

Rates of convergence for density estimation with generative adversarial networks

Nikita Puchkin

NPUCHKIN@HSE.RU

HSE University, Russian Federation

Institute for Information transmission Problems, Russian Federation

Sergey Samsonov

SVSAMSONOV@HSE.RU

HSE University, Russian Federation

Institute for Information transmission Problems, Russian Federation

Denis Belomestny

DENIS.BELOMESTNY@UNI-DUE.DE

Duisburg-Essen University, Germany

HSE University, Russian Federation

Eric Moulines

ERIC.MOULINES@POLYTECHNIQUE.EDU

École Polytechnique, France

Mohamed Bin Zayed University of AI, United Arab Emirates

Alexey Naumov

ANAUMOV@HSE.RU

HSE University, Russian Federation

Steklov Mathematical Institute of Russian Academy of Sciences, Russian Federation

Editor: Aarti Singh

Abstract

In this work we undertake a thorough study of the non-asymptotic properties of the vanilla generative adversarial networks (GANs). We prove an oracle inequality for the Jensen-Shannon (JS) divergence between the underlying density \mathbf{p}^* and the GAN estimate with a significantly better statistical error term compared to the previously known results. The advantage of our bound becomes clear in application to nonparametric density estimation. We show that the JS-divergence between the GAN estimate and \mathbf{p}^* decays as fast as $(\log n/n)^{2\beta/(2\beta+d)}$, where n is the sample size and β determines the smoothness of \mathbf{p}^* . This rate of convergence coincides (up to logarithmic factors) with minimax optimal for the considered class of densities.

Keywords: generative model, oracle inequality, Jensen-Shannon risk, minimax rates, nonparametric density estimation.

1. Introduction

Let X_1, \dots, X_n be i.i.d. random elements with values in $\mathsf{X} \subseteq \mathbb{R}^d$ drawn from a distribution P^* . We assume that P^* admits a density \mathbf{p}^* with respect to a dominating measure μ . The measure μ is not necessarily absolutely continuous with respect to the Lebesgue measure, it can be the counting measure or the Hausdorff measure on a low-dimensional manifold as well. Our goal is to estimate \mathbf{p}^* based on a finite sample. The problem of density estimation was extensively studied in the literature and encounters numerous approaches such as kernel (see, e.g., (Tsybakov, 2008, Section 1.2) and (McDonald, 2017)) and k-

nearest neighbors density estimators (Dasgupta and Kpotufe, 2014), wavelet thresholding (Donoho et al., 1996), and aggregation (Rakhlin et al., 2005; Rigollet and Tsybakov, 2006; Bunea et al., 2007; Butucea et al., 2017; Dalalyan and Sebbar, 2018). Recently, Goodfellow et al. (2014) have introduced a novel approach for a related problem of generative modeling called generative adversarial networks (or simply GANs). A generative adversarial network consists of a generator and a discriminator. Given a known easy-to-sample distribution on a latent space Y with a density ϕ , the generator $g : \mathsf{Y} \rightarrow \mathsf{X}$ takes i.i.d. samples Y_1, \dots, Y_n from ϕ and produces fake ones $g(Y_1), \dots, g(Y_n)$. The goal of the discriminator D is to distinguish between the real samples X_1, \dots, X_n and $g(Y_1), \dots, g(Y_n)$. Usually, D can be thought of as a map $\mathsf{X} \mapsto (0, 1)$, where larger values correspond to higher confidence that an input variable is drawn from \mathbf{p}^* . In practice, both generator and discriminator usually belong to some parametric families (for example, to classes of neural networks). Let us fix positive integers $d_{\mathcal{G}}$ and $d_{\mathcal{D}}$ and some compact sets $\mathsf{W} \subset \mathbb{R}^{d_{\mathcal{G}}}$ and $\Theta \subset \mathbb{R}^{d_{\mathcal{D}}}$. In our paper we assume that

$$g \in \mathcal{G} = \{g_w : w \in \mathsf{W}\} \quad \text{and} \quad D \in \mathcal{D} = \{D_\theta : \theta \in \Theta\}.$$

As a byproduct of the described generative approach, GANs also provide an implicit density estimate for \mathbf{p}^* . Indeed, if a statistician manages to find a good generator g_w , then the density of $g_w(Y)$ is a reasonable estimate of \mathbf{p}^* . In (Goodfellow et al., 2014), the authors suggested to solve the following minimax problem, also called *vanilla GAN*:

$$\min_{w \in \mathsf{W}} \max_{\theta \in \Theta} \left\{ L(w, \theta) := \frac{1}{2} \mathbb{E}_{X \sim \mathbf{p}^*} \log D_\theta(X) + \frac{1}{2} \mathbb{E}_{Y \sim \phi} \log(1 - D_\theta(g_w(Y))) \right\}. \quad (1.1)$$

Here and further in the paper, \log stands for the natural logarithm. The intuition behind (1.1) is that if \mathcal{D} contained all measurable functions on X with values in $(0, 1)$, the minimax problem (1.1) would reduce to (see (Goodfellow et al., 2014, Theorem 1))

$$\min_{w \in \mathsf{W}} [\text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2], \quad (1.2)$$

where JS is the Jensen-Shannon divergence (see (2.1) for the definition), and \mathbf{p}_w is the density of $g_w(Y)$ with $Y \sim \phi$. Unfortunately, since \mathcal{D} is a parametric class (and, hence, it cannot contain all the measurable functions), the actual value of $\max_{\theta \in \Theta} L(w, \theta)$ differs from $(\text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2)$. Fortunately, the gap between $\max_{\theta \in \Theta} L(w, \theta)$ and $(\text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2)$ may be rather small for a proper class of discriminators. Since the true distribution \mathbf{p}^* in (1.1) is unknown, we consider a plug-in estimate $\mathbf{p}_{\hat{w}}$, where \hat{w} is a solution of the optimization problem

$$\hat{w} \in \operatorname{argmin}_{w \in \mathsf{W}} \max_{\theta \in \Theta} L_n(w, \theta), \quad (1.3)$$

and

$$L_n(w, \theta) = \frac{1}{2n} \sum_{i=1}^n \log D_\theta(X_i) + \frac{1}{2n} \sum_{j=1}^n \log(1 - D_\theta(g_w(Y_j))) \quad (1.4)$$

is the empirical version of the functional $L(w, \theta)$ defined in (1.1).

GANs provide a flexible tool for sampling from an unknown distribution, and they have recently become extremely popular among practitioners. Using deep neural network classes

\mathcal{G} and \mathcal{D} , one can reach the state-of-the-art generative performance in many challenging tasks, including image super-resolution (Ledig et al., 2017), video synthesis (Kim et al., 2018), and many others. Various GAN formulations were later proposed by varying the divergence measure in (1.2). For instance, f-GAN (Nowozin et al., 2016) generalized vanilla GAN by minimizing a general f-divergence; Wasserstein GAN (WGAN) (Arjovsky et al., 2017) considered the first-order Wasserstein (Kantorovich) distance (W_1 distance); MMD-GAN (Dziugaite et al., 2015) was based on the maximum mean discrepancy; energy-based GAN (Zhao et al., 2016) minimizes the total variation distance as discussed in (Arjovsky et al., 2017); Quadratic GAN (Feizi et al., 2020) aimed to find the distribution minimizing the second-order Wasserstein (Kantorovich) distance.

The empirical success of GANs motivated many researchers to analyze their theoretical properties. For example, in (Biau et al., 2021; Schreuder et al., 2021), the authors carried out theoretical analysis of WGANs. In (Biau et al., 2021), the authors obtained upper bounds for the excess risk of GANs for parametric classes (including the class of neural networks). In contrast to Biau et al. (2021), Schreuder et al. (2021) considered generative models based on β times differentiable transformations of the d -dimensional unit hypercube and derived rates of convergence of order $O(n^{-\beta/d} \vee n^{-1/2})$ for the corresponding W_1 distance. Liang (2021) used results from the empirical process theory to prove upper bounds for Sobolev GANs (i.e., when generators and discriminators belong to Sobolev classes \mathcal{W}^α and \mathcal{W}^β , respectively), MMD GANs, WGANs, and vanilla GANs. In the case of Sobolev GANs, the obtained rate $n^{-(\alpha+\beta)/(2\beta+d)} \vee n^{-1/2}$ for the corresponding integral probability metric (IPM) is shown to be minimax optimal. The research of Liang (2021) was continued in the works of Singh and Póczos (2018); Uppal et al. (2019); Luise et al. (2020); Chen et al. (2022); Vardanyan et al. (2023) where the authors studied the performance of GANs in terms of different IPM losses and the Sinkhorn divergence. The vanilla GANs were studied in (Liang, 2021; Biau et al., 2020; Asatryan et al., 2023). However, the rates of convergence for them in terms of Jensen-Shannon divergence are not yet fully understood. Biau et al. (2020) and Asatryan et al. (2023) improved approximation terms as compared to Liang (2021), but it is not clear whether the rates obtained in (Biau et al., 2020; Asatryan et al., 2023) are minimax optimal. In this work, we provide a refined analysis of the theoretical properties of vanilla GANs and derive minimax optimal rates.

Contributions. Our contributions can be summarized as follows.

- We prove (Theorem 1) a sharp oracle inequality for the case when the classes \mathcal{G} and \mathcal{D} are general parametric classes, which significantly improves the existing inequalities from the works of Biau et al. (2020) and Asatryan et al. (2023).
- We apply the result of Theorem 1 to a nonparametric density estimation problem. Choosing \mathcal{G} and \mathcal{D} as classes of neural networks of appropriate architectures with ReQU activation functions, we derive the rates of convergence for the estimate $\mathbf{p}_{\hat{w}}$ to the true density \mathbf{p}^* in terms of the Jensen-Shannon divergence. Namely, we show that, with probability at least $1 - \delta$, it holds that

$$\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) \lesssim \left(\frac{\log n}{n}\right)^{\frac{2\beta}{2\beta+d}} + \frac{\log(1/\delta)}{n}, \quad (\text{Theorem 2})$$

provided that the actual density \mathbf{p}^* is the density of a random variable $g^*(Y)$ with Y being uniformly distributed on $[0, 1]^d$, for a smooth invertible transform g^* not necessary belonging to \mathcal{G} . We also discuss that previously known bounds do not yield this rate of convergence.

- We show that the result of Theorem 2 is minimax optimal up to a logarithmic factor. Namely, we prove that for any estimate $\hat{\mathbf{p}}$, it holds that

$$\sup_{\mathbf{p}^*} \mathbb{E} \text{JS}(\hat{\mathbf{p}}, \mathbf{p}^*) \gtrsim n^{-\frac{2\beta}{2\beta+d}}, \quad (\text{Theorem 3})$$

where the supremum is taken with respect to the densities \mathbf{p}^* satisfying the same regularity assumptions as in Theorem 2. Hence, our results imply the minimax optimality of GANs (up to logarithmic factors) in the context of nonparametric density estimation.

2. Preliminaries and notations

Kullback-Leibler and Jensen-Shannon divergences. Let $\Omega \subset \mathbb{R}^d$ be a bounded domain. For two probability measures on a measurable space $(\Omega, \mathcal{B}(\Omega))$ with Lebesgue densities \mathbf{p} and \mathbf{q} , respectively, we define the Kullback-Leibler divergence between them as

$$\text{KL}(\mathbf{p}, \mathbf{q}) = \begin{cases} \int \mathbf{p}(x) \log(\mathbf{p}(x)/\mathbf{q}(x)) d\mu, & \text{if } \mathbf{p} \ll \mathbf{q}, \\ +\infty, & \text{otherwise.} \end{cases}$$

Here and further in this paper, \log stands for the natural logarithm. By $\text{JS}(\mathbf{p}, \mathbf{q})$, we denote the Jensen-Shannon divergence

$$\text{JS}(\mathbf{p}, \mathbf{q}) = \frac{1}{2} \text{KL}\left(\mathbf{p}, \frac{\mathbf{p} + \mathbf{q}}{2}\right) + \frac{1}{2} \text{KL}\left(\mathbf{q}, \frac{\mathbf{p} + \mathbf{q}}{2}\right). \quad (2.1)$$

Norms. For a matrix A and a vector v , we denote by $\|A\|_\infty$ and $\|v\|_\infty$ the maximal absolute value of entries of A and v , respectively. $\|A\|_0$ and $\|v\|_0$ shall stand for the number of non-zero entries of A and v , respectively. Finally, the Frobenius norm and operator norm of A are denoted by $\|A\|_F$ and $\|A\|$, respectively, and the Euclidean norm of v is denoted by $\|v\|$. For $x \in \mathbb{R}^d$ and $r > 0$ we write $\mathcal{B}(x, r) = \{y \in \mathbb{R}^d, \|y - x\| \leq r\}$. For a function $f : \Omega \rightarrow \mathbb{R}^d$, we set

$$\begin{aligned} \|f\|_{L_\infty(\Omega)} &= \sup_{x \in \Omega} \|f(x)\|, \\ \|f\|_{L_2(\Omega)} &= \left\{ \int_{\Omega} \|f(x)\|^2 d\mu \right\}^{1/2}, \end{aligned}$$

and

$$\|f\|_{L_2(\mathbf{p}, \Omega)} = \left(\int_{\Omega} \|f(x)\|^2 \mathbf{p}(x) d\mu \right)^{1/2}.$$

Sometimes, we omit the domain Ω in the notations $L_\infty(\Omega)$, $L_2(\Omega)$, $L_2(\mathbf{p}, \Omega)$ and simply write L_∞ , L_2 , and $L_2(\mathbf{p})$, respectively, if there is no ambiguity.

Smoothness classes. For any $s \in \mathbb{N}$, the function space $C^s(\Omega)$ consists of those functions over the domain Ω which have partial derivatives up to order s in Ω , and these derivatives are bounded and continuous in Ω . Formally,

$$C^s(\Omega) = \{f : \Omega \rightarrow \mathbb{R}^m : \|f\|_{C^s} := \max_{|\gamma| \leq s} \|D^\gamma f\|_{L^\infty(\Omega)} < \infty\},$$

where, for any multi-index $\gamma = (\gamma_1, \dots, \gamma_d) \in \mathbb{N}_0^d$, the partial differential operator D^γ is defined as

$$D^\gamma f_i = \frac{\partial^{|\gamma|} f_i}{\partial x_1^{\gamma_1} \dots \partial x_d^{\gamma_d}}, \quad i \in \{1, \dots, m\}, \quad \text{and} \quad \|D^\gamma f\|_{L^\infty(\Omega)} = \max_{1 \leq i \leq m} \|D^\gamma f_i\|_{L^\infty(\Omega)}.$$

Here we have written $|\gamma| = \sum_{i=1}^d \gamma_i$ for the order of D^γ . To avoid confusion between multi-indices and scalars, we reserve the bold font for the former ones. For the matrix of first derivatives, we use the usual notation $\nabla f = (\partial f_i / \partial x_j)_{i=1, \dots, m, j=1, \dots, d}$. For a function $\varphi : \mathbb{R}^d \mapsto \mathbb{R}$, $\varphi \in C^2(\Omega)$, we write $\nabla^2 \varphi(x) \in \mathbb{R}^{d \times d}$ for its Hessian at point x . For a function $f : \Omega \rightarrow \mathbb{R}^m$ and any positive number $0 < \delta \leq 1$, the *Hölder constant* of order δ is given by

$$[f]_\delta := \max_{i \in \{1, \dots, m\}} \sup_{x \neq y \in \Omega} \frac{|f_i(x) - f_i(y)|}{\min\{1, \|x - y\|\}^\delta}. \quad (2.2)$$

Now, for any $\alpha > 0$, we set $s = \lfloor \alpha \rfloor$ and define the *Hölder ball* $\mathcal{H}^\alpha(\Omega, H)$ as

$$\mathcal{H}^\alpha(\Omega, H) = \{f \in C^s(\Omega) : \|f\|_{\mathcal{H}^\alpha} := \max\{\|f\|_{C^s}, \max_{|\gamma|=s} [D^\gamma f]_\delta\} \leq H\}.$$

Note that if $f \in \mathcal{H}^{1+\beta}(\Omega, H)$ for some $\beta > 0$, then, for any $i \in \{1, \dots, m\}$, $j \in \{1, \dots, d\}$, it holds that

$$\left| \frac{\partial f_i(x)}{\partial x_j} - \frac{\partial f_i(y)}{\partial x_j} \right| \leq \|f\|_{\mathcal{H}^{1+\beta}} \cdot \|x - y\|^{1 \wedge \beta} \leq H \cdot \|x - y\|^{1 \wedge \beta} \quad \text{for all } x, y \in \Omega,$$

since $\|f\|_{\mathcal{H}^{\beta_1}} \leq \|f\|_{\mathcal{H}^{\beta_2}}$ for any $\beta_2 \geq \beta_1$. We will also write $f \in \mathcal{H}^\alpha(\Omega)$ if $f \in \mathcal{H}^\alpha(\Omega, H)$ for some $H < \infty$. We also introduce a class of Λ -regular functions $\mathcal{H}_\Lambda^\alpha(\Omega, H)$, $\Lambda > 1$:

$$\mathcal{H}_\Lambda^\alpha(\Omega, H) = \left\{ f \in \mathcal{H}^\alpha(\Omega, H) : \Lambda^{-2} \mathbf{I}_{d \times d} \preceq \nabla f(x)^\top \nabla f(x) \preceq \Lambda^2 \mathbf{I}_{d \times d} \text{ for all } x \in \Omega \right\}, \quad (2.3)$$

where for symmetric matrices $A, B \in \mathbb{R}^{d \times d}$ we write $A \preceq B$ if $u^\top (B - A) u \geq 0$ for any $u \in \mathbb{R}^d$.

Neural networks. To give a formal definition of a neural network, we first fix an activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$. For a vector $v = (v_1, \dots, v_p) \in \mathbb{R}^p$, we define the shifted activation function $\sigma_v : \mathbb{R}^p \rightarrow \mathbb{R}^p$ as

$$\sigma_v(x) = (\sigma(x_1 - v_1), \dots, \sigma(x_p - v_p)), \quad x = (x_1, \dots, x_p) \in \mathbb{R}^p.$$

Given a positive integer N and a vector $\mathcal{A} = (p_0, p_1, \dots, p_{N+1}) \in \mathbb{N}^{N+2}$, a neural network of depth $N + 1$ (with N hidden layers) and architecture \mathcal{A} is a function of the form

$$f : \mathbb{R}^{p_0} \rightarrow \mathbb{R}^{p_{N+1}}, \quad f(x) = W_N \circ \sigma_{v_N} \circ W_{N-1} \circ \sigma_{v_{N-1}} \circ \dots \circ W_1 \circ \sigma_{v_1} \circ W_0 \circ x, \quad (2.4)$$

where $W_i \in \mathbb{R}^{p_{i+1} \times p_i}$ are weight matrices and $v_i \in \mathbb{R}^{p_i}$ are shift vectors. The numbers p, p_0, \dots, p_{N+1} should not be confused with the density $\mathfrak{p}(x)$ nor with $\mathfrak{p}^*(x)$, which are always displayed in serif. Next, we introduce a special subclass class of neural networks of depth $N + 1$ with architecture \mathcal{A} :

$$\text{NN}(N, \mathcal{A}) = \left\{ f \text{ of the form (2.4) : } \|W_0\|_\infty \vee \max_{1 \leq \ell \leq N} \{ \|W_\ell\|_\infty \vee \|v_\ell\|_\infty \} \leq 1 \right\}.$$

The maximum number of neurons in one layer $\|\mathcal{A}\|_\infty$ is called the width of the neural network. Similarly to (Schmidt-Hieber, 2020), we consider sparse neural networks assuming that only a few weights are not equal to zero. For this purpose, we introduce a class of neural networks of depth $N + 1$ with architecture \mathcal{A} and at most s non-zero weights:

$$\text{NN}(N, \mathcal{A}, s) = \left\{ f \in \text{NN}(N, \mathcal{A}) : \|W_0\|_0 + \sum_{\ell=1}^N (\|W_\ell\|_0 + \|v_\ell\|_0) \leq s \right\}.$$

3. Theoretical properties of vanilla GANs: a general oracle inequality

We begin with a sharp oracle inequality for general parametric classes of generators \mathcal{G} and discriminators \mathcal{D} . Following (Biau et al., 2020), we impose the next regularity assumptions on generators and discriminators.

Assumption AG. For all $g \in \mathcal{G}$, the image of the latent space Y is a subset of X , that is, $g(\mathsf{Y}) \subseteq \mathsf{X}$. Moreover, for all $y \in \mathsf{Y}$, the map $w \mapsto g_w(y)$ is Lipschitz on the parameter space W with a constant $L_{\mathcal{G}}$. That is, for any $y \in \mathsf{Y}$ and any $u, v \in \mathsf{W}$, it holds that

$$\|g_u(y) - g_v(y)\| \leq L_{\mathcal{G}} \|u - v\|_\infty.$$

Assumption AD. The maps $x \mapsto D_\theta(x)$ and $\theta \mapsto D_\theta(x)$ are Lipschitz on the ambient space X and on the parameter set Θ with constants L_{X} and L_Θ , respectively. More precisely, for any $x, x_1, x_2 \in \mathsf{X}$ and any $\theta, \theta_1, \theta_2 \in \Theta$, the following inequalities hold:

$$|D_\theta(x_1) - D_\theta(x_2)| \leq L_{\mathsf{X}} \|x_1 - x_2\| \quad \text{and} \quad |D_{\theta_1}(x) - D_{\theta_2}(x)| \leq L_\Theta \|\theta_1 - \theta_2\|_\infty.$$

Moreover, there exist constants $0 < D_{\min} \leq D_{\max} < 1$ such that

$$D_\theta(x) \in [D_{\min}, D_{\max}] \quad \text{for all } x \in \mathsf{X} \text{ and } \theta \in \Theta.$$

We would like to note that the requirement that all functions from \mathcal{D} are bounded away from 0 and 1 is needed for the $\log D_\theta$ and $\log(1 - D_\theta)$ to be well defined. Similar conditions appear in the literature for aggregation with the Kullback-Leibler loss (for instance, in (Polzehl and Spokoiny, 2006; Belomestny and Spokoiny, 2007; Rigollet, 2012; Butucea et al., 2017)). Finally, similarly to (Biau et al., 2020), we require the densities of fake random elements $\mathfrak{p}_w(x)$, $w \in \mathsf{W}$ to fulfil the following property.

Assumption Ap. For all $x \in \mathsf{X}$, the map $w \mapsto \mathfrak{p}_w(x)$ is Lipschitz on W with a constant $L_{\mathfrak{p}}$. That is, for any $x \in \mathsf{X}$ and any $u, v \in \mathsf{W}$, we have

$$|\mathfrak{p}_u(x) - \mathfrak{p}_v(x)| \leq L_{\mathfrak{p}} \|u - v\|_\infty.$$

Under Assumptions AG, AD, and Ap, we establish the following oracle inequality for the GAN estimate (1.3).

Theorem 1. *Assume AG, AD, and Ap. Let $\mathcal{W} \subseteq [-1, 1]^{d_{\mathcal{G}}}$ and $\Theta \subseteq [-1, 1]^{d_{\mathcal{D}}}$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, it holds that*

$$\begin{aligned} \text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} - \Delta_{\mathcal{D}} &\lesssim \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}})[(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\ &+ C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)}{n} \end{aligned} \quad (3.1)$$

where

$$\Delta_{\mathcal{G}} = \min_{w \in \mathcal{W}} \text{JS}(\mathbf{p}_w, \mathbf{p}^*), \quad \Delta_{\mathcal{D}} = \max_{w \in \mathcal{W}} \min_{\theta \in \Theta} [\text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta)],$$

and

$$C_{3.2} = \left(\frac{\log(1/D_{\min})}{D_{\min}^2} + \frac{\log(1/(1 - D_{\max}))}{(1 - D_{\max})^2} \right). \quad (3.2)$$

Here \lesssim stands for inequality up to an absolute multiplicative constant.

Remark 1. Let us recall that we assumed the number of fake samples m equal to the sample size n . In general, if $m \neq n$, n should be replaced by $(m \wedge n)$ in (3.1). This does not affect the bound much, because usually $m \geq n$.

Theoretical properties of vanilla GANs were studied in the works of Liang (2021); Biau et al. (2020); Asatryan et al. (2023). The results of Asatryan et al. (2023) mainly concern the case of highly smooth generators, so we postpone a comparison with their rates of convergence, and we will return to it after Theorem 2. To our knowledge, the first upper bound on the Jensen-Shannon divergence between the true density \mathbf{p}^* and the vanilla GAN estimate $\mathbf{p}_{\hat{w}}$ was obtained in (Liang, 2021, Theorem 13). In (Biau et al., 2020), the authors significantly improved the approximation terms in the oracle inequality of (Liang, 2021, Theorem 13). The closest result to our Theorem 1 in the literature is (Biau et al., 2020, Theorem 4.1), so let us focus on the comparison of these two results. First, in (Biau et al., 2020) the authors assumed that the true density \mathbf{p}^* is bounded away from 0 and $+\infty$ on its support and that \mathbf{p}_w is uniformly bounded over all $w \in \mathcal{W}$, while we avoid such requirements in our analysis. Second, the oracle inequality for the JS risk of the vanilla GAN estimate, established in (Biau et al., 2020) under similar assumptions is weaker than the bound from Theorem 1. The authors of the work (Biau et al., 2020) proved that

$$\mathbb{E} \text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} - \Delta_{\mathcal{D}} \lesssim \sqrt{\frac{d_{\mathcal{G}} + d_{\mathcal{D}}}{n}}. \quad (3.3)$$

One can also use McDiarmid's inequality (see, e.g. (Boucheron et al., 2004, Corollary 4)) to transform the in-expectation guarantee (3.3) into a large deviation bound on $\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*)$ of the form

$$\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} - \Delta_{\mathcal{D}} \lesssim \sqrt{\frac{d_{\mathcal{G}} + d_{\mathcal{D}}}{n}} + \sqrt{\frac{\log(1/\delta)}{n}}, \quad (3.4)$$

which holds with probability at least $1 - \delta$. If the classes \mathcal{G} , \mathcal{D} cannot approximate g^* and the respective optimal discriminator with high accuracy, then $\Delta_{\mathcal{G}}$ and $\Delta_{\mathcal{D}}$ are of order 1.

In this case, our rates show no improvements over (Biau et al., 2020). However, in practice one uses rather expressive classes of deep neural networks for \mathcal{G} and \mathcal{D} , and the bound (3.1) can be significantly better than (3.4). In Section 4, we give an example parametric families \mathcal{G}, \mathcal{D} such that $\Delta_{\mathcal{G}}$ and $\Delta_{\mathcal{D}}$ tend to 0 polynomially fast as n goes to ∞ . Finally, in contrast to (Biau et al., 2020), we specify the dependence of the rate on the constants $L_{\mathcal{G}}$ and L_{Θ} , which may be large, especially in the case of wide and deep networks.

To get further insights of the result of Theorem 1, let us elaborate on the properties of $\Delta_{\mathcal{D}}$. The next lemma shows that $\Delta_{\mathcal{D}}$ exhibits quadratic behaviour and its upper bound is closely related to approximation properties of the class of discriminators considered.

Lemma 1. *Under Assumption (AD), for any $w \in \mathbb{W}$ and any $\theta \in \Theta$, it holds that*

$$\text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta) \geq \left\| \frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} - D_{\theta} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}^2$$

and

$$\text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta) \leq \frac{C_{3.6}^2}{(C_{3.6} - 1)^2 D_{\min}(1 - D_{\max})} \left\| \frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} - D_{\theta} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}^2, \quad (3.5)$$

where

$$C_{3.6} = 1 + \sqrt{\frac{D_{\min}}{(1 - D_{\min}) \log(1/(1 - D_{\max}))} \wedge \frac{1 - D_{\max}}{D_{\max} \log(1/D_{\min})}}. \quad (3.6)$$

The proof of Lemma 1 is deferred to Appendix C.1. Lemma 1 plays a key role in derivation of faster rates of convergence. It shows that, for each $w \in \mathbb{W}$, $L(w, \theta)$ enjoys a similar curvature as the squared loss. This fact remained unnoticed in the literature. Though it is well known that $\text{JS}(\mathbf{p}_w, \mathbf{p}^*)$ has a quadratic behaviour with respect to \mathbf{p}_w (see, for instance, our Lemma 7 below), this fact alone is not enough to derive the uniform Bernstein-type inequality (3.1). Only the combination of Lemma 1 and Lemma 7 leads us to a new, significantly better result. In particular, if for any $w \in \mathbb{W}$ there exists $\theta(w) \in \Theta$ such that $\|D_{\theta(w)} - \mathbf{p}^*/(\mathbf{p}^* + \mathbf{p}_w)\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)} \leq \varepsilon$, then $\Delta_{\mathcal{D}} \lesssim \varepsilon^2$. Hence, in this case, Theorem 1 and the Cauchy-Schwarz inequality immediately yield that

$$\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} \lesssim \varepsilon^2 + \frac{d_{\mathcal{G}} + d_{\mathcal{D}} + \log(1/\delta)}{n}$$

with probability at least $1 - \delta$. In (Biau et al., 2020), the authors could not exploit the quadratic behaviour of $(\text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta))$ properly, and they only proved that

$$\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} \lesssim \varepsilon^2 + \sqrt{\frac{d_{\mathcal{G}} + d_{\mathcal{D}}}{n}} + \sqrt{\frac{\log(1/\delta)}{n}}$$

under more restrictive assumptions. Finally, we would like to note that, in contrary to the remark in (Singh et al., 2018, Section 10.1), the oracle inequality in Theorem 1 does not require the density \mathbf{p}^* to be bounded away from zero.

4. Example: deep nonparametric density estimation

The goal of this section is to show that GAN estimates achieve minimax rates of convergence in the problem of nonparametric density estimation. From now on, we assume that $\mathbf{X} = \mathbf{Y} = [0, 1]^d$, μ is the Lebesgue measure in \mathbb{R}^d , and the generators $g_w \in \mathcal{G}$ are non-degenerate maps, so the density of the fake samples with respect to μ is defined correctly. In this setup, if a latent random element Y is drawn according to the density ϕ supported on $[0, 1]^d$, then the corresponding density of $g_w(Y)$ is given by

$$\mathbf{p}_w(x) = |\det[\nabla g_w(g_w^{-1}(x))]|^{-1} \phi(g_w^{-1}(x)), \quad x \in [0, 1]^d. \quad (4.1)$$

For the ease of exposition, we assume that Y is distributed uniformly on $[0, 1]^d$, so that (4.1) simplifies to

$$\mathbf{p}_w(x) = |\det[\nabla g_w(g_w^{-1}(x))]|^{-1}, \quad x \in [0, 1]^d. \quad (4.2)$$

We also impose a structural assumption on the underlying density \mathbf{p}^* , assuming that it is the density of a random element $g^*(Y)$ where $g^* : [0, 1]^d \rightarrow [0, 1]^d$ is a smooth regular map and, as before, Y has a uniform distribution on $[0, 1]^d$.

Assumption Ap^ .* There exist constants $\beta > 2$, $H^* > 0$, and $\Lambda > 1$ such that \mathbf{p}^* is of the form

$$\mathbf{p}^*(x) = |\det[\nabla g^*((g^*)^{-1}(x))]|^{-1}, \quad x \in [0, 1]^d,$$

with $g^* \in \mathcal{H}_\Lambda^{1+\beta}([0, 1]^d, H^*)$.

In fact, Assumption Ap^* is not very restrictive and allows for a quite large class of densities. The celebrated Brenier's theorem (Brenier, 1991) implies that, for any density \mathbf{p} with a finite second moment, there exists a convex almost everywhere differentiable function φ such that $\nabla\varphi(Y) \sim \mathbf{p}$ where $Y \sim \mathcal{U}([0, 1]^d)$. The Caffarelli's regularity theory (Caffarelli, 1991, 1992a,b, 1996) extends the Brenier's result in the following way. If \mathbf{p} is bounded away from zero and infinity, $\Omega = \text{supp}(\mathbf{p})$ is convex, and \mathbf{p} is in $C^\beta(\text{Int}(\Omega))$, then the potential φ is in $C^{\beta+2}$.

Remark 2. Lemma 8 yields that \mathbf{p}^* is bounded away from zero and infinity in the considered model. This is a so-called strong density assumption (see, e.g., (Audibert and Tsybakov, 2007, Definition 2.2)), widely used in statistics.

When applying GANs to the problem of nonparametric density estimation, we shall take \mathcal{G} and \mathcal{D} to be classes of neural networks with ReQU (rectified quadratic unit) activation functions:

$$\sigma^{\text{ReQU}}(x) = (x \vee 0)^2.$$

While rectified linear unit (ReLU), defined as

$$\sigma^{\text{ReLU}}(x) = x \vee 0,$$

is the most common choice for the activation functions in neural networks, it is not suitable for our purposes. The reason is that we want to use neural networks as generators. The density of the fake random elements $g_w(Y_1), \dots, g_w(Y_n)$ is given by (4.2) and, to enforce differentiability, we use the ReQU activation function. Besides, a recent result of (Belomestny et al., 2023) on approximation properties of neural networks with ReQU activations can be

used to bound $\Delta_{\mathcal{G}}$ and $\Delta_{\mathcal{D}}$ from Theorem 1. Since the activation function is fixed, we will write $\sigma(x)$, instead of $\sigma^{\text{ReLU}}(x)$. In this section, we impose the following assumptions on the classes of generators and discriminators.

Assumption AG'. Fix $N_{\mathcal{G}}, d_{\mathcal{G}} \in \mathbb{N}$, and an architecture $\mathcal{A}_{\mathcal{G}} \in \mathbb{N}^{N_{\mathcal{G}}+2}$ with the first and the last component equal to d . There are constants $H_{\mathcal{G}} > 0$ and $\Lambda_{\mathcal{G}} > 1$ such that

$$\mathcal{G} = \mathcal{G}(\Lambda_{\mathcal{G}}, H_{\mathcal{G}}, N_{\mathcal{G}}, \mathcal{A}_{\mathcal{G}}, d_{\mathcal{G}}) = \mathcal{H}_{\Lambda_{\mathcal{G}}}^2([0, 1]^d, H_{\mathcal{G}}) \cap \text{NN}(N_{\mathcal{G}}, \mathcal{A}_{\mathcal{G}}, d_{\mathcal{G}}).$$

Besides, $g_w([0, 1]^d) \subseteq [0, 1]^d$ for all $g_w \in \mathcal{G}$.

Assumption AD'. Fix $N_{\mathcal{D}}, d_{\mathcal{D}} \in \mathbb{N}$, and an architecture $\mathcal{A}_{\mathcal{D}} \in \mathbb{N}^{N_{\mathcal{D}}+2}$ with the first and the last components equal to d and 1, respectively. There are constants $H_{\mathcal{D}} > 0$ and $0 < D_{\min} \leq D_{\max} < 1$ such that

$$\mathcal{D} = \mathcal{D}(D_{\min}, D_{\max}, H_{\mathcal{D}}, N_{\mathcal{D}}, \mathcal{A}_{\mathcal{D}}, d_{\mathcal{D}}) = \mathcal{H}^1([0, 1]^d, H_{\mathcal{D}}) \cap \text{NN}(N_{\mathcal{D}}, \mathcal{A}_{\mathcal{D}}, d_{\mathcal{D}}),$$

and each $D \in \mathcal{D}$ satisfies

$$D(x) \in [D_{\min}, D_{\max}] \subset [0, 1] \quad \text{for all } x \in [0, 1]^d.$$

According to Assumption AG' and the definition of the class $\text{NN}(N_{\mathcal{G}}, \mathcal{A}_{\mathcal{G}}, d_{\mathcal{G}})$, the generators are parametrized by vectors with at most $d_{\mathcal{G}}$ components taking values in $[-1, 1]$. Hence, in the context of Section 4, we have $\mathbb{W} = [-1, 1]^{d_{\mathcal{G}}}$. Similarly, we take $\Theta = [-1, 1]^{d_{\mathcal{D}}}$. Before applying Theorem 1 to the case of parametric families of neural networks, we first check that the conditions of Theorem 1 are fulfilled.

4.1 Towards the rates of convergence: verifying the conditions of Theorem 1

Let us start with Assumptions AG and AD. To this end we show that the maps $w \mapsto g_w(y)$ and $\theta \mapsto D_{\theta}(x)$ are Lipschitz on $[-1, 1]^{d_{\mathcal{G}}}$ and $[-1, 1]^{d_{\mathcal{D}}}$, respectively.

Lemma 2. *Let $N \in \mathbb{N}$ and fix an architecture $\mathcal{A} = (p_0, p_1, \dots, p_{N+1}) \in \mathbb{N}^{N+2}$. Let the matrices $W_i^{(1)}, W_i^{(2)} \in [-1, 1]^{p_{i+1} \times p_i}$, $0 \leq i \leq N$, and the vectors $v_i^{(1)}, v_i^{(2)} \in [-1, 1]^{p_i}$, $1 \leq i \leq N$, be such that*

$$\left\| W_i^{(1)} - W_i^{(2)} \right\|_{\infty} \leq \varepsilon \quad \text{for all } i \in \{0, \dots, L\}$$

and

$$\left\| v_i^{(1)} - v_i^{(2)} \right\|_{\infty} \leq \varepsilon \quad \text{for all } i \in \{1, \dots, L\}.$$

Then the neural networks

$$f^{(1)}(x) = W_N^{(1)} \circ \sigma_{v_N^{(1)}} \circ W_{N-1}^{(1)} \circ \sigma_{v_{N-1}^{(1)}} \circ \dots \circ \sigma_{v_1^{(1)}} \circ W_0^{(1)} \circ x,$$

$$f^{(2)}(x) = W_N^{(2)} \circ \sigma_{v_N^{(2)}} \circ W_{N-1}^{(2)} \circ \sigma_{v_{N-1}^{(2)}} \circ \dots \circ \sigma_{v_1^{(2)}} \circ W_0^{(2)} \circ x$$

satisfy the inequality

$$\left\| f^{(1)}(x) - f^{(2)}(x) \right\|_{\infty} \leq \varepsilon(N+1)2^N \prod_{\ell=0}^N (p_{\ell} + 1)^{2^N} \quad \text{for all } x \in [0, 1]^d. \quad (4.3)$$

The proof of Lemma 2 is moved to Appendix B.1. Lemma 2 and the fact that $\mathcal{D} \subset \mathcal{H}^1([0, 1]^d, H_{\mathcal{D}})$ (due to Assumption AD') imply that Assumptions AG and AD are fulfilled with the constants

$$\mathsf{L}_{\mathcal{G}} = (N_{\mathcal{G}} + 1)2^{N_{\mathcal{G}}} \prod_{p \in \mathcal{A}_{\mathcal{G}}} (p + 1)^{2^{N_{\mathcal{G}}}}, \quad \mathsf{L}_{\Theta} = (N_{\mathcal{D}} + 1)2^{N_{\mathcal{D}}} \prod_{p \in \mathcal{A}_{\mathcal{D}}} (p + 1)^{2^{N_{\mathcal{D}}}}, \quad \text{and} \quad \mathsf{L}_{\mathcal{X}} = H_{\mathcal{D}}.$$

It only remains to check that Assumption Ap is satisfied as well. We do it in two steps. First, we show that if two generators are close to each other with respect to the \mathcal{H}^1 -norm, then the corresponding densities are close as well.

Lemma 3. *Assume AG' and consider any $u, v \in [-1, 1]^{d_{\mathcal{G}}}$. Then the corresponding generators g_u, g_v and the densities $\mathfrak{p}_u, \mathfrak{p}_v$ fulfill*

$$\|\mathfrak{p}_u - \mathfrak{p}_v\|_{L_{\infty}([0, 1]^d)} \leq \mathsf{L}_2 \|g_u - g_v\|_{\mathcal{H}^1([0, 1]^d)}$$

with

$$\mathsf{L}_2 = d^{2+d/2} \Lambda^{3d} (1 + H_{\mathcal{G}} \Lambda \sqrt{d}). \quad (4.4)$$

We provide the proof of Lemma 3 in Appendix A.1. Finally, let us show that the \mathcal{H}^1 -norm of $(g_u - g_v)$ scales linearly with the norm of $(u - v)$. We need a counterpart of Lemma 2 for the Jacobi matrices $\nabla f^{(1)}(x)$ and $\nabla f^{(2)}(x)$.

Lemma 4. *Within the notations of Lemma 2, the neural networks $f^{(1)}$ and $f^{(2)}$ satisfy the inequality*

$$\left\| \nabla f^{(1)}(x) - \nabla f^{(2)}(x) \right\|_{\infty} \leq \varepsilon N(N + 1)2^{N+1} \prod_{\ell=0}^N (p_{\ell} + 1)^{2^{N+1}+1} \quad \text{for all } x \in [0, 1]^d. \quad (4.5)$$

One can find the proof of Lemma 4 in Appendix B.2. Lemmata 2 and 4 immediately yield that, under Assumption AG', we have

$$\|g_u - g_v\|_{\mathcal{H}^1([0, 1]^d)} \leq \|u - v\|_{\infty} \cdot N_{\mathcal{G}}(N_{\mathcal{G}} + 1)2^{N_{\mathcal{G}}+1} \prod_{p \in \mathcal{A}_{\mathcal{G}}} (p + 1)^{2^{N_{\mathcal{G}}+1}+1}.$$

Hence, Assumption Ap holds with

$$\mathsf{L}_{\mathfrak{p}} = d^{2+d/2} \Lambda^{3d} (1 + H_{\mathcal{G}} \Lambda \sqrt{d}) N_{\mathcal{G}}(N_{\mathcal{G}} + 1)2^{N_{\mathcal{G}}+1} \prod_{p \in \mathcal{A}_{\mathcal{G}}} (p + 1)^{2^{N_{\mathcal{G}}+1}+1}.$$

Thus, we proved that Assumptions AG' and AD' yield Assumptions AG, AD, and Ap.

4.2 Rates of convergence in nonparametric density estimation with GANS

The discussion in Section 4.1 implies that Theorem 1 can be applied to the setting described in the beginning of Section 4. In this section, we go further and provide upper bounds on $\Delta_{\mathcal{G}}$ and $\Delta_{\mathcal{D}}$ under assumptions AD', AG', and Ap*. The key ingredient of our analysis is the recent result of Belomestny et al. (2023) quantifying the expressiveness of neural networks with ReQU activations. It ensures that for any $f \in \mathcal{H}^{\beta+1}([0, 1]^d, H^*)$, $\beta > 1$, and any $\varepsilon > 0$,

there is a neural network with ReQU activation function from the class $\mathcal{H}^{[\beta]}([0, 1]^d, H^* + \varepsilon)$ which approximates f within the accuracy ε with respect to the norm in $\mathcal{H}^{[\beta]}([0, 1]^d)$. We would like to emphasize that, unlike many other results on approximation properties of neural networks, Belomestny et al. (2023) considers simultaneous approximation of a smooth function and its derivatives, which is crucial for our purposes. This result allows us to derive the optimal rates of convergence for GAN estimates in the nonparametric density estimation problem. However, we would like to emphasize that the proof of Theorem 2 works for any choice of activation function, which yields simultaneous approximation of a function and its derivatives by a neural network with weights taking their values in $[-1, 1]$.

Theorem 2. *Assume that Ap^* holds with $\beta > 2$. Choose $H_{\mathcal{G}} \geq 2H^*$, $\Lambda_{\mathcal{G}} \geq 2\Lambda$, $D_{\min} \leq \Lambda^{-d}/(\Lambda^d + \Lambda^{-d})$, $D_{\max} \geq \Lambda^d/(\Lambda^d + \Lambda^{-d})$, and $H_{\mathcal{D}}$ large enough. Then there are positive integers $N_{\mathcal{G}}$, $d_{\mathcal{G}}$, $N_{\mathcal{D}}$, $d_{\mathcal{D}}$ and the architectures $\mathcal{A}_{\mathcal{G}} \in \mathbb{N}^{N_{\mathcal{G}}+2}$, $\mathcal{A}_{\mathcal{D}} \in \mathbb{N}^{N_{\mathcal{D}}+2}$ such that AG' and AD' define nonempty sets \mathcal{G} and \mathcal{D} , respectively. Let us consider the estimator (1.3) with $\mathbf{W} = [-1, 1]^{d_{\mathcal{G}}}$, $\Theta = [-1, 1]^{d_{\mathcal{D}}}$, $g_w \in \mathcal{G}$ and $D_{\theta} \in \mathcal{D}$. For any $\delta \in (0, 1)$, with probability at least $1 - \delta$, $\mathbf{p}_{\hat{w}}$ satisfies the inequality*

$$\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) \lesssim \left(\frac{\log n}{n}\right)^{2\beta/(2\beta+d)} + \frac{\log(1/\delta)}{n}, \quad (4.6)$$

provided that $n \geq n_0$ with n_0 depending only on β, d, Λ, H^* , and $H_{\mathcal{G}}$. In (4.6), the notation \lesssim stands for an inequality up to a multiplicative constant depending on $d, \beta, H^*, H_{\mathcal{G}}$, and $H_{\mathcal{D}}$ only.

Remark 3. The dimensions $d_{\mathcal{G}}$ and $d_{\mathcal{D}}$ define the complexity of the optimization problem in (1.3) and depend on n . It follows from the proof of Theorem 2 that

$$d_{\mathcal{G}} \lesssim \left(\frac{n}{\log n}\right)^{d/(2\beta+d)}, \quad d_{\mathcal{D}} \lesssim \left(\frac{n}{\log n}\right)^{d/(2\beta+d)}.$$

Let us note that this dependence of $d_{\mathcal{G}}$ and $d_{\mathcal{D}}$ on n can not be avoided in general.

Theorem 2 improves the dependence on both n and δ in the existing bounds on the JS-divergence between the vanilla GAN estimate $\mathbf{p}_{\hat{w}}$ and \mathbf{p}^* . In (Asatryan et al., 2023, Theorem 3.13 and Theorem 4.4), the authors proved that $\mathbf{E} \text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*)$ decays as fast as $n^{-1/2}$ if $2\beta > d$. Our results show that in this case the rate of convergence can be much faster. The reason for suboptimality of (Asatryan et al., 2023, Theorem 4.4) is the use of the chaining technique to control the global supremum of the empirical process $L_n(w, \theta) - L(w, \theta)$. This approach was successfully applied to WGANs (see, e.g., (Liang, 2021)) but in the case of vanilla GANs, one can prove better upper bounds for the supremum of $L_n(w, \theta) - L(w, \theta)$ in a local vicinity of the saddle point. It is also worth mentioning that, if one uses the bound (3.3) from (Biau et al., 2020, Theorem 4.1) and (Belomestny et al., 2023, Theorem 2) to control $\Delta_{\mathcal{G}}$ and $\Delta_{\mathcal{D}}$, he will get a suboptimal rate of convergence

$$\mathbf{E} \text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) \lesssim \left(\frac{\log n}{n}\right)^{2\beta/(4\beta+d)}. \quad (4.7)$$

Finally, the second term $\log(1/\delta)/n$ in (4.6) significantly improves the standard rate $\sqrt{\log(1/\delta)/n}$ which follows from McDiarmid's inequality.

Theorem 2 also yields an upper bound on the squared L_2 -distance. Indeed, under the conditions of Theorem 2, Lemma 7 and Lemma 8 yield that

$$\|\mathbf{p}_{\hat{w}} - \mathbf{p}^*\|_{L_2(\mu)}^2 \lesssim \text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) \lesssim \left(\frac{\log n}{n}\right)^{2\beta/(2\beta+d)} + \frac{\log(1/\delta)}{n}.$$

This coincides (up to logarithmic factors) with the well-known minimax rate of convergence $(\log n/n)^{2\beta/(2\beta+d)}$ for estimation of a smooth bounded away from zero density under the squared L_2 -loss. It turns out that the bound on the JS-divergence from Theorem 2 is also minimax optimal up to some logarithmic factors, provided that \mathbf{p}^* belongs to the class of densities satisfying Assumption Ap^* .

Theorem 3. *Let (X_1, \dots, X_n) be a sample of i.i.d. observations generated from a density \mathbf{p}^* satisfying Ap^* . Then for any estimate $\hat{\mathbf{p}}$ of \mathbf{p}^* , that is, measurable function of X_1, \dots, X_n , it holds that*

$$\sup_{\mathbf{p}^*} \mathbb{E} \text{JS}(\hat{\mathbf{p}}, \mathbf{p}^*) \gtrsim n^{-2\beta/(2\beta+d)}$$

with a hidden constant depending on d only.

The proof of Theorem 3 relies on the van Trees inequality (see (Van Trees, 1968, p. 72) and (Gill and Levit, 1995)). Though under Assumption Ap^* $\text{JS}(\mathbf{p}_w, \mathbf{p}^*)$ is equivalent to $\|\mathbf{p}_w - \mathbf{p}^*\|_{L_2}^2$ (see Lemma 7 and Lemma 8 below), we would like to emphasize that Theorem 3 does not follow from the existing lower bounds in nonparametric density estimation (see, for instance, (Tsybakov, 2008, Exercise 2.10)). The reason is that the class of admissible densities in Assumption Ap^* is narrower, than $\mathcal{H}^\beta([0, 1]^d, H_0)$, $H_0 > 0$, due to the additional assumption that \mathbf{p}^* is the density of $g^*(Y)$ for some $g^* \in \mathcal{H}_\Lambda^{1+\beta}([0, 1]^d, H^*)$. If $\mathbf{p}^* \in \mathcal{H}^\beta([0, 1]^d, H_0)$, then, according to Brenier's theorem, there is $\varphi^* : [0, 1]^d \rightarrow \mathbb{R}^d$, such that $\nabla\varphi^*(Y) \sim \mathbf{p}^*$. However, Caffarelli's regularity theory does not guarantee that $\nabla^2\varphi^*$ satisfies the condition

$$\Lambda^{-2}\mathbf{I}_{d \times d} \preceq \nabla^2\varphi^*(x)^\top \nabla^2\varphi^*(x) \preceq \Lambda^2\mathbf{I}_{d \times d},$$

as required by Assumption Ap^* . Without it, the lower bound will be irrelevant. Moreover, the theory does not provide a uniform upper bound on the $(\beta + 2)$ -th derivative of φ^* . In contrary, it states that the $(\beta + 2)$ -th derivative of φ^* can tend to infinity at the border of $[0, 1]^d$. For these reasons, we find it necessary to prove an explicit lower bound in our setup.

5. Proofs of the main results

This section contains the proofs of our main results, Theorems 1, 2, and 3. The proof of the upper bounds relies on the uniform high probability bound on $L_n(w, \theta) - L(w, \theta)$, given below.

Proposition 1. *Grant Assumptions AG , AD , and Ap . Let $\mathbf{W} \subseteq [-1, 1]^{d_G}$ and $\Theta \subseteq [-1, 1]^{d_D}$. Then, for any $\delta \in (0, 1)$ and any $\varepsilon \in (0, 2]$, with probability at least $1 - \delta$, it*

holds that

$$\begin{aligned}
 |L_n(w, \theta) - L(w, \theta)| &\leq 3C_{5.1}\varepsilon + \sqrt{L_{\mathbf{p}}}\varepsilon^{1/2} \\
 &\quad + 4\sqrt{\frac{C_{5.2}\text{JS}(\mathbf{p}_w, \mathbf{p}^*)(d_{\mathcal{G}} \log(2/\varepsilon) + d_{\mathcal{D}} \log(2/\varepsilon) + \log(2/\delta))}{2n}} \\
 &\quad + \sqrt{\frac{C_{5.2}\Delta(w, \theta)(d_{\mathcal{G}} \log(2/\varepsilon) + d_{\mathcal{D}} \log(2/\varepsilon) + \log(2/\delta))}{2n}} \\
 &\quad + \frac{2(C_{5.2} + C_{5.3})(d_{\mathcal{G}} \log(2/\varepsilon) + d_{\mathcal{D}} \log(2/\varepsilon) + \log(2/\delta))}{3n}
 \end{aligned}$$

simultaneously for all $w \in \mathbb{W}$ and $\theta \in \Theta$. Here $\Delta(w, \theta) = \text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta)$ and the constants $C_{5.1}$, $C_{5.2}$, and $C_{5.3}$ are defined as follows:

$$C_{5.1} = \frac{L_{\mathcal{G}}L_{\mathcal{X}}}{2 - 2D_{\max}} + \frac{L_{\Theta}}{D_{\min} \wedge (1 - D_{\max})}, \quad (5.1)$$

$$\begin{aligned}
 C_{5.2} &= \left(\frac{\log^2(2D_{\min})}{1/2 - D_{\min}} \vee 2\log^2 2 \right) + \left(\frac{\log^2(2 - 2D_{\max})}{D_{\max} - 1/2} \vee 2\log^2 2 \right) \\
 &\quad + \left(\frac{\log(e/(2D_{\min}))}{2D_{\min}^2} + \frac{\log(e/(2 - 2D_{\max}))}{2(1 - D_{\max})^2} \right), \quad (5.2)
 \end{aligned}$$

$$C_{5.3} = \log \left(\frac{1}{D_{\max}} \vee \frac{1}{D_{\min}} \right) \vee \log \left(\frac{1}{1 - D_{\max}} \vee \frac{1}{1 - D_{\min}} \right) = \log \frac{1}{D_{\min} \wedge (1 - D_{\max})}. \quad (5.3)$$

The rest of this section is organized in the following way. Section 5.1 is devoted to the proof of Proposition 1. Then we prove Theorems 1, 2, and 3 in Sections 5.2, 5.3, and 5.4, respectively. In the beginning of these sections, we restate the main results for convenience.

5.1 Proof of Proposition 1

Let us start with the sketch of the proof. First, we study large deviations of $L_n(w, \theta) - L(w, \theta)$ for a fixed pair $(w, \theta) \in \mathbb{W} \times \Theta$. After that, we show that $L_n(w, \theta)$, $L(w, \theta)$, and $\sqrt{\text{JS}(\mathbf{p}_w, \mathbf{p}^*)}$ are Hölder with respect to $w \in \mathbb{W}$ and $\theta \in \Theta$. This and the ε -net argument allow us to derive a uniform large deviation bound on $L_n(w, \theta) - L(w, \theta)$. We split the proof into several steps for the sake of readability.

Step 1: large deviation bound on $L_n(w, \theta) - L(w, \theta)$. Under the conditions of Theorem 1, the functional $L(w, \theta)$ has a kind of curvature inherited from the properties of the cross-entropy. This special structure allows us to obtain tight large deviation bounds on $|L_n(w, \theta) - L(w, \theta)|$. To be more precise, we have the following concentration result.

Lemma 5. *Fix any $w \in \mathbb{W}$ and $\theta \in \Theta$. Assume the conditions of Theorem 1. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, it holds that*

$$|L_n(w, \theta) - L(w, \theta)| \leq \sqrt{\frac{C_{5.2}(9\text{JS}(\mathbf{p}_w, \mathbf{p}^*) + \Delta(w, \theta)) \log(2/\delta)}{2n}} + \frac{2C_{5.3} \log(2/\delta)}{3n}, \quad (5.4)$$

where $\Delta(w, \theta) = \text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta)$, and the constants $C_{5.2}$, $C_{5.3}$ are defined in (5.2) and (5.3), respectively.

The proof of Lemma 5 is moved to Appendix C.2. Our next goal is to convert the concentration inequality (5.4) into a uniform bound on $|L_n(w, \theta) - L(w, \theta)|$.

Step 2: towards a uniform bound on $L_n(w, \theta) - L(w, \theta)$. Let $\varepsilon > 0$ be a parameter to be specified later and let W_ε and Θ_ε be the minimal ε -nets of W and Θ , respectively, with respect to the norm $\|\cdot\|_\infty$. It is straightforward to check that the cardinalities of W_ε and Θ_ε satisfy the inequalities

$$|W_\varepsilon| \leq \left(\frac{2}{\varepsilon}\right)^{d_G} \quad \text{and} \quad |\Theta_\varepsilon| \leq \left(\frac{2}{\varepsilon}\right)^{d_D}.$$

Applying Lemma 5 and the union bound, we conclude that, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, it holds that

$$\begin{aligned} |L_n(w, \theta) - L(w, \theta)| &\leq \sqrt{\frac{C_{5.2}(9 \text{JS}(\mathbf{p}_w, \mathbf{p}^*) + \Delta(w, \theta)) \log(2|W_\varepsilon||\Theta_\varepsilon|/\delta)}{2n}} \\ &\quad + \frac{2C_{5.3} \log(2|W_\varepsilon||\Theta_\varepsilon|/\delta)}{3n} \tag{5.5} \\ &\leq \sqrt{\frac{C_{5.2}(9 \text{JS}(\mathbf{p}_w, \mathbf{p}^*) + \Delta(w, \theta))(d_G \log(2/\varepsilon) + d_D \log(2/\varepsilon) + \log(2/\delta))}{2n}} \\ &\quad + \frac{2C_{5.3}(d_G \log(2/\varepsilon) + d_D \log(2/\varepsilon) + \log(2/\delta))}{3n} \end{aligned}$$

simultaneously for all $(w, \theta) \in W_\varepsilon \times \Theta_\varepsilon$. If we show that $L_n(w, \theta)$, $L(w, \theta)$, and $\text{JS}(\mathbf{p}_w, \mathbf{p}^*)$ are Lipschitz or at least Hölder with respect to $w \in W$ and $\theta \in \Theta$, then we can easily extend the bound (5.5) from the ε -net $W_\varepsilon \times \Theta_\varepsilon$ to the whole set $W \times \Theta$. The Lipschitzness of $L_n(w, \theta)$ and $L(w, \theta)$ easily follows from Assumptions AG and AD.

Lemma 6. *Grant Assumptions AD and AG. Let $w, w_1, w_2 \in W$ and $\theta, \theta_1, \theta_2 \in \Theta$. Then it holds that*

$$|L_n(w_1, \theta) - L_n(w_2, \theta)| \leq \frac{L_G L_X \|w_1 - w_2\|_\infty}{2 - 2D_{\max}} \quad \text{almost surely.}$$

and

$$|L_n(w, \theta_1) - L_n(w, \theta_2)| \leq \frac{L_\Theta \|\theta_1 - \theta_2\|_\infty}{D_{\min} \wedge (1 - D_{\max})} \quad \text{almost surely.}$$

Moreover, under Assumptions AD and AG, we have

$$|L(w_1, \theta) - L(w_2, \theta)| \leq \frac{L_G L_X \|w_1 - w_2\|_\infty}{2 - 2D_{\max}}$$

and

$$|L(w, \theta_1) - L(w, \theta_2)| \leq \frac{L_\Theta \|\theta_1 - \theta_2\|_\infty}{D_{\min} \wedge (1 - D_{\max})}$$

for any $w, w_1, w_2 \in W$ and $\theta, \theta_1, \theta_2 \in \Theta$.

The proof of Lemma 6 is given in Appendix C.3. The next result plays a crucial role in the analysis of $\text{JS}(\mathbf{p}_w, \mathbf{p}^*)$.

Lemma 7. *Let \mathfrak{p} and \mathfrak{q} be the densities of two probability distributions with respect to a dominating measure μ . Then it holds that*

$$\frac{1}{4} \int \frac{(\mathfrak{p}(x) - \mathfrak{q}(x))^2}{\mathfrak{p}(x) + \mathfrak{q}(x)} d\mu \leq \text{JS}(\mathfrak{p}, \mathfrak{q}) \leq \frac{\log 2}{2} \int \frac{(\mathfrak{p}(x) - \mathfrak{q}(x))^2}{\mathfrak{p}(x) + \mathfrak{q}(x)} d\mu. \quad (5.6)$$

Lemma 7 yields the Hölderness of $\text{JS}^{1/2}(\mathfrak{p}_w, \mathfrak{p}^*)$ with respect to $w \in \mathbb{W}$, provided that Assumption Ap is fulfilled.

Corollary 1. *Grant Assumption Ap . Then, for any $u, v \in \mathbb{W}$, it holds that*

$$\left| \sqrt{\text{JS}(\mathfrak{p}_u, \mathfrak{p}^*)} - \sqrt{\text{JS}(\mathfrak{p}_v, \mathfrak{p}^*)} \right| \leq \frac{\sqrt{\mathbb{L}_\mathfrak{p} \log 2}}{2^{1/4}} \|u - v\|_\infty^{1/4} < \sqrt{\mathbb{L}_\mathfrak{p}} \|u - v\|_\infty^{1/4}.$$

The proofs of Lemma 7 and Corollary 1 are moved to Appendix C.4 and Appendix C.5, respectively. We are ready to prove a uniform bound on large deviations of $L_n(w, \theta) - L(w, \theta)$.

Step 3: uniform bound on $L_n(w, \theta) - L(w, \theta)$. Let E be the event where (5.5) holds. Choose any $(w, \theta) \in \mathbb{W} \times \Theta$ and denote the closest to (w, θ) element of $\mathbb{W}_\varepsilon \times \Theta_\varepsilon$ by $(w_\varepsilon, \theta_\varepsilon)$. Then, due to Lemma 6, the following holds on E :

$$\begin{aligned} & |L_n(w, \theta) - L(w, \theta)| \\ & \leq |L_n(w_\varepsilon, \theta_\varepsilon) - L(w_\varepsilon, \theta_\varepsilon)| + 2 \left(\frac{\mathbb{L}_\mathcal{G} \mathbb{L}_\mathcal{X}}{2 - 2D_{\max}} + \frac{\mathbb{L}_\Theta}{D_{\min} \vee (1 - D_{\max})} \right) \varepsilon \\ & \leq 2 \left(\frac{\mathbb{L}_\mathcal{G} \mathbb{L}_\mathcal{X}}{2 - 2D_{\max}} + \frac{\mathbb{L}_\Theta}{D_{\min} \vee (1 - D_{\max})} \right) \varepsilon \\ & \quad + \sqrt{\frac{C_{5.2}(9 \text{JS}(\mathfrak{p}_{w_\varepsilon}, \mathfrak{p}^*) + \Delta(w_\varepsilon, \theta_\varepsilon))(d_\mathcal{G} \log(2/\varepsilon) + d_\mathcal{D} \log(2/\varepsilon) + \log(2/\delta))}{2n}} \\ & \quad + \frac{2C_{5.3}(d_\mathcal{G} \log(2/\varepsilon) + d_\mathcal{D} \log(2/\varepsilon) + \log(2/\delta))}{3n}. \end{aligned}$$

According to Lemma 6 and Corollary 1,

$$\sqrt{\text{JS}(\mathfrak{p}_{w_\varepsilon}, \mathfrak{p}^*)} \leq \sqrt{\text{JS}(\mathfrak{p}_w, \mathfrak{p}^*)} + \sqrt{\mathbb{L}_\mathfrak{p}} \varepsilon^{1/4}$$

and

$$\begin{aligned} \sqrt{\Delta(w_\varepsilon, \theta_\varepsilon)} &= \sqrt{\text{JS}(\mathfrak{p}_{w_\varepsilon}, \mathfrak{p}^*) - \log 2 - L(w_\varepsilon, \theta_\varepsilon)} \\ &\leq \sqrt{2 \text{JS}(\mathfrak{p}_w, \mathfrak{p}^*) - \log 2 - L(w, \theta) + 2\mathbb{L}_\mathfrak{p} \varepsilon^{1/2} + C_{5.1} \varepsilon} \\ &\leq \sqrt{\text{JS}(\mathfrak{p}_w, \mathfrak{p}^*)} + \sqrt{\Delta(w, \theta)} + \sqrt{2\mathbb{L}_\mathfrak{p}} \varepsilon^{1/4} + \sqrt{C_{5.1} \varepsilon}. \end{aligned}$$

Thus, on E , it holds that

$$\begin{aligned}
 |L_n(w, \theta) - L(w, \theta)| &\leq 2C_{5.1}\varepsilon + 4\sqrt{\frac{C_{5.2}\text{JS}(\mathbf{p}_w, \mathbf{p}^*)(d_{\mathcal{G}}\log(2/\varepsilon) + d_{\mathcal{D}}\log(2/\varepsilon) + \log(2/\delta))}{2n}} \\
 &\quad + \sqrt{\frac{C_{5.2}\Delta(w, \theta)(d_{\mathcal{G}}\log(2/\varepsilon) + d_{\mathcal{D}}\log(2/\varepsilon) + \log(2/\delta))}{2n}} \\
 &\quad + \sqrt{\frac{C_{5.2}\mathbf{L}_p\varepsilon^{1/2}(d_{\mathcal{G}}\log(2/\varepsilon) + d_{\mathcal{D}}\log(2/\varepsilon) + \log(2/\delta))}{n}} \\
 &\quad + \sqrt{\frac{C_{5.2}C_{5.1}\varepsilon(d_{\mathcal{G}}\log(2/\varepsilon) + d_{\mathcal{D}}\log(2/\varepsilon) + \log(2/\delta))}{2n}} \\
 &\quad + \frac{2C_{5.3}(d_{\mathcal{G}}\log(2/\varepsilon) + d_{\mathcal{D}}\log(2/\varepsilon) + \log(2/\delta))}{3n}.
 \end{aligned}$$

where $C_{5.1}$ is defined in (5.1). Applying the Cauchy-Schwarz inequality to the fourth and the fifth terms, we get the desired bound.

5.2 Proof of Theorem 1

Theorem 4 (restatement of Theorem 1). *Assume AG, AD, and Ap. Let $\mathbf{W} \subseteq [-1, 1]^{d_{\mathcal{G}}}$ and $\Theta \subseteq [-1, 1]^{d_{\mathcal{D}}}$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, it holds that*

$$\begin{aligned}
 \text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} - \Delta_{\mathcal{D}} &\lesssim \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}})[(d_{\mathcal{G}} + d_{\mathcal{D}})\log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_p \vee 1)n) + \log(8/\delta)]}{n}} \\
 &\quad + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}})\log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_p \vee 1)n) + \log(8/\delta)}{n}
 \end{aligned}$$

where

$$\Delta_{\mathcal{G}} = \min_{w \in \mathbf{W}} \text{JS}(\mathbf{p}_w, \mathbf{p}^*), \quad \Delta_{\mathcal{D}} = \max_{w \in \mathbf{W}} \min_{\theta \in \Theta} [\text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta)],$$

and

$$C_{3.2} = \left(\frac{\log(1/D_{\min})}{D_{\min}^2} + \frac{\log(1/(1 - D_{\max}))}{(1 - D_{\max})^2} \right).$$

Here \lesssim stands for inequality up to an absolute multiplicative constant.

Let us introduce $\bar{w} \in \operatorname{argmin}_{w \in \mathbf{W}} \text{JS}(\mathbf{p}_w, \mathbf{p}^*)$. We begin with studying the excess risk

$$\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} = \text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \min_{w \in \mathbf{W}} \text{JS}(\mathbf{p}_w, \mathbf{p}^*) = \text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \text{JS}(\mathbf{p}_{\bar{w}}, \mathbf{p}^*), \quad (5.7)$$

where \hat{w} is given by (1.3). For any $w \in \mathbf{W}$, let θ_w^* denote the parameter, corresponding to the best discriminator in \mathcal{D} :

$$\theta_w^* \in \operatorname{argmax}_{\theta \in \Theta} L(w, \theta).$$

Then the definition of $\Delta_{\mathcal{D}}$ yields that

$$\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \text{JS}(\mathbf{p}_{\bar{w}}, \mathbf{p}^*) \leq \Delta_{\mathcal{D}} + L(\hat{w}, \theta_{\hat{w}}^*) - L(\bar{w}, \theta_{\bar{w}}^*).$$

Here we used the fact that for any $w \in \mathcal{W}$ and any $\theta \in \Theta$ it holds that $L(w, \theta) + \log 2 \leq \text{JS}(\mathbf{p}_w, \mathbf{p}^*)$ (see, e.g., (Goodfellow et al., 2014, Proposition 1) or (Biau et al., 2020, Section 2)) for the explanation. The expression in the right-hand side of the last inequality can be rewritten as follows:

$$\begin{aligned} \Delta_{\mathcal{D}} + & \underbrace{(L(\hat{w}, \theta_{\hat{w}}^*) - L_n(\hat{w}, \theta_{\hat{w}}^*))}_{T_1} \\ & + \underbrace{(L_n(\hat{w}, \theta_{\hat{w}}^*) - L_n(\bar{w}, \theta_{\bar{w}}^*))}_{T_2} + \underbrace{(L_n(\bar{w}, \theta_{\bar{w}}^*) - L(\bar{w}, \theta_{\bar{w}}^*))}_{T_3}. \end{aligned} \quad (5.8)$$

We split the rest of the proof into several steps for convenience.

Step 1: bounds on T_1 and T_3 . Let us take

$$\begin{aligned} \varepsilon = & \frac{C_{5.2}(d_{\mathcal{G}} \log(2n) + d_{\mathcal{D}} \log(2n) + \log(8/\delta))}{C_{5.1}n} \\ & \wedge \frac{C_{5.2}^2(d_{\mathcal{G}} \log(2n) + d_{\mathcal{D}} \log(2n) + \log(8/\delta))^2}{L_{\mathbf{p}}^2 n^2} \wedge 1. \end{aligned} \quad (5.9)$$

Then $C_{5.1}\varepsilon \leq 1/n$, $\sqrt{L_{\mathbf{p}}\varepsilon} \leq 1/n$, and

$$\log(1/\varepsilon) \leq \log n + \log \left(\frac{C_{5.1}}{C_{5.2}} \vee \frac{L_{\mathbf{p}}^2}{C_{5.2}^2} \right) \lesssim \log n + \log(L_{\mathcal{G}}L_{\mathcal{X}} \vee L_{\Theta} \vee L_{\mathbf{p}} \vee 1).$$

Applying Proposition 1 with ε defined in (5.9), we get that, with probability at least $1 - \delta/2$, we simultaneously have

$$\begin{aligned} T_1 & \lesssim \sqrt{\frac{(\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) + \Delta_{\mathcal{D}})[(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(L_{\mathcal{G}}L_{\mathcal{X}} \vee L_{\Theta} \vee L_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\ & + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(L_{\mathcal{G}}L_{\mathcal{X}} \vee L_{\Theta} \vee L_{\mathbf{p}} \vee 1)n) + \log(8/\delta)}{n} \\ & \leq \sqrt{\frac{(\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}})[(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(L_{\mathcal{G}}L_{\mathcal{X}} \vee L_{\Theta} \vee L_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\ & + \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}})[(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(L_{\mathcal{G}}L_{\mathcal{X}} \vee L_{\Theta} \vee L_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\ & + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(L_{\mathcal{G}}L_{\mathcal{X}} \vee L_{\Theta} \vee L_{\mathbf{p}} \vee 1)n) + \log(8/\delta)}{n} \end{aligned} \quad (5.10)$$

and

$$\begin{aligned} T_3 & \lesssim \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}})[(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(L_{\mathcal{G}}L_{\mathcal{X}} \vee L_{\Theta} \vee L_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\ & + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(L_{\mathcal{G}}L_{\mathcal{X}} \vee L_{\Theta} \vee L_{\mathbf{p}} \vee 1)n) + \log(8/\delta)}{n}. \end{aligned} \quad (5.11)$$

Here we used the fact that $\text{JS}(\mathbf{p}_{\bar{w}}, \mathbf{p}^*) = \Delta_{\mathcal{G}}$ by the definitions of \bar{w} and $\Delta_{\mathcal{G}}$. Also, it holds that $\Delta(w, \theta_w^*) \leq \Delta_{\mathcal{D}}$ due to the definitions of θ_w^* and $\Delta_{\mathcal{D}}$. Note that the hidden constants in (5.10) and (5.11) are absolute.

Step 2: a bound on T_2 . It remains to bound T_2 in (5.8). For any $w \in \mathcal{W}$, let us denote

$$\hat{\theta}_w \in \operatorname{argmax}_{\theta \in \Theta} L_n(w, \theta).$$

By the definition of \hat{w} , we have

$$L_n(\hat{w}, \hat{\theta}_{\hat{w}}) \leq L_n(\bar{w}, \hat{\theta}_{\bar{w}}).$$

Then $T_2 \leq L_n(\bar{w}, \hat{\theta}_{\bar{w}}) - L_n(\bar{w}, \theta_{\bar{w}}^*)$, and, applying Proposition 1 with ε from (5.9) again, we obtain that, with probability at least $1 - \delta/2$, it holds that

$$\begin{aligned} T_2 &\lesssim L(\bar{w}, \hat{\theta}_{\bar{w}}) - L(\bar{w}, \theta_{\bar{w}}^*) \\ &\quad + \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta(\bar{w}, \hat{\theta}_{\bar{w}})) [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\ &\quad + \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}}) [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\ &\quad + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)}{n}, \end{aligned}$$

where, as before, the hidden constants are absolute. Since

$$\Delta(\bar{w}, \hat{\theta}_{\bar{w}}) = \Delta(\bar{w}, \theta_{\bar{w}}^*) + L(\bar{w}, \theta_{\bar{w}}^*) - L(\bar{w}, \hat{\theta}_{\bar{w}}) \leq \Delta_{\mathcal{D}} + L(\bar{w}, \theta_{\bar{w}}^*) - L(\bar{w}, \hat{\theta}_{\bar{w}}),$$

the following inequality holds on the same event:

$$\begin{aligned} T_2 &\lesssim - \left(L(\bar{w}, \theta_{\bar{w}}^*) - L(\bar{w}, \hat{\theta}_{\bar{w}}) \right) \\ &\quad + \sqrt{\frac{\left(L(\bar{w}, \theta_{\bar{w}}^*) - L(\bar{w}, \hat{\theta}_{\bar{w}}) \right) [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\ &\quad + \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}}) [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\ &\quad + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)}{n}. \end{aligned}$$

Maximizing the right-hand side over $(L(\bar{w}, \theta_{\bar{w}}^*) - L(\bar{w}, \hat{\theta}_{\bar{w}}))^{1/2}$, we obtain that

$$\begin{aligned} T_2 &\lesssim \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}}) [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\ &\quad + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)}{n}. \end{aligned} \tag{5.12}$$

Step 3: final bound. Consider the union of events where (5.10), (5.11), and (5.12) hold. Note that the probability measure of this event is at least $1 - \delta$. Moreover, on this event, we have

$$\begin{aligned}
 & \text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} - \Delta_{\mathcal{D}} \leq T_1 + T_2 + T_3 \\
 & \lesssim \sqrt{\frac{(\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}}) [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}} \mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\
 & \quad + \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}}) [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}} \mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\
 & \quad + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}} \mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)}{n}. \\
 & \leq \sqrt{\frac{(\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} - \Delta_{\mathcal{D}})_+ [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}} \mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\
 & \quad + 2\sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}}) [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}} \mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\
 & \quad + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}} \mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)}{n}.
 \end{aligned}$$

Since the inequality $x \leq 2a\sqrt{x} + b$ yields $\sqrt{x} \leq a + \sqrt{a^2 + b}$ and, hence, $x \leq 4a^2 + 2b$, we obtain that

$$\begin{aligned}
 \text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} - \Delta_{\mathcal{D}} & \leq (\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} - \Delta_{\mathcal{D}})_+ \\
 & \lesssim \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}}) [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}} \mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\
 & \quad + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}} \mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)}{n}
 \end{aligned}$$

on the event of probability measure at least $1 - \delta$.

5.3 Proof of Theorem 2

Theorem 5 (restatement of Theorem 2). *Assume that Ap^* holds with $\beta > 2$. Choose $H_{\mathcal{G}} \geq 2H^*$, $\Lambda_{\mathcal{G}} \geq 2\Lambda$, $D_{\min} \leq \Lambda^{-d}/(\Lambda^d + \Lambda^{-d})$, $D_{\max} \geq \Lambda^d/(\Lambda^d + \Lambda^{-d})$, and $H_{\mathcal{D}}$ large enough. Then there are positive integers $N_{\mathcal{G}}$, $d_{\mathcal{G}}$, $N_{\mathcal{D}}$, $d_{\mathcal{D}}$ and the architectures $\mathcal{A}_{\mathcal{G}} \in \mathbb{N}^{N_{\mathcal{G}}+2}$, $\mathcal{A}_{\mathcal{D}} \in \mathbb{N}^{N_{\mathcal{D}}+2}$ such that AG' and AD' define nonempty sets \mathcal{G} and \mathcal{D} , respectively. Let us consider the estimator (1.3) with $\mathbf{W} = [-1, 1]^{d_{\mathcal{G}}}$, $\Theta = [-1, 1]^{d_{\mathcal{D}}}$, $g_w \in \mathcal{G}$ and $D_{\theta} \in \mathcal{D}$. For any $\delta \in (0, 1)$, with probability at least $1 - \delta$, $\mathbf{p}_{\hat{w}}$ satisfies the inequality*

$$\text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) \lesssim \left(\frac{\log n}{n} \right)^{2\beta/(2\beta+d)} + \frac{\log(1/\delta)}{n},$$

provided that $n \geq n_0$ with n_0 depending only on β, d, Λ, H^* , and $H_{\mathcal{G}}$. In (4.6), the notation \lesssim stands for an inequality up to a multiplicative constant depending on $d, \beta, H^*, H_{\mathcal{G}}$, and $H_{\mathcal{D}}$ only.

Applying Theorem 1, we obtain that, for any $\delta \in (0, 1)$, the GAN estimate $\mathbf{p}_{\hat{w}}$ satisfies the inequality

$$\begin{aligned} \text{JS}(\mathbf{p}_{\hat{w}}, \mathbf{p}^*) - \Delta_{\mathcal{G}} - \Delta_{\mathcal{D}} &\lesssim \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}})[(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)]}{n}} \\ &+ C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathbf{L}_{\mathcal{G}}\mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1)n) + \log(8/\delta)}{n} \end{aligned} \quad (5.13)$$

on an event with probability at least $1 - \delta$. Let us recall that

$$\Delta_{\mathcal{G}} = \min_{w \in \mathcal{W}} \text{JS}(\mathbf{p}_w, \mathbf{p}^*), \quad \text{and} \quad \Delta_{\mathcal{D}} = \max_{w \in \mathcal{W}} \min_{\theta \in \Theta} [\text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta)].$$

In the rest of the proof, we provide upper bounds on the approximation terms $\Delta_{\mathcal{G}}$ and $\Delta_{\mathcal{D}}$ and specify the architectures of neural networks and the parameters $d_{\mathcal{G}}$ and $d_{\mathcal{D}}$ as well. Our approach relies on the following result, concerning approximation properties of neural networks with ReQU activations.

Theorem 6 (Belomestny et al. (2023), Theorem 2). *Let $\beta > 2$ and let $p, d \in \mathbb{N}$. Then, for any $H > 0$, $f : [0, 1]^d \rightarrow \mathbb{R}^p$, $f \in \mathcal{H}^{\beta}([0, 1]^d, H)$ and any integer $K \geq 2$, there exists a neural network $h_f : [0, 1]^d \rightarrow \mathbb{R}^p$ of the width*

$$(4d(K + \lfloor \beta \rfloor)^d) \vee 12 ((K + 2\lfloor \beta \rfloor) + 1) \vee p$$

with

$$6 + 2(\lfloor \beta \rfloor - 2) + \lceil \log_2 d \rceil + 2(\lceil \log_2(2d\lfloor \beta \rfloor) + d \rceil \vee \log_2 \log_2 H \vee 1)$$

hidden layers and at most $p(K + \lfloor \beta \rfloor)^d C(\beta, d, H)$ non-zero weights taking their values in $[-1, 1]$, such that, for any $\ell \in \{0, \dots, \lfloor \beta \rfloor\}$,

$$\|f - h_f\|_{\mathcal{H}^{\ell}([0, 1]^d)} \leq \frac{(\sqrt{2ed})^{\beta} H}{K^{\beta - \ell}} + \frac{9^{d(\lfloor \beta \rfloor - 1)} (2\lfloor \beta \rfloor + 1)^{2d + \ell} (\sqrt{2ed})^{\beta} H}{K^{\beta - \ell}}. \quad (5.14)$$

The above constant $C(\beta, d, H)$ is given by

$$\begin{aligned} C(\beta, d, H) &= (60(\lceil \log_2(2d\lfloor \beta \rfloor) + d \rceil \vee \log_2 \log_2 H \vee 1) + 38) \\ &+ 20d^2 + 144d\lfloor \beta \rfloor + 8d. \end{aligned} \quad (5.15)$$

We split the proof of Theorem 2 into several steps for convenience.

Step 1: bounding $\Delta_{\mathcal{G}}$. According to Lemma 7, it holds that

$$\Delta_{\mathcal{G}} = \min_{w \in \mathcal{W}} \text{JS}(\mathbf{p}_w, \mathbf{p}^*) \leq \min_{w \in \mathcal{W}} \frac{\log 2}{2} \int \frac{(\mathbf{p}_w(x) - \mathbf{p}^*(x))^2}{\mathbf{p}_w(x) + \mathbf{p}^*(x)} d\mu.$$

Note that under Assumptions AG' and Ap*, the densities \mathbf{p}^* and \mathbf{p}_w are bounded away from zero and infinity.

Lemma 8. *Let a map $g : [0, 1]^d \rightarrow [0, 1]^d$ be from the class $\mathcal{H}_\Lambda^2([0, 1]^d, H_G)$ and let*

$$\mathfrak{p}(x) = |\det[\nabla g(g^{-1}(x))]|^{-1}, \quad x \in [0, 1]^d.$$

Then it holds that

$$\mathfrak{p}_{\min} \leq \mathfrak{p}(x) \leq \mathfrak{p}_{\max} \quad \text{for any } x \in [0, 1]^d,$$

where

$$\mathfrak{p}_{\min} = \Lambda^{-d}, \quad \mathfrak{p}_{\max} = \Lambda^d. \quad (5.16)$$

The proof of Lemma 8 is provided in Appendix A.2. Lemma 8 implies that

$$\Delta_{\mathcal{G}} \lesssim \Lambda^d \min_{w \in \mathcal{W}} \|\mathfrak{p}_w - \mathfrak{p}^*\|_{L^\infty([0,1]^d)}^2.$$

Combining this bound with the result of Lemma 3, we obtain that

$$\Delta_{\mathcal{G}} \lesssim \Lambda^{9d} \min_{w \in \mathcal{W}} \|g_w - g^*\|_{\mathcal{H}^1([0,1]^d)}^2.$$

Let us introduce an auxiliary parameter

$$K = \left\lceil \frac{n}{\log n} \right\rceil^{1/(2\beta+d)}$$

and set $N_{\mathcal{G}}$ and $d_{\mathcal{G}}$ as follows:

$$N_{\mathcal{G}} = 6 + 2(\lfloor \beta \rfloor - 1) + \lceil \log_2 d \rceil + 2(\lceil \log_2(2d\lfloor \beta \rfloor + 3d) \vee \log_2 \log_2 H^* \rceil \vee 1),$$

$$d_{\mathcal{G}} = d C(\beta + 1, d, H^*) (K + \lfloor \beta \rfloor + 1)^d \lesssim \left(\frac{n}{\log n} \right)^{d/(2\beta+d)}, \quad (5.17)$$

where the constant $C(\beta, d, H^*)$ is given by (5.15). According to Theorem 6, there is a neural network g_w of the width

$$(4d(K + \lfloor \beta \rfloor + 1)^d) \vee 12((K + 2\lfloor \beta \rfloor) + 3) \vee d,$$

with ReQU activations, $N_{\mathcal{G}}$ hidden layers and at most $d_{\mathcal{G}}$ non-zero weights taking their values in $[-1, 1]$, such that

$$\|g^* - g_w\|_{\mathcal{H}^1([0,1]^d)} \leq \frac{(\sqrt{2}ed)^{\beta+1} H^*}{K^\beta} + \frac{9^{d\lfloor \beta \rfloor} (2\lfloor \beta \rfloor + 3)^{2d+1} (\sqrt{2}ed)^{\beta+1} H^*}{K^\beta} \lesssim \left(\frac{\log n}{n} \right)^{\beta/(2\beta+d)}$$

and

$$\|g^* - g_w\|_{\mathcal{H}^2([0,1]^d)} \lesssim \left(\frac{\log n}{n} \right)^{(\beta-1)/(2\beta+d)}.$$

Hence, it holds that

$$\Delta_{\mathcal{G}} \lesssim \Lambda^{9d} \min_{w \in \mathcal{W}} \|g_w - g^*\|_{\mathcal{H}^1([0,1]^d)}^2 \lesssim \Lambda^{9d} \left(\frac{\log n}{n} \right)^{2\beta/(2\beta+d)}. \quad (5.18)$$

Step 2: bounding $\Delta_{\mathcal{D}}$. According to Lemma 1, it holds that

$$\Delta_{\mathcal{D}} \lesssim \max_{w \in \mathcal{W}} \min_{\theta \in \Theta} \left\| \frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} - D_{\theta} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}^2 \lesssim \max_{w \in \mathcal{W}} \min_{\theta \in \Theta} \left\| \frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} - D_{\theta} \right\|_{L_{\infty}([0,1]^d)}^2.$$

Our goal is to show that it is possible to approximate $\mathbf{p}^*/(\mathbf{p}^* + \mathbf{p}_w)$ with a neural network, taking its values in $[D_{\min}, D_{\max}]$. Using (4.2), Lemma 8, and the fact that $g_w \in \mathcal{H}^{\beta+1}(H_{\mathcal{G}}, [0, 1]^d)$, $g^* \in \mathcal{H}^{\beta+1}(H^*, [0, 1]^d)$, it is straightforward to check that, for any $w \in \mathcal{W}$,

$$\frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} \in \mathcal{H}^{\beta}(H^{\circ}, [0, 1]^d),$$

where H° is a constant, depending on $H_{\mathcal{G}}, H^*, \beta$, and d . Let us set the parameters $N_{\mathcal{D}}$, and $d_{\mathcal{D}}$ equal to

$$N_{\mathcal{D}} = 6 + 2(\lfloor \beta \rfloor - 2) + \lceil \log_2 d \rceil + 2(\lceil \log_2(2d\lfloor \beta \rfloor + d) \vee \log_2 \log_2 H^{\circ} \rceil \vee 1),$$

$$d_{\mathcal{D}} = d C(\beta, d, H^*) (K + \lfloor \beta \rfloor)^d \lesssim \left(\frac{n}{\log n} \right)^{d/(2\beta+d)}, \quad (5.19)$$

where the constant $C(\beta, d, H^{\circ})$ is defined in (5.15). Then, according to Theorem 6, for any $w \in \mathcal{W}$, there exists $\theta(w) \in \Theta$ and a neural network $D_{\theta(w)}$ of the width

$$(4d(K + \lfloor \beta \rfloor)^d) \vee 12((K + 2\lfloor \beta \rfloor) + 1),$$

with ReQU activations, $L_{\mathcal{D}}$ hidden layers and at most $d_{\mathcal{D}}$ non-zero weights taking their values in $[-1, 1]$, such that

$$\begin{aligned} \left\| \frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} - D_{\theta(w)} \right\|_{L_{\infty}([0,1]^d)} &\leq \frac{(\sqrt{2}ed)^{\beta} H^{\circ}}{K^{\beta}} + \frac{9^{d(\lfloor \beta \rfloor - 1)} (2\lfloor \beta \rfloor + 1)^{2d} (\sqrt{2}ed)^{\beta} H^{\circ}}{K^{\beta}} \\ &\lesssim \left(\frac{\log n}{n} \right)^{2\beta/(2\beta+d)}. \end{aligned}$$

Thus, we have

$$\Delta_{\mathcal{D}} \lesssim \max_{w \in \mathcal{W}} \min_{\theta \in \Theta} \left\| \frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} - D_{\theta} \right\|_{L_{\infty}([0,1]^d)}^2 \lesssim \left(\frac{\log n}{n} \right)^{2\beta/(2\beta+d)}. \quad (5.20)$$

Step 3: final bound. According to Lemma 2 and Lemma 4, it holds that

$$\log(\mathbf{L}_{\mathcal{G}} \mathbf{L}_{\mathcal{X}} \vee \mathbf{L}_{\Theta} \vee \mathbf{L}_{\mathbf{p}} \vee 1) \lesssim \log(K\Lambda).$$

This and the inequalities (5.13), (5.17), (5.18), (5.19), and (5.20) yield that

$$\begin{aligned}
 \text{JS}(\mathfrak{p}_{\hat{w}}, \mathfrak{p}^*) &\lesssim \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}}) [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathsf{L}_{\mathcal{G}}\mathsf{L}_{\mathcal{X}} \vee \mathsf{L}_{\Theta} \vee \mathsf{L}_{\mathsf{P}} \vee 1)n) + \log(8/\delta)]}{n}} \\
 &\quad + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2(\mathsf{L}_{\mathcal{G}}\mathsf{L}_{\mathcal{X}} \vee \mathsf{L}_{\Theta} \vee \mathsf{L}_{\mathsf{P}} \vee 1)n) + \log(8/\delta)}{n} \\
 &\lesssim \sqrt{\frac{(\Delta_{\mathcal{G}} + \Delta_{\mathcal{D}}) [(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2K\Lambda)]}{n}} + \log(8/\delta) \\
 &\quad + C_{3.2} \cdot \frac{(d_{\mathcal{G}} + d_{\mathcal{D}}) \log(2K\Lambda)}{n} + \log(8/\delta)n \\
 &\lesssim \Lambda^{9d/2} \left(\frac{\log n}{n} \right)^{\frac{2\beta}{2\beta+d}} + \frac{C_{3.2}}{n} \cdot \left(\log(2\Lambda n) \left(\frac{n}{\log n} \right)^{\frac{d}{2\beta+d}} + \log(8/\delta) \right) \\
 &\lesssim \left(\Lambda^{9d/2} + C_{3.2} \log \Lambda \right) \left(\frac{\log n}{n} \right)^{2\beta/(2\beta+d)} + C_{3.2} \cdot \frac{\log(1/\delta)}{n}
 \end{aligned}$$

with probability at least $1 - \delta$. The proof is finished.

5.4 Proof of Theorem 3

Theorem 7 (restatement of Theorem 3). *Let (X_1, \dots, X_n) be a sample of i.i.d. observations generated from a density \mathfrak{p}^* satisfying Ap^* . Then for any estimate $\hat{\mathfrak{p}}$ of \mathfrak{p}^* , that is, measurable function of X_1, \dots, X_n , it holds that*

$$\sup_{\mathfrak{p}^*} \mathbb{E} \text{JS}(\hat{\mathfrak{p}}, \mathfrak{p}^*) \gtrsim n^{-2\beta/(2\beta+d)}$$

with a hidden constant depending on d only.

Let

$$h = h(n) = \left(\frac{1}{n\Lambda^d d^2} \right)^{1/(2\beta+d)}.$$

and let M be the $(2h)$ -packing number of $[1/3, 2/3]^d$. It is clear that $M \gtrsim h^{-d}$. Let $\{x_1, \dots, x_M\} \subseteq [1/3, 2/3]^d$ be a $2h$ -separable set, that is, $\|x_i - x_j\| > 2h$ for all $i \neq j$.

Let $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ be a function from the class $\mathcal{H}^{\beta+2}(\mathbb{R}^d, H_\varphi)$ with some $H_\varphi > 0$, such that $\text{supp}(\varphi) = \mathcal{B}(0, 1) \subset \mathbb{R}^d$, φ attains its maximum at 0, and $-\nabla^2 \varphi(0) \succeq I_{d \times d}$. Consider a parametric class of generators $\{g_\theta : \theta \in \mathcal{B}(0, h) \subset \mathbb{R}^M\}$, where

$$g_\theta^{-1}(x) = x + h^\beta \sum_{j=1}^M \theta_j \nabla \varphi \left(\frac{x - x_j}{h} \right), \quad x \in [0, 1]^d.$$

Here and further in the proof, g^{-1} stands for the inverse map of g . For convenience, we split the rest of the proof into several steps.

Step 1: verifying the conditions. First, note that, according to the definition of g_θ , we have

$$\nabla g_\theta^{-1}(x) = I_{d \times d} + h^{\beta-1} \sum_{j=1}^M \theta_j \nabla^2 \varphi \left(\frac{x - x_j}{h} \right).$$

Since φ is supported on the unit ball $\mathcal{B}(0, 1)$, the last expression simplifies to

$$\nabla g_\theta^{-1}(x) = \begin{cases} I_{d \times d} + h^{\beta-1} \theta_j \nabla^2 \varphi \left(\frac{x-x_j}{h} \right), & \text{if } x \in \mathcal{B}(x_j, h) \text{ for some } j \in \{1, \dots, M\}, \\ I_{d \times d}, & \text{otherwise.} \end{cases}$$

The fact that $\varphi \in \mathcal{H}^{\beta+2}(\mathbb{R}^d, H_\varphi)$ immediately implies that

$$e^{-1/d} I_{d \times d} \preceq \left(1 - h^\beta H_\varphi\right) I_{d \times d} \preceq \nabla g_\theta^{-1}(x) \preceq \left(1 + h^\beta H_\varphi\right) I_{d \times d} \preceq e^{1/d} I_{d \times d}, \quad (5.21)$$

provided that

$$h^\beta \leq \frac{e^{1/d} - 1}{H_\varphi} \wedge \frac{1 - e^{-1/d}}{H_\varphi},$$

that is, if the sample size n is large enough. Hence, we checked that $g_\theta^{-1} \in \mathcal{H}_{e^{1/d}}^{\beta+1}([0, 1]^d, H_0)$ for some $H_0 > 0$ and each $\theta \in \mathcal{B}(0, h) \subset \mathbb{R}^M$. This yields that, for all $\theta \in \mathcal{B}(0, h) \subset \mathbb{R}^M$, g_θ belongs to a class $\mathcal{H}_{e^{1/d}}^{\beta+1}([0, 1]^d, H^\circ)$ with some $H^\circ > 0$.

Besides, for each $\theta \in \mathcal{B}(0, h)$, g_θ differs from $g_0(x) \equiv x$ only on $[1/3 - h, 2/3 + h]^d$. If n is sufficiently large, then h is small and $g_\theta^{-1}([1/3 - h, 2/3 + h]^d) \subset [0, 1]^d$. Then $g_\theta^{-1}([0, 1]^d) \subseteq [0, 1]^d$. Similarly, we can show that $g_\theta([0, 1]^d) \subseteq [0, 1]^d$. Hence, for any $\theta \in \mathcal{B}(0, h)$, g_θ is indeed a bijection between $[0, 1]^d$ and $[0, 1]^d$, provided that h is small enough.

Step 2: a minimax lower bound on the accuracy of parametric estimation. The next auxiliary result provides a lower bound on the accuracy of estimation of the parameter θ .

Lemma 9. *Let Y_1, \dots, Y_n be i.i.d. random elements on \mathcal{Y} , $\text{Vol}(\mathcal{Y}) = 1$, drawn from the uniform distribution on \mathcal{Y} and assume that a learner observes a sample (X_1, \dots, X_n) where $X_i = g_\theta(Y_i), i \in \{1, \dots, n\}$. Under the assumptions of Theorem 3, if n is large enough, for any estimate $\hat{\theta}$ taking its values in $\mathcal{B}(0, h) \subset \mathbb{R}^M$ it holds that*

$$\sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\hat{\theta} - \theta\|^2 \gtrsim \frac{Mh^2}{nh^{2\beta+d}d^2 + 1}.$$

Here we write \mathbf{E}_θ to emphasize that the expectation is taken with respect to the probability measure $\mathbb{P}_{g_\theta}^{\otimes n}$ where \mathbb{P}_{g_θ} is a distribution of random element $X = g_\theta(Y)$ with Y uniformly distributed on $[0, 1]^d$. The proof of Lemma 9 is based on the van Trees inequality and it is postponed to Appendix D.1 below.

Step 3: a minimax lower bound on the accuracy of density estimation. Denote a density of the measure \mathbb{P}_{g_θ} by \mathfrak{p}_{g_θ} . Our next goal is to convert the result of Lemma 9 to the lower bound on

$$\inf_{\hat{\theta}} \sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\mathfrak{p}_{g_{\hat{\theta}}} - \mathfrak{p}_{g_\theta}\|_{L_2(\mathcal{X})}^2$$

where the infimum is taken with respect to measurable functions $\hat{\theta}$ of the sample satisfying assumptions of Lemma 9. Recall that, for any $\theta \in \mathcal{B}(0, h)$, the log-density of r.v. $X = g_\theta(Y)$

is given by

$$\log \mathbf{p}_{g_\theta}(x) = \log \det \left(I_{d \times d} + h^{\beta-1} \sum_{j=1}^M \theta_j \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right). \quad (5.22)$$

Before obtaining a minimax lower bound on $\sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\mathbf{p}_{g_{\hat{\theta}}} - \mathbf{p}_{g_\theta}\|_{L_2(\mathcal{X})}^2$, we prove some properties of the map $\theta \mapsto \log \mathbf{p}_{g_\theta}$. The next lemma shows that, for any $j \in \{1, \dots, M\}$, the partial derivative

$$\frac{\partial \log \mathbf{p}_{g_\theta}(x)}{\partial \theta_j}$$

is bounded away from zero in a vicinity of x_j .

Lemma 10. *Let*

$$r_0 = \frac{h}{2(H_\varphi \vee 1)}.$$

Then, under the assumptions of Theorem 3, for a sufficiently small $h > 0$ and any estimate $\hat{\theta}$ taking its values in $\mathcal{B}(0, h) \subset \mathbb{R}^M$, it holds that

$$\left| \frac{\partial \log \mathbf{p}_{g_\theta}(x)}{\partial \theta_j} \right| \geq h^{\beta-1} d/2, \quad \text{for all } x \in \mathcal{B}(x_j, r_0).$$

The proof of Lemma 10 is postponed until Appendix D.2. Lemma 10 implies that, for any $j \in \{1, \dots, M\}$,

$$\left| \frac{\partial \mathbf{p}_{g_\theta}(x)}{\partial \theta_j} \right| = \frac{1}{\mathbf{p}_{g_\theta}(x)} \left| \frac{\partial \log \mathbf{p}_{g_\theta}(x)}{\partial \theta_j} \right| \geq \frac{h^{\beta-1} d}{2\mathbf{p}_{g_\theta}(x)}, \quad \text{for all } x \in \mathcal{B}(x_j, r_0).$$

The equality (5.22) implies that for $x \in \mathcal{B}(x_j, r_0)$, $\frac{\partial \mathbf{p}_{g_\theta}(x)}{\partial \theta_i} = 0$ for $i \neq j$. Hence, for $\theta, \theta' \in \mathcal{B}(0, h)$, it holds that

$$\begin{aligned} \|\mathbf{p}_{g_{\theta'}} - \mathbf{p}_{g_\theta}\|_{L_2(\mathcal{X})}^2 &\geq \sum_{j=1}^M \int_{\mathcal{B}(x_j, r_0)} (\mathbf{p}_{g_{\theta'}}(x) - \mathbf{p}_{g_\theta}(x))^2 dx \\ &\geq \sum_{j=1}^M \int_{\mathcal{B}(x_j, r_0)} \frac{h^{2\beta-2} d^2 (\theta'_j - \theta_j)^2}{4\mathbf{p}_{g_\theta(x)}(x)^2} dx \end{aligned}$$

for some $\vartheta(x) = t(x)\theta + (1-t(x))\theta'$, $t(x) \in [0, 1]$. The inequality (5.21) and Lemma 8 imply that

$$\begin{aligned} \|\mathbf{p}_{g_{\theta'}} - \mathbf{p}_{g_\theta}\|_{L_2(\mathcal{X})}^2 &\geq \sum_{j=1}^M \int_{\mathcal{B}(x_j, r_0)} \frac{h^{2\beta-2} d^2}{4e^2} (\theta'_j - \theta_j)^2 dx \\ &\gtrsim \sum_{j=1}^M \left(\frac{r_0}{h} \right)^d h^{2\beta-2+d} (\theta'_j - \theta_j)^2 \\ &\gtrsim h^{2\beta-2+d} \|\theta' - \theta\|^2, \end{aligned}$$

where we write \gtrsim for inequality up to an absolute constant in power d . Hence, Lemma 9 implies that

$$\begin{aligned} \sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\mathbf{p}_{g_{\hat{\theta}}} - \mathbf{p}_{g_\theta}\|_{L_2(\mathbf{X})}^2 &\gtrsim h^{2\beta-2+d} \sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\hat{\theta} - \theta\|^2 \\ &\gtrsim \frac{h^{2\beta}}{nh^{2\beta+d}d^2 + 1} \end{aligned}$$

for any estimate $\hat{\theta}$ taking its values in $\mathcal{B}(0, h)$. Since $h = (nd^2)^{-1/(2\beta+d)}$, we get

$$\sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\mathbf{p}_{g_{\hat{\theta}}} - \mathbf{p}_{g_\theta}\|_{L_2(\mathbf{X})}^2 \gtrsim n^{-2\beta/(2\beta+d)}, \quad (5.23)$$

where the hidden constant depends on d only.

Step 4: a minimax lower bound on the accuracy of nonparametric estimation.

We now prove that (5.23) yields

$$\inf_{\hat{\mathbf{p}}} \sup_{\mathbf{p}^*} \mathbf{E} \|\hat{\mathbf{p}} - \mathbf{p}^*\|_{L_2(\mathbf{X})}^2 \gtrsim n^{-2\beta/(2\beta+d)},$$

where \mathbf{p}^* satisfies Ap^* and $\hat{\mathbf{p}}$ is any estimate of \mathbf{p}^* . Let C_d be the hidden constant in (5.23), that is,

$$\sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\mathbf{p}_{g_{\hat{\theta}}} - \mathbf{p}_{g_\theta}\|_{L_2(\mathbf{X})}^2 \geq C_d n^{-2\beta/(2\beta+d)}. \quad (5.24)$$

Let $\hat{\mathbf{p}}$ be an arbitrary estimate. Since $\sup_{\mathbf{p}^*} \mathbf{E} \|\hat{\mathbf{p}} - \mathbf{p}^*\|_{L_2(\mathbf{X})}^2 \geq \sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\hat{\mathbf{p}} - \mathbf{p}_{g_\theta}\|_{L_2(\mathbf{X})}^2$, it is enough to show that

$$\sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\hat{\mathbf{p}} - \mathbf{p}_{g_\theta}\|_{L_2(\mathbf{X})}^2 \gtrsim n^{-2\beta/(2\beta+d)}.$$

Let us introduce

$$\tilde{\theta} \in \operatorname{argmin}_{\|\theta_0\| \leq h} \|\hat{\mathbf{p}} - \mathbf{p}_{g_{\theta_0}}\|_{L_2(\mathbf{X})}^2.$$

If

$$\sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\hat{\mathbf{p}} - \mathbf{p}_{g_{\tilde{\theta}}}\|_{L_2(\mathbf{X})}^2 \geq \frac{C_d}{4} n^{-2\beta/(2\beta+d)},$$

then, by the definition of $\tilde{\theta}$,

$$\begin{aligned} \sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\hat{\mathbf{p}} - \mathbf{p}_{g_\theta}\|_{L_2(\mathbf{X})}^2 &\geq \sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\hat{\mathbf{p}} - \mathbf{p}_{g_{\tilde{\theta}}}\|_{L_2(\mathbf{X})}^2 \\ &\geq \frac{C_d}{4} n^{-2\beta/(2\beta+d)}. \end{aligned}$$

On the other hand, if

$$\sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\hat{\mathbf{p}} - \mathbf{p}_{g_{\tilde{\theta}}}\|_{L_2(\mathbf{X})}^2 \leq \frac{C_d}{4} n^{-2\beta/(2\beta+d)},$$

then the Cauchy-Schwarz inequality and (5.24) yield

$$\begin{aligned} \sup_{\|\theta\| \leq h} \mathbb{E}_\theta \|\hat{\mathbf{p}} - \mathbf{p}_{g_\theta}\|_{L_2(\mathbf{X})}^2 &\geq \frac{1}{2} \sup_{\|\theta\| \leq h} \mathbb{E}_\theta \|\mathbf{p}_{g_\theta} - \mathbf{p}_{g_{\bar{\theta}}}\|_{L_2(\mathbf{X})}^2 - \sup_{\|\theta\| \leq h} \mathbb{E}_\theta \|\hat{\mathbf{p}} - \mathbf{p}_{g_{\bar{\theta}}}\|_{L_2(\mathbf{X})}^2 \\ &\geq \frac{C_d}{4} n^{-2\beta/(2\beta+d)}. \end{aligned}$$

Finally, Lemma 7, Lemma 8, and (5.21) yield that, for any $\hat{\mathbf{p}}$, it holds almost surely

$$\text{JS}(\hat{\mathbf{p}}, \mathbf{p}^*) \geq \frac{\|\hat{\mathbf{p}} - \mathbf{p}^*\|_{L_2(\mathbf{X})}^2}{8e}.$$

Hence, for any $\hat{\mathbf{p}}$, we obtain that

$$\sup_{\mathbf{p}^*} \mathbb{E} \text{JS}(\hat{\mathbf{p}}, \mathbf{p}^*) \geq \sup_{\mathbf{p}^*} \frac{\mathbb{E} \|\hat{\mathbf{p}} - \mathbf{p}^*\|_{L_2(\mathbf{X})}^2}{8e} \gtrsim C_d n^{-2\beta/(2\beta+d)},$$

where \gtrsim stands for an inequality up to an absolute constant.

6. Conclusion and future directions

Despite the huge recent interest to theoretical properties of generative adversarial networks, the existing papers mostly focus on the generalization ability of GANs, missing the issues occurring in their practical use. For instance, one of the most challenging problems in the GAN training process is the mode collapse phenomenon (Salimans et al., 2016; Che et al., 2016), which appears for various loss functions, including Wasserstein GANs, “vanilla” GANs, or GANs with other divergence measures. Moreover, in the present literature on Wasserstein GANs, the estimation error is mostly governed by the rates of convergence of the empirical measure to the population distribution. This leaves a question, whether WGANs are able to produce a distribution estimate, which is strictly better than just the empirical distribution. The first step in this direction was made in the recent work (Vardanyan et al., 2023), where the authors imposed regularity conditions on generators. This agrees with the papers on vanilla GANs in a sense that successful generalization requires smoothness of generators and discriminators. Practitioners often use regularization (see, e.g., Gulrajani et al. (2017)) for this purpose.

Acknowledgments

The publication was supported by the grant for research centers in the field of AI provided by the Analytical Center for the Government of the Russian Federation (ACRF) in accordance with the agreement on the provision of subsidies (identifier of the agreement 000000D730321P5Q0002) and the agreement with HSE University №70-2021-00139.

References

- 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, 2017.

- H. Asatryan, H. Gottschalk, M. Lippert, and M. Rottmann. A convenient infinite dimensional framework for generative adversarial learning. *Electronic Journal of Statistics*, 17(1):391–428, 2023.
- J.-Y. Audibert and A. B. Tsybakov. Fast learning rates for plug-in classifiers. *The Annals of Statistics*, 35(2):608 – 633, 2007.
- D. Belomestny and V. Spokoiny. Spatial aggregation of local likelihood estimates with applications to classification. *Ann. Statist.*, 35(5):2287–2311, 2007. ISSN 0090-5364. URL <https://doi.org/10.1214/009053607000000271>.
- D. Belomestny, A. Naumov, N. Puchkin, and S. Samsonov. Simultaneous approximation of a smooth function and its derivatives by deep neural networks with piecewise-polynomial activations. *Neural Networks*, 161:242–253, 2023.
- G. Biau, B. Cadre, M. Sangnier, and U. Tanielian. Some theoretical properties of GANs. *Annals of Statistics*, 48(3):1539–1566, 2020.
- G. Biau, M. Sangnier, and U. Tanielian. Some theoretical insights into Wasserstein GANs. *The Journal of Machine Learning Research*, 22(1):5287–5331, 2021.
- S. Boucheron, G. Lugosi, and O. Bousquet. Concentration inequalities. In O. Bousquet, U. von Luxburg, and G. Rätsch, editors, *Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2 - 14, 2003, Tübingen, Germany, August 4 - 16, 2003, Revised Lectures*, pages 208–240, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg. ISBN 978-3-540-28650-9. doi: 10.1007/978-3-540-28650-9_9. URL https://doi.org/10.1007/978-3-540-28650-9_9.
- Y. Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Comm. Pure Appl. Math.*, 44(4):375–417, 1991. ISSN 0010-3640. doi: 10.1002/cpa.3160440402. URL <https://doi.org/10.1002/cpa.3160440402>.
- F. Bunea, A. B. Tsybakov, and M. H. Wegkamp. Sparse density estimation with l1 penalties. In *International Conference on Computational Learning Theory*, pages 530–543. Springer, 2007.
- C. Butucea, J.-F. Delmas, A. Dutfoy, and R. Fischer. Optimal exponential bounds for aggregation of estimators for the Kullback-Leibler loss. *Electron. J. Stat.*, 11(1):2258–2294, 2017. ISSN 1935-7524. URL <https://doi.org/10.1214/17-EJS1269>.
- L. A. Caffarelli. Some regularity properties of solutions of Monge Ampère equation. *Comm. Pure Appl. Math.*, 44(8-9):965–969, 1991. ISSN 0010-3640. doi: 10.1002/cpa.3160440809. URL <https://doi.org/10.1002/cpa.3160440809>.
- L. A. Caffarelli. Boundary regularity of maps with convex potentials. *Comm. Pure Appl. Math.*, 45(9):1141–1151, 1992a. ISSN 0010-3640. doi: 10.1002/cpa.3160450905. URL <https://doi.org/10.1002/cpa.3160450905>.

- L. A. Caffarelli. The regularity of mappings with a convex potential. *J. Amer. Math. Soc.*, 5(1):99–104, 1992b. ISSN 0894-0347. doi: 10.2307/2152752. URL <https://doi.org/10.2307/2152752>.
- L. A. Caffarelli. Boundary regularity of maps with convex potentials. II. *Ann. of Math.*, 144(3):453–496, 1996. ISSN 0003-486X. doi: 10.2307/2118564. URL <https://doi.org/10.2307/2118564>.
- T. Che, Y. Li, A. P. Jacob, Y. Bengio, and W. Li. Mode regularized generative adversarial networks. *arXiv preprint arXiv:1612.02136*, 2016.
- M. Chen, W. Liao, H. Zha, and T. Zhao. Distribution approximation and statistical estimation guarantees of generative adversarial networks. Preprint, arXiv:2002.03938, 2022.
- A. S. Dalalyan and M. Sebbar. Optimal kullback–leibler aggregation in mixture density estimation by maximum likelihood. *Mathematical Statistics and Learning*, 1(1):1–35, 2018.
- S. Dasgupta and S. Kpotufe. Optimal rates for k-nn density and mode estimation. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL <https://proceedings.neurips.cc/paper/2014/file/bb7946e7d85c81a9e69fee1cea4a087c-Paper.pdf>.
- D. L. Donoho, I. M. Johnstone, G. Kerkycharian, and D. Picard. Density estimation by wavelet thresholding. *The Annals of Statistics*, 24(2):508 – 539, 1996. doi: 10.1214/aos/1032894451. URL <https://doi.org/10.1214/aos/1032894451>.
- G. K. Dziugaite, D. M. Roy, and Z. Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. *arXiv preprint arXiv:1505.03906*, 2015.
- D. Endres and J. Schindelin. A new metric for probability distributions. *IEEE Transactions on Information Theory*, 49(7):1858–1860, 2003.
- S. Feizi, F. Farnia, T. Ginart, and D. Tse. Understanding GANs in the LQG setting: Formulation, generalization and stability. *IEEE Journal on Selected Areas in Information Theory*, 2020.
- R. D. Gill and B. Y. Levit. Applications of the Van Trees inequality: a Bayesian Cramér–Rao bound. *Bernoulli*, 1(1-2):59–79, 1995. ISSN 1350-7265. doi: 10.2307/3318681. URL <https://doi.org/10.2307/3318681>.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

- H. Kim, P. Garrido, A. Tewari, W. Xu, J. Thies, M. Niessner, P. Pérez, C. Richardt, M. Zollhöfer, and C. Theobalt. Deep video portraits. *ACM Trans. Graph.*, 37(4), 2018. ISSN 0730-0301. doi: 10.1145/3197517.3201283. URL <https://doi.org/10.1145/3197517.3201283>.
- C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. *CVPR, 2017*, pages 105–114, 07 2017. doi: 10.1109/CVPR.2017.19.
- T. Liang. How well generative adversarial networks learn distributions. *Journal of Machine Learning Research*, 22(228):1–41, 2021. URL <http://jmlr.org/papers/v22/20-911.html>.
- G. Luise, M. Pontil, and C. Ciliberto. Generalization properties of optimal transport GANs with latent distribution learning. *arXiv preprint arXiv:2007.14641*, 2020.
- D. McDonald. Minimax Density Estimation for Growing Dimension. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 194–203, 2017. URL <https://proceedings.mlr.press/v54/mcdonald17a.html>.
- S. Nowozin, B. Cseke, and R. Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in neural information processing systems*, pages 271–279, 2016.
- J. Polzehl and V. Spokoiny. Propagation-separation approach for local likelihood estimation. *Probab. Theory Related Fields*, 135(3):335–362, 2006. ISSN 0178-8051. URL <https://doi.org/10.1007/s00440-005-0464-1>.
- A. Rakhlin, D. Panchenko, and S. Mukherjee. Risk bounds for mixture density estimation. *ESAIM: PS*, 9:220–229, 2005.
- P. Rigollet. Kullback-Leibler aggregation and misspecified generalized linear models. *Ann. Statist.*, 40(2):639–665, 2012. ISSN 0090-5364. URL <https://doi.org/10.1214/11-AOS961>.
- P. Rigollet and A. Tsybakov. Linear and convex aggregation of density estimators. *Mathematical Methods of Statistics*, 16:260–280, 2006.
- T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training GANs. In *Advances in neural information processing systems*, pages 2234–2242, 2016.
- J. Schmidt-Hieber. Nonparametric regression using deep neural networks with relu activation function. *Ann. Statist.*, 48(4):1875–1897, 08 2020. 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 *Algorithmic Learning Theory*, pages 1051–1071. PMLR, 2021.

- S. Singh and B. Póczos. Minimax distribution estimation in Wasserstein distance. *arXiv preprint arXiv:1802.08855*, 2018.
- S. Singh, A. Uppal, B. Li, C.-L. Li, M. Zaheer, and B. Póczos. Nonparametric density estimation with adversarial losses. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 10246–10257. Curran Associates Inc., 2018.
- A. B. Tsybakov. *Introduction to nonparametric estimation*. Springer Science & Business Media, 2008.
- A. Uppal, S. Singh, and B. Póczos. Nonparametric density estimation: Convergence rates for GANs under Besov IPM losses. In *Advances in Neural Information Processing Systems 32*, pages 9089–9100. Curran Associates, Inc., 2019.
- H. L. Van Trees. *Detection, Estimation, and Modulation Theory, Part 1*, volume 1. John Wiley and Sons, Inc., New York, 1968.
- E. Vardanyan, A. Minasyan, S. Hunanyan, T. Galstyan, and A. Dalalyan. Guaranteed optimal generative modeling with maximum deviation from the empirical distribution. Preprint, arXiv:2307.16422, 2023.
- J. Zhao, M. Mathieu, and Y. LeCun. Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*, 2016.

Appendix A. Properties of smooth maps of random elements with smooth densities

A.1 Proof of Lemma 3

Due to AG¹, $\nabla g_w(y)$ is non-degenerate for all $w \in W$ and all $y \in [0, 1]^d$. Thus, $\det \nabla g_w$ does not change its sign, and, without loss of generality, we can assume that $\det \nabla g_u(g_u^{-1}(x))$ and $\det \nabla g_v(g_v^{-1}(x))$ are positive for all $x \in [0, 1]^d$. Let $z = g_u^{-1}(x) \in [0, 1]^d$. Then $x = g_u(z)$ and

$$\begin{aligned} \|g_u^{-1}(x) - g_w^{-1}(x)\| &= \|g_u^{-1}(g_u(z)) - g_v^{-1}(g_u(z))\| = \|g_v^{-1}(g_v(z)) - g_v^{-1}(g_u(z))\| \\ &\leq \Lambda \|g_v(z) - g_u(z)\| \leq \Lambda \sqrt{d} \|g_v - g_u\|_{L_\infty([0,1]^d)} \end{aligned}$$

by the mean value theorem for vector valued functions. Furthermore, we have

$$\begin{aligned} |\mathbf{p}_u(x) - \mathbf{p}_v(x)| &= |\det[\nabla g_u(g_u^{-1}(x))]^{-1} - \det[\nabla g_v(g_v^{-1}(x))]^{-1}| \\ &\leq \min \{ \det[\nabla g_u(g_u^{-1}(x))], \det[\nabla g_v(g_v^{-1}(x))] \}^{-2} \\ &\quad \cdot |\det \nabla g_u(g_u^{-1}(x)) - \det \nabla g_v(g_v^{-1}(x))| \\ &\leq \Lambda^{2d} |\det \nabla g_u(g_u^{-1}(x)) - \det \nabla g_v(g_v^{-1}(x))|. \end{aligned}$$

Here we used the fact that, due to AG¹,

$$\begin{aligned} \det[\nabla g_u(g_u^{-1}(x))]^{-1} &= \sqrt{\det([\nabla g_u(g_u^{-1}(x))]^{-\top} [\nabla g_u(g_u^{-1}(x))]^{-1})} \\ &\leq \sqrt{\det(\Lambda^2 \mathbf{I}_{d \times d})} = \Lambda^d, \end{aligned}$$

and, similarly, $\det[\nabla g_v(g_v^{-1}(x))]^{-1} \leq \Lambda^d$. Next, since for any $d \times d$ matrices A and B it holds that

$$\begin{aligned} |\det A - \det B| &\leq \|A - B\|_{\text{F}} \frac{\|A\|_{\text{F}}^d - \|B\|_{\text{F}}^d}{\|A\|_{\text{F}} - \|B\|_{\text{F}}} \\ &\leq d \max\{\|A\|_{\text{F}}^d, \|B\|_{\text{F}}^d\} \|A - B\|_{\text{F}} \end{aligned}$$

and

$$\begin{aligned} \|\nabla g_u(g_u^{-1}(x))\|_{\text{F}}^2 &= \text{Tr} \left(\nabla g_u(g_u^{-1}(x))^\top \nabla g_u(g_u^{-1}(x)) \right) \leq \Lambda^2 d, \\ \|\nabla g_v(g_v^{-1}(x))\|_{\text{F}}^2 &= \text{Tr} \left(\nabla g_v(g_v^{-1}(x))^\top \nabla g_v(g_v^{-1}(x)) \right) \leq \Lambda^2 d, \end{aligned}$$

we obtain that

$$\begin{aligned} |\det \nabla g_u(g_u^{-1}(x)) - \det \nabla g_v(g_v^{-1}(x))| &\leq \Lambda^d d^{1+d/2} \|\nabla g_u(g_u^{-1}(x)) - \nabla g_v(g_v^{-1}(x))\|_{\text{F}} \\ &\leq \Lambda^d d^{1+d/2} \|\nabla g_u(g_u^{-1}(x)) - \nabla g_v(g_u^{-1}(x))\|_{\text{F}} \\ &\quad + \Lambda^d d^{1+d/2} \|\nabla g_v(g_u^{-1}(x)) - \nabla g_v(g_v^{-1}(x))\|_{\text{F}} \\ &\leq \Lambda^d d^{2+d/2} \|g_u - g_v\|_{\mathcal{H}^1(\mathcal{Y})} \\ &\quad + \Lambda^d d^{2+d/2} H_{\mathcal{G}} \|g_u^{-1}(x) - g_v^{-1}(x)\| \\ &\leq \Lambda^d d^{2+d/2} (1 + H_{\mathcal{G}} \Lambda \sqrt{d}) \|g_u - g_v\|_{\mathcal{H}^1(\mathcal{Y})}. \end{aligned}$$

Hence,

$$\|\mathfrak{p}_u - \mathfrak{p}_v\|_{L_\infty(\mathcal{X})} \leq d^{2+d/2} \Lambda^{3d} (1 + H_{\mathcal{G}} \Lambda \sqrt{d}) \|g_u - g_v\|_{\mathcal{H}^1(\mathcal{Y})}.$$

A.2 Proof of Lemma 8

According to the definition of $\mathcal{H}_\Lambda^2([0, 1]^d, H_{\mathcal{G}})$, we have $\Lambda^{-2} I_{d \times d} \preceq \nabla g(y)^\top \nabla g(y) \preceq \Lambda^2 I_{d \times d}$. Hence,

$$\Lambda^{-d} = \sqrt{\det(\Lambda^{-2} I_{d \times d})} \leq \sqrt{\det(\nabla g(g^{-1}(x))^\top \nabla g(g^{-1}(x)))} \leq \sqrt{\det(\Lambda^2 I_{d \times d})} = \Lambda^d.$$

Then the equality

$$\mathfrak{p}(x) = |\det \nabla g(g^{-1}(x))|^{-1}$$

yields that

$$\Lambda^{-d} \leq \mathfrak{p}(x) \leq \Lambda^d.$$

Appendix B. Some properties of feed-forward neural networks with ReQU activations

Let $N \in \mathbb{N}$, $\mathcal{A} = (p_0, p_1, \dots, p_{N+1}) \in \mathbb{N}^{N+2}$, and $f(x) : \mathbb{R}^{p_0} \rightarrow \mathbb{R}^{p_{N+1}}$ be a neural network

$$f(x) = W_N \circ \sigma_{v_N} \circ W_{L-1} \circ \sigma_{v_{L-1}} \circ \dots \circ \sigma_{v_1} \circ W_0 \circ x,$$

where $W_i \in \mathbb{R}^{p_{i+1} \times p_i}$ are weight matrices. Similarly to (Schmidt-Hieber, 2020, Lemma 5), we introduce the following notations. For $k \in \{1, \dots, N+1\}$, $i \leq k$, $k \leq j \leq N$, we define functions $\mathbf{B}_{k,i}(x) : \mathbb{R}^{p_{i-1}} \rightarrow \mathbb{R}^{p_k}$ and $\mathbf{A}_{j,k}f(x) : \mathbb{R}^{p_{k-1}} \rightarrow \mathbb{R}^{p_{j+1}}$ as follows

$$\begin{aligned} \mathbf{B}_{k,i}(x) &= \sigma_{v_k} \circ W_{k-1} \circ \sigma_{v_{k-1}} \cdots \circ \sigma_{v_i} \circ W_{i-1} \circ x; \\ \mathbf{A}_{j,k}(x) &= W_j \circ \sigma_{v_j} \cdots \circ W_k \circ \sigma_{v_k} \circ W_{k-1} \circ x. \end{aligned} \quad (\text{B.1})$$

Set by convention $\mathbf{A}_{N,N+2}(x) = \mathbf{B}_{0,1}f(x) = x$. For notation simplicity, we write \mathbf{A}_j instead of $\mathbf{A}_{N,j}$, and \mathbf{B}_j instead of $\mathbf{B}_{j,1}$. Note that with this notation $f(x) = \mathbf{A}_1(x)$.

Let us introduce the functions

$$\begin{aligned} f^{(1)}(x) &= W_N^{(1)} \circ \sigma_{v_N^{(1)}} \circ W_{N-1}^{(1)} \circ \sigma_{v_{N-1}^{(1)}} \cdots \circ \sigma_{v_1^{(1)}} \circ W_0^{(1)} \circ x, \\ f^{(2)}(x) &= W_N^{(2)} \circ \sigma_{v_N^{(2)}} \circ W_{N-1}^{(2)} \circ \sigma_{v_{N-1}^{(2)}} \cdots \circ \sigma_{v_1^{(2)}} \circ W_0^{(2)} \circ x, \end{aligned} \quad (\text{B.2})$$

where the parameters $W_i^{(1)}, W_i^{(2)}, v_i^{(1)}, v_i^{(2)}$ satisfy

$$\left\| W_i^{(1)} - W_i^{(2)} \right\|_\infty \leq \varepsilon, \quad \left\| v_i^{(1)} - v_i^{(2)} \right\|_\infty \leq \varepsilon, \quad \text{for all } i \in \{0, \dots, N\}.$$

B.1 Proof of Lemma 2

Before we prove Lemma 2, we need the following auxiliary result.

Lemma 11. *Let $x \in \mathbb{R}^d$, $\|x\|_\infty \leq \mathcal{K}$. Then for $k, i \in \{1, \dots, N\}$, $k \geq i$*

$$\|\mathbf{B}_{k,i}(x)\|_\infty \leq \left\{ \prod_{\ell=1}^{k-i+1} (p_{k-\ell} + 1)^{2^\ell} \right\} (\mathcal{K} \vee 1)^{2^{k-i+1}}, \quad (\text{B.3})$$

where $\mathbf{B}_{k,i}(x)$ are defined in (B.1). Moreover, function $\mathbf{A}_{j,k}(x)$ is Lipschitz for $x, y \in \mathbb{R}^d$: $\|x\|_\infty \leq \mathcal{K}, \|y\|_\infty \leq \mathcal{K}$, that is,

$$\|\mathbf{A}_{j,k}(x) - \mathbf{A}_{j,k}(y)\|_\infty \leq 2^{j-k+1} \prod_{\ell=0}^{j-k+1} (p_{j-\ell} + 1)^{2^\ell} (\mathcal{K} \vee 1)^{2^{j-k+1}} \|x - y\|_\infty. \quad (\text{B.4})$$

Proof of Lemma 11 The inequality (B.3) follows from an easy induction in k . Indeed, if $k = i$, $\mathbf{B}_{i,i} = \sigma_{v(i)} \circ W_{i-1}x$, and

$$\|\mathbf{B}_{i,i}\|_\infty \leq (\mathcal{K}p_{i-1} + 1)^2 \leq (p_{i-1} + 1)^2 (\mathcal{K} \vee 1)^2.$$

Using $\|\mathbf{B}_{k,i}\|_\infty \leq (\|\mathbf{B}_{k-1,i}\|_\infty p_{k-1} + 1)^2$ completes the proof.

To prove (B.4), we use an induction in j . Assume that (B.4) holds for any $k \in \{1, \dots, N+1\}$ and $j-1 \geq k$. Then

$$\begin{aligned} \|\mathbf{A}_{j,k}(x) - \mathbf{A}_{j,k}(y)\|_\infty &\leq p_j \|\mathbf{B}_{j,k}(x) - \mathbf{B}_{j,k}(y)\|_\infty \\ &\leq 2p_j \|\mathbf{A}_{j-1,k}(x) - \mathbf{A}_{j-1,k}(y)\|_\infty (\|\mathbf{A}_{j-1,k}(x)\|_\infty \vee \|\mathbf{A}_{j-1,k}(y)\|_\infty) \\ &\leq 2(p_j + 1)(p_{j-1} + 1) \prod_{\ell=1}^{j-k} (p_{j-\ell-1} + 1)^{2^\ell} (\mathcal{K} \vee 1)^{2^{j-k}} \times \\ &\quad \|\mathbf{A}_{j-1,k}(x) - \mathbf{A}_{j-1,k}(y)\|_\infty, \end{aligned}$$

and the statement follows from the elementary bound

$$\|\mathbf{A}_{k-1,k}(x) - \mathbf{A}_{k-1,k}(y)\|_\infty \leq p_{k-1} \|x - y\|_\infty \leq p_{k-1} \|x - y\|_\infty (\mathcal{K} \vee 1).$$

■

Proof of Lemma 2 Denote by $\mathbf{A}_j^{(1)}, \mathbf{B}_j^{(1)}, \mathbf{A}_j^{(2)}, \mathbf{B}_j^{(2)}$ the corresponding functions in (B.1). Following (Schmidt-Hieber, 2020, Lemma 5), we write

$$\|f^{(1)}(x) - f^{(2)}(x)\|_\infty \leq \sum_{k=1}^{N+1} \left\| \mathbf{A}_{k+1}^{(1)} \circ \sigma_{v_k^{(1)}} \circ W_{k-1}^{(1)} \circ \mathbf{B}_{k-1}^{(2)}(x) - \mathbf{A}_{k+1}^{(1)} \circ \sigma_{v_k^{(2)}} \circ W_{k-1}^{(2)} \circ \mathbf{B}_{k-1}^{(2)}(x) \right\|_\infty.$$

Due to Lemma 11, functions $\mathbf{A}_{k+1}^{(1)} = \mathbf{A}_{L,k+1}^{(1)}$ are Lipschitz, and

$$\begin{aligned} \|f^{(1)}(x) - f^{(2)}(x)\|_\infty &\leq \sum_{k=1}^{N+1} 2^{N-k} \prod_{\ell=0}^{N-k} (p_{N-\ell} + 1)^{2^\ell} \left(\|\mathbf{B}_k^{(2)}(x)\|_\infty \vee 1 \right)^{2^{N-k}} \times \\ &\quad \left\| \sigma_{v_k^{(1)}} \circ W_{k-1}^{(1)} \circ \mathbf{B}_{k-1}^{(2)}(x) - \sigma_{v_k^{(2)}} \circ W_{k-1}^{(2)} \circ \mathbf{B}_{k-1}^{(2)}(x) \right\|_\infty \end{aligned}$$

Note that

$$\begin{aligned} \left\| \sigma_{v_k^{(1)}} \circ W_{k-1}^{(1)} \circ \mathbf{B}_{k-1}^{(2)}(x) - \sigma_{v_k^{(2)}} \circ W_{k-1}^{(2)} \circ \mathbf{B}_{k-1}^{(2)}(x) \right\|_\infty &\leq \varepsilon (p_{k-1} + 1) \left(\|\mathbf{B}_{k-1}^{(2)}(x)\|_\infty \vee 1 \right) \times \\ \left(\left\| W_{k-1}^{(1)} \circ \mathbf{B}_{k-1}^{(2)}(x) \right\|_\infty + 1 \vee \left\| \mathbf{B}_k^{(2)}(x) \right\|_\infty + 1 \right) &\leq 2\varepsilon (p_{k-1} + 1)^2 \left(\|\mathbf{B}_{k-1}^{(2)}(x)\|_\infty \vee 1 \right)^2. \end{aligned}$$

Combining the previous bounds and (B.3) yields

$$\begin{aligned} \|f^{(1)}(x) - f^{(2)}(x)\|_\infty &\leq \varepsilon \sum_{k=1}^{N+1} 2^{N-k+1} \left\{ \prod_{\ell=0}^{N-k} (p_{N-\ell} + 1)^{2^\ell} \right\} (p_{k-1} + 1)^2 \left(\|\mathbf{B}_k^{(2)}(x)\|_\infty^2 \vee 1 \right)^{2^{N-k}} \\ &\leq \varepsilon (N+1) 2^N \prod_{\ell=0}^N (p_\ell + 1)^{2^\ell}, \end{aligned}$$

and the statement follows.

■

B.2 Proof of Lemma 4

Note that for $f(x)$ defined in (B.2), it holds

$$\nabla f(x) = 2^N W_N \prod_{\ell=1}^N \{ \text{diag} [\mathbf{A}_{N-\ell,1}(x) + v_{N-\ell+1} \vee 0] W_{N-\ell} \}. \quad (\text{B.5})$$

Let us define for $j \in \{0, \dots, N-1\}$ and $i \in \{1, 2\}$ the quantities

$$\Delta_j^{(i)} = \mathbf{A}_{j,1}(x) + v_{j+1}^{(i)} \vee 0.$$

Note that $\Delta_j^{(1)}, \Delta_j^{(2)} \in \mathbb{R}^{p_j+1}$. With the triangular inequality,

$$\begin{aligned} \left\| \nabla f^{(1)}(x) - \nabla f^{(2)}(x) \right\|_{\infty} &\leq 2^N \left\| \left(W_N^{(1)} - W_N^{(2)} \right) \prod_{\ell=1}^N \left\{ \text{diag} \left[\Delta_{N-\ell}^{(1)} \right] W_{N-\ell}^{(1)} \right\} \right\|_{\infty} + 2^N (p_N + 1) \times \\ &\sum_{k=1}^N \left\| \prod_{\ell=1}^{k-1} \left\{ \text{diag} \left[\Delta_{N-\ell}^{(1)} \right] W_{N-\ell}^{(1)} \right\} \left(\text{diag} \left[\Delta_k^{(1)} \right] W_k^{(1)} - \text{diag} \left[\Delta_k^{(2)} \right] W_k^{(2)} \right) \prod_{\ell=k+1}^N \left\{ \text{diag} \left[\Delta_{N-\ell}^{(2)} \right] W_{N-\ell}^{(2)} \right\} \right\|_{\infty} \end{aligned}$$

Proceeding as in Lemma 2, we obtain

$$\begin{aligned} \left\| \text{diag} \left[\Delta_k^{(1)} \right] W_k^{(1)} - \text{diag} \left[\Delta_k^{(2)} \right] W_k^{(2)} \right\|_{\infty} &\leq \varepsilon (k+1) 2^k \prod_{\ell=0}^k (p_{\ell} + 1)^{2^k} + (p_k + 1) \left\| \Delta_k^{(1)} \right\|_{\infty} \\ &\leq 2\varepsilon (k+1) 2^k \prod_{\ell=0}^k (p_{\ell} + 1)^{2^k}, \end{aligned}$$

and, similarly,

$$\left\| \text{diag} \left[\Delta_{N-\ell}^{(1)} \right] W_{N-\ell}^{(1)} \right\|_{\infty} \leq \prod_{j=1}^{N-\ell} (p_{N-j} + 1)^{2^j}.$$

Combining the previous bounds, we obtain

$$\left\| \nabla f^{(1)}(x) - \nabla f^{(2)}(x) \right\|_{\infty} \leq \varepsilon 2^N \prod_{\ell=0}^N (p_{\ell} + 1)^{2^{N+1}} + \varepsilon N(N+1) 2^N \prod_{\ell=0}^N (p_{\ell} + 1)^{2^{N+1}+1},$$

and the statement follows.

Appendix C. Proofs of the auxiliary results

C.1 Proof of Lemma 1

Let $D^*(x) = \mathbf{p}^*(x)/(\mathbf{p}^*(x) + \mathbf{p}_w(x))$. By the definition of $\text{JS}(\mathbf{p}_w, \mathbf{p}^*)$ and $L(w, \theta)$, it holds that

$$\begin{aligned} \text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta) &= \frac{1}{2} \int \left[\mathbf{p}^*(x) \log \left(\frac{D^*(x)}{D_{\theta}(x)} \right) + \mathbf{p}_w \log \left(\frac{1 - D^*(x)}{1 - D_{\theta}(x)} \right) \right] d\mu \\ &= \int \left[D^*(x) \log \left(\frac{D^*(x)}{D_{\theta}(x)} \right) + (1 - D^*(x)) \log \left(\frac{1 - D^*(x)}{1 - D_{\theta}(x)} \right) \right] \frac{\mathbf{p}^*(x) + \mathbf{p}_w(x)}{2} d\mu. \end{aligned}$$

Let $a \in [D_{\min}, D_{\max}]$ and introduce a function

$$h_a(v) = (a + v) \log \left(1 + \frac{v}{a} \right) + (1 - a - v) \log \left(1 - \frac{v}{1 - a} \right).$$

Then it holds that

$$\text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta) = \int h_{D_{\theta}(x)}(D^*(x) - D_{\theta}(x)) \frac{\mathbf{p}^*(x) + \mathbf{p}_w(x)}{2} d\mu.$$

Fix any $a \in [D_{\min}, D_{\max}]$ and prove that $h_a(v) \geq v^2$ for any $v \in [-a, 1-a]$. For this purpose, compute the derivatives of $h_a(v)$:

$$\begin{aligned} h'_a(v) &= \log\left(1 + \frac{v}{a}\right) - \log\left(1 - \frac{v}{1-a}\right), \\ h''_a(v) &= \frac{1}{(a+v)(1-a-v)} \geq 4 \quad \text{for all } v \in (-a, 1-a). \end{aligned}$$

This yields that

$$h_a(v) \geq h_a(0) + h'_a(0)v + 2v^2 = 2v^2 \quad \text{for all } v \in [-a, 1-a].$$

Hence,

$$\begin{aligned} \text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta) &= \int h_{D_\theta(x)}(D^*(x) - D_\theta(x)) \frac{\mathbf{p}^*(x) + \mathbf{p}_w(x)}{2} \, d\mu \\ &\geq \int (D^*(x) - D_\theta(x))^2 (\mathbf{p}^*(x) + \mathbf{p}_w(x)) \, d\mu. \\ &= \left\| \frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} - D_\theta \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}^2. \end{aligned}$$

To prove (3.5), it is enough to show that the following inequality holds for any $a \in [D_{\min}, D_{\max}]$ and any $v \in [-a, 1-a]$:

$$h_a(v) \leq \frac{C_a^2 v^2}{(C_a - 1)^2 a (1-a)}, \quad \text{where } C_a = 1 + \sqrt{\frac{a}{(1-a) \log(1/(1-a))} \wedge \frac{1-a}{a \log(1/a)}}.$$

Then it is easy to observe that

$$C_a \geq C_{3.6} > 1 \quad \text{and} \quad a(1-a) \geq D_{\min}(1 - D_{\max}) \quad \text{for all } a \in [D_{\min}, D_{\max}],$$

which yields

$$h_a(v) \leq \frac{C_{3.6}^2 v^2}{(C_{3.6} - 1)^2 D_{\min}(1 - D_{\max})} \quad \text{for all } a \in [D_{\min}, D_{\max}] \text{ and } v \in [-a, 1-a].$$

Hence,

$$\begin{aligned} \text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta) &= \int h_{D_\theta(x)}(D^*(x) - D_\theta(x)) \frac{\mathbf{p}^*(x) + \mathbf{p}_w(x)}{2} \, d\mu \\ &\leq \frac{C_{3.6}^2}{(C_{3.6} - 1)^2 D_{\min}(1 - D_{\max})} \int (D^*(x) - D_\theta(x))^2 (\mathbf{p}^*(x) + \mathbf{p}_w(x)) \, d\mu \\ &= \frac{C_{3.6}^2}{(C_{3.6} - 1)^2 D_{\min}(1 - D_{\max})} \left\| \frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} - D_\theta \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}^2. \end{aligned}$$

Note that, for any $v \in [-a/C_a, (1-a)/C_a]$, we have

$$h''_a(v) = \frac{1}{(a+v)(1-a-v)} \leq \frac{1}{(1 - 1/C_a)^2 a(1-a)}.$$

Since $h_a(0) = h'_a(0) = 0$, this yields that

$$h_a(v) \leq \frac{C_a^2 v^2}{2(C_a - 1)^2 a(1 - a)}.$$

Fix any $v \in [-a, -a/C_a]$ and consider $h_a(v)$. Since h_a decreases on $[-a, 0]$, we obtain that

$$h_a(v) \leq h_a(-a) = \log \frac{1}{1 - a}.$$

On the other hand, if $v \in [-a, -a/C_a]$, then $v^2 \geq a^2/C_a^2$. Taking into account that

$$(C_a - 1)^2 \leq \frac{a}{(1 - a) \log(1/(1 - a))}$$

due to the definition of C_a , we obtain that, for any $v \in [-a, -a/C_a]$, the following inequality holds:

$$\frac{C_a^2 v^2}{(C_a - 1)^2 a(1 - a)} \geq \frac{a}{(C_a - 1)^2 (1 - a)} \geq \log \frac{1}{1 - a} \geq h_a(v).$$

Similarly, since h_a increases on $[0, 1 - a]$, it holds that

$$h_a(v) \leq h_a(1 - a) = \log \frac{1}{a} \quad \text{for all } v \in \left[\frac{1 - a}{C_a}, 1 - a \right].$$

At the same time, $v^2 \geq (1 - a)^2/C_a^2$ for any $v \in [(1 - a)/C_a, 1 - a]$. Using the inequality

$$(C_a - 1)^2 \leq \frac{1 - a}{a \log(1/a)},$$

which follows from the definition of C_a , we obtain that

$$\frac{C_a^2 v^2}{(C_a - 1)^2 a(1 - a)} \geq \frac{1 - a}{a(C_a - 1)^2} \geq \log \frac{1}{a} \geq h_a(v), \quad \text{for all } v \in \left[\frac{1 - a}{C_a}, 1 - a \right],$$

and the proof is finished.

C.2 Proof of Lemma 5

Proof of Lemma 5 The proof of the lemma is quite long, so we split it in several steps.

Step 1: Bernstein's bound. Let us recall that

$$L_n(w, \theta) = \frac{1}{2n} \sum_{i=1}^n \log D_\theta(X_i) + \frac{1}{2n} \sum_{i=1}^n \log(1 - D_\theta(g_w(Y_i))),$$

and $L(w, \theta)$ is the expectation of $L_n(w, \theta)$. Assumption AD yields that the absolute values of the random variables $\log D_\theta(X_i)$ and $\log(1 - D_\theta(g_w(Y_i)))$ do not exceed $3C_{5.3}/2$, where the constant $C_{5.3}$ is given by (5.3). Then the Bernstein inequality implies that, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, we have

$$|L_n(w, \theta) - L(w, \theta)| \leq \sqrt{2\text{Var}[L_n(w, \theta)] \log(2/\delta)} + \frac{2C_D \log(2/\delta)}{3n}. \quad (\text{C.1})$$

It remains to prove that

$$4n\text{Var}[L_n(w, \theta)] \leq C_{5.2}(9 \text{JS}(\mathbf{p}_w, \mathbf{p}^*) + \Delta(w, \theta)).$$

Step 2a: bounding the variance. The variance of $L_n(w, \theta_w^*)$ satisfies the inequality

$$\begin{aligned} 4n\text{Var}[L_n(w, \theta)] &= \text{Var}[\log D_\theta(X_i)] + \text{Var}[\log(1 - D_\theta(g_w(Y_i)))] \\ &= \text{Var}[\log(2D_\theta(X_i))] + \text{Var}[\log(2 - 2D_\theta(g_w(Y_i)))] \\ &\leq \mathbb{E} \log^2(2D_\theta(X_1)) + \log^2(2 - 2D_\theta(g_w(Y_1))) \\ &= \int [\mathbf{p}^*(x) \log^2(2D_\theta(x)) + \mathbf{p}_w(x) \log^2(2 - 2D_\theta(x))] d\mu. \end{aligned}$$

Note that

$$\begin{aligned} &\int [\mathbf{p}^*(x) \log^2(2D_\theta(x)) + \mathbf{p}_w(x) \log^2(2 - 2D_\theta(x))] d\mu \\ &= \int \frac{\mathbf{p}^*(x)}{\mathbf{p}^*(x) + \mathbf{p}_w(x)} \log^2(2D_\theta(x)) (\mathbf{p}^*(x) + \mathbf{p}_w(x)) d\mu \\ &\quad + \frac{\mathbf{p}_w(x)}{\mathbf{p}^*(x) + \mathbf{p}_w(x)} \log^2(2 - 2D_\theta(x)) (\mathbf{p}^*(x) + \mathbf{p}_w(x)) d\mu \\ &= \int \left[\frac{\mathbf{p}^*(x)}{\mathbf{p}^*(x) + \mathbf{p}_w(x)} - \frac{1}{2} \right] \log^2(2D_\theta(x)) (\mathbf{p}^*(x) + \mathbf{p}_w(x)) d\mu \tag{C.2} \\ &\quad + \int \left[\frac{\mathbf{p}_w(x)}{\mathbf{p}^*(x) + \mathbf{p}_w(x)} - \frac{1}{2} \right] \log^2(2 - 2D_\theta(x)) (\mathbf{p}^*(x) + \mathbf{p}_w(x)) d\mu \\ &\quad + \int [\log^2(2D_\theta(x)) + \log^2(2 - 2D_\theta(x))] \frac{\mathbf{p}^*(x) + \mathbf{p}_w(x)}{2} d\mu. \end{aligned}$$

Step 2b: bounding the variance, the first term. Consider the first term in the right-hand side. The Cauchy-Schwarz inequality yields that

$$\begin{aligned} &\int \left[\frac{\mathbf{p}^*(x)}{\mathbf{p}^*(x) + \mathbf{p}_w(x)} - \frac{1}{2} \right] \log^2(2D_\theta(x)) (\mathbf{p}^*(x) + \mathbf{p}_w(x)) d\mu \\ &\leq \left\| \frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)} \left\| \log^2(2D_\theta) \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}. \end{aligned}$$

Applying Lemma 7, we obtain that

$$\begin{aligned} &\int \left[\frac{\mathbf{p}^*(x)}{\mathbf{p}^*(x) + \mathbf{p}_w(x)} - \frac{1}{2} \right] \log^2(2D_\theta(x)) (\mathbf{p}^*(x) + \mathbf{p}_w(x)) d\mu \\ &\leq 2\sqrt{\text{JS}(\mathbf{p}_w, \mathbf{p}^*)} \left\| \log^2(2D_\theta) \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}. \end{aligned}$$

According to Assumption AD, $D_\theta(x) \in [D_{\min}, D_{\max}] \subset (0, 1)$. The map $g(u) = \log^2(2u)$ is convex on $[D_{\min}, 1]$. Thus, it holds that

$$g(u) \leq g(1/2) + \frac{g(1)(u - 1/2)}{1 - 1/2} \quad \text{for all } u \in [1/2, D_{\max}]$$

and

$$g(u) \leq g(1/2) + \frac{g(D_{\min})(u - 1/2)}{D_{\min} - 1/2} \quad \text{for all } u \in [D_{\min}, 1/2].$$

This implies that

$$\|\log^2(2D_\theta)\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)} \leq \left(\frac{\log^2(2D_{\min})}{1/2 - D_{\min}} \vee 2 \log^2 2 \right) \left\| D_\theta - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}.$$

Hence, we finally obtain that

$$\begin{aligned} & \int \left[\frac{\mathbf{p}^*(x)}{\mathbf{p}^*(x) + \mathbf{p}_w(x)} - \frac{1}{2} \right] \log^2(2D_\theta(x)) (\mathbf{p}^*(x) + \mathbf{p}_w(x)) \, d\mu \\ & \leq 2 \left(\frac{\log^2(2D_{\min})}{1/2 - D_{\min}} \vee 2 \log^2 2 \right) \sqrt{\text{JS}(\mathbf{p}_w, \mathbf{p}^*)} \left\| D_\theta - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}. \end{aligned}$$

Step 2c: bounding the variance, the second term. For the second term in the right-hand side of (C.2) we similarly have

$$\begin{aligned} & \int \left[\frac{\mathbf{p}_w(x)}{\mathbf{p}^*(x) + \mathbf{p}_w(x)} - \frac{1}{2} \right] \log^2(2 - 2D_{\theta_w}(x)) (\mathbf{p}^*(x) + \mathbf{p}_w(x)) \, d\mu \\ & \leq 2 \left(\frac{\log^2(2 - 2D_{\max})}{D_{\max} - 1/2} \vee 2 \log^2 2 \right) \sqrt{\text{JS}(\mathbf{p}_w, \mathbf{p}^*)} \left\| D_\theta - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}. \end{aligned}$$

Step 2d: bounding the variance, the third term. It remains to bound

$$\int [\log^2(2D_\theta(x)) + \log^2(2 - 2D_\theta(x))] \frac{\mathbf{p}^*(x) + \mathbf{p}_w(x)}{2} \, d\mu.$$

Consider the function $h(u) = \log^2(2u) + \log^2(2 - 2u)$, $u \in [D_{\min}, D_{\max}]$. It is easy to observe that $h(1/2) = 0$, $h'(1/2) = 0$, and

$$h''(u) = \frac{2 \log(e/(2u))}{u^2} + \frac{2 \log(e/(2 - 2u))}{(1 - u)^2} \leq \frac{2 \log(e/(2D_{\min}))}{D_{\min}^2} + \frac{2 \log(e/(2 - 2D_{\max}))}{(1 - D_{\max})^2}.$$

Hence, it holds that

$$h(u) \leq \left(\frac{\log(e/(2D_{\min}))}{D_{\min}^2} + \frac{\log(e/(2 - 2D_{\max}))}{(1 - D_{\max})^2} \right) (u - 1/2)^2 \quad \text{for all } u \in [D_{\min}, D_{\max}],$$

and, therefore,

$$\begin{aligned} & \int [\log^2(2D_\theta(x)) + \log^2(2 - 2D_\theta(x))] \frac{\mathbf{p}^*(x) + \mathbf{p}_w(x)}{2} \, d\mu \\ & \leq \left(\frac{\log(e/(2D_{\min}))}{2D_{\min}^2} + \frac{\log(e/(2 - 2D_{\max}))}{2(1 - D_{\max})^2} \right) \left\| D_\theta - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}^2. \end{aligned}$$

Step 3: final part. Steps 2a–d yield that

$$\begin{aligned}
 4n\text{Var}[L_n(w, \theta)] &\leq 2 \left(\frac{\log^2(2D_{\min})}{1/2 - D_{\min}} \vee 2 \log^2 2 \right) \sqrt{\text{JS}(\mathbf{p}_w, \mathbf{p}^*)} \left\| D_\theta - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)} \\
 &\quad + 2 \left(\frac{\log^2(2 - 2D_{\max})}{D_{\max} - 1/2} \vee 2 \log^2 2 \right) \sqrt{\text{JS}(\mathbf{p}_w, \mathbf{p}^*)} \left\| D_{\theta_w} - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)} \\
 &\quad + \left(\frac{\log(e/(2D_{\min}))}{2D_{\min}^2} + \frac{\log(e/(2 - 2D_{\max}))}{2(1 - D_{\max})^2} \right) \left\| D_\theta - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}^2.
 \end{aligned}$$

By the Cauchy-Schwarz inequality,

$$4n\text{Var}[L_n(w, \theta)] \leq C_{5.2} \text{JS}(\mathbf{p}_w, \mathbf{p}^*) + C_{5.2} \left\| D_\theta - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}^2,$$

where $C_{5.2}$ is given by (5.2). Consider the norm $\|D_{\theta_w} - 1/2\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}$. By the triangle inequality, it holds that

$$\left\| D_{\theta_w} - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)} \leq \left\| D_{\theta_w} - \frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)} + \left\| \frac{\mathbf{p}^*}{\mathbf{p}^* + \mathbf{p}_w} - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}.$$

Applying Lemma 1 and Lemma 7, we obtain that

$$\begin{aligned}
 \left\| D_\theta - \frac{1}{2} \right\|_{L_2(\mathbf{p}^* + \mathbf{p}_w)}^2 &\leq \left(\sqrt{\text{JS}(\mathbf{p}_w, \mathbf{p}^*) - \log 2 - L(w, \theta)} + 2\sqrt{\text{JS}(\mathbf{p}_w, \mathbf{p}^*)} \right)^2 \\
 &\leq \left(\sqrt{\Delta(w, \theta)} + 2\sqrt{\text{JS}(\mathbf{p}_w, \mathbf{p}^*)} \right)^2 \\
 &\leq \Delta(w, \theta) + 8 \text{JS}(\mathbf{p}_w, \mathbf{p}^*).
 \end{aligned}$$

Then it holds that

$$4n\text{Var}[L_n(w, \theta)] \leq 9C_{5.2} \text{JS}(\mathbf{p}_w, \mathbf{p}^*) + C_{5.2} \Delta(w, \theta),$$

and the claim of the lemma follows. ■

C.3 Proof of Lemma 6

Since the statement of the lemma contains four inequalities, we split the proof in several steps for the sake of clarity.

Step 1. Prove that, for any $w_1, w_2 \in \mathcal{W}$ and $\theta \in \Theta$, we have

$$|L_n(w_1, \theta) - L_n(w_2, \theta)| \leq \frac{\mathbf{L}_G \mathbf{L}_X \|w_1 - w_2\|_\infty}{2 - 2D_{\max}} \quad \text{almost surely.}$$

It holds that

$$L_n(w_1, \theta) - L_n(w_2, \theta) = \frac{1}{2n} \sum_{i=1}^n \log \frac{1 - D_\theta(g_{w_1}(Y_i))}{1 - D_\theta(g_{w_2}(Y_i))}.$$

Since the map $x \mapsto \log(1-x)$ is Lipschitz on $[D_{\min}, D_{\max}]$ with the constant $1/(1-D_{\max})$, Assumption AD yields that

$$\begin{aligned} |L_n(w_1, \theta) - L_n(w_2, \theta)| &\leq \frac{1}{2n} \sum_{i=1}^n \frac{|D_{\theta}(g_{w_1}(Y_i)) - D_{\theta}(g_{w_2}(Y_i))|}{1 - D_{\max}} \\ &\leq \frac{1}{2n} \sum_{i=1}^n \frac{\mathsf{L}_X \|g_{w_1}(Y_i) - g_{w_2}(Y_i)\|}{1 - D_{\max}}. \end{aligned}$$

Due to Assumption AG, we have

$$|L_n(w_1, \theta) - L_n(w_2, \theta)| \leq \frac{\mathsf{L}_G \mathsf{L}_X \|w_1 - w_2\|_{\infty}}{2 - 2D_{\max}} \quad \text{almost surely.}$$

Step 2. Let us show that, for any $w \in \mathsf{W}$ and $\theta_1, \theta_2 \in \Theta$, we have

$$|L_n(w, \theta_1) - L_n(w, \theta_2)| \leq \frac{\mathsf{L}_{\Theta} \|\theta_1 - \theta_2\|_{\infty}}{D_{\min} \wedge (1 - D_{\max})} \quad \text{almost surely.}$$

It holds that

$$L_n(w, \theta_1) - L_n(w, \theta_2) = \frac{1}{2n} \sum_{i=1}^n \log \frac{D_{\theta_1}(X_i)}{D_{\theta_2}(X_i)} + \frac{1}{2n} \sum_{i=1}^n \log \frac{1 - D_{\theta_1}(g_w(Y_i))}{1 - D_{\theta_2}(g_w(Y_i))}$$

The maps $x \mapsto \log x$ and $x \mapsto \log(1-x)$ are Lipschitz on $[D_{\min}, D_{\max}]$ with the constants $1/D_{\min}$ and $1/(1-D_{\max})$, respectively. Hence, we have

$$|L_n(w, \theta_1) - L_n(w, \theta_2)| \leq \frac{1}{2n} \sum_{i=1}^n \frac{|D_{\theta_1}(X_i) - D_{\theta_2}(X_i)|}{D_{\min}} + \frac{1}{2n} \sum_{i=1}^n \frac{|D_{\theta_1}(X_i) - D_{\theta_2}(X_i)|}{1 - D_{\max}}.$$

Then, due to Assumption AD, it holds that

$$|L_n(w, \theta_1) - L_n(w, \theta_2)| \leq \frac{\mathsf{L}_{\Theta} \|\theta_1 - \theta_2\|_{\infty}}{2D_{\min}} + \frac{\mathsf{L}_{\Theta} \|\theta_1 - \theta_2\|_{\infty}}{2 - 2D_{\max}} \leq \frac{\mathsf{L}_{\Theta} \|\theta_1 - \theta_2\|_{\infty}}{D_{\min} \wedge (1 - D_{\max})}$$

almost surely.

Step 3. Finally, due to the Jensen inequality, it holds that

$$\begin{aligned} |L(w_1, \theta) - L(w_2, \theta)| &= |\mathsf{E}L_n(w_1, \theta) - \mathsf{E}L_n(w_2, \theta)| \\ &\leq \mathsf{E}|L_n(w_1, \theta) - L_n(w_2, \theta)| \\ &\leq \frac{\mathsf{L}_G \mathsf{L}_X \|w_1 - w_2\|_{\infty}}{2 - 2D_{\max}} \end{aligned}$$

and, similarly,

$$\begin{aligned} |L(w, \theta_1) - L(w, \theta_2)| &= |\mathsf{E}L_n(w, \theta_1) - \mathsf{E}L_n(w, \theta_2)| \\ &\leq \mathsf{E}|L_n(w, \theta_1) - L_n(w, \theta_2)| \\ &\leq \frac{\mathsf{L}_{\Theta} \|\theta_1 - \theta_2\|_{\infty}}{D_{\min} \wedge (1 - D_{\max})}. \end{aligned}$$

C.4 Proof of Lemma 7

Let $h(u) = (1+u) \log(1+u) + (1-u) \log(1-u)$, $u \in [-1, 1]$, and rewrite the Jensen-Shannon divergence between \mathbf{p} and \mathbf{q} in the following form:

$$\begin{aligned} \text{JS}(\mathbf{p}, \mathbf{q}) &= \frac{1}{2} \left[\int \log \left(\frac{2\mathbf{p}(x)}{\mathbf{p}(x) + \mathbf{q}(x)} \right) \mathbf{p}(x) d\mu + \int \log \frac{2\mathbf{q}(x)}{\mathbf{p}(x) + \mathbf{q}(x)} \mathbf{q}(x) d\mu \right] \\ &= \int \left[\frac{2\mathbf{p}(x)}{\mathbf{p}(x) + \mathbf{q}(x)} \log \frac{2\mathbf{p}(x)}{\mathbf{p}(x) + \mathbf{q}(x)} + \frac{2\mathbf{q}(x)}{\mathbf{p}(x) + \mathbf{q}(x)} \log \frac{2\mathbf{q}(x)}{\mathbf{p}(x) + \mathbf{q}(x)} \right] \frac{\mathbf{p}(x) + \mathbf{q}(x)}{4} d\mu \\ &= \int h \left(\frac{\mathbf{p}(x) - \mathbf{q}(x)}{\mathbf{p}(x) + \mathbf{q}(x)} \right) \frac{\mathbf{p}(x) + \mathbf{q}(x)}{4} d\mu. \end{aligned}$$

Note that the function $h(u)/u^2$ is even and increasing on $[0, 1]$. Hence, it attains its maximum on $[-1, 1]$ at the points $u = -1$ and $u = 1$ and its minimum at $u = 0$. This yields that

$$1 = \lim_{v \rightarrow 0} \frac{h(v)}{v^2} \leq \frac{h(u)}{u^2} \leq \frac{h(1)}{1^2} = 2 \log 2 \quad \text{for all } u \in [-1, 1].$$

Since $(\mathbf{p}(x) - \mathbf{q}(x))/(\mathbf{p}(x) + \mathbf{q}(x)) \in [-1, 1]$ for all x from $\text{supp}(\mathbf{p}) \cup \text{supp}(\mathbf{q})$, it holds that

$$\frac{1}{4} \int \frac{(\mathbf{p}^*(x) - \mathbf{p}_w(x))^2}{\mathbf{p}^*(x) + \mathbf{p}_w(x)} d\mu \leq \text{JS}(\mathbf{p}_w, \mathbf{p}^*) \leq \frac{\log 2}{2} \int \frac{(\mathbf{p}^*(x) - \mathbf{p}_w(x))^2}{\mathbf{p}^*(x) + \mathbf{p}_w(x)} d\mu.$$

C.5 Proof of Corollary 1

Fix any u and v from \mathbb{W} . The fact that the square root of the JS-divergence is a metric (see, e.g., (Endres and Schindelin, 2003)) implies that

$$\left| \sqrt{\text{JS}(\mathbf{p}_u, \mathbf{p}^*)} - \sqrt{\text{JS}(\mathbf{p}_v, \mathbf{p}^*)} \right| \leq \sqrt{\text{JS}(\mathbf{p}_u, \mathbf{p}_v)}$$

Then the claim of the corollary follows from the inequalities

$$\begin{aligned} \text{JS}(\mathbf{p}_u, \mathbf{p}_v) &\leq \frac{\log 2}{2} \int \frac{(\mathbf{p}_u(x) - \mathbf{p}_v(x))^2}{\mathbf{p}_u(x) + \mathbf{p}_v(x)} d\mu \\ &\leq \frac{\log 2}{2} \int |\mathbf{p}_u(x) - \mathbf{p}_v(x)| d\mu \\ &\leq \frac{\log 2}{2} \|\mathbf{p}_u - \mathbf{p}_v\|_{L^\infty(\mathbb{X})}^{1/2} \int |\mathbf{p}_u(x) - \mathbf{p}_v(x)|^{1/2} d\mu \\ &\leq \frac{\log 2}{2} \|\mathbf{p}_u - \mathbf{p}_v\|_{L^\infty(\mathbb{X})}^{1/2} \sqrt{\int |\mathbf{p}_u(x) - \mathbf{p}_v(x)| d\mu} \\ &\leq \frac{L_{\mathbf{p}} \log 2}{\sqrt{2}} \|u - v\|_{\infty}^{1/2} < L_{\mathbf{p}} \|u - v\|_{\infty}^{1/2}. \end{aligned}$$

where we used Lemma 7 and Assumption $\text{A}_{\mathbf{p}}$.

Appendix D. Proofs related to Theorem 3

D.1 Proof of Lemma 9

The main ingredient we use is the multivariate van Trees inequality (see (Van Trees, 1968), p. 72, and (Gill and Levit, 1995)). Choose a density λ of a prior distribution on the set of parameters $\{\theta \in \mathbb{R}^M : \|\theta\| \leq h\}$ of the form

$$\lambda(\theta) = \frac{1}{h^M} \lambda_0\left(\frac{\theta}{h}\right),$$

where λ_0 is a smooth density supported on $\mathcal{B}(0, 1) \subset \mathbb{R}^M$. We write \mathbb{E}_λ for the expectation with respect to the density λ . Then, for any estimate $\hat{\theta}$, it holds that

$$\sup_{\|\theta\| \leq h} \mathbb{E}_\theta \|\hat{\theta} - \theta\|^2 \geq \mathbb{E}_\lambda \mathbb{E}_\theta \|\hat{\theta} - \theta\|^2 \geq \frac{M^2}{n \sum_{j=1}^M \mathbb{E}_\lambda \mathcal{I}_j(\theta) + \mathcal{J}(\lambda)}, \quad (\text{D.1})$$

where

$$\mathcal{I}_j(\theta) = \mathbb{E}_\theta \left(\frac{\partial \log p_{g_\theta}(X)}{\partial \theta_j} \right)^2$$

is the Fischer information of one observation, $p_{g_\theta}(x)$ is the density of X , and

$$\mathcal{J}(\lambda) = \int_{\|\theta\| \leq h} \sum_{j=1}^M \left(\frac{\partial \lambda(\theta)}{\partial \theta_j} \right)^2 \frac{d\theta}{\lambda(\theta)}.$$

First, let us bound $\mathcal{J}(\lambda)$.

$$\mathcal{J}(\lambda) = \sum_{j=1}^M \int_{\|\theta\| \leq h} \frac{1}{h^{2M}} \left(\frac{\partial \lambda_0(\theta/h)}{\partial \theta_j} \right)^2 \frac{h^M}{\lambda_0(\theta/h)} d\theta$$

Substituting $\theta = h\nu$, $\nu \in \mathcal{B}(0, 1)$, we obtain

$$\mathcal{J}(\lambda) = \sum_{j=1}^M \int_{\|\nu\| \leq 1} h^{-2} \left(\frac{\partial \lambda_0(\nu)}{\partial \nu_j} \right)^2 \frac{d\nu}{\lambda_0(\nu)}.$$

Since $\int_{\|\nu\| \leq 1} \left(\frac{\partial \lambda_0(\nu)}{\partial \nu_j} \right)^2 \frac{d\nu}{\lambda_0(\nu)}$ is finite, we conclude that

$$\mathcal{J}(\lambda) \lesssim \frac{M}{h^2}.$$

Now, we focus on $\mathcal{I}_j(\theta)$:

$$\log p_{g_\theta}(x) = \log \det \left(I_{d \times d} + h^{\beta-1} \sum_{j=1}^M \theta_j \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right).$$

For any $j \in \{1, \dots, M\}$ the partial derivative of $\log \mathbf{p}_{g_\theta}$ with respect to θ_j is equal to

$$\frac{\partial \log \mathbf{p}_{g_\theta}(x)}{\partial \theta_j} = h^{\beta-1} \text{Tr} \left(A_\theta(x)^{-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right),$$

where

$$A_\theta(x) = I_{d \times d} + \sum_{j=1}^M \theta_j h^{\beta-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right).$$

Consider

$$\begin{aligned} & \left| \text{Tr} \left(A_\theta(x)^{-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right) \right| \\ &= \left| \text{Tr} \left(\left(I_{d \times d} + h^{\beta-1} \sum_{k=1}^M \theta_k \nabla^2 \varphi \left(\frac{x - x_k}{h} \right) \right)^{-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right) \right|. \end{aligned}$$

If $x \notin \mathcal{B}(x_j, h)$ then $\nabla^2 \varphi((x - x_j)/h) = O_{d \times d}$ and, consequently, $\text{Tr}(A_\theta(x)^{-1} \nabla^2 \varphi((x - x_j)/h)) = 0$. Otherwise, we have $\nabla^2 \varphi((x - x_k)/h) = O_{d \times d}$ for all $k \neq j$. Hence,

$$\begin{aligned} & \left| \text{Tr} \left(\left(I_{d \times d} + h^{\beta-1} \sum_{k=1}^M \theta_k \nabla^2 \varphi \left(\frac{x - x_k}{h} \right) \right)^{-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right) \right| \\ &= \left| \text{Tr} \left(\left(I_{d \times d} + h^{\beta-1} \theta_j \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right)^{-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right) \right| \\ &= \left| \text{Tr} \left(\sum_{k=0}^{\infty} (h^{\beta-1} \theta_j)^k \left(\nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right)^{k+1} \right) \right| \\ &\lesssim \left| \text{Tr} \left(\nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right) \right| \lesssim d. \end{aligned}$$

Here we used the fact that $h^{\beta-1} |\theta_j| \leq h^\beta$ is small, provided that the sample size is large enough. Consider

$$\mathcal{I}_j(\theta) = \int_{\mathbf{X}} \left(h^{\beta-1} \text{Tr} \left(A_\theta(x)^{-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right) \right)^2 \mathbf{p}_{g_\theta}(x) dx.$$

Since φ is supported on $\mathcal{B}(0, 1)$, the last expression is equal to

$$\begin{aligned} \mathcal{I}_j(\theta) &= h^{2\beta-2} \int_{\mathcal{B}(0, h)} \left(\text{Tr} \left(A_\theta(x)^{-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right) \right)^2 \mathbf{p}_{g_\theta}(x) dx \\ &\lesssim h^{2\beta-2} d^2 \int_{\mathcal{B}(0, h)} \mathbf{p}_{g_\theta}(x) dx \lesssim h^{2\beta+d-2} d^2. \end{aligned}$$

Here the last inequality follows from (5.21) and Lemma 8. Thus, it holds that

$$\sup_{\|\theta\| \leq h} \mathbf{E}_\theta \|\hat{\theta} - \theta\|^2 \geq \mathbf{E}_\lambda \mathbf{E}_\theta \|\hat{\theta} - \theta\|^2 \geq \frac{M^2}{n \sum_{j=1}^M \mathbf{E}_\lambda \mathcal{I}_j(\theta) + \mathcal{J}(\lambda)} \gtrsim \frac{Mh^2}{nh^{2\beta+d}d^2 + 1}.$$

D.2 Proof of Lemma 10

Recall that

$$\log \mathfrak{p}_{g_\theta}(x) = \log \det \left(I_{d \times d} + h^{\beta-1} \sum_{j=1}^M \theta_j \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right).$$

Then, for any $j \in \{1, \dots, M\}$ the partial derivative of $\log \mathfrak{p}_{g_\theta}$ with respect to θ_j is equal to

$$\frac{\partial \log \mathfrak{p}_{g_\theta}(x)}{\partial \theta_j} = h^{\beta-1} \text{Tr} \left(A_\theta(x)^{-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right),$$

where

$$A_\theta(x) = I_{d \times d} + \sum_{j=1}^M \theta_j h^{\beta-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right).$$

Since, for any $j \in \{1, \dots, J\}$, ψ_j is supported on $\mathcal{B}(0, 1)$ and $\{x_1, \dots, x_M\}$ is a $(2h)$ -packing on \mathbf{X} , then, for any $j \in \{1, \dots, J\}$ and for any $x \in \mathcal{B}(x_j, h/2)$,

$$A_\theta(x) = I_{d \times d} + \theta_j h^{\beta-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right).$$

Consider $\text{Tr} (A_\theta(x)^{-1} \nabla^2 \varphi((x - x_j)/h))$. We recall that, by the construction,

$$-\nabla^2 \varphi(0) \succeq I_{d \times d}.$$

Since $\varphi \in \mathcal{H}^{\beta+2}(\mathbf{X}, L)$, $\beta > 1$ and r_0 is such that $r_0 \leq h/(2H_\varphi)$, we have

$$-\nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \succeq 0.5 I_{d \times d}, \quad \text{for all } x \in \mathcal{B}(x_j, r_0).$$

Taking into account that

$$\left\| \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right\| \leq H_\varphi,$$

we obtain

$$\left\| \theta_j h^{\beta-1} \nabla^2 \varphi \left(\frac{x - x_j}{h} \right) \right\| \leq H_\varphi h^\beta < 1 \quad \text{for any } h < H_\varphi^{-1/\beta}.$$

Then

$$\begin{aligned} & \left(I_{d \times d} + \theta_j h^{\beta-1} \nabla g_0(g_0^{-1}(x)) \nabla \psi_j \left(\frac{x - x_j}{h} \right) \right)^{-1} \\ &= \sum_{k=0}^{\infty} \theta_j^k h^{k\beta-k} \left(-\nabla g_0(g_0^{-1}(x)) \nabla \psi_j \left(\frac{x - x_j}{h} \right) \right)^k. \end{aligned}$$

This implies

$$\begin{aligned}
 & -\operatorname{Tr}\left(\left(I_{d \times d} + \theta_j h^{\beta-1} \nabla^2 \varphi\left(\frac{x-x_j}{h}\right)\right)^{-1} \nabla^2 \varphi\left(\frac{x-x_j}{h}\right)\right) \\
 &= \operatorname{Tr}\left(\sum_{k=0}^{\infty} \theta_j^k h^{k\beta-k} \left(-\nabla^2 \varphi\left(\frac{x-x_j}{h}\right)\right)^{k+1}\right) \\
 &\geq \sum_{k=0}^{\infty} \frac{\theta_j^k h^{k\beta-k}}{2^{k+1}} \operatorname{Tr}(I_{d \times d}) \geq d/2.
 \end{aligned}$$

Then

$$\left| \frac{\partial \log p_{g_\theta}(x)}{\partial \theta_j} \right| = h^{\beta-1} \left| \operatorname{Tr}\left(A_\theta(x)^{-1} \nabla^2 \varphi\left(\frac{x-x_j}{h}\right)\right) \right| \geq h^{\beta-1} d/2 \quad \text{for all } x \in \mathcal{B}(x_j, r_0).$$