



A generative model for fBm with deep ReLU neural networks

Michaël Allouche, Stéphane Girard, Emmanuel Gobet

► To cite this version:

Michaël Allouche, Stéphane Girard, Emmanuel Gobet. A generative model for fBm with deep ReLU neural networks. 2021. hal-03237854v2

HAL Id: hal-03237854

<https://hal.science/hal-03237854v2>

Preprint submitted on 16 Jun 2021 (v2), last revised 25 Apr 2022 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A generative model for fBm with deep ReLU neural networks

Michaël Allouche

michael.allouche@polytechnique.edu

CMAP, CNRS, Ecole Polytechnique, Institut Polytechnique de Paris

Route de Saclay, 91128 Palaiseau Cedex, France

Stéphane Girard

stephane.girard@inria.fr

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK

38000 Grenoble, France

Emmanuel Gobet

emmanuel.gobet@polytechnique.edu

CMAP, CNRS, Ecole Polytechnique, Institut Polytechnique de Paris

Route de Saclay, 91128 Palaiseau Cedex, France

June 14, 2021

Abstract

We provide a large probability bound on the uniform approximation of fractional Brownian motion with Hurst parameter H , by a deep-feedforward ReLU neural network fed with a N -dimensional Gaussian vector, with bounds on the network construction (number of hidden layers and total number of neurons). Essentially, up to log terms, achieving an uniform error of $\mathcal{O}(N^{-H})$ is possible with $\log(N)$ hidden layers and $\mathcal{O}(N)$ parameters. Our analysis relies, in the standard Brownian motion case ($H = 1/2$), on the Levy construction and in the general fractional Brownian motion case ($H \neq 1/2$), on the Lemarié-Meyer wavelet representation. This work gives theoretical support on new generative models based on neural networks for simulating continuous-time processes.

Keywords: fractional Brownian motion, Gaussian process, neural networks, generative models

MSC: 62M45, 60G15, 60G22

1 Introduction

Over last few years a new paradigm of generative model has emerged in the machine learning community with the goal of sampling high-dimensional complex objects (such as images, videos or natural language) from a data set of these objects. If X denotes the random variable taking values in a general metric space $(\mathcal{X}, d_{\mathcal{X}})$ from which we have observations (X_1, \dots, X_n, \dots) , the problem of generative model amounts to finding a function $G_{\theta} : \mathbb{R}^N \mapsto \mathcal{X}$ and a latent probability distribution μ on \mathbb{R}^N such that

$$X \stackrel{d}{=} G_{\theta}(Z) \text{ and } Z \sim \mu. \tag{1}$$

Usually, the choice of the dimension N (the so-called latent dimension) is part of the problem. The function G_θ belongs to a parametric family of functions $\mathcal{G} = \{G_\theta\}_{\theta \in \Theta}$, and it is common to consider neural networks: in this work, we follow this approach. Essentially, two main questions have to be addressed to obtain a generative model: a) how to choose \mathcal{G} to have a chance to get the equality in distribution (1), or at least a good approximation of it for some $G_\theta \in \mathcal{G}$? b) how to learn the parameter θ from the data set? The second question b) has been tackled by [19] in their seminal work of Generative Adversarial Network (GAN). We will not focus on that problematic in this work, there is a tremendous number of works (about 30,000 citations of [19] on Google Scholar at the date of writing this article). Instead, we are to focus on a), *i.e.* quantifying how to choose \mathcal{G} and the latent space (N, μ) when \mathcal{X} is the space of continuous functions indexed by time, equipped with the sup norm $d_{\mathcal{X}}$, and when the distribution of X is that of a stochastic process (infinite dimensional object), possibly non-Markovian.

Among the huge and expanding literature on GANs, lot of works studied the ability to generate time-series (in a discrete time), either in finance [38], in medicine [13] or in meteorology [21], for citing only some of them. However, to the best of our knowledge, none of them is dealing with continuous-time processes. Moreover, designing the architecture of a neural network G_θ with respect to its depth (number of hidden layers), size (number of neurons), type (feed-forward, recurrent, convolutional, etc.) and activation functions (sigmoid, ReLU, etc.), is a very difficult question and therefore often left to empirical grid search. In this work, we aim at tackling these aspects and providing precise quantitative guidelines on \mathcal{G} in the case where X is a fractional Brownian motion (fBm) with Hurst parameter $H \in (0, 1)$ including standard Brownian motion ($H = 1/2$) as a particular case.

A fBm is a centered Gaussian process with a specific covariance function [29], detailed definition and properties are given in Section 2. The motivation in choosing such a model for our study is threefold. First, its stochastic simulation is known to be quite delicate (at least for $H \neq 1/2$), especially when the number of time points gets larger and larger – see [4, 11] for a review and [5, 23] for recent contributions – hence having at hand a generative model for the full path is really appealing for practical use. Second, it is widely used in various real-life modelings: uni and bipedal postural standing in biomechanics [3]; volatility of financial assets [7, 16]; vortex filament structures observed in 3D fluids [15]; prices of electricity in a liberated market [2]; solar cycle [34]; for other fractional-based modeling, see [6]. Third, understanding the right design of \mathcal{G} for generating the fBm distribution may well open the way to handle more complicated stochastic models written as a Stochastic Differential Equation (SDE) driven by fBm for instance: indeed, as we will see, the design of the current \mathcal{G} inherits much from the time-regularity of X ¹ and this property is lifted to SDE driven by X . This part is left to further investigation.

In this work we study the required depth and complexity (size and number of computations) of a deep-feedforward neural network (NN) for \mathcal{G} , with a Rectified Linear Unit (ReLU) for the activation function [18, Chapter 6]: it is referred to as ReLU NN in the sequel. For the latent distribution μ , we consider N independent components and without loss of generality for the simulation purpose, each of them is taken as a standard Gaussian random variable. Essentially, our results state (Theorems 2 and 3) that for a given latent dimension N , there is a $G_\theta \in \mathcal{G}$ such that the equality (1) holds with an error $N^{-H} (1 + \log(N))^{1/2}$ in sup norm with probability $1 - p$. Moreover, focusing on the rates with respect to $N \rightarrow +\infty$, the depth of G_θ is at most

$$\mathcal{O}(\log N)$$

and its global complexity is

$$\mathcal{O}(N^{1+\zeta} \log N),$$

where ζ is a positive parameter that can be taken as small as desired, and where the $\mathcal{O}(\cdot)$ depend on

¹remind that the parameter H mostly coincides with the Hölder exponent of the paths.

p, ζ and H . In particular for the Brownian motion ($H = 1/2$) we can take $\zeta = 0$. A more detailed dependence on p, ζ and H is given latter.

These results are original to the best of our knowledge, and should play a key role in tuning GAN-based methods in the choice of the parametric family of NN for generating fractional stochastic processes in continuous time. These results make a clear connection between the time-regularity of the path (that could be measured on the real observed data) and the architecture of the parameterization to set up.

This work is organized as follows. In Section 2, we recall few properties of fBm. Our approximations are based on wavelet decompositions and we will provide appropriate materials. Then we state our main quantitative results about depth and complexity of deep ReLU neural networks for generating fBm. Section 3 is devoted to the proofs. For pedagogical and technical reasons, we start with the case $H = 1/2$ (standard Brownian motion) in Subsection 3.1; then we handle the general case of fBm in Subsection 3.2. A few technical proofs and numerical illustrations are postponed to Appendix.

Notations: The set of naturals without zero $\mathbb{N}_0 := \{1, 2, \dots, n, \dots\}$ and $\mathbb{N} := \mathbb{N}_0 \cup \{0\}$; the set $\mathcal{M} := \{2^{n+1}, n \in \mathbb{N}\}$; the vector of N standard Gaussian random variables G_1, \dots, G_N is denoted by $G_{1:N}$; the imaginary number $\mathbf{i}^2 = -1$. We write $x = \mathcal{O}_c(y)$ if $|x| \leq c|y|$ for some positive constant c which, in the context where it is used, does not depend neither on the latent dimension N nor on the accuracy ε ; usually y will be a not-small quantity ($y \geq 1$) as a polynomial or logarithmic function of N or/and ε^{-1} according to the context. Finally, we write $a_N \asymp b_N$ if there exists a constant $c \geq 1$ such that $\forall N \in \mathbb{N}_0, c^{-1} \leq a_N/b_N \leq c$.

2 Preliminaries and main results

2.1 About Fractional Brownian motion

Fractional Brownian motion (fBm) $\{B^H(t)\}_{t \in \mathbb{R}}$ with a Hurst parameter $H \in (0, 1)$ is a Gaussian process, centered ($\mathbb{E}[B^H(t)] = 0$), with covariance function

$$\mathbb{Cov}(B^H(t), B^H(s)) = \frac{1}{2} \left(|t|^{2H} + |s|^{2H} - |t - s|^{2H} \right), \quad \text{for any } s, t \geq 0.$$

When $H = 1/2$, we will simply write B instead of $B^{1/2}$. Our aim is to approximate the distribution of B^H on a finite interval: owing to the self-similarity property of fBm [32, Proposition 2.1], we can consider, without loss of generality, the interval $[0, 1]$, which is our setting from now on.

As B^H is a centered Gaussian process in a Banach space $(C^0([0, 1], \mathbb{R}), \|\cdot\|_\infty)$ (see [25, Proposition 3.6]), B^H admits almost sure (a.s.) series representation of the form

$$B^H(t) = \sum_{k=0}^{\infty} u_k(t) G_k, \quad \forall t \in [0, 1], \quad (2)$$

where $\{u_k\}_{k \in \mathbb{N}}$ is a sequence of continuous non-random functions, and $\{G_k\}_{k \in \mathbb{N}}$ is a sequence of independent standard Gaussian variables $\mathcal{N}(0, 1)$. The equality (2) holds in the sense that the series converges a.s. uniformly. Such representations for fBm are studied in [31] using wavelets.

Let $H \in (0, 1)$; [24] showed that there exists a sequence $\{u_k\}_k$ such that the L^2 -truncation error is

$$\left(\mathbb{E} \left[\sup_{t \in [0, 1]} \left| \sum_{k=N}^{\infty} u_k(t) G_k \right|^2 \right] \right)^{1/2} \asymp N^{-H} (1 + \log(N))^{1/2}; \quad (3)$$

in addition, the above convergence rate is optimal among all sequences $\{u_k\}_k$ for which (2) converges a.s. in sup-norm.

In [31] the authors focused on the a.s. uniform convergence on $[0, 1]$ for different wavelet representations series (2) using a specific mother wavelet function ψ , and the authors of [1, Theorem 5] showed their optimality in the sense of (3). Not only ψ has to generate an orthonormal basis $\{\psi_{j,k}(t) = 2^{j/2}\psi(2^j t - k)\}_{(j,k) \in \mathbb{Z}^2}$ of $L^2(\mathbb{R}, dx) = \left\{f : \int_{-\infty}^{\infty} |f(x)|^2 dx < \infty\right\}$ [30, Theorem 1], but also it must respect some other regularity properties discussed hereafter.

In the following, our convention is to write the Fourier transform and its inverse as

$$\widehat{f}(\xi) := \int_{-\infty}^{\infty} f(t) e^{-it\xi} dt, \quad f(t) := \frac{1}{2\pi} \int_{-\infty}^{\infty} \widehat{f}(\xi) e^{it\xi} d\xi.$$

2.2 Brownian motion: wavelet representation and main result for NN generative model

A first well-known series representation is the so-called Lévy construction of the Brownian motion obtained by plugging in (2) the basis function

$$\psi_{j,k}^{\text{FS}}(t) = 2^{j/2} \psi^{\text{FS}}(2^j t - k), \quad j \in \mathbb{N}, k = 0, \dots, 2^j - 1, \quad (4)$$

where $\psi^{\text{FS}}(x) = 2(x\mathbb{1}_{0 \leq x < 1/2} + (1-x)\mathbb{1}_{1/2 \leq x \leq 1})$ is twice the antiderivative of the Haar mother wavelet [20]. The set $\{\psi_{j,k}^{\text{FS}}\}_{j \in \mathbb{N}, k=0, \dots, 2^j-1}$ defines the Faber-Schauder (F-S) system [14, 35] and forms an orthogonal basis of $L^2(\mathbb{R}, dx)$. Thus, given $\{G_1, G_{j,k}\}_{j \geq 0, 0 \leq k < 2^j}$ a sequence of independent standard Gaussian random variables $\mathcal{N}(0, 1)$, the Lévy construction of the Brownian motion states that a.s. the truncated series

$$B^{(n)}(t) := G_1 t + \sum_{j=0}^n \sum_{k=0}^{2^j-1} 2^{-(j+1)} \psi_{j,k}^{\text{FS}}(t) G_{j,k} \quad (5)$$

converges uniformly on $[0, 1]$ to a Brownian motion B as $n \rightarrow \infty$ (see [36, Section 3.4]). We write $B_N := B^{(n)}$ with $N = 2^{n+1}$, to emphasize that (5) contains N scalar Gaussian random variables, which is consistent with the latent dimension discussed above. The next result quantifies the a.s. convergence rate of B_N to B , the proof is postponed to Section A.1.

Lemma 1. *Let $N \in \mathcal{M}$. Then, there exists a finite random variable $C_{(6)} \geq 0$ such that almost surely*

$$\sup_{t \in [0, 1]} |B(t) - B_N(t)| \leq C_{(6)} N^{-1/2} (1 + \log(N))^{1/2}. \quad (6)$$

The above result is somehow well-known and shows that it is enough to approximate with a high probability the function $t \mapsto B_N(t)$ by a ReLU NN with suitable architecture (see the Appendix for a proof), which is the purpose of the following statement.

Theorem 2. *Let $N \geq 2$ and $(\Omega^N, \mathcal{F}^N, \mathbb{P}^N)$ be a probability space supporting N i.i.d. standard Gaussian random variables $G_{1:N}$. Therefore, there exists an extension $(\Omega, \mathcal{F}, \mathbb{P})$ supporting a Brownian motion B such that $\forall p \in (0, 1]$, there exist a ReLU neural network*

$$\tilde{B}_{N,p} : \begin{cases} \mathbb{R}^N & \rightarrow C^0([0, 1], \mathbb{R}) \\ G_{1:N} := (G_1, \dots, G_N) & \mapsto \tilde{B}_{N,p}(\cdot, G_{1:N}) \end{cases}$$

and a finite random variable $C \geq 0$ (independent from N and p) such that

$$\mathbb{P} \left(\sup_{t \in [0, 1]} |B(t) - \tilde{B}_{N,p}(t, G_{1:N})| \leq C N^{-1/2} (1 + \log N)^{1/2} \right) \geq 1 - p. \quad (7)$$

Additionally, $\tilde{B}_{N,p}$ is composed at most by

1. $\mathcal{O}_c \left(\log \left(\frac{N \rho_N}{(1 + \log N)^{1/2}} \right) \right)$ hidden layers,
2. $\mathcal{O}_c \left(N \log \left(\frac{N \rho_N}{(1 + \log N)^{1/2}} \right) \right)$ neurons and parameters,

with $\rho_N = -\Phi^{-1}(\frac{p}{2N})$ and Φ^{-1} the quantile function of the normal distribution.

The proof is postponed to Subsection 3.1.

Remark. It is known that $\Phi^{-1}(u) \sim -\sqrt{-2 \log u}$ as $u \rightarrow 0^+$, see [12]. Therefore we shall get equivalents of the architecture size, either as $p \rightarrow 0$ or as $N \rightarrow \infty$ (which results in $\rho_N \rightarrow \infty$ anyhow):

1. For a fixed p and as $N \rightarrow \infty$, $\rho_N/(1 + \log N)^{1/2}$ tends to a constant so the depth and the complexity are respectively of order $\mathcal{O}_c(\log(N))$ and $\mathcal{O}_c(N \log(N))$;
2. For a fixed N and as $p \rightarrow 0$, the impact on the network size is moderate since both depth and complexity are of order $\mathcal{O}_c(\log \log(1/p))$.

As a complement to the previous marginal asymptotics, the estimates of Theorem 2 allow to have p dependent on N : for instance, building a ReLU NN with an error tolerance of order $N^{-1/2} (1 + \log N)^{1/2}$ with probability $1 - N^{-k}$ (for any given $k > 0$) can be achieved using a depth $\mathcal{O}_c(\log N)$ and a complexity $\mathcal{O}_c(N \log N)$.

The next two remarks apply both to the current Brownian motion case and to the fBM studied in next subsection.

Remark. One way wonder about the optimality of Theorem 2 in terms of complexity. Actually, the convergence rate in (7) is optimal, as in Lemma 1: see for instance [24]. Then, once the latent dimension N is fixed, one cannot use less than $\mathcal{O}_c(N)$ parameters (no matter which NN is used), which is exactly what we propose (up to the logarithmic term). The remaining question may be related to the depth that we choose as $\log(N)$, showing a slowly-increasing depth regarding the latent dimension: although cheap, there is no reason that this $\log(N)$ -rate is the smallest possible one.

Taking $p = 0$ is not possible within our method of proof, it is an open question if another construction would allow $p = 0$.

Remark. From the GAN point of view, it is enough to know (like in Theorem 2) that such a ReLU NN can generate a Brownian motion without the knowledge of the NN parameters explicitly; indeed, the GAN optimization algorithm will retrieve the parameters.

Regarding the practical use of this generative model, notice that sampling a Brownian motion path boils down to sample $G_{1:N}$ (with N independent standard Gaussian variables), and then compute $\tilde{B}_{N,p}(t, G_{1:N})$ for all the times t required by the situation at hand.

2.3 Fractional Brownian motion: wavelet representation and main result for NN generative model

Among the wavelet fBm series representations proposed in [31], we will focus on the following one

$$B^H(t) = \sum_{j=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} 2^{-jH} (\Psi_H(2^j t - k) - \Psi_H(-k)) G_{j,k}, \quad (8)$$

with

$$\hat{\Psi}_H(\xi) := \frac{\hat{\psi}(\xi)}{(i\xi)^{H+1/2}}. \quad (9)$$

One choice for the wavelet ψ is the Lemarié-Meyer wavelet [26] (see [28, Equations (7.52)-(7.53)-(7.85) and Example 7.10] for more details on its construction) defined by its Fourier transform

$$\widehat{\psi^M}(\xi) := e^{-i\frac{\xi}{2}} \begin{cases} \sin\left(\frac{\pi}{2}\nu\left(\frac{3|\xi|}{2\pi} - 1\right)\right), & \frac{2}{3}\pi \leq |\xi| \leq \frac{4}{3}\pi, \\ \cos\left(\frac{\pi}{2}\nu\left(\frac{3|\xi|}{4\pi} - 1\right)\right), & \frac{4}{3}\pi \leq |\xi| \leq \frac{8}{3}\pi, \\ 0, & \text{otherwise,} \end{cases} \quad (10)$$

where $\nu : \mathbb{R} \rightarrow [0, 1]$ is a smooth function satisfying

$$\nu(x) = \begin{cases} 0, & \text{if } x \leq 0, \\ 1, & \text{if } x \geq 1, \end{cases} \quad \text{and} \quad \nu(x) + \nu(1-x) = 1. \quad (11)$$

Such properties allow to satisfy the quadrature conditions of the conjugate mirror filter [28, Subsection 7.1.3 p. 270] which specifies the scaling function in the construction of wavelet bases (see [28, Chapter 7] for a complete overview of wavelet bases analysis). Considering the truncated series of (8) over a specific set \mathcal{I}_N containing at most N indices (j, k) , the authors of [1, Section 5 p. 469] have showed that there exists a finite r.v. $C_{(12)} \geq 0$ such that

$$\sup_{t \in [0, 1]} |B^H(t) - B_N^H(t)| \leq C_{(12)} N^{-H} (1 + \log(N))^{1/2}. \quad (12)$$

In other words, if Ψ_H is well chosen such that ψ satisfies conditions $(\mathcal{A}_1), (\mathcal{A}_2), (\mathcal{A}_3)$ listed in [1, p. 456], then the wavelet decomposition (8) is optimal [1, Theorem 5] in the sense of (3).

Back to the construction of (10), a classical example of ν due to Daubechies [9, p. 119] is

$$\nu(x) = x^4 (35 - 84x + 70x^2 - 20x^3),$$

which entails $\widehat{\psi^M}$ having 3 vanishing derivatives at $|\xi| = 2\pi/3, 4\pi/3, 8\pi/3$. Below, we will propose another example of ν with higher order vanishing derivatives at the boundaries in order to get (see the proof in Subsection 3.2.1) a fast decay rate of all the derivatives $(\psi^M)^{(k)}$ at infinity, which results in reducing the complexity cost of the ReLU NN architecture that we will build in Theorem 3. Note that the construction below might not be numerically optimal among the large literature on wavelets and applications in signal processing, however it is a concrete example on which we can base our theoretical result.

Construction of ψ^M Let q_{st} be the quantile function of a Student distribution with β degrees of freedom. Thus, considering the function

$$\nu(u) := \frac{1}{1 + \exp(-q_{\text{st}}(u))}, \quad u \in [0, 1] \quad (13)$$

the conditions (11) are easily satisfied. We now briefly verify $(\mathcal{A}_1), (\mathcal{A}_2), (\mathcal{A}_3)$ from [1, p. 456] to validate the use of such ν . Condition (\mathcal{A}_1) clearly holds since the wavelets $\left\{\psi_{j,k}^M\right\}_{(j,k) \in \mathbb{Z}^2}$ generate an orthonormal basis of $L^2(\mathbb{R}, dx)$. Condition (\mathcal{A}_3) is straightforwardly satisfied because $\widehat{\psi^M}$ vanishes around 0. Last, consider (\mathcal{A}_2) : the decay of $\widehat{\psi^M}$ and its derivatives at infinity is straightforward since it has compact support. What really needs to be checked is the smoothness property of $\widehat{\psi^M}$: let us justify that it is C^∞ . Observe that this follows from the tentative property $\nu^{(q)}(0^+) = \nu^{(q)}(1^-) = 0$ for all $q \in \mathbb{N}_0$. To see this, remind that the Student distribution belongs to the Fréchet maximum domain of attraction [10, Theorem. 1.2.1(1.)], therefore q_{st} increases as a power function with a exponent

$\gamma = 1/\beta > 0$ called the tail-index. Second, for all $q \in \mathbb{N}$, $q_{\text{st}}^{(q)}(u)$ increases at most as a power function around 0 and 1, see [17, Lemma 15]. Moreover, the sigmoid function $\Delta(x) := (1 + \exp(-x))^{-1}$ clearly satisfies $\Delta^{(q)}(x) = \mathcal{O}_c(\exp(-|x|))$ for $q \geq 1$. Hence, applying the Faà di Bruno formula [22, p. 224-226] for expanding the derivative of the composition of $\Delta(\cdot)$ and $q_{\text{st}}(\cdot)$ gives

$$\nu^{(q)}(u) = \sum_{l=1}^q \frac{1}{l!} \Delta^{(l)}(q_{\text{st}}(u)) \sum_{i_s \in \mathbb{N}_0: i_1 + \dots + i_l = q} \frac{q!}{i_1! i_2! \dots i_l!} \prod_{s=1}^l q_{\text{st}}^{(i_s)}(u), \quad q \in \mathbb{N}_0,$$

which readily leads to $\nu^{(q)}(0^+) = \nu^{(q)}(1^-) = 0$ since the exponential function decays faster than any polynomials. See Figure 7 for an illustration of Ψ_H built with (10) and (13).

We are now in a position to state our second main result.

Theorem 3. *Let $N \geq 2$ and $(\Omega^N, \mathcal{F}^N, \mathbb{P}^N)$ be a probability space supporting N i.i.d. standard Gaussian random variables $G_{1:N}$. Therefore, there exists an extension $(\Omega, \mathcal{F}, \mathbb{P})$ supporting a fractional Brownian motion B^H such that $\forall p \in (0, 1]$, for all $r \in \mathbb{N}_0$ there exist a ReLU neural network*

$$\tilde{B}_{N,p}^H : \begin{cases} \mathbb{R}^N & \rightarrow C^0([0, 1], \mathbb{R}) \\ G_{1:N} := (G_1, \dots, G_N) & \mapsto \tilde{B}_{N,p}^H(\cdot, G_{1:N}) \end{cases}$$

and a finite random variable $C \geq 0$ (independent from N and p) such that

$$\mathbb{P} \left(\sup_{t \in [0, 1]} \left| B^H(t) - \tilde{B}_{N,p}^H(t, G_{1:N}) \right| \leq C N^{-H} (1 + \log(N))^{1/2} \right) \geq 1 - p.$$

Additionally, $\tilde{B}_{N,p}^H$ is composed by

1. $\mathcal{O}_c \left(\log \left(\frac{N \rho_N}{(1 + \log(N))^{1/2}} \right) \right)$ hidden layers,
2. $\mathcal{O}_c \left(N^{1 + \frac{H+1}{2r}} \log \left(\frac{N \rho_N}{(1 + \log(N))^{1/2}} \right) \left(\frac{\rho_N}{(1 + \log(N))^{1/2}} \right)^{\frac{1}{2r}} \right)$ neurons and parameters,

where ρ_N is defined in Theorem 2. The constants in $\mathcal{O}_c(\cdot)$ may depend on r and H .

Observe that Remarks 2.2-2.2-2.2 apply similarly to the above Theorem.

2.4 Discussion

In Table 1 we compare the asymptotic architecture bounds between a BM and a fBm. Note that the BM benefits from a natural construction of the F-S wavelet through ReLU functions. In comparison, the fBm construction suffers from 1) an additional approximation of the wavelet Ψ_H and 2) a larger bound on the sum over \mathcal{I}_N , which has only a log impact in the asymptotic NN architecture (see details in the proof in Subsection ou 3.2.2). Therefore, both models have the same asymptotic depth (with a constant depending on r) and a very close complexity in terms of the latent dimension N .

The takeaway message from these results is that a NN with N Gaussian r.v. as inputs for approximating a process with a time regularity H (and an approximation error N^{-H} up to log-term) may have at most a depth $\log N$ and a complexity $N \log N$. Although the set \mathcal{I}_N is not explicit for finding the optimal fBm NN parameters, this part can be achieved through the optimization of the GAN model with the appropriate architecture detailed in Subsection 3.2.

	BM	fBm
error tolerance (TOL)	$N^{-1/2}$	N^{-H}
depth	$\log(\text{TOL}^{-1})$	$\log(\text{TOL}^{-1})$
complexity	$\text{TOL}^{-2} \log(\text{TOL}^{-1})$	$\text{TOL}^{-(\frac{1}{H}+\zeta)} \log(\text{TOL}^{-1})$

Table 1: For a given confidence probability p , asymptotic complexity rates with respect to tolerance error (TOL). The parameter ζ can be taken arbitrary small, constants depending on H , p and $\zeta > 0$ are omitted.

3 Proofs

In this section we will discuss the constructive proofs of the ReLU NN that appear in the main results. We start with the case $H = 1/2$ (standard Brownian motion) in Subsection 3.1 and we then handle the general case of fBm in Subsection 3.2. Before that, recall the output expression of a 1-hidden layer NN given some input $x \in \mathbb{R}$ and parameterized by $\theta = \left\{ w_k^{(1)}, w_k^{(2)}, b_k^{(1)}, b_k^{(2)} \right\}_{k=1}^K$ is

$$\sum_{k=1}^K w_k^{(2)} \sigma \left(w_k^{(1)} x + b_k^{(1)} \right) + b_k^{(2)}, \quad (14)$$

with $\sigma(x) := \max(0, x)$ the ReLU function. Similarly, a multi-layer NN is just multiple compositions of σ with (14) between different hidden layers. For readers interested in having references on approximation properties of NN, we may refer to [27, Theorem 1 p. 70] for L_2 error using single hidden layer NN, to [33, Corollary 6.4 p. 170] for uniform approximation, and to a more recent paper [39] which has shown some uniform convergence rate for multi-layers NN.

3.1 NN representation of BM

In the following proof of Theorem 2 we will restrict to $N \in \mathcal{M} = \{2^{n+1}, n \in \mathbb{N}\}$. However note that if one wants to choose a $N \notin \mathcal{M}$, it will neither impact the error nor the complexity bounds in Theorem 2. Indeed, it suffices to take $n = \left\lfloor \frac{\log(N)}{\log(2)} - 1 \right\rfloor$ and $N' = 2^{n+1}$ such that $N' \in (\frac{N}{2}, N]$, and then set $\tilde{B}_{p,N}(t, G_{1:N}) := \tilde{B}_{p,N'}(t, G_{1:N'})$. Regarding the error bound and complexity w.r.t. N , use those for N' by easily adjusting constants: indeed, since $N' \leq N$, it follows that $\rho_{N'} \leq \rho_N$, $\frac{1}{(1+\log(N'))} \leq \frac{1}{1-\log(2)} \frac{1}{(1+\log(N))}$ for the complexity bound and $N'^{-1/2}(1+\log(N'))^{1/2} \leq \sqrt{2}N^{-1/2}(1+\log(N))^{1/2}$ for the error bound.

From now on, $N = 2^{n+1}$. For ease of notation, let

$$s_{j,k}(t) := \frac{\psi_{j,k}(t)}{2^{j/2}} = \psi(2^j t - k) \in [0, 1], \quad (15)$$

be the normalized F-S wavelet, where $\psi = \psi^{\text{FS}}$ in this section. Then, in view of (5) and Lemma 1, the objective is to find a ReLU NN with N standard Gaussian variables and the time t as inputs, that can approximate with uniform error and high probability

$$B_N(t) = G_1 t + \sum_{j=0}^n \sum_{k=0}^{2^j-1} 2^{-(j/2+1)} s_{j,k}(t) G_{j,k}. \quad (16)$$

The key advantage with the F-S wavelet (4) is that the mother wavelet ψ can be built easily with 3 ReLUs and 9 parameters such as

$$\psi(x) = 2 \left(\sigma(x) - 2\sigma\left(x - \frac{1}{2}\right) + \sigma(x-1) \right). \quad (17)$$

Clearly a product operation in (16) is required between the inputs $G_{j,k}$ (i.e. the latent space in a GAN setting) and the normalized wavelets $s_{j,k}$ just built (see Figure 1). Since such an operation is not natively done in a feedforward network, let us study how to approximate it.




Figure 1: Neural network construction of a normalized Faber-Schauder basis function (15). The circles filled with σ represent a ReLU function, while the ones with a $/$ represent the identity function.

3.1.1 How to make a product with a NN

Let $h(x) = x^2$. The key observation in [39, Proposition 2] based on [37] is that h can be approximated by piece-wise linear interpolation

$$\tilde{h}_\ell(x) = x - \sum_{j=1}^{\ell} \frac{\psi^{[\circ j]}(x)}{2^{2j}}, \quad (18)$$

with

$$\psi^{[\circ j]}(x) := \underbrace{\psi \circ \dots \circ \psi}_{j \text{ times}}(x), \quad (19)$$

such that

$$\sup_{x \in [0,1]} |h(x) - \tilde{h}_\ell(x)| = 2^{-2(\ell+1)}.$$

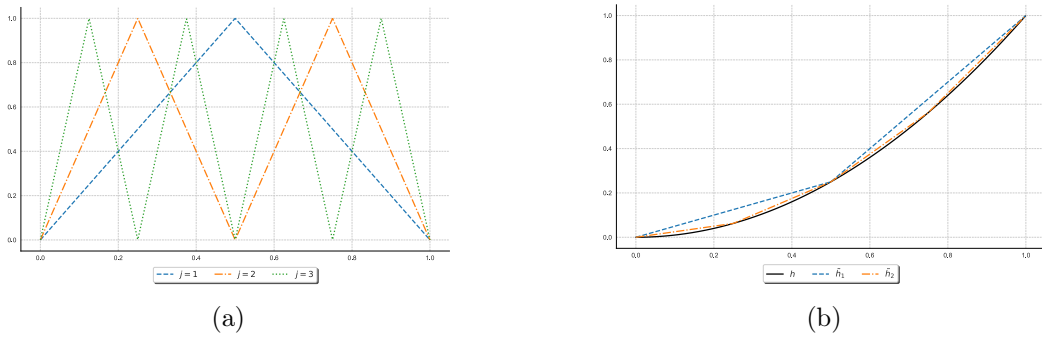


Figure 2: Plot of $\psi^{[\circ j]}$ for $j = \{1, 2, 3\}$ (a); approximation of $h(x) = x^2$ with \tilde{h}_ℓ for $\ell = \{1, 2\}$

The expression (18) can be interpreted as a NN approximation with ℓ hidden layers, where each composition in (19) is just the sum of all translated positions of a F-S wavelet, i.e. $\forall j \geq 0, \psi^{[\circ j+1]}(x) = \sum_{k=0}^{2^j-1} \psi(2^j x - k)$. Therefore, instead of making a linear combination of such functions built through a long single hidden layer, the benefit of increasing the depth of the network allows to increase at a geometric rate the number of wavelets and to reduce the complexity cost from $3 \times 2^{\ell-1}$ to 3ℓ neurons.

Additionally, one shall be aware that \tilde{h}_ℓ does approximate the square function only inside the interval $[0, 1]$ (see Figure 2b). Therefore we introduce a new function

$$\check{h}_\ell(x) = \tilde{h}_\ell(|x|) = \tilde{h}_\ell(\sigma(x) + \sigma(-x)), \quad (20)$$

which applies a ReLU absolute value on the input. Obviously \check{h}_ℓ extends the approximation of h on $[-1, 1]$ such that

$$\sup_{|x| \leq 1} |h(x) - \check{h}_\ell(x)| = 2^{-2(\ell+1)}. \quad (21)$$

The NN construction of (20) requires $\ell + 1$ hidden layers and $\mathcal{O}_c(\ell)$ neurons and parameters (see Figure 3).




Figure 3: Neural network architecture of \check{h}_ℓ with $\ell = 2$. Lighter arrows refer to similar parameters which can easily be inferred from (17). For implementation purpose, one can obviously bypass the identity function in the middle of the network which is put here for the sake of clarity.

Once the square operation is approximately synthesized through a ReLU NN, we can leverage the polarization identity to get the product operation $(x, y) \mapsto xy$. Because the above approximation (21) is valid only on the interval $[-1, 1]$, it is useful to use a polarization identity with some flexible rescalings of x and y . It writes, for any $a, b > 0$,

$$xy = ab \left(-\left(\frac{x}{2a} - \frac{y}{2b} \right)^2 + \left(\frac{x}{2a} + \frac{y}{2b} \right)^2 \right).$$

The following Proposition provides a uniform error bound on the approximation of the product with a ReLU NN.

Proposition 4. *Let $k(x, y) := xy$. Then, for any $\ell \in \mathbb{N}_0$, for given $a > 0$ and $b > 0$,*

1. *there exists a NN $\tilde{k}_\ell^{a,b} : \mathbb{R}^2 \rightarrow \mathbb{R}$ with $\ell + 1$ hidden layers such that*

$$\sup_{x, y: |x| \leq a, |y| \leq b} |k(x, y) - \tilde{k}_\ell^{a,b}(x, y)| \leq ab 2^{-(2\ell+1)}; \quad (22)$$

2. *if $x = 0$ or $y = 0$, then $\tilde{k}_\ell^{a,b}(x, y) = 0$;*
3. *the ReLU NN $\tilde{k}_\ell^{a,b}$ can be implemented with no more than $\mathcal{O}_c(\ell)$ complexity and a depth $\ell + 1 := \left\lceil \frac{1}{2 \log(2)} \log \left(\frac{ab}{\varepsilon} \right) - \frac{1}{2} \right\rceil + 1$, where ε is the error tolerance in sup norm.*

Proof. It is enough to set

$$\tilde{k}_\ell^{a,b}(x, y) := ab \left(-\check{h}_\ell \left(\frac{x}{2a} - \frac{y}{2b} \right) + \check{h}_\ell \left(\frac{x}{2a} + \frac{y}{2b} \right) \right)$$

and to apply (21), while observing that when $|x| \leq a$ and $|y| \leq b$, $\frac{x}{2a} \pm \frac{y}{2b} \in [-1, 1]$. \square

3.1.2 Final approximation of B_N

Based on Proposition 4, it seems that we can deduce a uniform bound on the product $s_{j,k}(t)G_{j,k}$ by a linear combination of composition functions of ReLUs (i.e. a multi-layer NN). Nevertheless, recall that (22) only holds for $|x| \leq a$ and $|y| \leq b$. Thus, although it is clear from (15) that for all $t \in [0, 1]$ we have $s_{j,k}(t) \in [0, 1]$, the random variables $G_{j,k}$ need however to be bounded in order to use Proposition 4: it can be made only with some probability.

Proposition 5. *Let $N \in \mathbb{N}_0$ and $p \in (0, 1]$, set*

$$\rho_N = -\Phi^{-1}\left(\frac{p}{2N}\right) \geq 0,$$

with Φ^{-1} the quantile function of the normal distribution, and let $G_{1:N}$ be i.i.d. standard Gaussian r.v.. Then

$$\mathbb{P}(\forall i = 1, \dots, N : |G_i| \leq \rho_N) \geq 1 - p.$$

Proof. Clearly, the probability on the above left hand side equals

$$1 - \mathbb{P}\left(\bigcup_{i=1}^N \left\{|G_i| \geq \rho_N\right\}\right) \geq 1 - 2N\Phi(-\rho_N) = 1 - p.$$

□

Therefore combining Propositions 4 and 5 with $a = 1$ and $b = \rho_N$, we can define

$$\tilde{k}_\ell^{1,\rho_N}(s_{j,k}(t), G_{j,k}) := \rho_N \left(-\check{h}_\ell \left(\frac{s_{j,k}(t)}{2} - \frac{G_{j,k}}{2\rho_N} \right) + \check{h}_\ell \left(\frac{s_{j,k}(t)}{2} + \frac{G_{j,k}}{2\rho_N} \right) \right), \quad (23)$$

which can be implemented with $\ell + 2$ hidden layers (since we need an additional one to build the $s_{j,k}$) and $\mathcal{O}_c(\ell)$ neurons and parameters (see Figure 4).

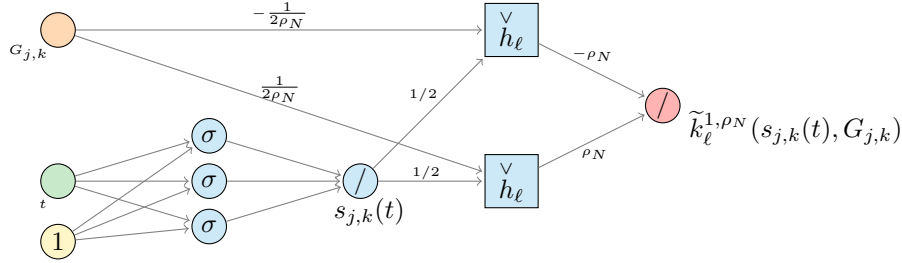


Figure 4: Neural network architecture of $\tilde{k}_\ell^{1,\rho_N}(s_{j,k}(t), G_{j,k})$.

Remark. *The key advantage of (23) over the polarization identity in [39, Equation (3) p. 8] is that ρ_N can be directly parameterized inside the NN instead of pre-normalizing $G_{j,k}$ for the approximation of $G_{j,k}^2$.*

Let \tilde{B}_N be the NN approximation of (16) such that

$$\tilde{B}_N(t) = \tilde{k}_\ell^{1,\rho_N}(t, G_1) + \sum_{j=0}^n \sum_{k=0}^{2^j-1} 2^{-(j/2+1)} \tilde{k}_\ell^{1,\rho_N}(s_{j,k}(t), G_{j,k}), \quad (24)$$

with $\tilde{k}_\ell^{1,\rho_N}$ defined in (23). Therefore, on the event $\{|G_i| \leq \rho_N : i = 1, \dots, N\}$ which has a probability greater than $1 - p$, one has

$$\begin{aligned} \sup_{t \in [0,1]} |B_N(t) - \tilde{B}_N(t)| &\leq \sup_{t \in [0,1]} |tG_1 - \tilde{k}_\ell^{1,\rho_N}(t, G_1)| \\ &\quad + \sup_{t \in [0,1]} \left| \sum_{j=0}^n \sum_{k=0}^{2^j-1} 2^{-(j/2+1)} \left(s_{j,k}(t) G_{j,k} - \tilde{k}_\ell^{1,\rho_N}(s_{j,k}(t), G_{j,k}) \right) \right| \end{aligned}$$

(from Proposition 4 and since $s_{j,k}(t) \in [0, 1]$ and $|G_{j,k}| \leq \rho_N$)

$$\begin{aligned} &\leq \rho_N 2^{-(2\ell+1)} \left(1 + \sum_{j=0}^n \sum_{k=0}^{2^j-1} 2^{-(j/2+1)} \right) \\ &= \rho_N 2^{-(2\ell+1)} \left(1 + \frac{N^{1/2} - 1}{2(\sqrt{2} - 1)} \right) \quad (\text{recall } N = 2^{n+1}) \\ &\leq \rho_N 2^{-2\ell} N^{1/2}. \end{aligned}$$

Hence, with probability at least $1 - p$, combining Lemma 1 with the above yields

$$\sup_{t \in [0,1]} |B(t) - \tilde{B}_N(t)| \leq C_{(6)} N^{-1/2} (1 + \log(N))^{1/2} + \rho_N 2^{-2\ell} N^{1/2}.$$

It follows that, if we choose

$$\ell = \left\lceil \frac{1}{2 \log(2)} \log \left(\frac{N \rho_N}{(1 + \log(N))^{1/2}} \right) \right\rceil \vee 1, \quad (25)$$

then (7) is proved with $\tilde{B}_{N,p}(t, G_{1:N}) := \tilde{B}_N(t)$. All in all, based on Figure 4, the architecture required for the N products in (24), *i.e.* N sub-networks, yields a total of at most $\ell + 2$ hidden layers and a complexity $\mathcal{O}_c(N\ell)$. Replacing with (25) gives the stated bounds. \square

3.2 NN representation of fBm

Now that we are acquainted with the case of BM, we can move on to the more general case which requires additional arguments. In view of (12) with a fixed $\gamma > 0$, the goal here is to prove that there exists a ReLU NN approximating uniformly

$$B_N^H(t) = \sum_{(j,k) \in \mathcal{I}_N} 2^{-jH} (\Psi_H(2^j t - k) - \Psi_H(-k)) G_{j,k}, \quad (26)$$

with $\text{Card}(\mathcal{I}_N) \leq N$ and $G_{j,k} \sim \mathcal{N}(0, 1)$. The proof will be composed in two parts. First we will discuss how Ψ_H can be approximated by ReLU basis functions in \mathbb{R} . Second, we will see how to control the error on the product with Gaussians in (26). In this section we will write $\psi^M(\cdot) = \psi(\cdot)$ for the Lemarié-Meyer wavelet (10) with $\nu(\cdot)$ as in (13).

3.2.1 Approximation of Ψ_H

We want to show that for all $\varepsilon \in (0, 1)$ there exists a ReLU NN \tilde{g} such that

$$\sup_{u \in \mathbb{R}} |\Psi_H(u) - \tilde{g}(u)| \leq \varepsilon. \quad (27)$$

Note that we cannot apply the universal approximation theorem [8, Theorem 1] which holds for continuous functions with compact support. To tackle the infinite support, the strategy will consist of

first approximating Ψ_H in some interval $[-u_{\max}, u_{\max}]$, and then using the fast decay rate of $|\Psi_H(u)|$ for $|u| > u_{\max}$. Indeed, since by construction $\widehat{\psi}$ and its derivatives vanish in the neighborhood of $\xi = 0$, $\widehat{\Psi}_H$ defined in (9) is \mathcal{C}^∞ with compact support for any parameter $H \in \mathbb{R}$. So for all $(m, q) \in \mathbb{N}^2$, we claim that

$$\left| \Psi_H^{(q)}(u) \right| \leq \frac{C_{H-q,m}}{1 + |u|^{m+1}}, \quad (28)$$

where $C_{H-q,m}$ is a constant depending on $H-q$ and m . The property for $q = 0$ is clear: use the inverse Fourier transform and $m+1$ integration by parts, taking advantage that the derivatives of $\widehat{\Psi}_H$ vanish at the boundary of its support (see discussion after (13)). For $q \neq 0$, observe that $\Psi_H^{(q)}(u) = \Psi_{H-q}(u)$ and the property follows. Now we proceed to (27), by following the ideas of [39, Theorem 1] with some variations. In (28), we have a degree of freedom with the choice of the parameter m , it will be fixed at the end of the proof.

Consider a uniform grid of M points $\{u_i = (i-1)\delta - u_{\max}\}_{i=1}^M$ with $M > 1$ and $\delta = \frac{2u_{\max}}{M-1}$ on the domain $[-u_{\max}, u_{\max}]$, assuming $\delta \leq 1/2$. The parameter $u_{\max} > 0$ will be fixed later. Additionally, for $i = 1, \dots, M$, we define a triangular function

$$\phi_i(u) := \phi\left(\frac{u - u_i}{\delta}\right),$$

where

$$\phi(t) := \sigma(t+1) + \sigma(t-1) - 2\sigma(t),$$

and with the following (obvious) properties:

1. $\phi_i(\cdot)$ is symmetric around u_i ,
2. $\sup_{u \in \mathbb{R}} |\phi_i(u)| = \phi_i(u_i) = 1$,
3. $\text{supp}(\phi_i) \in [u_i - \delta, u_i + \delta]$,
4. $\sum_{i=1}^M \phi_i(u) \equiv 1$, for $u \in [-u_{\max}, u_{\max}]$.

The function ϕ_i is nothing else than another FS wavelet $\psi_{j,k}^{\text{FS}}$ with slightly different scaling and position parameters. Now let $r \in \mathbb{N}_0$ and consider a localized Taylor polynomial function

$$g_1(u) := \sum_{i=1}^M \phi_i(u) P_i(u), \quad (29)$$

where P_i is the Taylor polynomial of degree $(r-1)$ of $\Psi_H \in \mathcal{C}^\infty$ at the point u_i given by

$$P_i(u) := \sum_{q=0}^{r-1} \frac{\Psi_H^{(q)}(u_i)}{q!} (u - u_i)^q.$$

To approximate the q -power function, we will need the following result.

Proposition 6. *Let $\ell \in \mathbb{N}_0$, $a > 0$ and $b > 0$. For any $q \in \mathbb{N}$, define recursively the ReLU NN with at most $(q-1)(\ell+1)$ hidden layers by*

$$y \mapsto \widetilde{y}^q := \widetilde{k}_\ell^{b, b_q} \left(y, \widetilde{y}^{q-1} \right), \quad q \geq 2,$$

with by convention $\widetilde{y}^0 := 1$, $\widetilde{y}^1 := y$, where $\widetilde{k}_\ell^{a,b}$ is defined in Proposition 4 and where

$$b_q := b^{q-1} \left(1 + 2^{-(2\ell+1)} \right)^{q-2}. \quad (30)$$

It is such that

$$\sup_{y: |y| \leq b} |y^q - \tilde{y}^q| \leq b^q \left(\left(1 + 2^{-(2\ell+1)} \right)^{q-1} - 1 \right), \quad (31)$$

$$\sup_{x, y: |x| \leq a, |y| \leq b} \left| xy^q - \tilde{k}_\ell^{a, b_{q+1}}(x, \tilde{y}^q) \right| \leq ab^q \left(\left(1 + 2^{-(2\ell+1)} \right)^q - 1 \right). \quad (32)$$

Proof. We set $\eta := 2^{-(2\ell+1)}$ and we proceed by induction. The inequality (31) holds for $q = 2$ thanks to Proposition 4. Now take $q \geq 3$, assume (31) holds for $q - 1$. Clearly, this implies

$$\sup_{|y| \leq b} \left| y^{q-1} \right| \leq b^{q-1} (1 + \eta)^{q-2} = b_q. \quad (33)$$

Therefore,

$$\begin{aligned} \sup_{|y| \leq b} |y^q - \tilde{y}^q| &\leq \sup_{|y| \leq b} |y^q - y \widetilde{y^{q-1}}| + \sup_{|y| \leq b} |y \widetilde{y^{q-1}} - \tilde{y}^q| \\ &\leq b \sup_{|y| \leq b} |y^{q-1} - \widetilde{y^{q-1}}| + \sup_{|y| \leq b} |y \widetilde{y^{q-1}} - \tilde{k}_\ell^{b, b_q}(y, \widetilde{y^{q-1}})| \\ &\leq b b^{q-1} ((1 + \eta)^{q-2} - 1) + b b^{q-1} (1 + \eta)^{q-2} \eta \\ &= b^q ((1 + \eta)^{q-1} - 1) \end{aligned}$$

where, at the last inequality, we have used Proposition 4 combined with the bound (33). We are done with (31). Similarly for (32) we get

$$\sup_{|x| \leq a, |y| \leq b} \left| xy^q - \tilde{k}_\ell^{a, b_{q+1}}(x, \tilde{y}^q) \right| \leq a \sup_{|y| \leq b} |y^q - \tilde{y}^q| + \sup_{|x| \leq a, |y| \leq b} \left| x \tilde{y}^q - \tilde{k}_\ell^{a, b_{q+1}}(x, \tilde{y}^q) \right|.$$

Combining (31) and Proposition 4 with (33), we get (32). \square




Figure 5: Neural network architecture of all power functions of $(u - u_i)^q$ for $q \in \{2, \dots, r - 1\}$ with $b = 2\delta$ and b_q defined in (30).

We are now in a position to prove (27). Given the support property of ϕ_i , the strategy consists of splitting the error approximation in three pieces:

1. A classical Taylor bound on the main interval yields

$$\begin{aligned} \sup_{|u| \leq u_{\max}} |\Psi_H(u) - g_1(u)| &= \sup_{|u| \leq u_{\max}} \left| \sum_{i=1}^M \phi_i(u) (\Psi_H(u) - P_i(u)) \right| \\ &\leq 2 \max_{i=1, \dots, M} \sup_{u \in \text{supp}(\phi_i)} |\Psi_H(u) - P_i(u)| \end{aligned}$$

since u is in the support of at most two ϕ_i 's and $|\phi_i(u)| \leq 1$,

$$\begin{aligned} &\leq 2 \max_{i=1,\dots,M} \sup_{u \in \text{supp}(\phi_i)} \frac{|\Psi_H^{(r)}(u)|}{r!} (2\delta)^r \\ &\leq \frac{2}{r!} C_{H-r,m} (2\delta)^r, \end{aligned}$$

using (28). Let $\widetilde{g}_{i,q}(\cdot)$ be the ReLU NN approximation of $u \mapsto \phi_i(u)(u - u_i)^q$ using (32) with $|\phi_i(u)| \leq 1 = a$ and $|u - u_i| \leq 2\delta = b$ (see Figure 5). In view of (29), set

$$\widetilde{g}(u) := \sum_{q=0}^{r-1} \frac{1}{q!} \sum_{i=1}^M \Psi_H^{(q)}(u_i) \widetilde{g}_{i,q}(u). \quad (34)$$

Observe, from statement (2) of Proposition 4, that we have $\widetilde{g}_{i,q}(u) = 0$ for $u \notin \text{supp}(\phi_i)$. So using (32) (setting $\eta := 2^{-(2\ell_1+1)}$ with $\ell_1 \in \mathbb{N}_0$) leads to

$$\begin{aligned} \sup_{|u| \leq u_{\max}} |g_1(u) - \widetilde{g}(u)| &\leq \sum_{q=0}^{r-1} \sup_{|u| \leq u_{\max}} \frac{|\Psi_H^{(q)}(u)|}{q!} \sum_{i=1}^M \sup_{|u| \leq u_{\max}} |\phi_i(u)(u - u_i)^q - \widetilde{g}_{i,q}(u)| \\ &\leq 2r \max_{q=0,\dots,r-1} C_{H-q,m} \sum_{q=1}^{\infty} (2\delta)^q \frac{((1+\eta)^q - 1)}{q!}. \end{aligned} \quad (35)$$

Using that $2\delta \leq 1$ and $\eta \leq 1/2$, we easily get that the above right hand side is bounded by $r \delta \eta C_e \max_{q=0,\dots,r-1} C_{H-q,m}$ for some universal constant C_e . To sum up, we have proved

$$\sup_{|u| \leq u_{\max}} |\Psi_H(u) - \widetilde{g}(u)| \leq C_{H,r,m} (\delta^r + \delta\eta)$$

where, here and in what follows, $C_{H,r,m}$ stands for a finite positive constant depending on H, r, m , which value may change from line to line, without changing its name. By taking

$$\delta = \left(\frac{\varepsilon}{6C_{H,r,m}} \right)^{\frac{1}{r}} \wedge \frac{1}{2} = \mathcal{O}_c \left(\varepsilon^{\frac{1}{r}} \right), \quad \eta \leq \frac{\varepsilon}{6C_{H,r,m}\delta} = \mathcal{O}_c \left(\varepsilon^{1-\frac{1}{r}} \right), \quad (36)$$

we have

$$\sup_{|u| \leq u_{\max}} |\Psi_H(u) - \widetilde{g}(u)| \leq \frac{\varepsilon}{3}.$$

The condition on η is satisfied for

$$\ell_1 = \left\lceil \frac{1}{2 \log(2)} \log \left(\frac{6C_{H-r,m}\delta}{\varepsilon} \right) - \frac{1}{2} \right\rceil \vee 1 = \mathcal{O}_c(\log(\varepsilon)). \quad (37)$$

2. Focusing on the short interval $|u| \in [u_{\max}, u_{\max} + \delta]$ where u belongs to $\text{supp}(\phi_M)$ only, write

$$\begin{aligned} &\sup_{|u| \in [u_{\max}, u_{\max} + \delta]} |\Psi_H(u) - g_1(u)| \\ &= \sup_{|u| \in [u_{\max}, u_{\max} + \delta]} |\Psi_H(u) - \phi_M(u) P_M(u)| \\ &\leq \sup_{|u| \in [u_{\max}, u_{\max} + \delta]} |\Psi_H(u)| + \sup_{|u| \in [u_{\max}, u_{\max} + \delta]} |\Psi_H(u) - P_M(u)| \\ &\stackrel{(28)}{\leq} \frac{C_{H,m}}{1 + u_{\max}^{m+1}} + \frac{C_{H-r,m}}{1 + u_{\max}^{m+1}} \frac{\delta^r}{r!} \end{aligned}$$

$$\stackrel{(36)}{\leq} \frac{\mathcal{C}_{H,r,m}}{1 + u_{\max}^{m+1}}.$$

Similarly to the bound (35) but taking advantage of the fast decay of $\sup_{u \in [u_{\max}, u_{\max} + \delta]} |\Psi_H^{(q)}(u)|$ yields to

$$\sup_{|u| \in [u_{\max}, u_{\max} + \delta]} |g_1(u) - \tilde{g}(u)| \leq \mathcal{C}_{H,r,m} \frac{\delta \eta}{1 + u_{\max}^{m+1}}.$$

All in all, and using $\delta \eta \leq 1/4$,

$$\sup_{|u| \in [u_{\max}, u_{\max} + \delta]} |\Psi_H(u) - \tilde{g}(u)| \leq \frac{\mathcal{C}_{H,r,m}}{1 + u_{\max}^{m+1}} \leq \frac{\varepsilon}{3}$$

for a new constant $\mathcal{C}_{H,r,m}$ at the first inequality and from the choice

$$u_{\max} := \left(\frac{3\mathcal{C}_{H,r,m}}{\varepsilon} \right)^{\frac{1}{m+1}} \quad (38)$$

at the last inequality.

3. Finally, on the last interval $|u| \in [u_{\max} + \delta, +\infty)$, $\tilde{g}(\cdot)$ vanishes and from (28)-(38), we readily get

$$\sup_{|u| \in [u_{\max} + \delta, +\infty)} |\Psi_H(u) - \tilde{g}(u)| \leq \frac{\mathcal{C}_{H,m}}{1 + u_{\max}^{m+1}} \leq \frac{\varepsilon}{3}.$$

All in all, (27) is proved with the ReLU NN (34). Collecting previous asymptotics, we get

$$M = \frac{2u_{\max}}{\delta} + 1 = \mathcal{O}_c \left(\varepsilon^{-\frac{1}{m+1}} \varepsilon^{-\frac{1}{r}} \right). \quad (39)$$

3.2.2 Error control including Gaussian variables

We are back to the approximation of (26). For $(j, k) \in \mathcal{I}_N$ we set

$$Y_{j,k}(t) := \Psi_H(2^j t - k) - \Psi_H(-k) \quad \text{and} \quad \tilde{Y}_{j,k}(t) := \tilde{g}(2^j t - k) - \tilde{g}(-k)$$

for its ReLU NN approximation. In view of (26), let us derive an error bound of the product $Y_{j,k}(t)G_{j,k}$ for $t \in [0, 1]$ and $G_{j,k}$ a standard Gaussian random variable. From (27) with $\varepsilon \leq 1$ and (28), we get

$$\sup_{t \in [0, 1]} \left| \tilde{Y}_{j,k}(t) \right| \vee \sup_{t \in [0, 1]} |Y_{j,k}(t)| \leq 2\varepsilon + 2 \sup_{u \in \mathbb{R}} |\Psi_H(u)| \leq 2(1 + C_{H,m}) := \bar{C}_H. \quad (40)$$

Similarly to (23), we can rewrite for $t \in [0, 1]$ and $(j, k) \in \mathcal{I}_N$ the NN product approximation of $\tilde{Y}_{j,k}(t)G_{j,k}$ with $\ell_2 \in \mathbb{N}_0$ as

$$\tilde{k}_{\ell_2}^{\bar{C}_H, \rho_N} \left(\tilde{Y}_{j,k}(t), G_{j,k} \right) = \bar{C}_H \rho_N \left(-\check{h}_{\ell_2} \left(\frac{\tilde{Y}_{j,k}(t)}{2\bar{C}_H} - \frac{G_{j,k}}{2\rho_N} \right) + \check{h}_{\ell_2} \left(\frac{\tilde{Y}_{j,k}(t)}{2\bar{C}_H} + \frac{G_{j,k}}{2\rho_N} \right) \right). \quad (41)$$

Let us work on the event $\{|G_{j,k}| \leq \rho_N : (j, k) \in \mathcal{I}_N\}$ which has a probability greater than $1 - p$ and let us focus on the approximation error of the first term on the right-hand side of (41):

$$\begin{aligned} & \sup_{t \in [0, 1]} \left| \check{h}_{\ell_2} \left(\frac{\tilde{Y}_{j,k}(t)}{2\bar{C}_H} - \frac{G_{j,k}}{2\rho_N} \right) - \left(\frac{Y_{j,k}(t)}{2\bar{C}_H} - \frac{G_{j,k}}{2\rho_N} \right)^2 \right| \\ & \leq \sup_{t \in [0, 1]} \left| \check{h}_{\ell_2} \left(\frac{\tilde{Y}_{j,k}(t)}{2\bar{C}_H} - \frac{G_{j,k}}{2\rho_N} \right) - \left(\frac{\tilde{Y}_{j,k}(t)}{2\bar{C}_H} - \frac{G_{j,k}}{2\rho_N} \right)^2 \right| \end{aligned}$$

$$\begin{aligned}
& + \sup_{t \in [0,1]} \left| \left(\frac{\tilde{Y}_{j,k}(t)}{2\bar{C}_H} - \frac{G_{j,k}}{2\rho_N} \right)^2 - \left(\frac{Y_{j,k}(t)}{2\bar{C}_H} - \frac{G_{j,k}}{2\rho_N} \right)^2 \right| \\
& \stackrel{(21)-(40)}{\leq} 2^{-2(\ell_2+1)} + \frac{\sup_{t \in [0,1]} |\tilde{Y}_{j,k}(t) - Y_{j,k}(t)|}{2\bar{C}_H} \\
& \quad \times \left(\frac{\sup_{t \in [0,1]} |\tilde{Y}_{j,k}(t)| + \sup_{t \in [0,1]} |Y_{j,k}(t)|}{2\bar{C}_H} + \frac{|G_{j,k}|}{\rho_N} \right) \\
& \stackrel{(27)}{\leq} 2^{-2(\ell_2+1)} + \frac{2\varepsilon}{\bar{C}_H}.
\end{aligned}$$

So replacing in (41) and similarly for the second term, it entails

$$\sup_{(j,k) \in \mathcal{I}_N} \sup_{t \in [0,1]} \left| Y_{j,k}(t) G_{j,k} - \tilde{k}_{\ell_2}^{\bar{C}_H, \rho_N} \left(\tilde{Y}_{j,k}(t), G_{j,k} \right) \right| \leq 2^{-(2\ell_2+1)} \bar{C}_H \rho_N + 4\varepsilon \rho_N.$$

For the final ReLU NN approximation of (26), define \tilde{B}_N^H as

$$\tilde{B}_N^H(t) := \sum_{(j,k) \in \mathcal{I}_N} 2^{-jH} \tilde{k}_{\ell_2}^{\bar{C}_H, \rho_N} (\tilde{g}(2^j t - k) - \tilde{g}(-k), G_{j,k}). \quad (42)$$

Combining (26)-(12)-(42) gives (still on the event $\{|G_{j,k}| \leq \rho_N : (j,k) \in \mathcal{I}_N\}$)

$$\begin{aligned}
\sup_{t \in [0,1]} |B^H(t) - \tilde{B}_N^H(t)| & \leq \sup_{t \in [0,1]} |B^H(t) - B_N^H(t)| + \sup_{t \in [0,1]} |B_N^H(t) - \tilde{B}_N^H(t)| \\
& \leq C_{(12)} N^{-H} (1 + \log(N))^{1/2} + N \left(2^{-(2\ell_2+1)} \bar{C}_H \rho_N + 4\varepsilon \rho_N \right)
\end{aligned}$$

recalling that $\text{Card}(\mathcal{I}_N) \leq N$. It suffices to ensure that the second term at the right-hand side is bounded by $2N^{-H} (1 + \log(N))^{1/2}$ thanks to the choice

$$\varepsilon = \frac{(1 + \log(N))^{1/2}}{4\rho_N N^{H+1}} \wedge 1, \quad (43)$$

$$\ell_2 = \left\lceil \frac{1}{2\log(2)} \log \left(\frac{\bar{C}_H \rho_N N^{H+1}}{(1 + \log(N))^{1/2}} \right) \right\rceil \vee 1 = \mathcal{O}_c(\log(\varepsilon^{-1})). \quad (44)$$

3.2.3 Architecture

The total architecture of \tilde{g} is composed by M sub-networks, where each $\tilde{g}_{i,q}$ is built as a cascade of q NN with $(\ell_1 + 1)$ hidden layers, i.e. $(q - 1)$ NN from Proposition 6 and 1 more from the product with ϕ_i . Therefore, \tilde{g} requires at most a depth $\mathcal{O}_c(\ell_1)$ and a complexity $\mathcal{O}_c(M\ell_1)$, with constants clearly depending on r . Using M in (39) and ℓ_1 in (37), we get the architecture bounds as a function of the accuracy ε , for just one approximation of Ψ_H with \tilde{g} .

As mentioned above, $\tilde{Y}_{j,k}$ is composed of $2\tilde{g}$ NN and so it has the same depth but twice the complexity (number of neurons and parameters) of \tilde{g} . Additionally, (42) requires $(\ell_2 + 1)$ hidden layers to perform the multiplications with the Gaussian variables. Finally, all these operations are computed for N different scaling/transition parameters (j, k) . All in all, the total architecture of \tilde{B}_N^H (see Figure 6) is composed of at most an order

1. $\mathcal{O}_c(\ell_1 + \ell_2)$ hidden layers,
2. $\mathcal{O}_c\left(N\left(M\ell_1 + \ell_2\right)\right)$ neurons and parameters.

Replacing with (36), (39) and (44) gives the architecture bounds with respect to ε , i.e.




Figure 6: Neural network architecture of $\tilde{B}_N^H(t, G_{1:N})$.

1. hidden layers:

$$\mathcal{O}_c(\ell_1 + \ell_2) = \mathcal{O}_c\left(\log\left(\varepsilon^{-(1-\frac{1}{r})}\right) + \log(\varepsilon^{-1})\right) = \mathcal{O}_c(\log(\varepsilon^{-1})),$$

where we have observed that the exponent inside the log term can be put in the \mathcal{O}_c since the constants are allowed to depend on r in our notation;

2. neurons and parameters:

$$\begin{aligned} \mathcal{O}_c\left(N(M\ell_1 + \ell_2)\right) &= \mathcal{O}_c\left(N\left(\varepsilon^{-\frac{1}{2r}} \log\left(\varepsilon^{-(1-\frac{1}{r})}\right) + \log(\varepsilon^{-1})\right)\right) \\ &= \mathcal{O}_c\left(N\varepsilon^{-\frac{1}{2r}} \log(\varepsilon^{-1})\right), \end{aligned}$$

with equilibrium $r = m + 1$ in (39). Remembering the choice (43) of ε w.r.t. N gives the announced result.

A Complements

A.1 Proof of Lemma 1

Let us bound the truncated approximation of (5) for all $n \in \mathbb{N}$,


$$\begin{aligned} &\sup_{t \in [0,1]} \left| B(t) - \left(G_1 t + \sum_{j=0}^n \sum_{k=0}^{2^j-1} 2^{-(j+1)} \psi_{j,k}^{\text{FS}}(t) G_{j,k} \right) \right| \\ &= \sup_{t \in [0,1]} \left| \sum_{j=n+1}^{\infty} \sum_{k=0}^{2^j-1} 2^{-(j+1)} \psi_{j,k}^{\text{FS}}(t) G_{j,k} \right| \\ &\leq \sum_{j=n+1}^{\infty} 2^{-(j/2+1)} \sup_{0 \leq k \leq 2^j-1} |G_{j,k}| \\ &\leq C \sum_{j=n+1}^{\infty} 2^{-(j/2+1)} (\log(j + 2^j + 1))^{1/2} \quad \text{a.s.} \\ &\leq C \sum_{j=n+1}^{\infty} 2^{-\frac{1}{2}(j+1)} (1 + j)^{1/2} \\ &\leq C 2^{-\frac{1}{2}(n+1)} (n + 1)^{1/2} \end{aligned}$$

$$\leq C N^{-1/2} (1 + \log(N))^{1/2},$$


where C is a non-negative random variable which may change from line to line. In the third line, use the fact that the wavelets have disjoint support in k and so for fixed j , any t belongs to the support of at most one $\psi_{j,k}^{\text{FS}}$, with $\|\psi_{j,k}^{\text{FS}}\|_{\infty} \leq 2^{j/2}$. In the fourth, invoke [1, Lemma 2]; in the fifth the inequality holds for j large enough; in the sixth use a classical integral test, and lastly replace with N . \square

A.2 Wavelet representation


Integrals are computed numerically using the function `quad` from the package `scipy` in `Python`.




(a) $u \mapsto \nu(u)$



(b) $\xi \mapsto |\hat{\psi}(\xi)|$



(c) $x \mapsto \psi(x)$



(d) $x \mapsto \Psi_H(x)$ with $H = 0.1$

Figure 7: Lemarié-Meyer wavelet constructed with (13) for $\gamma = 2$

Acknowledgements

This action benefited from the support of the Chair Stress Test, Risk Management and Financial Steering, led by the French Ecole polytechnique and its foundation and sponsored by BNP Paribas. This work has been partially supported by MIAI @ Grenoble Alpes, (ANR-19-P3IA-0003).

References

- [1] A. Ayache and M. S. Taqqu. Rate optimality of wavelet series approximations of fractional Brownian motion. *J. Fourier Anal. Appl.*, 9(5):451–471, 2003.

- [2] M. Bennedsen. A rough multi-factor model of electricity spot prices. *Energy Economics*, 63:301–313, 2017.
- [3] C. Burdet and P. Rougier. Analysis of center-of-pressure data during unipedal and bipedal standing using fractional Brownian motion modeling. *Journal of Applied Biomechanics*, 23(1):63–69, 2007.
- [4] J.-F. Coeurjolly. Simulation and identification of the fractional Brownian motion: A bibliographical and comparative study. *Journal of Statistical Software*, 5:1–53, 2000.
- [5] J.-F. Coeurjolly and E. Porcu. Fast and exact simulation of complex-valued stationary Gaussian processes through embedding circulant matrix. *J. Comput. Graph. Statist.*, 27(2):278–290, 2018.
- [6] S. Cohen and J. Istas. *Fractional fields and applications*, volume 73 of *Mathématiques & Applications*. Springer, Heidelberg, 2013.
- [7] F. Comte and E. Renault. Long memory in continuous-time stochastic volatility models. *Mathematical Finance*, 8(4):291–323, 1998.
- [8] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Math. Control Signals Systems*, 2(4):303–314, 1989.
- [9] I. Daubechies. *Ten lectures on wavelets*, volume 61 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1992.
- [10] L. de Haan and A. Ferreira. *Extreme value theory*. Springer Series in Operations Research and Financial Engineering. Springer, New York, 2006.
- [11] T. Dieker. Simulation of fractional Brownian motion. Master’s thesis, Department of Mathematical Sciences, University of Twente, <http://www.columbia.edu/~ad3217/fbm/thesis.pdf>, 2004.
- [12] D. E. Dominici. The inverse of the cumulative standard normal probability function. *Integral Transforms Spec. Funct.*, 14(4):281–292, 2003.
- [13] C. Esteban, S. L. Hyland, and G. Räscher. Real-valued (medical) time series generation with recurrent conditional gans. *arXiv preprint arXiv:1706.02633*, 2017.
- [14] G. Faber. Über stetige Funktionen. *Math. Ann.*, 66(1):81–94, 1908.
- [15] F. Flandoli and M. Gubinelli. Random currents and probabilistic models of vortex filaments. In *Seminar on Stochastic Analysis, Random Fields and Applications IV*, volume 58 of *Progr. Probab.*, pages 129–139. Birkhäuser, Basel, 2004.
- [16] J. Gatheral, T. Jaisson, and M. Rosenbaum. Volatility is rough. *Quant. Finance*, 18(6):933–949, 2018.
- [17] E. Gobet, J. G. López-Salas, and C. Vázquez. Quasi-regression Monte-Carlo scheme for semi-linear PDEs and BSDEs with large scale parallelization on GPUs. *Arch. Comput. Methods Eng.*, 27(3):889–921, 2020.
- [18] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, 2016.
- [19] 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.
- [20] A. Haar. Zur Theorie der orthogonalen Funktionensysteme. *Math. Ann.*, 69(3):331–371, 1910.

- [21] M. Haas and S. Richter. Statistical analysis of Wasserstein GANs with applications to time series forecasting. *arXiv preprint arXiv:2011.03074*, 2020.
- [22] W. P. Johnson. The curious history of Faà di Bruno’s formula. *Amer. Math. Monthly*, 109(3):217–234, 2002.
- [23] Y. Kozachenko, A. Pashko, and O. Vasylyk. Simulation of generalized fractional Brownian motion in $C([0, T])$. *Monte Carlo Methods Appl.*, 24(3):179–192, 2018.
- [24] T. Kühn and W. Linde. Optimal series representation of fractional Brownian sheets. *Bernoulli*, 8(5):669–696, 2002.
- [25] M. Ledoux and M. Talagrand. *Probability in Banach spaces*, volume 23 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1991.
- [26] P. G. Lemarié and Y. Meyer. Ondelettes et bases Hilbertiennes. *Rev. Mat. Iberoamericana*, 2(1-2):1–18, 1986.
- [27] V. E. Maiorov. On best approximation by ridge functions. *J. Approx. Theory*, 99(1):68–94, 1999.
- [28] S. Mallat. *A wavelet tour of signal processing*. Elsevier/Academic Press, Amsterdam, third edition, 2009.
- [29] B. B. Mandelbrot and J. W. Van Ness. Fractional Brownian motions, fractional noises and applications. *SIAM Rev.*, 10:422–437, 1968.
- [30] Y. Meyer. Principe d’incertitude, bases Hilbertiennes et algèbres d’opérateurs. In *Astérisque*, number 145-146, pages 209–223. Soc. Math. France, 1987. Séminaire Bourbaki, Vol. 1985/86.
- [31] Y. Meyer, F. Sellan, and M. S. Taqqu. Wavelets, generalized white noise and fractional integration: the synthesis of fractional Brownian motion. *J. Fourier Anal. Appl.*, 5(5):465–494, 1999.
- [32] I. Nourdin. *Selected aspects of fractional Brownian motion*, volume 4 of *Bocconi & Springer Series*. Springer ; Bocconi University Press, Milan, 2012.
- [33] A. Pinkus. Approximation theory of the MLP model in neural networks. In *Acta numerica, 1999*, volume 8 of *Acta Numer.*, pages 143–195. Cambridge Univ. Press, Cambridge, 1999.
- [34] A. Pontieri, F. Lepreti, L. Sorriso-Valvo, A. Vecchio, and V. Carbone. A simple model for the solar cycle. *Solar Physics*, 213(1):195–201, 2003.
- [35] J. Schauder. Eine Eigenschaft des Haarschen Orthogonalsystems. *Math. Z.*, 28(1):317–320, 1928.
- [36] J. M. Steele. *Stochastic calculus and financial applications*, volume 45 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, 2001.
- [37] M. Telgarsky. Representation benefits of deep feedforward networks. *arXiv preprint arXiv:1509.08101*, 2015.
- [38] M. Wiese, R. Knobloch, R. Korn, and P. Kretschmer. Quant GANs: deep generation of financial time series. *Quant. Finance*, 20(9):1419–1440, 2020.
- [39] D. Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103–114, 2017.