for Hölder-continuous g and $d \leq D$, respectively. Zhang et al. (2024) use a more classical kernel estimation and truncation strategy instead of neural network approximations. Stéphanovitch et al. (2025) exploit deeper results on the space-time regularity of the score function of an Ornstein–Uhlenbeck process for direct approximation with tanh activation function, thereby avoiding the need to approximate p_t and ∇p_t separately. Finally Yakovlev and Puchkin (2025) exploit the structure of the score induced by their data assumption $p_0 \sim g_{\#}\mathcal{U}[0, 1]^d$ and the Gaussian Ornstein–Uhlenbeck forward transition densities to construct their ReLU neural network based approximation

class in a very specific way. These approaches do not translate directly to our non-Gaussian reflected setting and we leave the statistical study of reflected diffusion models without lower bound assumptions to future work.

In general, our analytical approach in DRDMS differs significantly from that in unconstrained diffusion models, as the bounded domain prevents the explicit Gaussian transition densities commonly used for error control. The semi-explicit nature of these densities in the reflected setting means that, rather than relying on straightforward Gaussian approximations, we implement general spectral decompositions and Sobolev-based bounds informed by the Sobolev smoothness of p_0 . The score approximation here presents additional technical challenges, which we address through an innovative polynomial-time interpolation procedure that proves crucial to achieving feasible convergence rates. This technique introduces new and effective methods for controlling error contributions in generative models with bounded state spaces and, as outlined in the introduction, may also provide a versatile tool for statistical analysis of generalised denoising Markov models (Benton et al., 2024; Ren et al., 2025) beyond the scope of this paper.

Finally, it should be noted that our analysis does not address sampling issues, in particular those arising from simulating the backward reflected process with an estimated drift. While this aspect is important in practical implementations, our current focus on theoretical convergence guarantees serves to isolate and address the approximation errors inherent to the reflection-based model setup itself. Future studies could incorporate error analysis for sampling methodologies specific to reflected processes to extend the results presented here.

Acknowledgments

Financial support from the Villum Synergy Programme, Project No. 50099, is gratefully acknowledged.

Appendix A. Girsanov’s theorem for reflected diffusions

The following result is a version of Girsanov’s theorem for reflected diffusions, which is a correction of Theorem 7.1 and Theorem A.6 in Lou and Ermon (2023), where the influence of the diffusion coefficient on the KL divergence has been overlooked.

Theorem 14 *Let $(X_t)_{t \in [0, T]}$ and $(\tilde{X}_t)_{t \in [0, T]}$ be solutions of the normally reflected SDEs*

$$\begin{aligned} dX_t &= b(t, X_t) dt + \sigma(t, X_t) dW_t + \nu(X_t) d\ell_t, \\ d\tilde{X}_t &= \tilde{b}(t, \tilde{X}_t) dt + \sigma(t, \tilde{X}_t) d\tilde{W}_t + \nu(\tilde{X}_t) d\tilde{\ell}_t, \quad \tilde{X}_0 \stackrel{d}{=} X_0, \end{aligned} \tag{28}$$

where b, \tilde{b} are bounded on $[0, T] \times \bar{D}$ and $\sigma(t, \cdot)$ is bounded, Lipschitz and uniformly elliptic in the sense $a(t, \cdot) := \sigma(t, \cdot)\sigma(t, \cdot)^\top \succeq \underline{\lambda}\mathbb{I}$ for some $\underline{\lambda} > 0$ and all $t \in [0, T]$. Denote by \mathbb{P}_{X^T} and $\mathbb{P}_{\tilde{X}^T}$ their respective path measures on $\mathcal{C}([0, T], \bar{D})$. Then, for $L_t = \int_0^t \nu(X_s) d\ell_s$,

$$\begin{aligned} &\log \frac{d\mathbb{P}_{\tilde{X}^T}}{d\mathbb{P}_{X^T}}(X^T) \\ &= \int_0^T (\tilde{b}(t, X_t) - b(t, X_t))^\top a^{-1}(t, X_t) d(X_t - L_t) \end{aligned}$$

$$- \frac{1}{2} \int_0^T (\tilde{b}(t, X_t) - b(t, X_t))^\top a^{-1}(t, X_t) (\tilde{b}(t, X_t) + b(t, X_t)) dt,$$

a.s., and

$$\text{KL}(\mathbb{P}_{X^T} \parallel \mathbb{P}_{\tilde{X}^T}) = \frac{1}{2} \mathbb{E} \left[\int_0^T \|\sigma^{-1}(t, X_t) (\tilde{b}(t, X_t) - b(t, X_t))\|^2 dt \right].$$

Proof Let

$$\begin{aligned} Z_T := & \exp \left(\int_0^T (\tilde{b}(t, X_t) - b(t, X_t))^\top (\sigma^{-1}(t, X_t))^\top dW_t \right. \\ & \left. - \frac{1}{2} \int_0^T \|\sigma^{-1}(t, X_t) (\tilde{b}(t, X_t) - b(t, X_t))\|^2 dt \right), \end{aligned}$$

and define $d\mathbb{Q}_T := Z_T d\mathbb{P}$. Since $\bar{\beta} := \sup_{t \in [0, T], x \in \bar{D}} \|\beta(t, x)\| < \infty$ for $\beta \in \{b, \tilde{b}\}$, we have

$$\sup_{t \in [0, T], x \in \bar{D}} \|\sigma^{-1}(t, x) \beta(t, x)\| \leq \frac{\bar{\beta}}{\lambda} < \infty.$$

Thus, Novikov's condition is satisfied, making \mathbb{Q}_T a probability measure equivalent to \mathbb{P} , and Girsanov's theorem implies that

$$\tilde{W}_t := W_t - \int_0^t \sigma^{-1}(s, X_s) (\tilde{b}(s, X_s) - b(s, X_s)) ds, \quad t \in [0, T],$$

is a \mathbb{Q}_T -Brownian motion, see Karatzas and Shreve (1991, Chapter 3, Theorem 5.1, Corollary 5.13). Thus,

$$X_t = X_0 + \int_0^t \tilde{b}(s, X_s) ds + \int_0^t \sigma(s, X_s) d\tilde{W}_s + \int_0^t \nu(X_s) d\ell_s, \quad t \in [0, T],$$

where $\ell_t = \int_0^t \mathbf{1}_{\{X_s \in \partial D\}} d\ell_s$, for any $t \in [0, T]$, \mathbb{P} -a.s. and hence also \mathbb{Q}_T -a.s. Since $X_0 \stackrel{d}{=} \tilde{X}_0$, it follows that X^T is a weak solution to the reflected SDE (28) and analogously to Karatzas and Shreve (1991, Chapter 5, Proposition 3.10), it holds that under the given assumptions on \tilde{b}, σ , weak solutions to (28) are unique in law. Thus, X^T under \mathbb{Q}_T has the same law as \tilde{X}^T under \mathbb{P} , such that, upon noting that

$$\begin{aligned} \log Z_T = & \int_0^T (\tilde{b}(t, X_t) - b(t, X_t))^\top a^{-1}(t, X_t) d(X_t - L_t) \\ & - \frac{1}{2} \int_0^T (\tilde{b}(t, X_t) - b(t, X_t))^\top a^{-1}(t, X_t) (\tilde{b}(t, X_t) + b(t, X_t)) dt, \end{aligned}$$

the first assertion follows. This yields immediately

$$\begin{aligned} \text{KL}(\mathbb{P}_{X^T} \parallel \mathbb{P}_{\tilde{X}^T}) &= -\mathbb{E} \left[\log \frac{d\mathbb{P}_{\tilde{X}^T}}{d\mathbb{P}_{X^T}}(X^T) \right] = -\mathbb{E}[\log Z_T] \\ &= \frac{1}{2} \mathbb{E} \left[\int_0^T \|\sigma^{-1}(t, X_t) (\tilde{b}(t, X_t) - b(t, X_t))\|^2 dt \right]. \end{aligned}$$

■

Appendix B. Bernstein condition for the denoising score matching excess loss and generalisation error

The goal of this section is to prove (16) and thereby justify the generalisation error decomposition from Theorem 4 in a generalised Markovian framework that in particular includes our reflected forward model, as well as the unconstrained Ornstein–Uhlenbeck models used in other works.

Let $Y_1, \dots, Y_n \stackrel{\text{i.i.d.}}{\sim} \mu$ be a collection of random variables with state space $\mathcal{X} \subset \mathbb{R}^d$. Let also $(X_t)_{t \geq 0}$ be a Markov process with state space \mathcal{X} and transition densities $(q_t)_{t \geq 0}$ w.r.t. some reference measure ν , that is, $\mathbb{P}^x(X_t \in dy) = q_t(x, y) \nu(dy)$, where $\mathbb{P}^x = \mathbb{P}(\cdot | X_0 = x)$. Denote by p_t the density of X_t started in μ at time $t > 0$, i.e., $p_t(y) \nu(dy) = \mathbb{P}^\mu(X_t \in dy) = \int_{\mathcal{X}} q_t(x, y) \mu(dx) \nu(dy)$. We assume that the *score* s° of this Markov process given by $s^\circ(x, t) = \nabla \log p_t(x)$ as well as $\nabla \log q_t(x, y)$ are well defined for all $t > 0, x, y \in \mathcal{X}$, and define for some measurable candidate function s the *denoising score matching loss* over an interval $[\underline{T}, \bar{T}] \subset (0, \infty)$ by

$$\begin{aligned} L_s(x) &= \int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(y, t) - \nabla \log q_t(x, y)|^2 q_t(x, y) \nu(dy) dt \\ &= \mathbb{E}^x \left[\int_{\underline{T}}^{\bar{T}} |s(X_t, t) - \nabla \log q_t(x, X_t)|^2 dt \right], \end{aligned}$$

which gives

$$\mathbb{E}^\mu[L_s(X_0)] = \int_{\underline{T}}^{\bar{T}} \iint_{\mathcal{X} \times \mathcal{X}} |s(y, t) - \nabla \log q_t(x, y)|^2 q_t(x, y) \nu(dy) \mu(dx) dt.$$

Then, given sufficient integrability properties that allow interchanging the order of differentiation and integration, a few simple lines of calculation establish the equivalence between denoising and *explicit* score matching Vincent (2011), which is the first equality in

$$\begin{aligned} \mathbb{E}^\mu[L_s(X_0)] &= \int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(y, t) - \nabla \log p_t(y)|^2 p_t(y) \nu(dy) dt + C_{\underline{T}, \bar{T}} \\ &= \int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(y, t) - s^\circ(y, t)|^2 p_t(y) \nu(dy) dt + C_{\underline{T}, \bar{T}}, \end{aligned}$$

where $C_{\underline{T}, \bar{T}} = \mathbb{E}^\mu[L_{s^\circ}(X_0)]$ is independent of \mathcal{S} . In particular, we obtain

$$\int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(y, t) - \nabla \log p_t(y)|^2 p_t(y) \nu(dy) dt = \mathbb{E}^\mu[L_s(X_0)] - \mathbb{E}^\mu[L_{s^\circ}(X_0)]. \quad (29)$$

As discussed before, a natural and tractable choice for the score s° given the data Y_1, \dots, Y_n is therefore given by the empirical risk minimiser

$$\hat{s} \in \arg \min_{s \in \mathcal{S}} \frac{1}{n} \sum_{i=1}^n L_s(Y_i),$$

where \mathcal{S} is some appropriate class of candidate functions. The following result shows that the excess risk $L_s - L_{s^\circ}$ satisfies a Bernstein condition w.r.t. μ , which is essential for the denoising score loss to work well as an empirical risk minimisation objective.

Lemma 15 *Suppose that $\sup_{s \in \mathcal{S} \cup \{s^\circ\}} \|L_s\|_{\mathcal{X}} \leq C(\mathcal{L}) < \infty$. Then,*

$$\mathbb{E}^\mu [(L_s(X_0) - L_{s^\circ}(X_0))^2] \leq 4C(\mathcal{L})\mathbb{E}^\mu [L_s(X_0) - L_{s^\circ}(X_0)].$$

Proof Using the elementary equality $|a|^2 - |b|^2 = \langle a + b, a - b \rangle$ for vectors a, b , and the Cauchy–Schwarz inequality several times, we find for any $x \in \mathcal{X}$,

$$\begin{aligned} & |L_s(x) - L_{s^\circ}(x)| \\ & \leq \int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |\langle s(y, t) - s^\circ(y, t), s(y, t) + s^\circ(y, t) - 2\nabla \log q_t(x, y) \rangle| q_t(x, y) \nu(dy) dt \\ & \leq \int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(y, t) - s^\circ(y, t)| |s(y, t) + s^\circ(y, t) - 2\nabla \log q_t(x, y)| q_t(x, y) \nu(dy) dt \\ & \leq \int_{\underline{T}}^{\bar{T}} \left(\int_{\mathcal{X}} |s(y, t) - s^\circ(y, t)|^2 q_t(x, y) \nu(dy) \right)^{1/2} \\ & \quad \times \left(\int_{\mathcal{X}} |s(y, t) + s^\circ(y, t) - 2\nabla \log q_t(x, y)|^2 q_t(x, y) \nu(dy) \right)^{1/2} dt \\ & \leq \left(\int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(y, t) - s^\circ(y, t)|^2 q_t(x, y) \nu(dy) dt \right)^{1/2} \\ & \quad \times \left(\int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(y, t) + s^\circ(y, t) - 2\nabla \log q_t(x, y)|^2 q_t(x, y) \nu(dy) dt \right)^{1/2} \\ & \leq \sqrt{2(L_s(x) + L_{s^\circ}(x))} \left(\int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(y, t) - s^\circ(y, t)|^2 q_t(x, y) \nu(dy) dt \right)^{1/2} \\ & \leq 2\sqrt{C(\mathcal{L})} \left(\int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(y, t) - s^\circ(y, t)|^2 q_t(x, y) \nu(dy) dt \right)^{1/2} \end{aligned}$$

Thus, using (29),

$$\begin{aligned} \mathbb{E}^\mu [(L_s(X_0) - L_{s^\circ}(X_0))^2] & \leq 4C(\mathcal{L}) \int_{\mathcal{X}} \int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(y, t) - s^\circ(y, t)|^2 q_t(x, y) \nu(dy) dt \mu(dx) \\ & = 4C(\mathcal{L}) \int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(y, t) - s^\circ(y, t)|^2 p_t(y) \nu(dy) dt \\ & = 4C(\mathcal{L})\mathbb{E}^\mu [L_s(X_0) - L_{s^\circ}(X_0)]. \end{aligned}$$

■

This result renders the reasoning in Oko et al. (2023, Theorem C.4) valid up to a multiplicative factor that only results in a minor change of constants in their generalisation error upper bound. For completeness we state a corrected version of Oko et al. (2023, Theorem C.4) in the general context of this section with corrected constants.

Theorem 16 *Suppose that $\sup_{s \in \mathcal{S}} \|L_s\|_{\mathcal{X}} \leq C(\mathcal{L})$, where $C(\mathcal{L}) < \infty$. Then, for any $\delta > 0$ such that $\mathcal{N}(\mathcal{L}, \|\cdot\|_{\mathcal{X}}, \delta) \geq 3$, it holds that*

$$\begin{aligned} & \mathbb{E} \left[\int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |\widehat{s}(x, t) - \nabla \log p_t(x)|^2 p_t(x) \nu(dx) dt \right] \\ & \leq 2 \inf_{s \in \mathcal{S}} \int_{\underline{T}}^{\bar{T}} \int_{\mathcal{X}} |s(x, t) - \nabla \log p_t(x)|^2 p_t(x) \nu(dx) dt + \frac{2C(\mathcal{L})}{n} \left(\frac{145}{9} \log \mathcal{N}(\mathcal{L}, \|\cdot\|_{\mathcal{X}}, \delta) + 160 \right) \\ & \quad + 5\delta. \end{aligned}$$

Appendix C. Technical results on neural network approximations

We first consider neural network approximations of products.

Lemma 17 *For any $C \geq 1$ and $m, d \in \mathbb{N}$, there exist neural networks $\varphi_m^{\text{mult}} \in \Phi(L, W, S, B)$ and $\varphi_m^{\text{mult}, d} \in \Phi(L, W_d, d \cdot S, B)$ satisfying*

$$|\varphi_m^{\text{mult}}(x, y) - xy| \leq C2^{-m}, \quad x \in [0, 1], y \in [-C, C].$$

and

$$|\varphi_m^{\text{mult}, d}(x, y) - xy| \leq \sqrt{d}C2^{-m}, \quad x \in [0, 1], y \in [-C, C]^d.$$

Furthermore, $\varphi_m^{\text{mult}}(0, y) = \varphi_m^{\text{mult}}(x, 0) = 0$. The sizes of the networks are evaluated as $L = m + 8$, $S = 58 + 16m$, $B = C$ and

$$W = [2 \quad 3 \quad 3 \quad \underbrace{12 \quad \cdots \quad 12}_{m+2 \text{ times}} \quad 2 \quad 2 \quad 1]^\top, \quad (W_d)_i = \begin{cases} d+1, & \text{if } i = 1 \\ d \cdot W_i & \text{otherwise.} \end{cases}$$

In particular, we have $L, S \lesssim m$, $\|W\|_\infty \lesssim 1$, $\|W_d\|_\infty \lesssim d$ and $B \lesssim C$.

Proof We proceed as in the proof of Oko et al. (2023, Lemma F.6), but adapted to this specific setting. Thus, let $m \in \mathbb{N}$ and $C \geq 0$ be fixed. Then, by Schmidt-Hieber (2020, Lemma A.2), there exists a neural network $\overline{\varphi}_m^{\text{mult}} \in \Phi(m+4, W_0, 24+16m, 1)$, where

$$W_0 = [2 \quad \underbrace{6 \quad \cdots \quad 6}_{m+2 \text{ times}} \quad 1]^\top,$$

satisfying $|\overline{\varphi}_m^{\text{mult}}(x, y) - xy| \leq 2^{-m}$ and $\overline{\varphi}_m^{\text{mult}}(x, 0) = \overline{\varphi}_m^{\text{mult}}(0, y) = 0$ for all $x, y \in [0, 1]$. Thus, for $y \in [-1, 1]$, we have $|\text{sgn}(y)\overline{\varphi}_m^{\text{mult}}(x, |y|) - xy| \leq 2^{-m}$. Note then that

$$\begin{aligned} \text{sgn}(y)\overline{\varphi}_m^{\text{mult}}(x, |y|) &= \sigma_R(\overline{\varphi}_m^{\text{mult}}(x, \sigma_R(y))) - \sigma_R(\overline{\varphi}_m^{\text{mult}}(x, \sigma_R(-y))) \\ &= \overline{\varphi}_m^{\text{mult}, 3} \circ \overline{\varphi}_m^{\text{mult}, 2} \circ \overline{\varphi}_m^{\text{mult}, 1}(x, y), \end{aligned}$$

where

$$\overline{\varphi}_m^{\text{mult}, 1} = I_3 \sigma \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & -1 \end{bmatrix} \in \Phi(1, [2 \quad 3 \quad 3]^\top, 6, 1),$$

$$\bar{\varphi}_m^{\text{mult},3} = [1 \quad -1] \sigma I_2 \in \Phi(1, [2 \quad 2 \quad 1]^\top, 4, 1),$$

and $\bar{\varphi}_m^{\text{mult},2}(x, y, z) = [\bar{\varphi}_m^{\text{mult}}(x, y) \quad \bar{\varphi}_m^{\text{mult}}(x, z)]^\top$. Thus, $\bar{\varphi}_m^{\text{mult},2} \in \Phi(m+4, W_1, 48+32m, 1)$ where W_1 is the parallelization of W_0 with itself, and hence $\text{sgn}(y)\bar{\varphi}_m^{\text{mult}}(x, |y|) := \tilde{\varphi}_m^{\text{mult}}(x, y) \in \Phi(m+8, W, 58+16m, 1)$, where W is as specified in the statement of the lemma. Finally, if $y \in [-C, C]$, we have similarly $|C\tilde{\varphi}_m^{\text{mult}}(x, y/C) - xy| \leq C2^{-m}$, and since these are simply linear transformations (i.e., neural networks of depth 0), it follows that $C\tilde{\varphi}_m^{\text{mult}}(x, y/C) := \varphi_m^{\text{mult}}(x, y) \in \Phi(m+8, W, 58+16m, C)$. Parallelization of φ_m^{mult} a total of d times yields a network $\varphi_m^{\text{mult},d} \in \Phi(m+8, W_d, d(58+16m), C)$ with W_d as specified above, which satisfies

$$\max_{i \in \{1, \dots, d\}} |(\varphi_m^{\text{mult},d}(x, y))_i - xy_i| \leq C2^{-m}, \quad x \in [0, 1], y \in [-C, C]^d,$$

implying the desired result since $|y| \leq \sqrt{d}\|y\|_\infty$ for $y \in \mathbb{R}^d$. \blacksquare

Next, we consider the reciprocals. Note that by not relying on the comparatively slow converging Taylor series of the reciprocal, we achieve a substantially smaller network size than Oko et al. (2023, Lemma F.7), who achieve a network size of $\tilde{\Phi}(m^2, m^3, m^4, 2^{2m})$ (assuming, as they do, that $m = \underline{k} = \bar{k}$).

Lemma 18 *For any $\underline{k}, \bar{k}, m \in \mathbb{N}$ with $m > \bar{k}$, there exists a neural network $\varphi_m^{\text{rec}} \in \Phi(L, W, S, B)$ satisfying*

$$|\varphi_m^{\text{rec}}(x) - \frac{1}{x}| \leq 2^{-m}, \quad x \in [2^{-\underline{k}}, 2^{\bar{k}}].$$

The size of φ_m^{rec} is evaluated as

$$\begin{aligned} L &= (4(\underline{k} + \bar{k}) + 2m + 17) \lceil \log_2(\underline{k} + m + 2) \rceil + 2 \lceil \log_2(\underline{k} + \bar{k}) \rceil + 1, \\ S &= (260 + 68(\underline{k} + \bar{k} + m)) \lceil \log_2(\underline{k} + m + 2) \rceil + 8(\underline{k} + \bar{k}), \end{aligned}$$

$B = 2^{2(\underline{k} + \bar{k})}$ and W is the concatenation of $\lceil \log_2(\underline{k} + m + 2) \rceil$ copies of W_{iter} and W_{init} , where

$$W_{\text{iter}} = \left[2 \quad \underbrace{7 \quad \dots \quad 7}_{2(\underline{k} + \bar{k}) + m + 3 \text{ times}} \quad 2 \quad \underbrace{7 \quad \dots \quad 7}_{2(\underline{k} + \bar{k}) + m + 3 \text{ times}} \quad 2 \right]^\top,$$

and

$$W_{\text{init}} = \left[1 \quad \underline{k} + \bar{k} \quad 2^{\lceil \log_2(\underline{k} + \bar{k}) \rceil} \quad 2^{\lceil \log_2(\underline{k} + \bar{k}) \rceil - 1} \quad 2^{\lceil \log_2(\underline{k} + \bar{k}) \rceil - 1} \quad \dots \quad 2 \quad 2 \quad 1 \right]^\top.$$

In particular, we have $L, S \lesssim (k+m)(\log(k+m))$, $\|W\|_\infty \lesssim k$, where $k = \underline{k} + \bar{k}$.

Proof Let $\underline{k}, \bar{k}, m \in \mathbb{N}$ with $m > \bar{k}$ be fixed. We approximate the reciprocal by Newton–Raphson iterations, adapted to neural networks. In a usual Newton–Raphson scheme approximating x^{-1} for some $x > 0$, we would take x_0 to be some initial approximation and set $x_n = x_{n-1}(2 - xx_{n-1})$ for $n \in \mathbb{N}$. This, however, involves two multiplications by non-constants in each iteration, and as such is not directly accessible by neural networks. To overcome this, suppose that we have access to a neural network φ_m^{iter}

such that $|\varphi_m^{\text{iter}}(x, y) - x(2 - xy)| \leq 2^{-(m+1)}$ as well as a neural network φ^{init} such that $|\varphi^{\text{init}} - x^{-1}| \leq \frac{1}{2}x^{-1}$. Then, setting $\tilde{x}_0 = \varphi^{\text{init}}(x)$ and $\tilde{x}_n = \varphi_m^{\text{iter}}(\tilde{x}_{n-1}, x)$ for $n \in \mathbb{N}$, we claim that $|\tilde{x}_{n_0} - x^{-1}| \leq 2^{-m}$ where $n_0 = \lceil \log_2(\underline{k} + m + 2) \rceil$. To show this, we first find by the recursive definition of \tilde{x}_n that

$$|\tilde{x}_{n+1} - x^{-1}| = |\varphi_m^{\text{iter}}(\tilde{x}_n, x) - x^{-1}| \leq |\tilde{x}_n(2 - x\tilde{x}_n) - x^{-1}| + 2^{-(m+1)} = x(\tilde{x}_n - x^{-1})^2 + 2^{-(m+1)}.$$

Hence, setting $e_n(x) := xe_{n-1}^2(x) + 2^{-(m+1)}$ with $e_0(x) := \frac{1}{2}x^{-1}$, we find that $|\tilde{x}_n - x^{-1}| \leq e_n(x)$ for all $n \in \mathbb{N}$. To show that $e_{n_0}(x) \leq 2^{-m}$ on $[2^{-\underline{k}}, 2^{\bar{k}}]$, we introduce $d_n(x) := e_n(x) - \frac{1}{x}2^{-2^n}$ for $n \in \mathbb{N}$ and $x \in [2^{-\underline{k}}, 2^{\bar{k}}]$. Since $\frac{1}{x}2^{-2^n} = x(\frac{1}{x}2^{-2^{n-1}})^2$, d_n satisfies the recursion $d_n(x) = xd_{n-1}^2(x) + 2^{-2^{n+1}}d_{n-1}(x) + 2^{-(m+1)}$ with $d_0 \equiv 0$. Furthermore, we claim that d_n is non-decreasing and convex. Clearly, this is true for d_0 , and by the recursion above we have

$$d'_n(x) = \frac{d}{dx}[xd_{n-1}^2(x) + 2^{-2^{n+1}}d_{n-1}(x)] = d_{n-1}^2(x) + 2xd'_{n-1}(x)d_{n-1}(x) + 2^{-2^{n+1}}d'_{n-1}(x).$$

By induction, this is non-negative and non-decreasing, so d_n is non-decreasing and convex as claimed. Since $e_n(x) = \frac{1}{x}2^{-2^n} + d_n(x)$, it follows that e_n is also convex, and we thus only need to check that $e_{n_0}(2^{-\underline{k}}) \vee e_{n_0}(2^{\bar{k}}) \leq 2^{-m}$. To this end, it is clear that for $n \geq 2$, we have $d_n(2^{-\underline{k}}) \geq d_{n+1}(2^{-\underline{k}})$, and some tedious but straightforward calculations show that

$$d_3(2^{-\underline{k}}) = \left(1 + 2^{-\underline{k}}2^{-(m+1)}\right)\left(2^{-\underline{k}}2^{-(m+1)} + \frac{17}{16}\right)^2 + \frac{1}{256}\left(2^{-\underline{k}}2^{-(m+1)} + \frac{17}{16}\right)2^{-(m+1)}.$$

A very rough estimate yields from this that $d_n(2^{-\underline{k}}) \leq \frac{3}{2}2^{-(m+1)} = 2^{-m} - 2^{-(m+2)}$ for $n \geq 3$. Now, since $\underline{k} + m + 2 > 4$, we have $n_0 \geq 3$, and so

$$e_{n_0}(2^{-\underline{k}}) = 2^{\underline{k}}2^{-2^{n_0}} + d_{n_0}(2^{-\underline{k}}) \leq 2^{\underline{k}}2^{-(\underline{k}+m+2)} + 2^{-m} - 2^{-(m+2)} = 2^{-m}.$$

As for $e_{n_0}(2^{\bar{k}})$, since $\bar{k} \leq m - 1$, the same convexity argument implies that it suffices to check that $e_{n_0}(2^{m-1}) \leq 2^{-m}$. But we clearly have that $e_0(2^{m-1}) = 2^{-m}$, and since

$$2^{m-1}(2^{-m})^2 + 2^{-(m+1)} = 2 \cdot 2^{-(m+1)} = 2^{-m},$$

it follows that $e_n(2^{m-1}) = 2^{-m}$ for all $n \in \mathbb{N}$. Now, let us construct this network. For notation, let $k = \underline{k} + \bar{k}$. We begin by constructing φ^{init} as a piecewise linear function. In particular, for $i = 0, \dots, k$, let $p_i = 2^{i-\underline{k}}$, and for $i = 1, \dots, k$ let ℓ_i be the linear interpolation between (p_{i-1}, p_{i-1}^{-1}) and (p_i, p_i^{-1}) , i.e.,

$$\ell_i(x) = \frac{p_i^{-1} - p_{i-1}^{-1}}{p_i - p_{i-1}}x - \left(\frac{p_i^{-1} - p_{i-1}^{-1}}{p_i - p_{i-1}}p_i - \frac{1}{p_i}\right) = \frac{1}{p_i}\left(3 - \frac{2}{p_i}x\right), \quad x \in [p_{i-1}, p_i].$$

Note that by convexity of the reciprocal function, we have $\ell_i(x) \geq x^{-1}$ for all $x \in [p_{i-1}, p_i]$, and hence $E_i(x) := |\ell_i(x) - \frac{1}{x}| = \frac{1}{p_i}\left(3 - \frac{2}{p_i}x\right) - \frac{1}{x}$. We then have that $E'_i(x) = \frac{1}{x^2} - \frac{2}{p_i^2}$, and since E_i has exactly one extremal point and this is a maximum, we have that the error is maximized when $x = \frac{p_i}{\sqrt{2}}$ and that the corresponding error is $E_i\left(\frac{p_i}{\sqrt{2}}\right) = \frac{3-2\sqrt{2}}{p_i} \leq \frac{1}{2x}$. Thus, since ℓ_i is merely an affine transformation and hence clearly representable as a neural

network, ℓ_i satisfies the properties of φ^{init} on $[p_{i-1}, p_i]$. Also, if $x \in [p_{i-1}, p_i]$, then $\ell_j(x) < x^{-1}$ for all $j \neq i$, and so $\varphi^{\text{init}}(x) := \max_{j=1, \dots, k} \ell_j(x)$ works. Since $a \vee b = a + \sigma(b - a)$ for $a, b \geq 0$, the function $(y_1, \dots, y_k) \mapsto \max_{j=1, \dots, k} y_j$ is realized by a network $\varphi^{\text{max}} \in \Phi(2n_k, W_{\text{max}}, 5k, 1)$, where $n_k = \lceil \log_2(k) \rceil$ and

$$W_{\text{max}} = [k \quad 2^{n_k} \quad 2^{n_k-1} \quad 2^{n_k-1} \quad \dots \quad 2 \quad 2 \quad 1]^\top.$$

Thus, by parallelizing, we have $\varphi^{\text{init}} \in \Phi(2(n_k + 1), [1 \quad k \quad W_{\text{max}}^\top]^\top, 8k, 2^{2k})$. Next, we construct φ_m^{iter} . Similarly to the proof of Lemma 17, a small modification of Schmidt-Hieber (2020, Lemma A.2) yields a network $\varphi_l^{\text{mult}} \in \Phi(l + 4, W_{\text{mult}}, 24 + 16l, 2^k)$ with

$$W_{\text{mult}} = [2 \quad \underbrace{6 \quad \dots \quad 6}_{l+2 \text{ times}} \quad 1]^\top,$$

satisfying $|\varphi_l^{\text{mult}}(x, y) - xy| \leq 2^{k+1-l}$ for all $x \in [0, 2]$ and $y \in [0, 2^k]$. Since we always have $x \wedge \tilde{x}_n \leq 2$ and $|2 - x\tilde{x}_n| \leq 2$, setting $\varphi_m^{\text{iter}}(x, y) = \varphi_l^{\text{mult}}(x, 2 - \varphi_l^{\text{mult}}(x, y))$ with $l \geq 2k + m + 3$ works. Parallelizing with the identity, we can now simply chain these together without altering the size of the network other than increasing the width of each layer by 1. That is, setting $\tilde{\varphi}^{\text{init}}(x) := [\varphi^{\text{init}}(x) \quad x]^\top$ and $\tilde{\varphi}_m^{\text{iter}}(x, y) := [\varphi_m^{\text{iter}}(x, y) \quad x]^\top$,

$$\varphi_m^{\text{rec}}(x) := \varphi_m^{\text{iter}} \circ \underbrace{\tilde{\varphi}_m^{\text{iter}} \circ \dots \circ \tilde{\varphi}_m^{\text{iter}}}_{n_0-1 \text{ times}} \circ \tilde{\varphi}^{\text{init}}(x),$$

yields the desired network. ■

Proof of Lemma 12 We first approximate $t \mapsto \sqrt{t}$ by a piecewise linear function. To this end, let $t_i = 2^{-m} + i/m$ for $i = 0, \dots, m$ such that $t_m \geq 1$. Next, for $j = 1, \dots, m$ and $t \in [t_{j-1}, t_j]$, set

$$\ell_j(t) = m(\sqrt{t_j} - \sqrt{t_{j-1}})(t - t_{j-1}) + \sqrt{t_{j-1}},$$

such that ℓ_j is the linear interpolation between the points $(t_{j-1}, \sqrt{t_{j-1}})$ and $(t_j, \sqrt{t_j})$. Since $t \mapsto \sqrt{t}$ is concave, it follows that, for $t \in [t_{i-1}, t_i]$, we have $\ell_i(t) \leq \sqrt{t}$ while $\ell_j(t) \geq \sqrt{t}$ for all $j \neq i$. Hence, setting $\ell(t) = \min_{j \in [m]} \ell_j(t)$ yields the desired piecewise linear approximation, and since both affine functions and minima are exactly representable as neural networks, we find that $\ell \in \tilde{\Phi}(\log m, m, m, \sqrt{m})$. We claim that setting $\varphi^{\text{cap}} = \varphi_1^{\text{rec}} \circ \ell$ yields the desired network. Indeed, for the size of the network, note that since $\ell(t) \in [2^{-m/2}, 1]$ for all $t \in [2^{-m}, 1]$, we have $\varphi_1^{\text{rec}} \in \tilde{\Phi}(m \log m, m, m \log m, 2^{m/2})$, and since this dominates the size of ℓ , it follows that φ^{cap} is of the same size as φ_1^{rec} .

As for the claim that $\varphi^{\text{cap}}(t) \asymp \frac{1}{\sqrt{t}}$, first note that

$$\varphi^{\text{cap}}(t) = \frac{1}{\sqrt{t}} + \left(\varphi_1^{\text{rec}} \circ \ell(t) - \frac{1}{\ell(t)} \right) + \left(\frac{1}{\ell(t)} - \frac{1}{\sqrt{t}} \right),$$

where the second term is in $[-1/2, 1/2]$, whence

$$\frac{1}{2\sqrt{t}} \leq \frac{1}{\sqrt{t}} + \left(\varphi_1^{\text{rec}} \circ \ell(t) - \frac{1}{\ell(t)} \right) \leq \frac{3}{2\sqrt{t}}.$$

For the last term, note that since $\ell(t_i) = \sqrt{t_i}$ for $i = 0, \dots, m$, we have

$$\left| \frac{1}{\ell(t)} - \frac{1}{\sqrt{t}} \right| \leq \frac{1}{2^{m/2}} - \frac{1}{\sqrt{2^m + m}} < \frac{1}{4},$$

whereby

$$\frac{1}{4\sqrt{t}} \leq \varphi^{\text{cap}}(t) \leq \frac{7}{4\sqrt{t}},$$

as desired. ■

References

- B. D. O. Anderson. Reverse-time diffusion equation models. *Stochastic Process. Appl.*, 12(3):313–326, 1982. ISSN 0304-4149,1879-209X. doi: 10.1016/0304-4149(82)90051-5. URL [https://doi.org/10.1016/0304-4149\(82\)90051-5](https://doi.org/10.1016/0304-4149(82)90051-5).
- M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 214–223. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/arjovsky17a.html>.
- I. Azangulov, G. Deligiannidis, and J. Rousseau. Convergence of diffusion models under the manifold hypothesis in high-dimensions, 2024. URL <https://arxiv.org/abs/2409.18804>.
- D. Bakry and Z. M. Qian. Harnack inequalities on a manifold with positive or negative Ricci curvature. *Rev. Mat. Iberoamericana*, 15(1):143–179, 1999. ISSN 0213-2230. doi: 10.4171/RMI/253. URL <https://doi.org/10.4171/RMI/253>.
- P. L. Bartlett and S. Mendelson. Empirical minimization. *Probab. Theory Related Fields*, 135(3):311–334, 2006. ISSN 0178-8051,1432-2064. doi: 10.1007/s00440-005-0462-3. URL <https://doi.org/10.1007/s00440-005-0462-3>.
- J. Benton, Y. Shi, V. De Bortoli, G. Deligiannidis, and A. Doucet. From denoising diffusions to denoising Markov models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 86(2):286–301, 2024. ISSN 1369-7412,1467-9868. doi: 10.1093/jrsssb/qkae005. URL <https://doi.org/10.1093/jrsssb/qkae005>.
- P. Cattiaux. Time reversal of diffusion processes with a boundary condition. *Stochastic Process. Appl.*, 28(2):275–292, 1988. ISSN 0304-4149,1879-209X. doi: 10.1016/0304-4149(88)90101-9. URL [https://doi.org/10.1016/0304-4149\(88\)90101-9](https://doi.org/10.1016/0304-4149(88)90101-9).
- M. Chae. Rates of convergence for nonparametric estimation of singular distributions using generative adversarial networks, 2022. URL <https://arxiv.org/abs/2202.02890>.
- M. Chen, W. Liao, H. Zha, and T. Zhao. Distribution approximation and statistical estimation guarantees of generative adversarial networks, 2022. URL <https://arxiv.org/abs/2002.03938>.

- M. Chen, K. Huang, T. Zhao, and M. Wang. Score approximation, estimation and distribution recovery of diffusion models on low-dimensional data. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 4672–4712. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/chen23o.html>.
- S. Chewi, J. Niles-Weed, and P. Rigollet. *Statistical Optimal Transport*, volume 2364 of *Lecture Notes in Mathematics*. Springer, Cham, 2025. ISBN 978-3-031-85159-9; 978-3-031-85160-5. doi: 10.1007/978-3-031-85160-5. URL <https://doi.org/10.1007/978-3-031-85160-5>. École d’Été de Probabilités de Saint-Flour XLIX – 2019.
- E. B. Davies. *Heat kernels and spectral theory*, volume 92 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1989. ISBN 0-521-36136-2. doi: 10.1017/CBO9780511566158. URL <https://doi.org/10.1017/CBO9780511566158>.
- N. Fishman, L. Klarner, V. D. Bortoli, E. Mathieu, and M. J. Hutchinson. Diffusion models for constrained domains. *Transactions on Machine Learning Research*, 2023a. ISSN 2835-8856. URL <https://openreview.net/forum?id=xuWTFQ4VG0>.
- N. Fishman, L. Klarner, E. Mathieu, M. Hutchinson, and V. De Bortoli. Metropolis sampling for constrained diffusion models. In *Advances in Neural Information Processing Systems*, volume 36, pages 62296–62331, 2023b. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/c47bfcc8e2eccdc540fad1e25f13aa4d-Paper-Conference.pdf.
- M. Giordano and S. Wang. Statistical algorithms for low-frequency diffusion data: a PDE approach. *Ann. Statist.*, 53(3):1150–1175, 2025. ISSN 0090-5364,2168-8966. doi: 10.1214/25-aos2496. URL <https://doi.org/10.1214/25-aos2496>.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144, 2020.
- U. G. Haussmann and E. Pardoux. Time reversal of diffusions. *Ann. Probab.*, 14(4):1188–1205, 1986. ISSN 0091-1798,2168-894X. doi: 10.1214/aop/1176992362. URL <https://doi.org/10.1214/aop/1176992362>.
- J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf.
- M. Hoffmann and K. Ray. Nonparametric Bayesian estimation in a multidimensional diffusion model with high frequency data. *Probab. Theory Relat. Fields*, 2024. doi: 10.1007/s00440-024-01317-w. URL <https://doi.org/10.1007/s00440-024-01317-w>.
- S. Itô. *Diffusion equations*, volume 114 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 1992. ISBN 0-8218-4570-5. doi: 10.1090/mmono/114.

- I. Karatzas and S. E. Shreve. *Brownian motion and stochastic calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1991. ISBN 0-387-97655-8. doi: 10.1007/978-1-4612-0949-2. URL <https://doi.org/10.1007/978-1-4612-0949-2>.
- D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations*, 2014.
- P. Li and S.-T. Yau. On the parabolic kernel of the Schrödinger operator. *Acta Math.*, 156(3-4):153–201, 1986. ISSN 0001-5962,1871-2509. doi: 10.1007/BF02399203. URL <https://doi.org/10.1007/BF02399203>.
- T. Liang. How well generative adversarial networks learn distributions. *J. Mach. Learn. Res.*, 22:Paper No. 228, 41, 2021. ISSN 1532-4435,1533-7928.
- A. Lou and S. Ermon. Reflected diffusion models. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 22675–22701. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/lou23a.html>.
- R. Nickl. Consistent inference for diffusions from low frequency measurements. *Ann. Statist.*, 52(2):519–549, 2024. ISSN 0090-5364,2168-8966. doi: 10.1214/24-aos2357. URL <https://doi.org/10.1214/24-aos2357>.
- S. Nowozin, B. Cseke, and R. Tomioka. f-GAN: training generative neural samplers using variational divergence minimization. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 271–279, Red Hook, NY, USA, 2016. Curran Associates Inc. ISBN 9781510838819.
- K. Oko, S. Akiyama, and T. Suzuki. Diffusion models are minimax optimal distribution estimators. In *International Conference on Machine Learning*, 2023. URL <https://openreview.net/forum?id=0Ryo7fxcIA>.
- E. M. Ouhabaz. *Analysis of heat equations on domains*, volume 31 of *London Mathematical Society Monographs Series*. Princeton University Press, Princeton, NJ, 2005. ISBN 0-691-12016-1.
- G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021. URL <http://jmlr.org/papers/v22/19-1028.html>.
- N. Puchkin, S. Samsonov, D. Belomestny, E. Moulines, and A. Naumov. Rates of convergence for density estimation with generative adversarial networks. *Journal of Machine Learning Research*, 25(29):1–47, 2024. URL <http://jmlr.org/papers/v25/23-0062.html>.
- Z. M. Qian. Gradient estimates and heat kernel estimate. *Proc. Roy. Soc. Edinburgh Sect. A*, 125(5):975–990, 1995. ISSN 0308-2105,1473-7124. doi: 10.1017/S0308210500022599. URL <https://doi.org/10.1017/S0308210500022599>.

- Y. Ren, G. M. Rotskoff, and L. Ying. A unified approach to analysis and design of denoising Markov models, 2025. URL <https://arxiv.org/abs/2504.01938>. arXiv preprint, arXiv:2504.01938.
- D. Rezende and S. Mohamed. Variational inference with normalizing flows. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1530–1538, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/rezende15.html>.
- J. Schmidt-Hieber. Nonparametric regression using deep neural networks with ReLU activation function. *Ann. Statist.*, 48(4):1875–1897, 2020. ISSN 0090-5364,2168-8966. doi: 10.1214/19-AOS1875. URL <https://doi.org/10.1214/19-AOS1875>.
- N. Schreuder, V.-E. Brunel, and A. Dalalyan. Statistical guarantees for generative models without domination. In *Proceedings of the 32nd International Conference on Algorithmic Learning Theory*, volume 132 of *Proceedings of Machine Learning Research*, pages 1051–1071. PMLR, 16–19 Mar 2021. URL <https://proceedings.mlr.press/v132/schreuder21a.html>.
- J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.
- Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/3001ef257407d5a371a96dcd947c7d93-Paper.pdf.
- Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=PxtTIG12RRHS>.
- A. Stéphanovitch, E. Aamari, and C. Levrard. Wasserstein generative adversarial networks are minimax optimal distribution estimators. *Ann. Statist.*, 52(5):2167–2193, 2024a. ISSN 0090-5364,2168-8966. doi: 10.1214/24-aos2430. URL <https://doi.org/10.1214/24-aos2430>.
- A. Stéphanovitch, U. Tanielian, B. Cadre, N. Klutchnikoff, and G. Biau. Optimal 1-Wasserstein distance for WGANs. *Bernoulli*, 30(4):2955–2978, 2024b. ISSN 1350-7265,1573-9759. doi: 10.3150/23-bej1701. URL <https://doi.org/10.3150/23-bej1701>.
- A. Stéphanovitch, E. Aamari, and C. Levrard. Generalization bounds for score-based generative models: a synthetic proof, 2025. URL <https://arxiv.org/abs/2507.04794>. arXiv preprint, arXiv:2507.04794.

- T. Suzuki. Adaptivity of deep ReLU network for learning in Besov and mixed smooth Besov spaces: optimal rate and curse of dimensionality. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=H1ebTsActm>.
- E. G. Tabak and C. V. Turner. A family of nonparametric density estimation algorithms. *Comm. Pure Appl. Math.*, 66(2):145–164, 2013. ISSN 0010-3640,1097-0312. doi: 10.1002/cpa.21423. URL <https://doi.org/10.1002/cpa.21423>.
- R. Tang and Y. Yang. Minimax rate of distribution estimation on unknown submanifolds under adversarial losses. *Ann. Statist.*, 51(3):1282–1308, 2023. ISSN 0090-5364,2168-8966. doi: 10.1214/23-aos2291. URL <https://doi.org/10.1214/23-aos2291>.
- R. Tang and Y. Yang. Adaptivity of diffusion models to manifold structures. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 1648–1656. PMLR, 02–04 May 2024. URL <https://proceedings.mlr.press/v238/tang24a.html>.
- M. Telgarsky. Neural networks and rational functions. In *Proceedings of the 34th International Conference on Machine Learning*, Proceedings of Machine Learning Research, pages 3387–3393. PMLR, 2017.
- L. N. Trefethen. *Approximation theory and approximation practice*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, extended ed. edition, 2020. ISBN 978-1-611975-93-2. doi: 10.1137/1.978161197594.
- H. Triebel. *Theory of function spaces*. Modern Birkhäuser Classics. Birkhäuser/Springer Basel AG, Basel, 2010. ISBN 978-3-0346-0415-4; 978-3-0346-0416-1.
- C. Villani. *Optimal transport*, volume 338 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2009. ISBN 978-3-540-71049-3. doi: 10.1007/978-3-540-71050-9. URL <https://doi.org/10.1007/978-3-540-71050-9>. Old and new.
- P. Vincent. A connection between score matching and denoising autoencoders. *Neural Comput.*, 23(7):1661–1674, 2011. ISSN 0899-7667,1530-888X. doi: 10.1162/NECO_a_00142. URL https://doi.org/10.1162/NECO_a_00142.
- K. Yakovlev and N. Puchkin. Generalization error bound for denoising score matching under relaxed manifold assumption. In *Proceedings of 38th Conference on Learning Theory*, volume 291 of *Proceedings of Machine Learning Research*, pages 5824–5891. PMLR, 2025.
- Y. Yang and A. Barron. Information-theoretic determination of minimax rates of convergence. *Ann. Statist.*, 27(5):1564–1599, 1999. ISSN 0090-5364,2168-8966. doi: 10.1214/aos/1017939142. URL <https://doi.org/10.1214/aos/1017939142>.
- D. Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Netw.*, 94:103–114, 2017. ISSN 0893-6080. doi: <https://doi.org/10.1016/j.neunet.2017.07.002>. URL <https://www.sciencedirect.com/science/article/pii/S0893608017301545>.

- K. Zhang, C. H. Yin, F. Liang, and J. Liu. Minimax optimality of score-based diffusion models: beyond the density lower bound assumptions. In *Proceedings of the 41st International Conference on Machine Learning*, Proceedings of Machine Learning Research. PMLR, 2024.
- J. Zhao, M. Mathieu, and Y. LeCun. Energy-based generative adversarial networks. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=ryh9pmcee>.