

This is a repository copy of *Support Collapse of Deep Gaussian Processes with Polynomial Kernels for a Wide Regime of Hyperparameters*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/224572/>

Preprint:

Chernobrovkina, Daryna and Grunewalder, Steffen (2025) Support Collapse of Deep Gaussian Processes with Polynomial Kernels for a Wide Regime of Hyperparameters. [Preprint]

<https://doi.org/10.48550/arXiv.2503.12266>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Support Collapse of Deep Gaussian Processes with Polynomial Kernels for a Wide Regime of Hyperparameters

Daryna Chernobrovkina and Steffen Grünewälder
University of York
Department of Mathematics
York, UK

Abstract

We analyze the prior that a Deep Gaussian Process with polynomial kernels induces. We observe that, even for relatively small depths, averaging effects occur within such a Deep Gaussian Process and that the prior can be analyzed and approximated effectively by means of the Berry-Esseen Theorem. One of the key findings of this analysis is that, in the absence of careful hyper-parameter tuning, the prior of a Deep Gaussian Process either collapses rapidly towards zero as the depth increases or places negligible mass on low norm functions. This aligns well with experimental findings and mirrors known results for convolution based Deep Gaussian Processes.

1 Introduction

Deep Gaussian processes (DGPs) have been introduced by [1] as a natural extension of Gaussian processes (GPs) that has been inspired by deep neural networks. Like deep neural networks, DGPs have multiple layers and each layer corresponds to an individual GP. It has recently been noted by [2] that traditional GPs attain for certain compositional regression problems a strictly slower rate of convergence than the minimax optimal rate. This is demonstrated in [2] by showing that for a class of generalized additive models any GP will be suboptimal, independently of the kernel function that is used. Generalized additive models can be regarded as a simple form of a compositional model with two layers. In contrast, [3] have shown that DGPs can attain for such problems the minimax optimal rate of convergence (up to logarithmic factors) when the DGPs are *carefully tuned*. In fact, they show that DGPs are able to attain optimal rates of convergence for many compositional problems. Along similar lines, [4] show that for nonlinear inverse problems DGPs can attain a rate of convergence that is polynomially faster than the rate that GPs with Matérn kernel functions can attain when the unknown parameter has a compositional structure. One well known downside of DGPs is the difficulty of sampling from the posterior distribution. [5] approach this problem by providing a particularly *simple prior which facilitates*

posterior calculations while guaranteeing adaptivity in the context of regression to both the smoothness of the unknown regression function and the compositional structure.

In this paper, we focus on the prior that a DGP places on a function space. We work in the context of polynomial kernels and we study the behavior of the priors as the depth of the DGP increases. We show that the prior is very sensitive to the hyperparameters that are used for the individual GPs and that small deviations of the ‘correct regime’ of hyperparameters would either lead to an extremely tight concentration at zero or would result in prior measures that place negligible mass on functions of low norm. In earlier work, [6] have observed in experiments that the prior of a DGP with Gaussian kernels shows pathological behavior. They also analyzed the derivative of a DGP to get insight into this pathological behavior, but did not provide an analysis of the behavior of the prior itself. [7] provide a deeper analysis by phrasing a DGP as a Markov chain and studying its ergodic behavior. In particular, [7, Thm 4] states that the output of a DGP becomes constant (in a form of point-wise convergence) as the depth increases when a Gaussian kernel is used and a condition on the parameters of the kernel is satisfied. They also study a DGP where instead of a composition of GPs a convolution of GPs is used. This form of a DGP differs from the DGPs that are commonly used in the literature [1, 6, 2, 3, 5], but has the advantage that it is amenable to a convolution and Fourier theory based argument. This allows the authors to get deep insights into this type of DGPs. They find that for a convolutional DGP, Fourier coefficients associated with the DGP converge either to zero or diverge (almost surely). Furthermore, the eigenvalues of a covariance operator associated with the DGP control if the coefficients converge to zero or diverge [7, Thm 16].

One of the key research challenges in the area of DGPs is to gain deeper insight into the behavior of standard DGPs. Such insight is crucial to make sense of the ‘contradictory’ observations in the literature: on the one hand, DGPs are often used successfully in practice [1] and DGPs outperform GPs in a variety of statistical tasks in terms of rate of convergence [3] while, on the other hand, there is the pathological behavior of DGPs that has been observed in experiments and in convolutional DGPs [6, 7]. This research challenge is also far from trivial since the convolutional structure allows for significant simplifications in the analysis of [7], and it is unclear how to get tight control of the behavior of a DGP in its absence.

In the context of polynomial kernels, we develop an alternative approach that does not rely on convolutions and applies to standard DGPs. Our approach makes use of the fact that for polynomial kernel functions the sample paths of GPs lie within the reproducing kernel Hilbert space associated to that kernel function. Combining this fact with a Karhunen-Loève type decomposition of the GPs allows us to write the composition of GPs as a product of normally distributed vectors. We study these products then with the help of the Berry-Esseen Theorem. It is worth highlighting that earlier works focused on Law of Large Numbers and Ergodic type results which provide neither rates of convergence nor finite sample bounds. In contrast to that, the Berry-Esseen approach that we develop provides both rates of convergence and finite sample bounds.

Our main result is Theorem 1, which provides a bound on the approximation of a DGP $g_\ell \circ \dots \circ g_1(x)$, where g_1 has covariance $k_1(x, y) = (xy + c)^{d_1}$, $c \geq 0$, and the g_i ’s have covariance $k_i(x, y) = \sigma_i^2(xy)^{d_i}$, where $\sigma_i > 0$ and the d_i ’s are non-zero integers. For such

a process we find a normally distributed random variable Y and a random sign S such that

$$\sup_{x,t \in \mathbb{R}} |\Pr(g_\ell \circ \dots \circ g_1(x) \leq t) - \Pr(Se^Y(g_1(x))^{c_1} \leq t)| \leq 0.56 \left(\sum_{i=2}^{\ell} \sigma_{i,\log}^2 \right)^{-3/2} \sum_{i=2}^{\ell} \rho_{i,\log},$$

where $\sigma_{\log}^2, \rho_{\log}$ is the variance and absolute third moment of certain log-normal random variables. The constant c_1 is equal to $d_2 + \dots + d_\ell$ and will generally be very large. It is worth highlighting that we have here an approximation of a DGP that consists of a product of a single GP, a random sign, and a log-normal random variable.

Another important result that can be derived from the theorem is that when $\sigma_2 = \dots = \sigma_\ell =: \sigma$, then the median of the DGP converges rapidly to zero in ℓ if $\sigma < \exp((\gamma + \log 2)/2)$, where γ is the Euler-Mascheroni constant, and diverges when $\sigma > \exp((\gamma + \log 2)/2)$. This is the same threshold that was found by [7] in the context of convolutional DGPs.

The remainder of the paper is organized as follows: in Section 1.1 we provide key definitions and results that we use throughout. In Section 2, we start with the simple case of products of Gaussian random variables; the motivation for this is that the main averaging effects that are at play are very transparent in this simple setting. Section 3 is our main section. We start with the simple case of a product of GPs with linear kernels before approaching the case of polynomial kernels. In Section 4 we provide then a discussion of the results and we put these in perspective. In particular, we highlight challenges that need to be overcome to extend our results beyond the polynomial kernel case. There are also two appendices with technical results. In Appendix A we provide a variety of closed form expressions for moments of log-normal random variables that we use throughout, and Appendix B contains a variety of auxiliary results for DGPs that we use.

1.1 Preliminaries

A zero mean GP g on \mathbb{R} is a stochastic process which is fully specified by its covariance function $k(x, y), x, y \in \mathbb{R}$. The covariance function k is positive semi-definite. The function k is called the covariance function since $\text{Var}(g(x)) = k(x, x)$ and $\text{Cov}(g(x)g(y)) = k(x, y)$ for all $x, y \in \mathbb{R}$. In the context of kernel methods, one also calls k the kernel function and we use the two terms interchangeably. To each covariance function there corresponds a reproducing kernel Hilbert space (RKHS) \mathcal{H}_k . In case that \mathcal{H}_k is finite dimensional it is known that the GP g attains values in \mathcal{H}_k . In other words, the sample paths are RKHS functions when \mathcal{H}_k is finite dimensional. If \mathcal{H}_k is infinite dimensional then the sample paths lie almost surely not in \mathcal{H}_k . Often it is convenient to work with a so called feature map $\phi : \mathbb{R} \rightarrow \mathcal{H}_k$ which satisfies $\langle \phi(x), \phi(y) \rangle = k(x, y)$, where the inner product is here the inner product of \mathcal{H}_k . In the finite dimensional case, we also write $\phi(x)^\top \phi(y) = k(x, y)$. A DGP of depth ℓ on \mathbb{R} is a composition of ℓ zero mean GPs $g_\ell \circ \dots \circ g_1$ with corresponding covariance functions k_ℓ, \dots, k_1 .

We make frequent use of the Central Limit Theorem (CLT) and different versions of the Berry-Esseen Theorem. In particular, we use the following two versions of the Berry-Esseen Theorem, which guarantee uniform convergence of certain normalized sums to a Gaussian limit: (1) The first version that we use applies to zero mean i.i.d. random variables

X_1, \dots, X_n with variance $\text{Var}(X_1) = \sigma^2$ and absolute third moment $\rho = E(|X_1|^3)$. Let $S_n = X_1 + \dots + X_n$ then this version of the Berry-Esseen Theorem states that

$$\sup_{x \in \mathbb{R}} |\Pr(n^{-1/2} \sigma^{-1} S_n \leq x) - \Phi(x)| \leq \frac{0.336(\rho + 0.415\sigma^3)}{\sigma^3 n^{1/2}},$$

where Φ denotes the cumulative distribution (CDF) function of a standard normal random variable.

(2) The second version avoids the need of identically distributed random variables at the cost of a slightly more conflated theorem statement. Consider again independent zero mean random variables X_1, \dots, X_n but now with individual variances $\text{Var}(X_i) = \sigma_i^2$ and absolute third moments $\rho_i = E(|X_i|^3)$, $i \leq n$. The second version of the Berry-Esseen Theorem states that

$$\sup_{x \in \mathbb{R}} |\Pr\left(\frac{S_n}{\sqrt{\sigma_1^2 + \dots + \sigma_n^2}} \leq x\right) - \Phi(x)| \leq 0.56 \left(\sum_{i=1}^n \sigma_i^2\right)^{-3/2} \sum_{i=1}^n \rho_i.$$

Note that the $n^{-1/2}$ factors that appear in the first version are subsumed in the variance terms; i.e. when $\sigma_1 = \dots = \sigma_n = \sigma$ then $\sqrt{\sigma_1^2 + \dots + \sigma_n^2} = \sqrt{n}\sigma$ and when additionally $\rho_1 = \dots = \rho_n = \rho$ then $(\sum_{i=1}^n \sigma_i^2)^{-3/2} \sum_{i=1}^n \rho_i = \rho / \sqrt{n}\sigma^3$.

Besides the definition of a DGP, all of the above results are classical and can be found in textbooks such as [8].

2 Products of Gaussian Random Variables

We start by analyzing the products of Gaussian random variables before approaching DGPs in the following section. We will see that such products are closely related to compositions of GPs with linear kernels. Let X_1, \dots, X_ℓ be i.i.d. standard random variables with variance $\sigma^2 > 0$ and consider their product $\prod_{i=1}^\ell X_i$. Figure 1 plots the density of the product in dependence of ℓ . Notice that the left plot uses $\sigma = 1$ and that the density rapidly concentrates around zero in this case, as ℓ increases. The right plot considers larger values of σ (the values are 2, 2.5 and 3) and we can notice the opposite effect: the probability for the absolute value of the product to attain values below 1/2 rapidly falls as ℓ increases. We will observe this effect repeatedly in other contexts.

We will now aim to characterize the distribution of the product as ℓ increases. In order to do that, we apply the CLT to the product. We write the product as

$$\prod_{i=1}^\ell X_i = \left(\prod_{i=1}^\ell S_i\right) \left(\prod_{i=1}^\ell |X_i|\right),$$

where S_i is the sign of X_i ,

$$S_i = \begin{cases} 1 & \text{if } X_i \geq 0, \\ -1 & \text{otherwise.} \end{cases}$$

Note that S_i is independent of $|X_i|$ (Appendix A.2) and that $\prod_{i=1}^\ell S_i$ attains values 1 and -1 each with probability 1/2. If we take the logarithm of $\prod_{i=1}^\ell |X_i|$ then the CLT is applicable




Figure 1: (a) The densities of the product of $\ell = 1, 10, 30$ normally distributed random variables with mean $\mu = 0$ and variance $\sigma^2 = 1$ are shown. (b) The probability of the product attaining values around zero for larger σ is shown.

if the variance of $\log |X_i|$ is finite. The variance of $\log |X_i|$ is, in fact, finite (See (12) in the Appendix) and the CLT can be applied. Under the assumption X_1, \dots, X_ℓ are i.i.d. we can infer that

$$\ell^{-1/2} \sum_{i=1}^{\ell} (\log |X_i| - E(\log |X_i|)) \xrightarrow{d} N(0, \text{Var}(\log |X_1|)),$$

where d denotes convergence in distribution. In particular, for large ℓ the sum $\ell^{-1/2} \sum_{i=1}^{\ell} \log |X_i|$ has approximately the distribution $N(\sqrt{\ell} E(\log |X_i|), \text{Var}(\log |X_i|))$. Furthermore, the continuous mapping theorem [9, Thm 2.3] can be applied since the exponential function is continuous, and it follows that a normalized version of the product converges in distribution,

$$\left(\prod_{i=1}^{\ell} \frac{|X_i|}{\exp(E(\log |X_i|))} \right)^{1/\sqrt{\ell}} = \exp\left(\frac{1}{\sqrt{\ell}} \sum_{i=1}^{\ell} (\log |X_i| - E(\log |X_i|))\right) \xrightarrow{d} e^Z,$$

where Z is normally distributed with mean zero and variance $\text{Var}(\log |X_1|)$. For large enough ℓ we then have the approximation,

$$\prod_{i=1}^{\ell} |X_i|^{1/\sqrt{\ell}} \approx e^{Z + \sqrt{\ell} E(\log |X_1|)} \quad (\text{in distribution}).$$

In other words, $\prod_{i=1}^{\ell} |X_i|^{1/\sqrt{\ell}}$ is approximately log normally distributed with mean parameter $\sqrt{\ell} E(\log |X_i|)$ and variance parameter $\text{Var}(\log |X_i|)$. Figure 2 shows a comparison of this approximation and the corresponding distribution of the scaled product (gained by sampling).

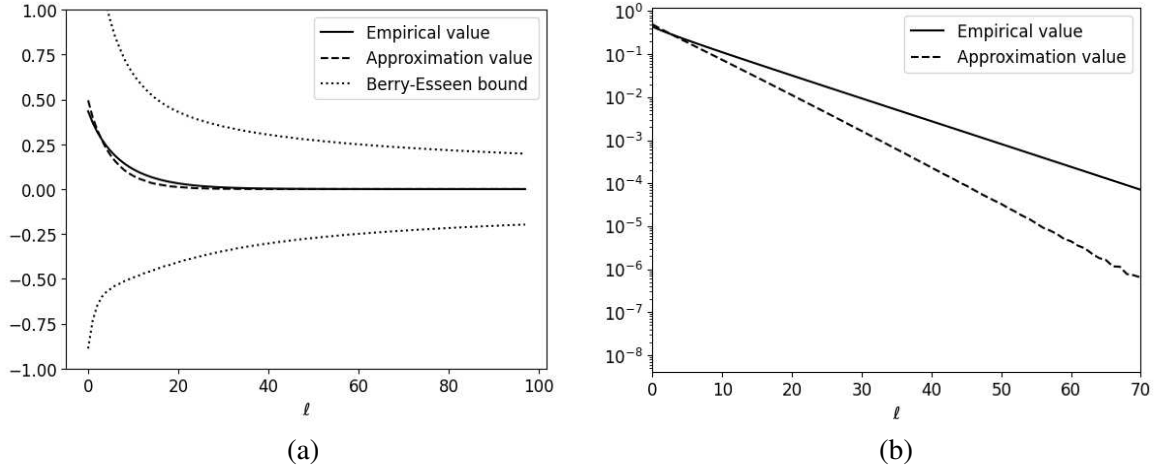


Figure 2: (a) The probability of the scaled product and the log-normal approximation to attain values above $1/2$ are compared ($\sigma = 1$). The plot is complemented by an error bound. (b) The same quantities are compared but on a logarithmic scale (the error bound is omitted).

One might wonder if the normalizing factor $\ell^{-1/2}$ can be incorporated into the variance of the X_i so that we can say something about the product $\prod_{i=1}^{\ell} \tilde{X}_i$ of suitably normalized Gaussian random variables \tilde{X}_i . This does not work, however, since $\text{Var}(\log |X_i|) = \pi^2/8$ independently of the variance parameter $\sigma > 0$. This is a consequence of $\text{Var}(\log |aX_i|) = \text{Var}(\log |a| + \log |X_i|) = \text{Var}(\log |X_i|)$, which holds for any $a \in \mathbb{R}$.

2.1 An Application of the Berry-Esseen Theorem

Ideally, we want to be able to infer properties of $\prod_{i=1}^{\ell} X_i$ or of $\prod_{i=1}^{\ell} X_i^{\alpha}$, $\alpha \in \mathbb{N}$. This will ultimately be useful for understanding how products of Gaussian processes with polynomial kernels behave. When following the earlier approach, we are led to expressions of the form $\sum_{i=1}^{\ell} \log |X_i|$ and $\alpha \sum_{i=1}^{\ell} \log |X_i|$. The leading α term in the latter expression is of minor importance. However, the lack of the normalizing factor $\ell^{-1/2}$ in front of the sums is a significant problem, since we cannot apply the CLT directly. This problem can be understood in terms of point-wise convergence. The CLT tells us that for any $x \in \mathbb{R}$, $\lim_{\ell \rightarrow \infty} \Pr(\ell^{-1/2} \sum_{i=1}^{\ell} (\log |X_i| - E(\log |X_i|)) \leq x) = \Phi(x/\sigma_{\log})$, where Φ denotes the CDF of a standard normal random variable and $\sigma_{\log}^2 = \text{Var}(\log |X_1|) = \pi^2/8$. To control the difference between the CDF of the unnormalized sum and Φ we can try

$$\Pr\left(\frac{1}{\sigma_{\log}} \sum_{i=1}^{\ell} (\log |X_i| - E(\log |X_i|)) \leq x\right) = \Pr\left(\frac{\ell^{-1/2}}{\sigma_{\log}} \sum_{i=1}^{\ell} (\log |X_i| - E(\log |X_i|)) \leq \ell^{-1/2} x\right)$$

and hope that the latter expression gets close to $\Phi(\ell^{-1/2} x)$. However, the CLT does not allow us to infer this convergence since the location $\ell^{-1/2} x$ changes with ℓ .

One way to address this nettle is to move from point-wise convergence to uniform convergence. This can be achieved by using the Berry-Esseen Theorem instead of the CLT.

The Berry-Esseen Theorem guarantees that

$$\sup_{x \in \mathbb{R}} \left| \Pr \left(\ell^{-1/2} \sigma_{\log}^{-1} \sum_{i=1}^{\ell} (\log |X_i| - E(\log |X_i|)) \leq x \right) - \Phi(x) \right| \leq \frac{0.336(\rho_{\log}^{\sigma} + 0.415\sigma_{\log}^3)}{\sqrt{\ell}\sigma_{\log}^3},$$

where we assume that our Gaussian variables X_1, \dots, X_{ℓ} are i.i.d., centered, and have variance σ_{\log}^2 and where we use the definition $\rho_{\log}^{\sigma} = E(|\log^3 |X_i||)$. We provide a closed-form expression of ρ_{\log}^{σ} in Appendix A, (13), as well as an easier to interpret bound (14). It is also easy to get very accurate approximations of ρ_{\log}^{σ} through sampling. We can now infer that, uniformly in $x \in \mathbb{R}$,

$$\begin{aligned} & \left| \Pr \left(\sigma_{\log}^{-1} \sum_{i=1}^{\ell} (\log |X_i| - E(\log |X_i|)) \leq x \right) - \Phi(\ell^{-1/2}x) \right| \\ &= \left| \Pr \left(\ell^{-1/2} \sigma_{\log}^{-1} \sum_{i=1}^{\ell} (\log |X_i| - E(\log |X_i|)) \leq \ell^{-1/2}x \right) - \Phi(\ell^{-1/2}x) \right| \\ &\leq \sup_{y \in \mathbb{R}} \left| \Pr \left(\ell^{-1/2} \sigma_{\log}^{-1} \sum_{i=1}^{\ell} (\log |X_i| - E(\log |X_i|)) \leq y \right) - \Phi(y) \right| \\ &\leq \frac{0.336(\rho_{\log}^{\sigma} + 0.415\sigma_{\log}^3)}{\sqrt{\ell}\sigma_{\log}^3}. \end{aligned}$$

We can rewrite this further to get an approximation of the law of $\sum_{i=1}^{\ell} \log |X_i|$,

$$\begin{aligned} \Pr \left(\sum_{i=1}^{\ell} \log |X_i| \leq \sigma_{\log}x + \ell E(\log |X_1|) \right) &= \Pr \left(\sigma_{\log}^{-1} \sum_{i=1}^{\ell} (\log |X_i| - E(\log |X_i|)) \leq x \right) \\ &\approx \Phi(\ell^{-1/2}x). \end{aligned}$$

In other words, with $y = \sigma_{\log}x + \ell E(\log |X_1|)$,

$$\Pr \left(\sum_{i=1}^{\ell} \log |X_i| \leq y \right) \approx \Phi(\ell^{-1/2} \sigma_{\log}^{-1}(y - \ell E(\log |X_1|))). \quad (1)$$

If we let $Z \sim N(\ell E(\log |X_1|), \ell \sigma_{\log}^2)$ then (1) implies

$$\sup_{x \in \mathbb{R}} \left| \Pr \left(\sum_{i=1}^{\ell} \log |X_i| \leq x \right) - \Pr(Z \leq x) \right| \leq \frac{0.336(\rho_{\log}^{\sigma} + 0.415\sigma_{\log}^3)}{\sqrt{\ell}\sigma_{\log}^3}.$$

Since $\Pr(Z \leq x) = \Pr(e^Z \leq e^x)$, and similarly for $\sum_{i=1}^{\ell} \log |X_i|$, we find that

$$\begin{aligned} \sup_{x \in \mathbb{R}} \left| \Pr \left(\prod_{i=1}^{\ell} |X_i| \leq x \right) - \Pr(e^Z \leq x) \right| &= \sup_{x \in \mathbb{R}} \left| \Pr \left(\prod_{i=1}^{\ell} |X_i| \leq e^x \right) - \Pr(e^Z \leq e^x) \right| \\ &\leq \frac{0.336(\rho_{\log}^{\sigma} + 0.415\sigma_{\log}^3)}{\sqrt{\ell}\sigma_{\log}^3}. \end{aligned} \quad (2)$$

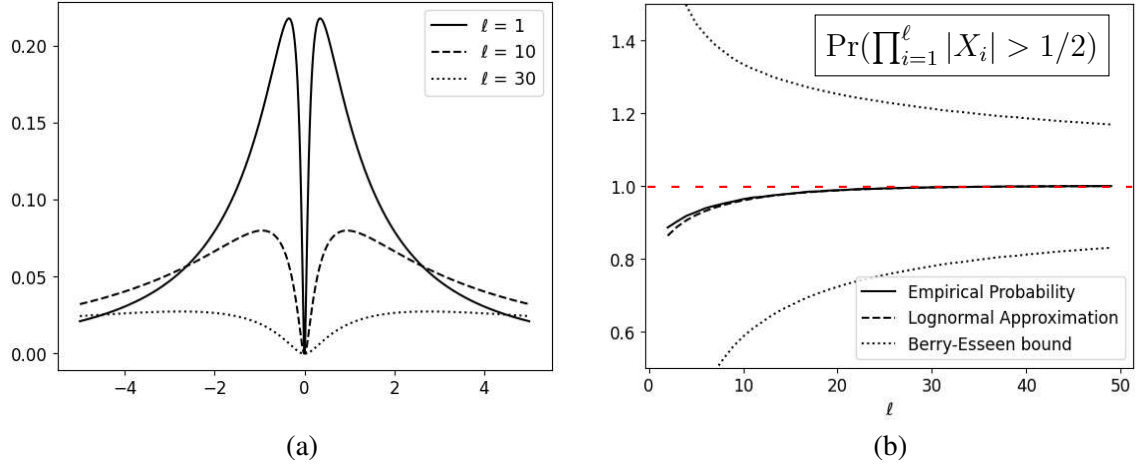


Figure 3: (a) The distribution of the product of $\ell = 1, 10$ and 30 log-normal random variables with $\sigma = 3$ is shown. (b) The probability for the product and the log-normal approximation to attain values above $1/2$ is shown ($\sigma = 3$).

The median of the log-normal random variable e^Z is approximately $\exp(\ell(\log(\sigma) - 0.63))$. In particular, when $\sigma < e^{0.63} \approx 1.87$, the median approaches exponentially fast 0, while when $\sigma > e^{0.63}$, the median diverges to infinity at an exponential rate in ℓ . We demonstrate this divergence effect in Figure 3 for $\prod_{i=1}^{\ell} X_i$ and the approximation $S_Z e^Z$, where S_Z is a random variable that is independent of Z , and which attains values $+1$ and -1 with probability $1/2$ each.

2.2 Convergence & Approximation for Powers of X

We aim to generalize the above approach to products of the form $\prod_{i=1}^{\ell} X_i^{\alpha}$. When α is even then this product will always be positive and will be equal to $\prod_{i=1}^{\ell} |X_i|^{\alpha}$. When α is odd then $\prod_{i=1}^{\ell} X_i = S_{\alpha} \prod_{i=1}^{\ell} |X_i|$, where S_{α} attains values $+1$ and -1 with equal probability. We will apply again the the Berry-Esseen Theorem to approximate the distribution of $\prod_{i=1}^{\ell} |X_i|^{\alpha}$. To this end, note that

$$\begin{aligned} & \sup_{x \in \mathbb{R}} \left| \Pr \left(\ell^{-1/2} (\alpha \sigma_{\log})^{-1} \sum_{i=1}^{\ell} (\alpha \log |X_i| - \alpha E(\log |X_i|)) \leq x \right) - \Phi(x) \right| \\ & \leq \frac{0.336(\alpha^3 \rho_{\log}^{\sigma} + 0.415(\alpha \sigma_{\log})^3)}{(\alpha \sigma_{\log})^3 \sqrt{\ell}}, \end{aligned}$$

where the α 's on the right side cancel. As in the previous section, after introducing the random variable $Z_\alpha \sim N(\ell\alpha E(\log |X_1|), \ell(\alpha\sigma_{\log})^2)$, we can infer that

$$\begin{aligned} \sup_{x \in \mathbb{R}} |\Pr\left(\prod_{i=1}^{\ell} |X_i|^\alpha \leq x\right) - \Pr(e^{Z_\alpha} \leq x)| &= \sup_{x \in \mathbb{R}} |\Pr\left(\prod_{i=1}^{\ell} |X_i|^\alpha \leq e^x\right) - \Pr(e^{Z_\alpha} \leq e^x)| \\ &\leq \frac{0.336(\rho_{\log}^\sigma + 0.415\sigma_{\log}^3)}{\sigma_{\log}^3 \sqrt{\ell}}. \end{aligned}$$

Note that α plays a similar role as ℓ and the distribution of Z_α for, say, $\alpha = 5$ and $\ell = 10$ has the same mean as the random variable Z_1 with $\ell = 50$. This implies that larger values of α lead to a more rapid collapse of the support of e^{Z_α} . For example, the median of the variable e^{Z_α} is $\ell\alpha(\log(\sigma) - 0.63)$. As the factor α increases, we have an even faster convergence of the median to 0 for $\sigma < 1.87$ and divergence to infinity when $\sigma > 1.88$.




Figure 4: (a) Five draws from a DGP with $\ell = 30$ layers, a linear kernel, and $\sigma = 1$ is shown. Notice the scale of the y -axis. (b) As for (a) but with $\sigma = 2.5$.

3 Limit Distributions of Deep Gaussian Processes

The results in the last section on the distribution of $\prod_{i=1}^{\ell} X_i$, where X_i are normally distributed, have natural links to the distribution of deep GPs. To see this, let us start with the simple case where the kernel functions of the different layers are all $k(x, y) = \sigma^2 xy$ and $x, y \in \mathbb{R}$. Now, consider ℓ independent Gaussian processes g_1, \dots, g_ℓ each of which has zero mean and kernel function k . The process $g_\ell \circ \dots \circ g_1$ is a Deep GP on \mathbb{R} . The kernel $k(x, y)$ can also be written as an inner product of a feature map ϕ . In particular, if we use $\phi(x) = \sigma x$ then $k(x, y) = \phi(x)\phi(y)$. Each of the individual GPs attains values in the RKHS \mathcal{H}_k that corresponds to k . Note that this RKHS can be written as $\{h : \mathbb{R} \rightarrow \mathbb{R} : h(x) = \alpha\phi(y)\phi(x), \alpha, y \in \mathbb{R}\} = \{h : \mathbb{R} \rightarrow \mathbb{R} : h(x) = \alpha\phi(1)\phi(x), y \in \mathbb{R}\}$ since ϕ is linear in this context. Since $\phi(1) = \sigma$ we know that each path drawn from

the GP will be of the form $\alpha\sigma\phi(x)$, where the slope α changes depending on our draw from the GP. Let us use the notation $\alpha_\omega^{(i)}$ to denote the slope corresponding to the draw from GP g_i for experiment $\omega \in \Omega$, i.e. $\alpha_\omega^{(i)}$ is the random variable that corresponds to the slope of the paths drawn from g_i . Let us now take a look at g_1 and $\alpha_\omega^{(1)}$. Fixing some $x \in \mathbb{R}$, we know that $g_1(x)$ is a zero mean Gaussian random variable with variance $k(x, x) = \sigma^2 x^2$. We also know that $g_1(x) = \alpha_\omega^{(1)} \sigma \phi(x) = \alpha_\omega^{(1)} \sigma^2 x$ and $\alpha_\omega^{(1)} = g_1(x) / \sigma^2 x$. In other words, $\alpha_\omega^{(1)}$ is a Gaussian random variable with zero mean and variance $1/\sigma^2$. Hence, with $U_1 = \sigma^2 \alpha_\omega^{(1)} \sim N(0, \sigma^2)$ it follows that $g_1(x) = \alpha_\omega^{(1)} \sigma \phi(x) = U_1 x$. By induction we can generalize this to

$$g_\ell \circ \dots \circ g_1(x) = \prod_{i=1}^{\ell} U_i x,$$

where U_1, \dots, U_ℓ are i.i.d. with distribution $N(0, \sigma^2)$. The independence of U_1, \dots, U_ℓ follows right away from the independence of g_1, \dots, g_ℓ since each U_i is a function of g_i .

3.1 Approximation of DGPs with Linear Kernels

From the previous section, we know that $g_\ell \circ \dots \circ g_1(x)$ can be written as $xS \prod_{i=1}^{\ell} |U_i|$, where S is independent of the U_i 's and attains values $+1$ and -1 with equal probability. The statement about S follows by the same argument as in Section 2. We can now approximate the distribution of the product by means of the Berry-Esseen Theorem. In particular, for large ℓ the product $\prod_{i=1}^{\ell} |U_i|$ will be close in distribution to e^Z , where $Z \sim N(\ell E(\log |U_1|), \ell \sigma_{\log}^2)$, with $\sigma_{\log}^2 = \text{Var}(\log |U_1|)$ and $\rho_{\log}^\sigma = E(|\log |U_1||^3)$. In particular, by following the same argument as in Section 2, and by incorporating the random signs and x , we find that

$$\begin{aligned} \sup_{x, c \in \mathbb{R}} \left| \Pr(g_\ell \circ \dots \circ g_1(x) \leq c) - \Pr(Sxe^Z \leq c) \right| &= \sup_{x, c \in \mathbb{R}} \left| \Pr\left(Sx \prod_{i=1}^{\ell} |U_i| \leq c\right) - \Pr(Sxe^Z \leq c) \right| \\ &\leq \frac{0.336(\rho_{\log}^\sigma + 0.415\sigma_{\log}^3)}{\sqrt{\ell}\sigma_{\log}^3}. \end{aligned}$$

It is worth highlighting that this bound holds uniformly over all values x . In fact, in the linear case, this follows right away since the Berry-Esseen bound is uniform in c and we can use a simple substitution from c/x to c to infer that the bound also holds uniformly in x .

3.2 Approximation of DGPs with Polynomial Kernels

We will now extend the above results to DGPs of the form $g_\ell \circ \dots \circ g_1$, where the GP g_1 has a polynomial kernel of order d_1 , that is $k_1(x, y) = (xy + c)^{d_1}$, where $d_1 > 0$ is some integer and $c \geq 0$, and the successive GPs have kernels $k_i(x, y) = \sigma_i^2(xy)^{d_i}$, $d_i \geq 1, \sigma_i > 0$

and $i \geq 2$. The GP g_1 can be written in the following way (see Appendix B.1 on p. 25),

$$g_1(x) = \sum_{i=1}^{d+1} Z_i \phi_i(x) = \sum_{i=0}^d \binom{d}{i}^{1/2} Z_{i+1} x^{d-i} c^{i/2},$$

where $(Z_1, \dots, Z_{d+1})^\top \sim N(0, I)$ and $\phi_i(x) = \binom{d}{i-1}^{1/2} x^{d-i+1} c^{(i-1)/2}$, $i \leq d+1$. Similarly, there are independent random variables $Y_i \sim N(0, \sigma_i^2)$, which are also independent of Z_1, \dots, Z_{d+1} , and such that $g_i(x) = Y_i x^{d_i}$, for all $2 \leq i \leq \ell$. We can therefore write the DGP as

$$g_\ell \circ \dots \circ g_1(x) = Y_\ell (Y_{\ell-1})^{d_1^\downarrow} (Y_{\ell-2})^{d_2^\downarrow} \times \dots \times (Y_2)^{d_{\ell-2}^\downarrow} \left(\sum_{i=1}^{d+1} Z_i \phi_i(x) \right)^{d_{\ell-1}^\downarrow},$$

where we use the notation $d_i^\downarrow = \sum_{j=0}^{i-1} d_{\ell-j}$, for $i = 1, \dots, \ell-1$. Taking the logarithm of the product of the absolute values of the Y -terms gives us $d_{\ell-2}^\downarrow \log |Y_2| + \dots + d_1^\downarrow \log |Y_{\ell-1}| + \log |Y_\ell| = \sum_{j=2}^\ell c_j \log |Y_j|$, where $c_j = d_{\ell-j}^\downarrow$ for $j = 2, \dots, \ell-1$, and $c_\ell = 1$. We are now in a position to apply the Berry-Esseen Theorem for *non-identically distributed* random variables. To set this up, let $\sigma_{i,\log}^2 = c_i^2 \text{Var}(\log |Y_i|)$ and $\rho_{i,\log} = c_i^3 E(|\log |Y_i||^3)$ for $2 \leq i \leq \ell$. Then

$$\begin{aligned} & \sup_{x \in \mathbb{R}} |\Pr((\sigma_{2,\log}^2 + \dots + \sigma_{\ell,\log}^2)^{-1/2} \sum_{j=2}^\ell (c_j \log |Y_j| - c_j E(\log |Y_j|)) \leq x) - \Phi(x)| \\ & \leq 0.56 \left(\sum_{i=1}^n \sigma_{i,\log}^2 \right)^{-3/2} \sum_{i=1}^n \rho_{i,\log}. \end{aligned} \quad (3)$$

As earlier, we can translate this statement into a statement about $\sum_{j=2}^\ell c_j \log |Y_j|$ by means of substitution. For a given x , let $y = (\sigma_{2,\log}^2 + \dots + \sigma_{\ell,\log}^2)^{1/2} x + \sum_{j=2}^\ell c_j E(\log |Y_j|)$, then

$$|\Pr\left(\sum_{j=2}^\ell c_j \log |Y_j| \leq y\right) - \Phi\left((\sigma_{2,\log}^2 + \dots + \sigma_{\ell,\log}^2)^{-1/2} (y - \sum_{j=2}^\ell c_j E(\log |Y_j|))\right)|$$

is also upper bounded by the right side of (3). In fact, this bound holds uniformly over all $y \in \mathbb{R}$. Let us introduce a random variable Y that is independent of $Z_1, \dots, Z_{d+1}, S_2, \dots, S_\ell$ and which has the law $N(\sum_{i=2}^\ell c_i E(\log |Y_i|), \sum_{i=2}^\ell \sigma_{i,\log}^2)$, then

$$\sup_{y \in \mathbb{R}} |\Pr\left(|Y_\ell| \prod_{i=2}^{\ell-1} |Y_i|^{c_i} \leq y\right) - \Pr(e^Y \leq y)| = \sup_{y \in \mathbb{R}} |\Pr\left(\sum_{j=2}^\ell c_j \log |Y_j| \leq y\right) - \Pr(Y \leq y)|$$

and the latter term is again upper bounded by the right side of (3). We can translate this into a statement about the product of the Y_i 's by observing that $Y_\ell \prod_{i=2}^{\ell-1} Y_i^{c_i} = S_\ell \prod_{i \in I} S_i |Y_\ell| \prod_{i=2}^{\ell-1} |Y_i|^{c_i}$, where $I \subset \{2, \dots, \ell\}$ are the indices which correspond to odd c_i values and S_i is the sign of Y_i for all $i \leq \ell$. Since the different S_i 's are independent of each other and independent of the $|Y_i|$'s it follows that $S = S_\ell \prod_{i \in I} S_i$ is a random variable that is independent of the

$|Y_i|$'s (in fact, S is also independent of Z_1, \dots, Z_{d+1}) and it attains values $+1$ and -1 with equal probability. There is one final technical hurdle in the way to an approximation of the DGP. We need to multiply both the product and the approximation by the random variable $(g_1(x))^{c_1}$ (strictly speaking, we can have two different probability spaces and the new space needs to contain a copy of $(g_1(x))^{c_1}$). In any case, we can relate the two distributions that include $(g_1(x))^{c_1}$ by means of a conditional expectation argument, which we provide in Appendix B.3. From here we get directly to an approximation of the DGP. We summarize this statement in the following theorem.

Theorem 1. *Given a DGP $g_\ell \circ \dots \circ g_1$ on \mathbb{R} with ℓ -layers and corresponding independent GPs g_1, \dots, g_ℓ with covariance functions $k_1(x, y) = (xy + c)^{d_1}$, $c \geq 0$, and $k_i(x, y) = \sigma_i^2(xy)^{d_i}$ where $\sigma_i > 0$, $2 \leq i \leq \ell$, and $d_1, \dots, d_\ell \geq 1$ are integers. There exist independent Y_2, \dots, Y_ℓ , such that each Y_i has distribution $N(0, \sigma_i^2)$ and $g_i(x) = Y_i x^{d_i}$. For $2 \leq i \leq \ell$, let $\sigma_{i,\log}^2 = c_i^2 \text{Var}(\log |Y_i|)$ and $\rho_{i,\log} = c_i^3 E(|\log |Y_i||^3)$. We have the following approximation of the DGP:*

$$\sup_{x, t \in \mathbb{R}} |\Pr(g_\ell \circ \dots \circ g_1(x) \leq t) - \Pr(Se^Y (g_1(x))^{c_1} \leq t)| \leq 0.56 \left(\sum_{i=2}^{\ell} \sigma_{i,\log}^2 \right)^{-3/2} \sum_{i=2}^{\ell} \rho_{i,\log},$$

where

$$Y \sim N\left(\sum_{i=2}^{\ell} c_i E(\log |Y_i|), \sum_{i=2}^{\ell} c_i^2 \text{Var}(\log |Y_i|)\right),$$

and $c_\ell = 1$, $c_i = \sum_{j=i+1}^{\ell} d_j$, for $1 \leq i \leq \ell - 1$. The random sign S attains values $+1$ and -1 with equal probability. Furthermore, g_1 , S and Y are independent and we can write

$$g_1(x) = \sum_{i=0}^{d_1} \binom{d_1}{i}^{1/2} Z_{i+1} x^{d_1-i} c^{i/2},$$

with Z_1, \dots, Z_n independent standard normal random variables that are independent of S and Y .

Example ($d_2 = \dots = d_\ell = 2$): It is instructive to analyze the distribution of $S(g_1(x))^{c_1} e^Y$ and the Berry-Esseen bound in a concrete setting. Assume that $d_2 = \dots = d_\ell = 2$ and $\sigma_2 = \dots = \sigma_\ell = \sigma$, for some σ that we will vary, and let $\ell \geq 2$. We show in Appendix B.2 (Eq. (15)) that the *Berry Esseen bound* takes in this setting the form

$$0.56 \left(\sum_{i=2}^{\ell} \sigma_{i,\log}^2 \right)^{-3/2} \sum_{i=2}^{\ell} \rho_{i,\log} \leq 3\ell^{-1/2} \frac{E(|\log |Y_1||^3)}{(\text{Var}(\log |Y_1|))^{3/2}}.$$

Note that the bound improves with the familiar $\ell^{-1/2}$ rate. In terms of the dependence on σ , recall that $\text{Var}(\log |Y_1|)$ is independent of σ and the bound in Appendix A.1 on p. 22 shows that the dependence of $E(|\log |Y_1||^3)$ on σ is at most logarithmical.

In terms of the distribution of Y , first note that for $2 \leq i \leq \ell - 1$ the coefficients become $c_i = 2(\ell - i - 1)$ and the *mean* of Y becomes (App. B.2, Eq. (16)),

$$E\left(\sum_{i=2}^{\ell} c_i E(\log |Y_i|)\right) = (\ell(\ell - 1) - 1)E(\log |Y_1|) \approx ((\ell - 1)^2 + \ell)(\log(\sigma) - 0.63).$$

Similarly, the *variance* becomes (App B.2, Eq. (17))

$$\sum_{i=2}^{\ell} c_i^2 \text{Var}(\log |Y_i|) = \left(\frac{2\ell(\ell - 1)(2\ell - 1)}{3} - 3\right) \text{Var}(\log |Y_1|) = \frac{\pi^2}{8} \left(\frac{2\ell(\ell - 1)(2\ell - 1)}{3} - 3\right).$$

In terms of the log-normal random variable e^Y we have a similar effect as in the earlier settings: the *median* of e^Y is $\exp((1/2)((\ell - 1)^2 + \ell)(\log(\sigma) - (\gamma + \log 2)/2))$ which approaches rapidly zero when $\sigma < \exp((\gamma + \log 2)/2) \approx 1.88$ and, otherwise, diverges to infinity as ℓ increases. In terms of the approximation $Se^Y(g_1(x))^{c_1}$ of the DGP, note that $c_1 = 2(\ell - 1)$ and

$$(g_1(x))^{c_1} = \left(\sum_{i=0}^{d_1} \binom{d_1}{i}^{1/2} Z_{i+1} x^{d_1-i} c^{i/2}\right)^{2(\ell-1)}.$$

Hence, we have two terms of large order that interact multiplicatively. If $\sigma < 1.88$ and the sum $\sum_{i=0}^{d_1} \binom{d_1}{i}^{1/2} Z_{i+1} x^{d_1-i} c^{i/2}$ is strictly smaller than one then we expect the DGP to attain tiny values. Similarly, when $\sigma > 1.88$ and the sum attains value above one then we expect the DGP to attain huge values. The most interesting case is the case where the two multiplicative terms compete: for simplicity consider the median value of e^Y , let $\sigma = 1$, and assume that the sum attains a value of $a > 1$, then the product of these two terms is approximately of the order $e^{-\ell^2} \times a^{2\ell}$ and the term e^Y will dominate when ℓ growth. The probability for attaining large values depends obviously on x ; the larger $|x|$ the higher the probability to observe large values a and the longer it will take for $e^{-\ell^2}$ to dominate.

4 Discussion

In this paper, we demonstrated that, for a range of polynomial kernels, a DGP will either concentrate around zero or it will generate functions that have increasingly large norms. The regime where these two pathologies do not occur becomes increasingly smaller as the size of the DGP increases. One can also note that kernels with higher order polynomials will likely give rise to DGPs that show this pathological behavior earlier (we discuss this phenomenon in the context of products of normally distributed random variables in Section 2.2). Another important observation is that these effects arise due to averaging effects across the layers. One can break these averaging effects in the context of polynomial kernels by choosing the degree of some kernels significantly higher than of others. Then only a few GPs will dominate. This, however, will be similar to a DGP with fewer layers, which defeats the purpose of DGPs.

There is a wide range of open questions, that are important to address. One of the most important questions is how to generalize our approach to a larger class of kernel functions.

A major hurdle that needs to be overcome to be able to generalize the results is that GPs with infinite dimensional RKHSs do not attain values within their RKHS. In all likelihood, the fact that the RKHS lies dense in the space of functions in which the GP attains values will be crucial for extensions. Another important question is why DGPs show near optimal asymptotic rates of convergence in various problems given their pathological behavior. Is this simply the Bernstein-von-Mises Theorem at work, or is there a deeper reason?

A Closed-Form Solutions & Bounds

A.1 Closed form expressions for $E(\log |X|)$, $E(\log^2 |X|)$, $\text{Var}(\log |X|)$ and $E(|\log^3 |X||)$

It is known that the values $E(\log |X|)$ and $E(\log^2 |X|)$ can be derived analytically when X is normally distributed. With some more work one can also derive a closed-form expression for $E(|\log^3 |X||)$ whenever the variance σ^2 of X satisfies $\sigma^2 > 1/2$. We complement this last result with the simple bound $E(|\log^3 |X||) \leq (E(\log^4 |X|))^{3/4}$ for the case where $\sigma^2 \leq 1/2$. For the reader's convenience, we present the various derivations in detail below.

Applications of the Γ function and the incomplete Γ functions. The first step is to derive closed-form expressions for $\int_0^\infty e^{-x^2} \log |x| dx$, $\int_0^\infty e^{-x^2} \log^2 |x| dx$ and $\int_0^\infty e^{-x^2} |\log^3 |x|| dx$. An easy argument to derive these uses the Γ function

$$\Gamma(s) = \int_0^\infty x^{s-1} e^{-x} dx, \quad (4)$$

as well as the upper incomplete Γ function

$$\Gamma(s, x) = \int_x^\infty t^{s-1} e^{-t} dt,$$

and the lower incomplete Γ function

$$\gamma(s, x) = \int_0^x t^{s-1} e^{-t} dt,$$

where for us it suffices to consider real valued $s > 0$, and where $x > 0$.

Taking derivatives on both sides of (4), where we can exchange the derivative and the integral since the integral is defined for all $s > 0$, the derivative with respect to s exists in a neighborhood around $s = 1/2$, and there exists an integrable function that bounds the absolute value of the derivative from above ([10], Th.123D). Moreover, using the di-gamma function ψ , which has the property that $\psi(s) = \Gamma'(s)/\Gamma(s)$, for all $s > 0$, gives

$$\Gamma(s)\psi(s) = \int_0^\infty x^{s-1} e^{-x} \log(x) dx.$$

In particular, for $s = 1/2$,

$$\Gamma(1/2)\psi(1/2) = \int_0^\infty x^{-1/2} e^{-x} \log(x) dx = 2 \int_0^\infty e^{-x^2} \log(x^2) dx = 4 \int_0^\infty e^{-x^2} \log(x) dx$$

and

$$\int_0^\infty e^{-x^2} \log(x) dx = \frac{\Gamma(1/2)\psi(1/2)}{4} = -\frac{\sqrt{\pi}(\gamma + 2 \log 2)}{4} = \frac{\Gamma'(1/2)}{4}. \quad (5)$$

Similarly, we can approach the integral for $\log^2 |x|$. Analogously to the calculation above, we can exchange the derivative and the integral by ([10], Th.123D). Taking the second derivative of the Γ function, yields

$$\frac{d}{ds}\Gamma(s)\psi(s) = \int_0^\infty x^{s-1}e^{-x} \log^2(x) dx.$$

On the left we have

$$\frac{d}{ds}\Gamma(s)\psi(s) = \Gamma'(s)\psi(s) + \Gamma(s)\psi'(s) = \Gamma(s)\psi^2(s) + \Gamma(s)\psi'(s).$$

The derivative ψ' can be found under the name polygamma function. Evaluating the integral at $s = 1/2$ yields

$$2 \int_0^\infty (1/2)x^{-1/2}e^{-x} \log^2(x) dx = 2 \int_0^\infty e^{-x^2} \log^2(x^2) dx = 8 \int_0^\infty e^{-x^2} \log^2(x) dx$$

and

$$\int_0^\infty e^{-x^2} \log^2(x) dx = \frac{\Gamma(1/2)\psi^2(1/2) + \Gamma(1/2)\psi'(1/2)}{8} = \frac{\sqrt{\pi}((2 \log 2 + \gamma)^2 + 3\zeta(2))}{8},$$

where ζ is the Riemman-Zeta function. It is known that $\zeta(2) = \pi^2/6$. For later use we note that

$$\Gamma''(1/2) = \sqrt{\pi}((2 \log 2 + \gamma)^2 + 3\zeta(2)). \quad (6)$$

Finding an analytic solution for $\rho = E(|\log^3 |X||)$ when $\sigma^2 > 1/2$ requires the use of the incomplete Γ functions. With the same arguments as for the Γ function, one can show that:

$$\begin{aligned} \frac{d}{ds}\Gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} &= 4 \int_{1/\sqrt{2}\sigma}^\infty \log(x)e^{-x^2} dx, \\ \frac{d}{ds}\gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} &= 4 \int_0^{1/\sqrt{2}\sigma} \log(x)e^{-x^2} dx, \\ \frac{d^2}{ds^2}\Gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} &= 8 \int_{1/\sqrt{2}\sigma}^\infty \log^2(x)e^{-x^2} dx, \\ \frac{d^2}{ds^2}\gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} &= 8 \int_0^{1/\sqrt{2}\sigma} \log^2(x)e^{-x^2} dx, \\ \frac{d^3}{ds^3}\Gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} &= 16 \int_{1/\sqrt{2}\sigma}^\infty \log^3(x)e^{-x^2} dx, \\ \frac{d^3}{ds^3}\gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} &= 16 \int_0^{1/\sqrt{2}\sigma} \log^3(x)e^{-x^2} dx. \end{aligned}$$

The challenge is to find closed-form expressions for the derivatives of the incomplete Γ functions. Because, $\Gamma(s) = \Gamma(s, 1/\sqrt{2}\sigma) + \gamma(s, 1/\sqrt{2}\sigma)$ it suffices to find the derivatives

of $\Gamma(s)$ and $\Gamma(s, 1/\sqrt{2}\sigma)$. We calculated the first and second derivatives of $\Gamma(s)$ already and the third and the fourth derivatives are also not hard to derive:

$$\begin{aligned} \frac{d^3}{ds^3}\Gamma(s)\Big|_{s=1/2} &= \frac{d}{ds}(\Gamma(s)\psi^2(s) + \Gamma(s)\psi'(s))\Big|_{s=1/2} = \Gamma(s)(\psi^3(s) + 3\psi(s)\psi'(s) + \psi''(s))\Big|_{s=1/2} \\ &= -\sqrt{\pi}((2\log 2 + \gamma)^3 + 9(2\log 2 + \gamma)\zeta(2) + 14\zeta(3)). \end{aligned} \quad (7)$$

$$\begin{aligned} \frac{d^4}{ds^4}\Gamma(s)\Big|_{s=1/2} &= \frac{d}{ds}(\Gamma(s)(\psi^3(s) + 3\psi(s)\psi'(s) + \psi''(s)))\Big|_{s=1/2} \\ &= \Gamma(s)(\psi^4(s) + 6\psi'(s)\psi^2(s) + \psi'''(s) + 4\psi''(s)\psi(s) + 3(\psi'(s))^2)\Big|_{s=1/2} \\ &= \sqrt{\pi}((2\log 2 + \gamma)^4 + 18(2\log 2 + \gamma)^2\zeta(2) \\ &\quad + 56(2\log 2 + \gamma)\zeta(3) + 27\zeta^2(2) + 90\zeta(4)). \end{aligned} \quad (8)$$

The derivatives of $\Gamma(s, 1/\sqrt{2}\sigma)$ are more challenging to derive and we are using results from [11, p.156/157] which apply when $\sigma^2 > 1/2$. The first derivative is

$$\frac{d}{ds}\Gamma(s, 1/\sqrt{2}\sigma)\Big|_{s=1/2} = -\log(\sqrt{2}\sigma)\Gamma(1/2, 1/\sqrt{2}\sigma) + \frac{1}{\sqrt{2}\sigma}T(3, 1/2, 1/\sqrt{2}\sigma), \quad (9)$$

where the function T is defined in [11]. Further below, we derive the particular values of T that we need. The second derivative is

$$\begin{aligned} \frac{d^2}{ds^2}\Gamma\left(s, \frac{1}{\sqrt{2}\sigma}\right)\Big|_{s=1/2} &= \log^2(\sqrt{2}\sigma)\Gamma\left(\frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right) \\ &\quad + \frac{2}{\sqrt{2}\sigma}\left(-\log(\sqrt{2}\sigma)T\left(3, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right) + T\left(4, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right)\right), \end{aligned} \quad (10)$$

and the third derivative is

$$\begin{aligned} \frac{d^3}{ds^3}\Gamma\left(s, \frac{1}{\sqrt{2}\sigma}\right)\Big|_{s=1/2} &= -\log^3(\sqrt{2}\sigma)\Gamma\left(s, \frac{1}{\sqrt{2}\sigma}\right) + \frac{3}{\sqrt{2}\sigma}\left(\log^2(\sqrt{2}\sigma)T\left(3, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right) \right. \\ &\quad \left. - 2\log(\sqrt{2}\sigma)T\left(4, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right) + 2T\left(5, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right)\right). \end{aligned} \quad (11)$$

Expansions of $T(n, 1/2, 1/\sqrt{2}\sigma)$. We need the values $T(3, 1/2, 1/\sqrt{2}\sigma)$, $T(5, 1/2, 1/\sqrt{2}\sigma)$ as well as $T(5, 1/2, 1/\sqrt{2}\sigma)$. Under the assumption that $\sigma^2 > 1/2$, we can follow the arguments in [11, (31), p.156; (36), p.157] and [12, p.144]. For any $n \in \mathbb{N}$,

$$T\left(n, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right) = \frac{1}{2\pi i} \oint_C \left(\frac{-1}{s+1}\right)^{n-1} \Gamma(-1/2 - s)(\sqrt{2}\sigma)^{-s} ds,$$

where i denotes the imaginary unit and C is a contour given in [12, (4), p.144] which surrounds the poles at $s = -1$ and at $s = 1/2 + k$, for all $k \in \mathbb{N}$. The pole at $s = -1$ is of order $n - 1$ and the residue for $n = 3$ is

$$\begin{aligned}
c_3 &:= \frac{1}{2} \lim_{z \rightarrow -1} \frac{d^2}{dz^2} \frac{(z+1)^3}{(z+1)^3} \Gamma(-1/2 - z) (\sqrt{2}\sigma)^{-z} \\
&= \frac{\sigma}{\sqrt{2}} \left(\frac{d^2}{dz^2} \Gamma\left(-\frac{1}{2} - z\right) \Big|_{z=-1} - 2 \log(\sqrt{2}\sigma) \frac{d}{dz} \Gamma\left(-\frac{1}{2} - z\right) \Big|_{z=-1} + \Gamma\left(\frac{1}{2}\right) \log^2(\sqrt{2}\sigma) \right) \\
&= \frac{\sigma}{\sqrt{2}} \left(\Gamma\left(\frac{1}{2}\right) \psi^2\left(\frac{1}{2}\right) + \Gamma\left(\frac{1}{2}\right) \psi'\left(\frac{1}{2}\right) - 2 \log(\sqrt{2}\sigma) \Gamma\left(\frac{1}{2}\right) \psi\left(\frac{1}{2}\right) + \Gamma\left(\frac{1}{2}\right) \log^2(\sqrt{2}\sigma) \right) \\
&= \sigma \sqrt{\frac{\pi}{2}} \left((2 \log 2 + \gamma + \log(\sqrt{2}\sigma))^2 + 3\zeta(2) \right) \approx \sigma \sqrt{\frac{\pi}{2}} ((2.31 + \log \sigma)^2 + \pi^2/2).
\end{aligned}$$

For $n = 4$ the residue is

$$\begin{aligned}
c_4 &:= -\frac{1}{6} \frac{d^3}{dz^3} \Gamma(-1/2 - z) (\sqrt{2}\sigma)^{-z} \Big|_{z=-1} \\
&= -\frac{\sqrt{2}\sigma}{6} \left(\frac{d^3}{dz^3} \Gamma(-1/2 - z) \Big|_{z=-1} - 3 \log(\sqrt{2}\sigma) \frac{d^2}{dz^2} \Gamma(-1/2 - z) \Big|_{z=-1} \right. \\
&\quad \left. + 3 \log^2(\sqrt{2}\sigma) \frac{d}{dz} \Gamma(-1/2 - z) \Big|_{z=-1} - \log^3(\sqrt{2}\sigma) \Gamma(1/2) \right) \\
&= -\frac{\sqrt{2}\sigma}{6} \left(\Gamma\left(\frac{1}{2}\right) \left(\psi^3\left(\frac{1}{2}\right) + 3\psi\left(\frac{1}{2}\right) \psi'\left(\frac{1}{2}\right) + \psi''\left(\frac{1}{2}\right) \right) \right. \\
&\quad \left. - 3 \log(\sqrt{2}\sigma) \left(\Gamma\left(\frac{1}{2}\right) \psi^2\left(\frac{1}{2}\right) + \Gamma\left(\frac{1}{2}\right) \psi'\left(\frac{1}{2}\right) \right) \right. \\
&\quad \left. + 3 \log^2(\sqrt{2}\sigma) \Gamma\left(\frac{1}{2}\right) \psi\left(\frac{1}{2}\right) - \log^3(\sqrt{2}\sigma) \Gamma\left(\frac{1}{2}\right) \right) \\
&= \frac{\sigma \sqrt{2\pi}}{6} \left((2 \log 2 + \gamma + \log(\sqrt{2}\sigma))^3 + 9\zeta(2)(2 \log 2 + \gamma + \log(\sqrt{2}\sigma)) + 14\zeta(3) \right) \\
&\approx \frac{\sigma \sqrt{2\pi}}{6} ((2.31 + \log \sigma)^3 + \frac{3\pi^2}{2}(2.31 + \log \sigma) + 16.828).
\end{aligned}$$

Finally, for $n = 5$, the residue is

$$\begin{aligned}
c_5 &:= \frac{1}{2} \frac{d^4}{dz^4} \Gamma(-1/2 - z) (\sqrt{2}\sigma)^{-z} \Big|_{z=-1} \\
&= \frac{\sqrt{2}\sigma}{24} \left(\frac{d^4}{dz^4} \Gamma(-1/2 - z) \Big|_{z=-1} - 4 \log(\sqrt{2}\sigma) \frac{d^3}{dz^3} \Gamma(-1/2 - z) \Big|_{z=-1} \right. \\
&\quad + 6 \log^2(\sqrt{2}\sigma) \frac{d^2}{dz^2} \Gamma(-1/2 - z) \Big|_{z=-1} - 4 \log^3(\sqrt{2}\sigma) \frac{d}{dz} \Gamma(-1/2 - z) \Big|_{z=-1} \\
&\quad \left. + \log^4(\sqrt{2}\sigma) \Gamma(1/2) \right) \\
&= \frac{\sigma \sqrt{2\pi}}{24} \left((2 \log 2 + \gamma + \log(\sqrt{2}\sigma))^4 + 18\zeta(2)(2 \log 2 + \gamma + \log(\sqrt{2}\sigma))^2 \right. \\
&\quad \left. + 56\zeta(3)(2 \log 2 + \gamma + \log(\sqrt{2}\sigma)) + 27\zeta^2(2) + 90\zeta(4) \right) \\
&\approx \frac{\sigma \sqrt{2\pi}}{24} ((2.31 + \log \sigma)^4 + 3\pi^2(2.31 + \log \sigma)^2 + 67.312 \log \sigma + 155.49 + \frac{7\pi^4}{4})
\end{aligned}$$

The residues at $z = 1/2 + k$, $k \in \mathbb{N}$, are

$$\left(\frac{-1}{3/2 + k} \right)^{n-1} \frac{(-1)^k}{k!} (\sqrt{2}\sigma)^{-(1/2+k)}.$$

Hence,

$$\begin{aligned}
T\left(3, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right) &= c_3 + \sum_{k=0}^{\infty} \frac{(-1)^k}{k!(3/2 + k)^2 (\sqrt{2}\sigma)^{k+1/2}}, \\
T\left(4, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right) &= c_4 - \sum_{k=0}^{\infty} \frac{(-1)^k}{k!(3/2 + k)^3 (\sqrt{2}\sigma)^{k+1/2}}, \\
T\left(5, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right) &= c_5 + \sum_{k=0}^{\infty} \frac{(-1)^k}{k!(3/2 + k)^4 (\sqrt{2}\sigma)^{k+1/2}}.
\end{aligned}$$

A good approximation of these values can be attained by summing up to 10. For this case we get

$$\begin{aligned}
T\left(3, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right) &\approx c_3 + 0.374\sigma^{-0.5} - 0.095\sigma^{-1.5} + 0.017\sigma^{-2.5} - 0.0024\sigma^{-3.5}, \\
T\left(4, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right) &\approx c_4 - 0.374\sigma^{-0.5} + 0.095\sigma^{-1.5} - 0.017\sigma^{-2.5} + 0.0024\sigma^{-3.5}, \\
T\left(5, \frac{1}{2}, \frac{1}{\sqrt{2}\sigma}\right) &\approx c_5 + 0.374\sigma^{-0.5} - 0.095\sigma^{-1.5} + 0.017\sigma^{-2.5} - 0.0024\sigma^{-3.5}.
\end{aligned}$$

Derivation of $E(\log |X|)$, $E(\log^2 |X|)$, $\text{Var}(\log |X|)$ and $E(|\log^3 |X||)$. Now, integration by substitution allows us to derive the expected value, the variance, and the absolute

third moment. In the following, assume that X is a zero mean Gaussian random variable with variance $\sigma^2 > 0$. Then

$$\begin{aligned}
E(\log |X|) &= \sqrt{2/\pi\sigma^2} \int_0^\infty \log(x) e^{-x^2/2\sigma^2} dx \\
&= \frac{2}{\sqrt{\pi}} \int_0^\infty \log(x) e^{-x^2} dx + \frac{2}{\sqrt{\pi}} \log(\sqrt{2}\sigma) \int_0^\infty e^{-x^2} dx \\
&= -\frac{\gamma + 2\log 2}{2} + \log \sigma + \frac{1}{2} \log 2 \\
&= \log \sigma - \frac{\gamma + \log 2}{2} \approx \log \sigma - 0.63.
\end{aligned}$$

Similarly, we can derive $E(\log^2 |X|)$. In detail,

$$\begin{aligned}
E(\log^2 |X|) &= \sqrt{2/\pi\sigma^2} \int_0^\infty \log^2(x) e^{-x^2/2\sigma^2} dx \\
&= (2/\sqrt{\pi}) \int_0^\infty (\log(\sqrt{2}\sigma) + \log(x))^2 e^{-x^2} dx \\
&= (2/\sqrt{\pi}) \log^2(\sqrt{2}\sigma) \int_0^\infty e^{-x^2} dx + (4/\sqrt{\pi}) \log(\sqrt{2}\sigma) \int_0^\infty \log(x) e^{-x^2} dx \\
&\quad + (2/\sqrt{\pi}) \int_0^\infty \log^2(x) e^{-x^2} dx \\
&= \log^2(\sqrt{2}\sigma) - \log(\sqrt{2}\sigma)(\gamma + 2\log 2) + \frac{((2\log 2 + \gamma)^2 + \pi^2/2)}{4}.
\end{aligned}$$

Combining these we find the variance of $\log |X|$ to be

$$\begin{aligned}
\text{Var}(\log |X|) &= E(\log^2 |X|) - E(\log |X|)^2 \\
&= \log^2(\sqrt{2}\sigma) - \log(\sqrt{2}\sigma)(\gamma + 2\log 2) + \frac{((2\log 2 + \gamma)^2 + \pi^2/2)}{4} \\
&\quad - \frac{(2\log \sigma - \gamma - \log 2)^2}{4} \\
&= \pi^2/8,
\end{aligned} \tag{12}$$

which is independent of $\sigma > 0$.

To derive the third absolute moment of $\log |X|$, we expand a third-order polynomial below and we apply the earlier derived results on the incomplete Γ functions.

$$\begin{aligned}
E(|\log^3 |X||) &= \sqrt{2/\pi\sigma^2} \int_0^\infty |\log^3(x)| e^{-x^2/2\sigma^2} dx \\
&= \sqrt{2/\pi\sigma^2} \left(\int_1^\infty \log^3(x) e^{-x^2/2\sigma^2} dx - \int_0^1 \log^3(x) e^{-x^2/2\sigma^2} dx \right) \\
&= \frac{2}{\sqrt{\pi}} \left(\int_{1/\sqrt{2}\sigma}^\infty (\log(\sqrt{2}\sigma) + \log(x))^3 e^{-x^2} dx - \int_0^{1/\sqrt{2}\sigma} (\log(\sqrt{2}\sigma) + \log(x))^3 e^{-x^2} dx \right).
\end{aligned}$$

The third-order polynomial inside the integrals is

$$\log^3(\sqrt{2}\sigma) + 3\log^2(\sqrt{2}\sigma)\log(x) + 3\log(\sqrt{2}\sigma)\log^2(x) + \log^3(x).$$

We will now address each of the terms in turn. The first term corresponds to

$$\log^3(\sqrt{2}\sigma) \left(\frac{2}{\sqrt{\pi}} \int_{1/\sqrt{2}\sigma}^{\infty} e^{-x^2} dx - \frac{2}{\sqrt{\pi}} \int_0^{1/\sqrt{2}\sigma} e^{-x^2} dx \right) = \log^3(\sqrt{2}\sigma)(1 - 2\text{erf}(1/\sqrt{2}\sigma)),$$

where erf is the error function. The second term corresponds to

$$\begin{aligned} & \frac{6}{\sqrt{\pi}} \log^2(\sqrt{2}\sigma) \left(\int_{1/\sqrt{2}\sigma}^{\infty} \log(x) e^{-x^2} dx - \int_0^{1/\sqrt{2}\sigma} \log(x) e^{-x^2} dx \right) \\ &= \frac{6}{\sqrt{\pi}} \log^2(\sqrt{2}\sigma) \left(\frac{d}{ds} \Gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} - \frac{d}{ds} \gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} \right) \\ &= \frac{6}{\sqrt{\pi}} \log^2(\sqrt{2}\sigma) \left(2 \frac{d}{ds} \Gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} - \frac{d}{ds} \Gamma(s) \Big|_{s=1/2} \right). \end{aligned}$$

The third term is

$$\begin{aligned} & \frac{6}{\sqrt{\pi}} \log(\sqrt{2}\sigma) \left(\int_{1/\sqrt{2}\sigma}^{\infty} \log^2(x) e^{-x^2} dx - \int_0^{1/\sqrt{2}\sigma} \log^2(x) e^{-x^2} dx \right) \\ &= \frac{6}{\sqrt{\pi}} \log(\sqrt{2}\sigma) \left(2 \frac{d^2}{ds^2} \Gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} - \frac{d^2}{ds^2} \Gamma(s) \Big|_{s=1/2} \right). \end{aligned}$$

Similarly, the last term is equal to

$$\frac{2}{\sqrt{\pi}} \left(2 \frac{d^3}{ds^3} \Gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} - \frac{d^3}{ds^3} \Gamma(s) \Big|_{s=1/2} \right).$$

Combining these, we find that under the assumption that $\sigma^2 > 1/2$, $\rho = E(|\log^3 X|)$ is equal to

$$\begin{aligned} & \log^3(\sqrt{2}\sigma)(1 - 2\text{erf}(1/\sqrt{2}\sigma)) + \frac{6}{\sqrt{\pi}} \log^2(\sqrt{2}\sigma) \left(2 \frac{d}{ds} \Gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} - \frac{d}{ds} \Gamma(s) \Big|_{s=1/2} \right) \\ &+ \frac{6}{\sqrt{\pi}} \log(\sqrt{2}\sigma) \left(2 \frac{d^2}{ds^2} \Gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} - \frac{d^2}{ds^2} \Gamma(s) \Big|_{s=1/2} \right) \\ &+ \frac{2}{\sqrt{\pi}} \left(2 \frac{d^3}{ds^3} \Gamma(s, 1/\sqrt{2}\sigma) \Big|_{s=1/2} - \frac{d^3}{ds^3} \Gamma(s) \Big|_{s=1/2} \right). \end{aligned} \tag{13}$$

The derivatives of the Γ function and incomplete Γ function are given in Equations (5), (6), (7), (9), (10) and (11).

A bound on ρ . The above approach does not work when $\sigma^2 \leq 1/2$ and we need an alternative approach for that. Even when the above approach works it might be convenient to have a simple expression that bounds ρ . By an application of Hölder's inequality we get such a simple bound that holds for any $\sigma > 0$. The bound is based on the 4'th moment of $\log |X|$ which is

$$\begin{aligned}
E(\log^4 |X|) &= \sqrt{\frac{2}{\pi\sigma^2}} \int_0^\infty \log^4(x) e^{-x^2/2\sigma^2} dx \\
&= \frac{2}{\sqrt{\pi}} \int_0^\infty \log^4(\sqrt{2}\sigma x) e^{-x^2} dx \\
&= \log^4(\sqrt{2}\sigma) - 2\log^3(\sqrt{2}\sigma)(\gamma + 2\log 2) + \frac{3}{2}\log^2(\sqrt{2}\sigma)((2\log 2 + \gamma)^2 + \pi^2/2) \\
&\quad - \frac{1}{2}\log(\sqrt{2}\sigma)((2\log 2 + \gamma)^3 + \frac{3\pi^2}{2}(2\log 2 + \gamma) + 14\zeta(3)) \\
&\quad + \frac{1}{16}((2\log 2 + \gamma)^4 + 3\pi^2(2\log 2 + \gamma)^2 + 56(2\log 2 + \gamma)\zeta(3) + 7\pi^4/4).
\end{aligned}$$

and

$$\rho = E(\log^3 |X|) \leq (E(\log^4 |X|))^{3/4}. \quad (14)$$

A.2 Independence of S and $|X|$

It is well known that the sign S of a centered normal distributed random variable with variance $\sigma^2 > 0$ and its absolute value $|X|$ are independent. One way to verify this is to recall that S and $|X|$ are independent if for all $a, b \in \mathbb{R}$, $\Pr(S \leq a, |X| \leq b) = \Pr(S \leq a) \Pr(|X| \leq b)$. It is easy to verify this: for any $b \in \mathbb{R}$ and $a < -1$

$$\Pr(S_i \leq a, |X_i| \leq b) = 0 = \Pr(S_i \leq a) \Pr(|X_i| \leq b)$$

and for $a \geq 1$,

$$\Pr(S_i \leq a, |X_i| \leq b) = \Pr(|X_i| \leq b) = \Pr(S_i \leq a) \Pr(|X_i| \leq b).$$

Finally, for $-1 \leq a < 1$, and $b \geq 0$ ($b < 0$ is trivial), we have due to the symmetry of X that

$$\begin{aligned}
\Pr(S_i \leq a) \Pr(|X_i| \leq b) &= \Pr(|X_i| \leq b)/2 \\
&= \Pr(-b \leq X_i \leq b)/2 = \Pr(0 \leq X_i \leq b) = \Pr(-b \leq X_i \leq 0) \\
&= \Pr(X_i < 0, -X_i \leq b) = \Pr(X_i < 0, |X_i| \leq b) = \Pr(S_i \leq a, |X_i| \leq b).
\end{aligned}$$

B Further Results on GPs and DGPs

B.1 Representing a GP with Quadratic Kernel by a Multivariate Normal RV

GPs can often be written as a (potentially infinite) linear combination of normally distributed random variables. That is the Karhunen-Loève expansion which is based on an

eigendecomposition of the kernel-integral operator corresponding to the covariance function (Mercer's theorem). For polynomial kernels a more direct approach is possible. We demonstrate this approach here. We start with a simple 2-dimensional example, for which the key steps are transparent, before approaching the more abstract general case.

2-Dimensional feature map. Consider the following kernel function on \mathbb{R} ,

$$k_0(x, y) = x^2y^2 + xy = \begin{pmatrix} x^2 \\ x \end{pmatrix}^\top \begin{pmatrix} y^2 \\ y \end{pmatrix} = \psi(x)^\top \psi(y),$$

where we will represent $\psi(x)$ as $\psi(x) = (\psi_1(x), \psi_2(x))^\top$. The RKHS \mathcal{H} corresponding to k_0 is 2-dimensional since x^2 and x are linearly independent.

We start by taking a closer look at \mathcal{H}_0 . Similarly as for the linear kernel in Section 3, we can write the RKHS in the form

$$\begin{aligned} \mathcal{H}_0 &= \{h(x) = \alpha_1\psi(-1)^\top \psi(x) + \alpha_2\psi(-1/2)^\top \psi(x) : \alpha_1, \alpha_2 \in \mathbb{R}\} \\ &= \{h(x) = \alpha_1(x^2 - x) + \alpha_2(\frac{1}{4}x^2 - \frac{1}{2}x) : \alpha_1, \alpha_2 \in \mathbb{R}\}, \end{aligned}$$

since $\psi(-1)^\top \psi(x)$ and $\psi(-1/2)^\top \psi(x)$ are linearly independent and \mathcal{H}_0 is 2-dimensional. Let g_0 be the GP corresponding to kernel k_0 then g_0 attains values in \mathcal{H}_0 and there are random variables $\alpha_{1,\omega}$ and $\alpha_{2,\omega}$ such that for all $x \in \mathbb{R}$,

$$g_0(x) = \alpha_{1,\omega}(x^2 - x) + \alpha_{2,\omega}(\frac{1}{4}x^2 - \frac{1}{2}x).$$

Consider $x_1 = -1$ then $g_0(-1) = 2\alpha_{1,\omega} + (3/4)\alpha_{2,\omega}$. Furthermore, for $x_2 = -1/2$ we find that $g_0(-1/2) = (3/4)\alpha_{1,\omega} + (5/16)\alpha_{2,\omega}$. Let C be a covariance matrix of a gaussian vector $(g_0(x_1), g_0(x_2))$, and is given by

$$C = \begin{pmatrix} k(x_1, x_1) & k(x_1, x_2) \\ k(x_2, x_1) & k(x_2, x_2) \end{pmatrix} = \begin{pmatrix} 2 & 3/4 \\ 3/4 & 5/16 \end{pmatrix}$$

since $k(-1, -1) = 1 + 1 = 2$, $k(-1, -1/2) = 3/4$, $k(-1/2, -1/2) = 5/16$. Therefore, we get

$$\begin{pmatrix} g_0(-1) \\ g_0(-1/2) \end{pmatrix} = C \begin{pmatrix} \alpha_{1,\omega} \\ \alpha_{2,\omega} \end{pmatrix} = \begin{pmatrix} 2 & 3/4 \\ 3/4 & 5/16 \end{pmatrix} \begin{pmatrix} \alpha_{1,\omega} \\ \alpha_{2,\omega} \end{pmatrix}.$$

The covariance between $\alpha_{1,\omega}$ and $\alpha_{2,\omega}$ is given by

$$\begin{aligned} Cov\left(\begin{pmatrix} \alpha_{1,\omega} \\ \alpha_{2,\omega} \end{pmatrix}\right) &= C^{-1}Cov\left(\begin{pmatrix} g_0(-1) \\ g_0(-1/2) \end{pmatrix}\right)C^{-1} \\ &= \begin{pmatrix} 5 & -12 \\ -12 & 32 \end{pmatrix} \begin{pmatrix} 2 & 3/4 \\ 3/4 & 5/16 \end{pmatrix} \begin{pmatrix} 5 & -12 \\ -12 & 32 \end{pmatrix} \\ &= \begin{pmatrix} 5 & -12 \\ -12 & 32 \end{pmatrix} = C^{-1}. \end{aligned}$$

Hence, we get

$$E(\alpha_{1,\omega}\alpha_{2,\omega}) = -12,$$

and

$$\begin{pmatrix} \alpha_{1,\omega} \\ \alpha_{2,\omega} \end{pmatrix} \sim N\left(0, \begin{pmatrix} 5 & -12 \\ -12 & 32 \end{pmatrix}\right).$$

Rearranging the terms gives

$$\begin{aligned} g_0(x) &= x^2(\alpha_{1,\omega} + (1/4)\alpha_{2,\omega}) + x(\alpha_{1,\omega} + (1/2)\alpha_{2,\omega}) \\ &= \psi_1(x)(\alpha_{1,\omega}\psi_1(x_1) + \alpha_{2,\omega}\psi_1(x_2)) + \psi_2(x)(\alpha_{1,\omega}\psi_2(x_1) + \alpha_{2,\omega}\psi_2(x_2)) \\ &= \begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix}^\top \psi(x), \end{aligned}$$

where

$$\begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix} = \begin{pmatrix} \psi_1(x_1) & \psi_1(x_2) \\ \psi_2(x_1) & \psi_2(x_2) \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix}.$$

It is easy to see that

$$\begin{pmatrix} \psi_1(x_1) & \psi_1(x_2) \\ \psi_2(x_1) & \psi_2(x_2) \end{pmatrix}^\top \begin{pmatrix} \psi_1(x_1) & \psi_1(x_2) \\ \psi_2(x_1) & \psi_2(x_2) \end{pmatrix} = C.$$

In the following, let

$$B = \begin{pmatrix} \psi_1(x_1) & \psi_1(x_2) \\ \psi_2(x_1) & \psi_2(x_2) \end{pmatrix} \quad \text{and note that} \quad \begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix} = B \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix}.$$

The covariance between Y_1 and Y_2 is

$$\text{Cov}\left(\begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix}\right) = B \text{Cov}\left(\begin{pmatrix} \alpha_{1,\omega} \\ \alpha_{2,\omega} \end{pmatrix}\right) B^\top.$$

Following that, we multiply both sides by B on the right, which gives

$$\text{Cov}\left(\begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix}\right) B = B C^{-1} B^\top B = B.$$

Therefore,

$$\text{Cov}\left(\begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix}\right) = I,$$

and

$$\begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix} \sim N\left(0, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\right).$$

The general case. We consider now the polynomial kernel function. The argument below generalizes, however, right away to any other kernel function which has a finite dimensional feature vector. Denote the parameters of the polynomial kernel function with integer parameter $d \geq 1$, and real valued $c > 0$ on \mathbb{R} ,

$$k(x, y) = (xy + c)^d = \phi(x)^\top \phi(y),$$

where

$$\phi(x) = \left(x^d, \binom{d}{1}^{1/2} x^{d-1} c^{1/2}, \binom{d}{2}^{1/2} x^{d-2} (c^2)^{1/2}, \dots, c^{d/2} \right)^\top.$$

The functions $x^d, x^{d-1} c^{1/2}, \dots, c^{d/2}$ are linearly independent and [13, Ex 3.7] shows that these functions all lie in \mathcal{H} . Hence, \mathcal{H} is at least $d + 1$ -dimensional. In fact, it follows from [13, Thm 2.10] that \mathcal{H} is $d + 1$ -dimensional. This implies that there are $x_1, \dots, x_{d+1} \in \mathbb{R}$ such that

$$\mathcal{H} = \left\{ \sum_{i=1}^{d+1} \alpha_i k(x_i, \cdot) : \alpha_i \in \mathbb{R}, i \leq d + 1 \right\}.$$

Let g be a GP with kernel k then g attains values in \mathcal{H} and

$$g(x) = \sum_{i=1}^{d+1} \alpha_{i,\omega} k(x_i, x)$$

for $d + 1$ stochastic real valued coefficients $\alpha_{1,\omega}, \dots, \alpha_{d+1,\omega}$. In particular,

$$\begin{pmatrix} g(x_1) \\ \vdots \\ g(x_{d+1}) \end{pmatrix} = C \boldsymbol{\alpha}_\omega \quad \text{and} \quad C^{-1} \begin{pmatrix} g(x_1) \\ \vdots \\ g(x_{d+1}) \end{pmatrix} = \boldsymbol{\alpha}_\omega,$$

where $\boldsymbol{\alpha}_\omega = (\alpha_{1,\omega}, \dots, \alpha_{d+1,\omega})^\top$ and $C = (k(x_i, x_j))_{i,j \leq d+1}$. That C is invertible can be seen in the following way: the functions $k(x_1, \cdot), \dots, k(x_{d+1}, \cdot)$ are linearly independent since they span the $d + 1$ -dimensional space \mathcal{H} . In particular, for any h there exists a_1, \dots, a_{d+1} such that $h = \sum_{i=1}^{d+1} a_i k(x_i, \cdot)$ and

$$\|h\|^2 = \mathbf{a}^\top C \mathbf{a},$$

where $\mathbf{a} = (a_1, \dots, a_{d+1})^\top$. If C would not be of full rank then there would exist an eigenvector \mathbf{e} of C with eigenvalue 0. The function h corresponding to \mathbf{e} would not be the constant 0 function since \mathbf{e} would not be zero. However, in this case

$$0 \neq \|h\|^2 = \mathbf{e}^\top C \mathbf{e} = 0,$$

and C has to be of full rank.

Because $\boldsymbol{\alpha}_\omega$ is a linear transformation of a zero mean Gaussian vector it follows that $\boldsymbol{\alpha}_\omega$ is also a zero mean Gaussian vector. The matrix C is the covariance matrix of the $d + 1$ -dimensional Gaussian vector $(g(x_1), \dots, g(x_{d+1}))^\top$ and

$$\text{Cov}(\boldsymbol{\alpha}_\omega) = C^{-1} C C^{-1} = C^{-1}.$$

Writing the Gaussian process in terms of the feature map ϕ ,

$$g(x) = \sum_{i=1}^{d+1} \alpha_{i,\omega} \phi(x_i)^\top \phi(x)$$

and denoting the different entries of $\phi(x)$ by $\phi_1(x), \dots, \phi_{d+1}(x)$, we define $d+1$ zero mean Gaussian random variables Y_1, \dots, Y_{d+1} through

$$\sum_{i=1}^{d+1} \alpha_{i,\omega} \phi(x_i) = \underbrace{(\phi(x_1) \ \dots \ \phi(x_{d+1}))}_B \boldsymbol{\alpha}_\omega = \begin{pmatrix} Y_1 \\ \vdots \\ Y_{d+1} \end{pmatrix}$$

Note that $g(x) = Y^\top \phi(x)$. The random vector $Y = (Y_1, \dots, Y_{d+1})^\top$ has covariance

$$\text{Cov} \begin{pmatrix} Y_1 \\ \vdots \\ Y_{d+1} \end{pmatrix} = B \text{Cov}(\boldsymbol{\alpha}_\omega) B^\top.$$

Multiplying by B on the right yields

$$\text{Cov} \begin{pmatrix} Y_1 \\ \vdots \\ Y_{d+1} \end{pmatrix} B = B \text{Cov}(\boldsymbol{\alpha}_\omega) B^\top B = B,$$

since $B^\top B = C$. The matrix B is invertible ... and multiplying the above from the right by B^{-1} leads us to

$$\text{Cov} \begin{pmatrix} Y_1 \\ \vdots \\ Y_{d+1} \end{pmatrix} = I, \quad \text{and} \quad \begin{pmatrix} Y_1 \\ \vdots \\ Y_{d+1} \end{pmatrix} \sim N(0, I).$$

B.2 A Berry-Esseen Bound for the case that $d_1 = \dots = d_{\ell-1} = 2$

Before considering a general case, we first examine the approximation of DGPs, where the successive layers use a kernel $k_i(x, y) = (xy)^{d_i}$ with $d_i = 2$. We will also define $d_i^\downarrow = \sum_{j=0}^{i-1} 2$ for all $i = 1, \dots, \ell - 1$. As it was shown above, the GP g_1 can be represented as

$$g_1(x) = \begin{pmatrix} Z_1 \\ \vdots \\ Z_{d+1} \end{pmatrix}^\top \begin{pmatrix} \phi_1(x) \\ \vdots \\ \phi_{d+1}(x) \end{pmatrix},$$

where $(Z_1, \dots, Z_{d+1})^\top \sim N(0, I)$. Also, let $Y_i \sim N(0, \sigma_i^2)$ be i.i.d, and independent of Z_1, \dots, Z_{d+1} , and such that $g_i(x) = Y_i x^2$ for all $2 \leq i \leq \ell$. Then the DGP can be written as

$$\begin{aligned}
g_\ell \circ \dots \circ g_1(x) &= Y_\ell(Y_{\ell-1})^{d_1^\downarrow}(Y_{\ell-2})^{d_2^\downarrow} \times \dots \times (Y_2)^{d_{\ell-2}^\downarrow} \left(\sum_{i=1}^{d+1} Z_i \phi_i(x) \right)^{d_{\ell-1}^\downarrow} \\
&= Y_\ell Y_{\ell-1}^2 Y_{\ell-2}^4 \times \dots \times Y_2^{2^{(\ell-2)}} \left(\sum_{i=1}^{d+1} Z_i \phi_i(x) \right)^{2^{(\ell-1)}}.
\end{aligned}$$

Taking the logarithm of the absolute values of the product of Y -terms yields $2(\ell-2) \log |Y_2| + \dots + 2 \log |Y_{\ell-1}| + \log |Y_\ell| = \sum_{j=2}^{\ell} c_j \log |Y_j|$, where $c_j = 2(\ell - j)$ for $j = 2, \dots, \ell - 1$ and $c_\ell = 1$. We then want to apply the Berry-Esseen Theorem, and in order to do this we first define $\sigma_{i,\log} = c_i^2 \text{Var}(\log |Y_i|)$ and $\rho_{i,\log} = c_i^3 E(|\log |Y_i||^3)$ for all $i = 2, \dots, \ell$. Furthermore, to derive the expression for the Berry-Esseen bound, we first have

$$\left(\sum_{i=2}^{\ell} \sigma_{i,\log}^2 \right)^{-3/2} \sum_{i=2}^{\ell} \rho_{i,\log} = \frac{\sum_{i=2}^{\ell} c_i^3 E(|\log |Y_i||^3)}{(\sum_{i=2}^{\ell} c_i^2)^{3/2} (\text{Var}(\log |Y_i|))^{3/2}},$$

where the sums of the coefficients $\sum_{i=2}^{\ell-1} c_i^2$ and $\sum_{i=2}^{\ell-1} c_i^3$ can be found as

$$\sum_{i=2}^{\ell-1} c_i^2 = 4 \sum_{i=2}^{\ell-1} (\ell - i + 1)^2 = 4 \sum_{i=1}^{\ell-2} (\ell - i)^2 = 4 \sum_{i=1}^{\ell-1} i^2 - 4 = \frac{2\ell(\ell-1)(2\ell-1)}{3} - 4,$$

$$\begin{aligned}
\sum_{i=2}^{\ell-1} c_i^3 &= 8 \sum_{i=2}^{\ell-1} (\ell - i + 1)^3 = 8 \sum_{i=1}^{\ell-2} (\ell - i)^3 \\
&= 8 \left(-\frac{1}{2} \ell^3 (\ell - 1) + \frac{1}{2} \ell^2 (\ell - 1)(2\ell - 1) - \sum_{i=1}^{\ell-1} i^3 - 1 \right) \\
&= 8 \left(-\frac{1}{2} \ell^3 (\ell - 1) + \frac{1}{2} \ell^2 (\ell - 1)(2\ell - 1) - \frac{\ell^2 (\ell - 1)^2}{4} - 1 \right) \\
&= 2\ell^2 (\ell - 1)^2 - 8.
\end{aligned}$$

For later use we also find

$$\sum_{i=2}^{\ell} c_i = \sum_{i=2}^{\ell-1} c_i + 1 = 2 \sum_{i=2}^{\ell-1} (\ell - i + 1) + 1 = 2 \sum_{i=2}^{\ell-2} (\ell - i) + 1 = \ell(\ell - 1) - 2 + 1 = \ell(\ell - 1) - 1$$

Hence, we get $\sum_{i=2}^{\ell} c_i^2 = \frac{2\ell(\ell-1)(2\ell-1)}{3} - 3$, and $\sum_{i=2}^{\ell} c_i^3 = 2\ell^2(\ell - 1)^2 - 7$, and this gives

us the expressions for the bound, the mean and the variance of Y

$$0.56 \left(\sum_{i=2}^{\ell} \sigma_{i,\log}^2 \right)^{-3/2} \sum_{i=2}^{\ell} \rho_{i,\log} = 0.56 \frac{E(|\log |Y_1||^3)}{(\text{Var}(\log |Y_1|))^{3/2}} \frac{(2\ell^2(\ell-1)^2 - 7)3^{3/2}}{(2\ell(\ell-1)(2\ell-1) - 9)^{3/2}} \quad (15)$$

$$\leq 3\ell^{-1/2} \frac{E(|\log |Y_1||^3)}{(\text{Var}(\log |Y_1|))^{3/2}},$$

$$\sum_{i=2}^{\ell} c_i E(\log |Y_i|) = \sum_{i=2}^{\ell} c_i E(\log |Y_i|) \approx (\ell(\ell-1) - 1)(\log \sigma - 0.63) \quad (16)$$

$$= ((\ell-1)^2 + \ell)(\log \sigma - 0.63),$$

$$\text{Var}\left(\sum_{i=2}^{\ell} c_i^2 \log |Y_i|\right) = \sum_{i=2}^{\ell} c_i^2 \text{Var}(\log |Y_i|) = \left(\frac{2\ell(\ell-1)(2\ell-1)}{3} - 3 \right) \text{Var}(\log |Y_1|)$$

$$= \frac{\pi^2}{8} \left(\frac{2\ell(\ell-1)(2\ell-1)}{3} - 3 \right). \quad (17)$$

The Berry-Esseen Theorem guarantees that

$$\sup_{x \in \mathbb{R}} |\Pr((\sigma_{2,\log}^2 + \dots + \sigma_{\ell,\log}^2)^{-1/2} \sum_{j=2}^{\ell} (c_j \log |Y_j| - c_j E(\log |Y_j|)) \leq x) - \Phi(x)|$$

$$\leq 0.56 \left(\sum_{i=1}^n \sigma_{i,\log}^2 \right)^{-3/2} \sum_{i=1}^n \rho_{i,\log}.$$

Similarly, to the derivation in Section 3.2 by substitution,

$$\sup_{x \in \mathbb{R}} |\Pr\left(|Y_{\ell}| \prod_{i=2}^{\ell-1} |Y_i|^{2(\ell-i)} \leq x\right) - \Pr(e^Y \leq x)| = \sup_{x \in \mathbb{R}} \left| \sum_{j=2}^{\ell} c_j \log |Y_j| - \Pr(Y \leq x) \right|,$$

where $Y \sim N(\sum_{i=2}^{\ell} c_i E(\log |Y_i|), \sum_{i=2}^{\ell} \sigma_{i,\log}^2)$ and the last term is upper bounded by the bound given in (15). Note that $Y_{\ell} \prod_{i=2}^{\ell-1} Y_i^{2(\ell-i)} = S_{\ell} \prod_{i \in I} S_i |Y_{\ell}| \prod_{i=2}^{\ell-1} |Y_i|^{2(\ell-i)}$, where we can write $S = S_{\ell} \prod_{i \in I} S_i$, since S has the same distribution as the product on the right hand side, and attains values -1 and 1 with probability $1/2$.

B.3 Incorporating $(g_1(x))^{c_1}$ in the Approximation

In Section 3.2 we derive an approximation that leads us to a bound on

$$\sup_{y \in \mathbb{R}} |\Pr\left(S \prod_{i=2}^{\ell} Y_i \leq y\right) - \Pr(Se^Y \leq y)|.$$

The involved random variables are specified in Section 3.2. We also have a random variable $(g_1(x))^{c_1}$ that we like to incorporate on both sides. In detail, we want a bound on

$$\sup_{y \in \mathbb{R}} |\Pr\left(S(g_1(x))^{c_1} \prod_{i=2}^{\ell} Y_i \leq y\right) - \Pr(S(g_1(x))^{c_1} e^Y \leq y)|.$$

We can attain such a bound by using a conditional expectation argument. Note that, due to the towering rule of conditional expectations,

$$\begin{aligned} \Pr\left(S(g_1(x))^{c_1} \sum_{j=2}^{\ell} c_j \log |Y_j| \leq y\right) &= E\left(\mathbf{1}\left\{S(g_1(x))^{c_1} \sum_{j=2}^{\ell} c_j \log |Y_j| \leq y\right\}\right) \\ &= E\left(E\left(\mathbf{1}\left\{S(g_1(x))^{c_1} \sum_{j=2}^{\ell} c_j \log |Y_j| \leq y\right\}\right) \middle| (g_1(x))^{c_1}\right), \end{aligned}$$

where $\mathbf{1}$ denotes the indicator function. Since the indicator function is real-valued, there exists a measurable function h such that the conditional expectation is equal to $h((g_1(x))^{c_1})$ (see, for example, [14, Eq. 10, p.220]). In fact, one can observe that we can choose

$$h(z) = E\left(\mathbf{1}\left\{Sz \sum_{j=2}^{\ell} c_j \log |Y_j| \leq y\right\}\right) = \Pr\left(Sz \sum_{j=2}^{\ell} c_j \log |Y_j| \leq y\right).$$

The same approach leads us to a function

$$\tilde{h}(z) = \Pr\left(Sze^Y \leq y\right)$$

and

$$\sup_{z \in \mathbb{R}} |h(z) - \tilde{h}(z)| = \sup_{y \in \mathbb{R}} \left| \Pr\left(S \prod_{i=2}^{\ell} Y_i \leq y\right) - \Pr(Se^Y \leq y) \right|.$$

Hence,

$$\begin{aligned} &\sup_{y \in \mathbb{R}} \left| \Pr\left(S(g_1(x))^{c_1} \prod_{i=2}^{\ell} Y_i \leq y\right) - \Pr(S(g_1(x))^{c_1} e^Y \leq y) \right| \\ &= \sup_{y \in \mathbb{R}} \left| E(h((g_1(x))^{c_1})) - E(\tilde{h}((g_1(x))^{c_1})) \right| \\ &\leq \sup_{z \in \mathbb{R}} |h(z) - \tilde{h}(z)| \\ &= \sup_{y \in \mathbb{R}} \left| \Pr\left(S \prod_{i=2}^{\ell} Y_i \leq y\right) - \Pr(Se^Y \leq y) \right| \end{aligned}$$

and bound is not affected by the introduction of $(g_1(x))^{c_1}$.

References

- [1] A. Damianou and N. D. Lawrence. Deep Gaussian processes. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, volume 31, 2013.
- [2] M. Giordano, K. Ray, and J. Schmidt-Hieber. On the inability of gaussian process regression to optimally learn compositional functions. In *Advances in Neural Information Processing Systems*, volume 35, 2022.

- [3] G. Finocchio and J. Schmidt-Hieber. Posterior contraction for deep gaussian process priors. *Journal of Machine Learning Research*, 24(66):1–49, 2023.
- [4] K. Abraham and N. Deo. Deep gaussian process priors for bayesian inference in nonlinear inverse problems, 2023.
- [5] I. Castillo and T. Randrianarisoa. Deep horseshoe gaussian processes, 2024.
- [6] D. Duvenaud, O. Rippel, R. Adams, and Z. Ghahramani. Avoiding pathologies in very deep networks. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, 2014.
- [7] M. M. Dunlop, M. A. Girolami, A. M. Stuart, and A. L. Teckentrup. How deep are deep gaussian processes? *Journal of Machine Learning Research*, 19(1), 2018.
- [8] O. Kallenberg. *Foundations of Modern Probability*. Springer, third edition, 2021.
- [9] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.
- [10] D.H. Fremlin. *Measure Theory, Volume 1*. Torres Fremlin, 2000.
- [11] K.O. Geddes, M.L. Glasser, R.A. Moore, and T.C. Scott. Evaluation of classes of definite integrals involving elementary functions via differentiation of special functions. *Applicable Algebra in Engineering, Communication and Computing*, 1, 1990.
- [12] Y. L. Luke. *The special functions and their approximations*, volume 1. London: Academic Press, 1969.
- [13] V. Paulsen and M. Raghupathi. *An Introduction to the Theory of Reproducing Kernel Hilbert Spaces*. Cambridge University Press, 2016.
- [14] A. N. Shiryaev. *Probability (2nd ed.)*. Springer-Verlag, Berlin, Heidelberg, 1995.