Spectral Diffusion Processes

Angus Phillips[*,†], **Thomas Seror**[*,‡], **Michael Hutchinson**[†],
Valentin De Bortoli[‡], **Arnaud Doucet**[†], **Émile Mathieu**[§]

Abstract

Score-based generative modelling (SGM) has proven to be a very effective method for modelling densities on finite-dimensional spaces. In this work we propose to extend this methodology to learn generative models over functional spaces. To do so, we represent functional data in spectral space to dissociate the stochastic part of the processes from their space-time part. Using dimensionality reduction techniques we then sample from their stochastic component using finite dimensional SGM. We demonstrate our method's effectiveness for modelling various multimodal datasets.

1 Introduction

Score-based Generative Models (SGMs) (Song and Ermon, 2019; Song et al., 2021; Ho et al., 2020; Dhariwal and Nichol, 2021) are a powerful class of generative models. Noise is progressively added to data following a Stochastic Differential Equation (SDE)—the forward *noising* process—until it is approximately Gaussian. The generative model is given by an approximation of the associated time-reversed process called the backward *denoising* process. This process is an SDE whose drift depends on the gradient of the logarithm of the densities of the forward process. The drift is estimated by leveraging techniques from deep learning and score matching (Hyvärinen, 2005; Vincent, 2011).

Probabilistic modelling over functional spaces is of high importance in the physical sciences where objects of interest are often *fields*, mapping (space-time) inputs to tensors like scalars or vectors. Aleatoric and epistemic uncertainties motivate for a *probabilistic* treatment of these phenomena, like weather forecasting (Ravuri et al., 2021) or electric potential prediction (Yang et al., 2020).

A well-established framework to model distributions over functions are Gaussian processes (GPs) (Rasmussen, 2003). To alleviate their scalability and flexibility limitations, neural processes (NPs) have been introduced and successfully applied on a wide variety of probabilistic prediction tasks. Yet, their Gaussian likelihood assumption limits their flexibility.

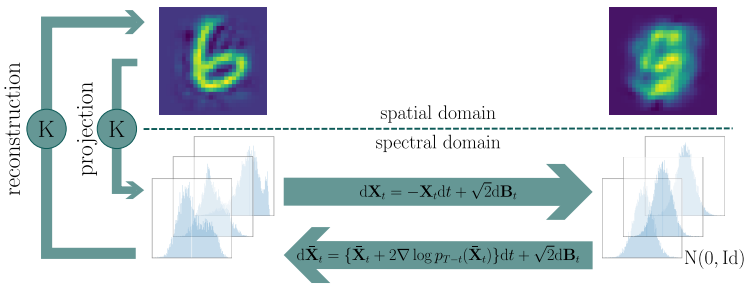

Figure 1: Illustration of our methodology. SGM is performed in a spectral space.

[*]Equal contribution. Correspondence to: angus.phillips@stats.ox.ac.uk.
[†]Dept. of Statistics, University of Oxford, Oxford, UK.
[‡]Dept. of Computer Science ENS, CNRS, PSL University Paris, France.
[§]Dept. of Engineering, University of Cambridge, Cambridge, UK.

Score-based Method Workshop at Conference on Neural Information Processing Systems (NeurIPS 2022).

In this work, we circumvent this problem by casting the modelling problem in a spectral space. Doing so, we decouple the *stochastic* part of the process from its *space-time* part. This approach turns an *uncountable* dimensional problem into a *countable* dimensional one, which is turned into a *finite* dimensional one by truncation. Finally, we deal with this last problem using standard SGM techniques. Our methodology is illustrated in Fig. 1 and presented in more detail in Sec. 3.

2 Background

Score-based generative modelling. We briefly recall the concepts behind SGMs on the Euclidean space $\mathbb{R}^d$. Let $(\mathbf{Y}_t)_{t\geq 0}$ be a *noising* process defined by the following SDE

$$\mathrm{d}\mathbf{Y}_t = -\mathbf{Y}_t \mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \mathbf{Y}_0 \sim p_0,$$

with $(\mathbf{B}_t)_{t\geq 0}$ a d-dimensional Brownian motion and p_0 the data distribution on $\mathbb{R}^d$. This process converges towards the unit Gaussian distribution $\mathrm{N}(0, \mathrm{Id})$. Under conditions on p_0, the time-reversal $(\bar{\mathbf{Y}}_t)_{t\in[0,T]} = (\mathbf{Y}_{T-t})_{t\in[0,T]}$ is given by (Cattiaux et al., 2021; Haussmann and Pardoux, 1986)

$$\mathrm{d}\bar{\mathbf{Y}}_t = \{\bar{\mathbf{Y}}_t + 2\nabla \log p_{T-t}(\bar{\mathbf{Y}}_t)\}\mathrm{d}t + \sqrt{2}\mathrm{d}\mathbf{B}_t, \quad \bar{\mathbf{Y}}_0 \sim p_T, \tag{1}$$

with p_t the density of $\mathbf{Y}_t$. Denoising diffusion models are defined to sample approximately from (1). First, we replace p_T by $\mathrm{N}(0, \mathrm{Id})$. Then $\nabla \log p_t$ is approximated by a parametric function $\boldsymbol{s}_\theta(t, \cdot)$ so as to minimise $\ell_t(\mathbf{s}) = \mathbb{E}[\|\mathbf{s}_\theta(t, \mathbf{Y}_t) - \nabla_{y_t} \log p_{t|0}(\mathbf{Y}_t|\mathbf{Y}_0)\|^2]$ (Hyvärinen, 2005). Finally, an Euler–Maruyama discretisation of Eq. (1) is considered using a discretisation step γ and $N = T/\gamma$

$$\bar{Y}_{n+1} = \bar{Y}_n + \gamma\{\bar{Y}_n + 2\mathbf{s}_\theta(T - n\gamma, \bar{Y}_n)\} + \sqrt{2\gamma}G_{n+1}, \quad \bar{Y}_0 \sim \mathrm{N}(0, \mathrm{Id}), \quad G_{n+1} \overset{\text{i.i.d.}}{\sim} \mathrm{N}(0, \mathrm{Id}).$$

Stochastic processes. Here we are interested in modelling distributions over functions. Given a compact input space $\mathcal{X}$, a $\mathbb{R}^d$-valued stochastic process $(\mathbf{Y}_x)_{x\in\mathcal{X}}$ is a collection of random variables $\mathbf{Y}_x$ [5]. Most existing methods which model stochastic processes rely on the parametrisation of finite dimensional marginals $\{\mathbf{Y}_{x_i} : i \in \{1, \dots, n\},\ x_i \in \mathcal{X}\}$ for every $n \in \mathbb{N}$. However, it is not clear *a priori* that a collection of probability distributions $\mathrm{S} = \{\pi_{x_1,\dots,x_n} : (x_1, \dots, x_n) \in \mathcal{X}^n,\ n \in \mathbb{N}\}$ defines a distribution $\pi \in \mathscr{P}((\mathbb{R}^d)^{\mathcal{X}})$[6]. This is the focus of the Kolmogorov extension theorem (Charalambos and Aliprantis, 2013, Thm 15.26) which ensures the existence of such a distribution under a set of conditions. The conditions required of the system are exchangeability and consistency. A system S is *exchangeable*, if for any $n \in \mathbb{N}$, n-permutation $\sigma \in \mathfrak{S}_n$, $x_1, \dots, x_n \in \mathcal{X}$ and continuous and bounded function $f \in \mathrm{C}_b((\mathbb{R}^d)^n)$, we have $\int f(y_{\sigma^{-1}(1)}, \dots, y_{\sigma^{-1}(n)})\mathrm{d}\pi_{x_{\sigma(1)},\dots,x_{\sigma(n)}}(y_1, \dots, y_n) = \int f(y_1, \dots, y_n)\mathrm{d}\pi_{x_1,\dots,x_n}(y_1, \dots, y_n)$. A system S is *consistent*, if for any $n_1 \leq n_2$, $x_1, \dots, x_{n_2} \in \mathcal{X}$ and $f \in \mathrm{C}_b((\mathbb{R}^d)^{n_1})$ we have $\int f(y_1, \dots, y_{n_1})\mathrm{d}\pi_{x_1,\dots,x_{n_1}}(y_1, \dots, y_{n_1}) = \int f(y_1, \dots, y_{n_1})\mathrm{d}\pi_{x_1,\dots,x_{n_2}}(y_1, \dots, y_{n_2})$. Existing generative modelling approaches often satisfy the exchangeability criterion but fall short on the consistency. In this work we aim at bridging this gap by developing a *spectral* method which immediately yields the consistency of the obtained system.

3 Spectral Score-Based Processes

Karhunen-Loève theorem. In order to derive a consistent method, our key insight is to dissociate the *stochastic* part of the process from its *space-time* component using a spectral decomposition. This approach stems from the Karhunen-Loève theorem which we recall below.

Theorem 1. *Let $\mathcal{X}$ be a compact space and $(\mathbf{Y}_x)_{x\in\mathcal{X}}$ a continuous stochastic process such that for any $x \in \mathcal{X}$, $\mathbb{E}[\mathbf{Y}_x] = 0$ and $\mathbb{E}[\|\mathbf{Y}_x\|^2] < +\infty$. For any $x_1, x_2 \in \mathcal{X}$, we denote the covariance kernel as $\mathrm{K}_{\mathbf{Y}}(x_1, x_2) = \mathbb{E}[\langle \mathbf{Y}_{x_1}, \mathbf{Y}_{x_2}\rangle]$ and define a $\mathrm{L}^2(\mathcal{X})$ operator $T_{\mathrm{K}_{\mathbf{Y}}}$ given for any $f \in \mathrm{L}^2(\mathcal{X})$ by $T_{\mathrm{K}_{\mathbf{Y}}} f(x_2) = \int_{\mathcal{X}} \mathrm{K}_{\mathbf{Y}}(x_1, x_2) f(x_1)\mathrm{d}x_1$. Denoting by $(e_m)_{m\in\mathbb{N}} \in (\mathrm{L}^2(\mathcal{X}))^{\mathbb{N}}$, $(\lambda_m)_{m\in\mathbb{N}} \in \mathbb{R}^{\mathbb{N}}$ the eigenfunctions and eigenvalues of $T_{\mathrm{K}_{\mathbf{Y}}}$, we have*

$$\lim_{M\to+\infty} \mathbb{E}[\sup_{x\in\mathcal{X}} \|\mathbf{Y}_x - \textstyle\sum_{m=0}^{M} {\lambda_m}^{1/2} Z_m e_m(x)\|^2] = 0,$$

where for any $m \in \mathbb{N}$, $Z_m = \lambda_m^{-1/2} \int_{\mathcal{X}} \langle \mathbf{Y}_x, e_m(x)\rangle \mathrm{d}x$.

[5] In this work, we focus on *continuous* stochastic processes, *i.e.* any trajectory of $(\mathbf{Y}_x)_{x\in\mathcal{X}}$ belongs to $\mathrm{C}(\mathcal{X}, \mathbb{R}^d)$. In that case $\omega \mapsto (\mathbf{Y}_x(\omega))_{x\in\mathcal{X}}$ is measurable and we denote $\pi \in \mathscr{P}(\mathrm{C}(\mathcal{X}, \mathbb{R}^d))$ its distribution.

[6] $\mathscr{P}((\mathbb{R}^d)^{\mathcal{X}})$ is the space of distributions over the product space $(\mathbb{R}^d)^{\mathcal{X}}$.

Therefore, a natural approximation of the process $(\mathbf{Y}_x)_{x\in\mathcal{X}}$ [7] is given by $F_M \triangleq \sum_{m=0}^{M} {\lambda_m}^{1/2} Z_m e_m$. In the previous theorem, we dissociated the *stochastic* part of the process—embedded in $\{Z_m\}_{m=0}^M$—from its *space-time* part—embedded in $\{e_m\}_{m=0}^M$. In fact, even though our approach is originally based on the Karhunen-Loève theorem, in practice we consider *arbitrary* kernels K (not necessarily based on the process $\mathbf{Y}$) using Mercer's theorem (Ferreira and Menegatto, 2009, Thm 1.1).

Proposition 2. *We denote by* π^M *the distribution of* $(\sum_{m=0}^{M} {\lambda_m}^{1/2} Z_m e_m(x))_{x\in\mathcal{X}}$ *and* $\mathrm{S} = \{\pi^M_{x_1,\dots,x_n} : (x_1,\dots,x_n) \in \mathcal{X}^n,\ n \in \mathbb{N}\}$. *Then we have that* S *is* exchangeable *and* consistent.

In order to sample from $(\sum_{m=0}^{M} {\lambda_m}^{1/2} Z_m e_m(x))_{x\in\mathcal{X}}$ one only needs to sample $\{Z_m\}_{m=0}^M$, effectively approximating the infinite dimensional modelling problem by a finite-dimensional one. We are now ready to present our method: Spectral Process Score-Based Generative Modelling (SP-SGM).

Spectral Process Score-Based Generative modelling. SP-SGM works by first choosing a kernel K and computing the associated system of orthonormal eigenfunctions $\{e_m\}_{m\in\mathbb{N}}$—see Appendix B.3. We assume we have access to a dataset $\mathcal{D} = \{\mathbf{Y}^i\}_{i=1}^L$ of *functions*[8] from the target distribution π. We set $M \in \mathbb{N}$ and compute a new $M+1$-dimensional dataset $\mathcal{D}^M = \{\{Z_m^i\}_{m=0}^M : i \in \{1,\dots,L\}\}$ where $Z_m^i = \lambda_m^{-1/2} \int_{\mathcal{X}} \langle \mathbf{Y}_x^i, e_m(x)\rangle \mathrm{d}x$. Then, we train a $M+1$-dimensional standard SGM using $\mathcal{D}^M$. The SGM model is initialised with $\{\bar{Z}_{m,0}\}_{m=0}^M \sim \mathrm{N}(0, \mathrm{Id})$ and outputs $\{\bar{Z}_m\}_{m=0}^M$ which is approximately distributed as $\{Z_m\}_{m=0}^M$. Finally, SP-SGM outputs the *functional* model $\bar{F}_M = \sum_{m=0}^{M} \lambda_m^{1/2} \bar{Z}_m e_m$. We remark that standard SGMs would only be able to deal with *finite* input space. Note that $\bar{F}_{M,0} = \sum_{m=0}^{M} \lambda_m^{1/2} \bar{Z}_{m,0} e_m$ is a Gaussian process (GP) and in particular—when $\mathrm{K} = \mathrm{K}_{\mathbf{Y}}$— it is the closest GP to the target distribution π in the following sense.

Proposition 3. *Let* π^0 *be the distribution of* $\sum_{m=0}^{+\infty} \lambda_m^{1/2} \bar{Z}_{m,0} e_m$ *and* π *the target distribution. Denote* $\mathrm{GP}(\mathcal{X})$ *the space of Gaussian processes on* $\mathcal{X}$ *and assume that* K *is the covariance kernel. Then,* $\pi^0 \in \arg\min_{\pi_{\mathrm{GP}}\in\mathrm{GP}(\mathcal{X})} \mathrm{KL}\left(\pi|\pi_{\mathrm{GP}}\right)$. *In addition,* $\sum_{m=0}^{M} \lambda_m^{1/2} \bar{Z}_{m,0} e_m$ *is the projection of* $\sum_{m=0}^{+\infty} \lambda_m^{1/2} \bar{Z}_{m,0} e_m$ *on the subspace of* $\mathrm{L}^2(\mathcal{X})$ *spanned by* $\{e_m\}_{m=0}^M$.

4 Related work

Gaussian processes and the Neural processes family. One standard and powerful framework to construct distributions over functional spaces are Gaussian processes (Rasmussen, 2003). However, they scale badly with the number of datapoints and are restricted in their modelling capacity. These problems can be partially alleviated by using neural processes (Kim et al., 2019; Garnelo et al., 2018b,a; Jha et al., 2022; Louizos et al., 2019; Singh et al., 2019), although they also assume a Gaussian likelihood and only satisfy predictive consistency, meaning one must condition on a context set to achieve consistency. Recently, denoising diffusion models have been combined with attention network architectures to build consistent models of stochastic processes (Dutordoir et al., 2022). Finally, Dupont et al. (2022) model weights of implicit neural representation using SGM.

Spectral approaches. Another possibility to circumvent the consistency problem is to consider spectral decompositions—such as the Karhunen-Loéve one— which shows that stochastic processes can be represented over a functional basis with random coefficients. Some prior works have explored this avenue using variational auto-encoders to model these coefficients (Mishra et al., 2020) whereas Lim et al. (2022) considered energy-based models. Another related line of work is 'inter-domain Gaussian processes' (Lázaro-Gredilla and Figueiras-Vidal, 2009), which rely on inducing variables that are obtained via inducing functions such as elements of a Fourier basis (Hensman et al., 2018).

Spatial structure in score-based generative models A variety of approaches have also been proposed to incorporate spatial correlation in the noising process of finite-dimensional diffusion models leveraging the multiscale structure of data (Jing et al., 2022; Guth et al., 2022; Ho et al., 2022; Saharia et al., 2021; Hoogeboom and Salimans, 2022; Rissanen et al., 2022). Our methodology can also be seen as a principled way to modify the forward dynamics in classical denoising diffusion models. Indeed by applying the diffusion in the spectral space the destruction process *blurs* the samples instead of *noising* them, see Fig. 3 for an illustration. Therefore, our contribution can be

[7] We assume that $\mathbb{E}[\|\mathbf{Y}_x\|^2] < +\infty$ and w.l.o.g. $\mathbb{E}[\mathbf{Y}_x] = 0$ as the mean can be subtracted otherwise.

[8] In practice we have access to these functions via a finite number of observations.

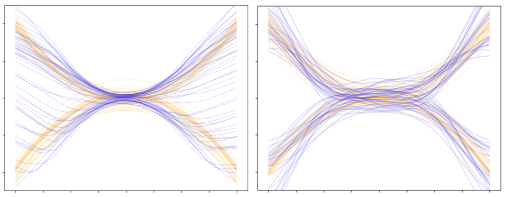

Figure 2: Samples from the Quadratic dataset (orange), and from a trained NP [Left, blue] and a trained SP-SGM [Right, blue].

	SP-SGM	NP	GP
Quadratic	$\mathbf{5.4}_{\pm \mathbf{0.7}}$	$8.6_{\pm 1.5}$	$100.0_{\pm 0.0}$
Melbourne	$\mathbf{5.3}_{\pm \mathbf{0.7}}$	$10.1_{\pm 1.9}$	$20.1_{\pm 4.0}$
Gridwatch	$\mathbf{4.7}_{\pm \mathbf{0.5}}$	$51.8_{\pm 15.1}$	$29.2_{\pm 5.5}$

Table 1: Power (percent) of a kernel two-sample hypothesis test on 1D datasets. Lower is better. Statistically significant best result is in bold.

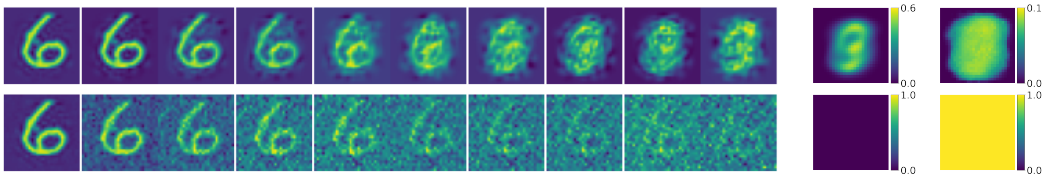

Figure 3: Forward process in SP-SGM [Top] vs standard SGM [Bottom] on MNIST digits. Pixel-wise mean and standard deviation of reference measure in rightmost columns respectively.

understood in the light of recent advances in generative modelling on cold and soft denoising diffusion models (Daras et al., 2022; Bansal et al., 2022; Hoogeboom and Salimans, 2022).

5 Experimental results

First, we assess our method's capacity to model bi-modality by qualitative comparison of samples on the synthetic Quadratic dataset. Figure 2 shows that our method is capable of expressing bi-modality while NPs are not. While SGMs alone exhibit multimodality, they only apply to finite input spaces, hence expressing *multimodality in stochastic processes* is indeed a feature of our approach. To provide a quantitative assessment we compute the power of a two-sample hypothesis test for functional data (Wynne and Duncan, 2022) which tries to discriminate between samples from the model and the dataset. Table 1 shows that our method outperforms GPs and NPs.

Then, we illustrate the behaviour of our proposed parametrisation of stochastic processes. We plot the forward noising process of both SP-SGM (reconstructed in the original image space) and of a standard SGM in Figure 3. We see that by performing the diffusion in the spectral space, we capture and incorporate spatial correlations governed by the spatial component of the Karhunen-Loève theorem.

Finally we conduct ablation studies on the MNIST dataset. We consider the effect of the truncation order M in Figure 5, where we find best performance for intermediate values of M. We decompose this effect into the reconstruction between spectral and spatial domain, which improves with M, and the quality of the stochastic model in the spectral domain, which decreases with M. These conflicting objectives are also visualised in Figure 5. We also assess different kernels in Figure 4. Qualitatively, eigenfunctions of the covariance kernel capture spatial correlations effectively and thus generates the best samples, while the RBF kernel imposes a smoothness which tends to produce blurry samples.

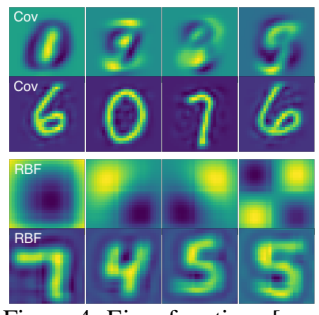

Figure 4: Eigenfunctions [row 1, 3] and samples from SP-SGM [row 2, 4] for covariance and RBF kernels.

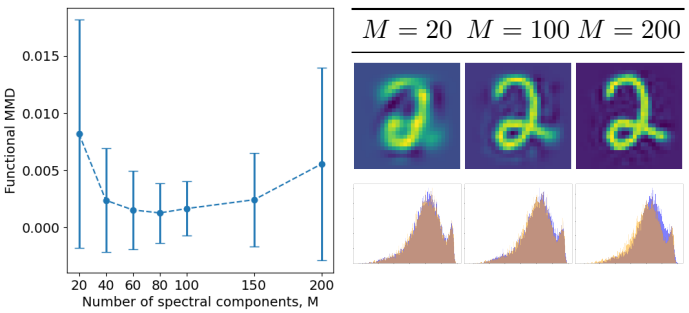

Figure 5: On the left, functional MMD vs M. Lower is better. On the right, Karhunen-Loève recompositions of MNIST samples [top row] and the distribution of the first spectral component from the dataset [orange] and from the SGM in spectral space [blue].

Acknowledgements

We thank the anonymous reviewers for their insightful and constructive comments. AP and MH are funded through the StatML CDT through grant EP/S023151/1. AD acknowledges support of the UK Defence Science and Technology Laboratory (Dstl) and Engineering and Physical Research Council (EPSRC) under grant EP/R013616/1. This is part of the collaboration between US DOD, UK MOD and UK EPSRC under the Multidisciplinary University Research Initiative. AD is also partially supported by the EPSRC grant EP/R034710/1 CoSines. EM is supported by an EPSRC Prosperity Partnership EP/T005386/1 between Microsoft Research and the University of Cambridge.

References

Ambrosio, L., Gigli, N., and Savaré, G. (2008). *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, second edition.

Baker, C. T. H. (1979). Numerical Integration in the Treatment of Integral Equations. In Hämmerlin, G., editor, *Numerische Integration: Tagung Im Mathematischen Forschungsinstitut Oberwolfach Vom 1. Bis 7. Oktober 1978*, pages 44–53. Birkhäuser Basel, Basel.

Bansal, A., Borgnia, E., Chu, H.-M., Li, J. S., Kazemi, H., Huang, F., Goldblum, M., Geiping, J., and Goldstein, T. (2022). Cold diffusion: Inverting arbitrary image transforms without noise. *arXiv preprint arXiv:2208.09392*.

Cattiaux, P., Conforti, G., Gentil, I., and Léonard, C. (2021). Time reversal of diffusion processes under a finite entropy condition. *arXiv preprint arXiv:2104.07708*.

Charalambos, D. and Aliprantis, B. (2013). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag Berlin and Heidelberg GmbH & Company KG.

Daras, G., Delbracio, M., Talebi, H., Dimakis, A. G., and Milanfar, P. (2022). Soft diffusion: Score matching for general corruptions. *arXiv preprint arXiv:2209.05442*.

Dhariwal, P. and Nichol, A. (2021). Diffusion models beat GAN on image synthesis. *arXiv preprint arXiv:2105.05233*.

Dupont, E., Kim, H., Eslami, S., Rezende, D., and Rosenbaum, D. (2022). From data to functa: Your data point is a function and you should treat it like one. *arXiv preprint arXiv:2201.12204*.

Dupuis, P. and Ellis, R. S. (2011). *A weak convergence approach to the theory of large deviations*. John Wiley & Sons.

Dutordoir, V., Saul, A., Ghahramani, Z., and Simpson, F. (2022). Neural diffusion processes. *arXiv preprint arXiv:2206.03992*.

Fernique, X. (1975). Regularité des trajectoires des fonctions aléatoires gaussiennes. In *Ecole d'Eté de Probabilités de Saint-Flour IV—1974*, pages 1–96. Springer.

Ferreira, J. and Menegatto, V. (2009). Eigenvalues of integral operators defined by smooth positive definite kernels. *Integral Equations and Operator Theory*, 64(1):61–81.

Garnelo, M., Rosenbaum, D., Maddison, C., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y. W., Rezende, D., and Eslami, S. M. A. (2018a). Conditional neural processes. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1704–1713. PMLR.

Garnelo, M., Schwarz, J., Rosenbaum, D., Viola, F., Rezende, D. J., Eslami, S., and Teh, Y. W. (2018b). Neural processes. *arXiv preprint arXiv:1807.01622*.

Guth, F., Coste, S., De Bortoli, V., and Mallat, S. (2022). Wavelet score-based generative modeling. *arXiv preprint arXiv:2208.05003*.

Haussmann, U. G. and Pardoux, E. (1986). Time reversal of diffusions. *The Annals of Probability*, 14(4):1188–1205.

Hensman, J., Durrande, N., and Solin, A. (2018). Variational fourier features for gaussian processes. *Journal of Machine Learning Research*, Volume 18, issue 1:52.

Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*.

Ho, J., Saharia, C., Chan, W., Fleet, D. J., Norouzi, M., and Salimans, T. (2022). Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.*, 23:47–1.

Hoogeboom, E. and Salimans, T. (2022). Blurring diffusion models. *arXiv preprint arXiv:2209.05557*.

Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4).

Jha, S., Gong, D., Wang, X., Turner, R. E., and Yao, L. (2022). The neural process family: Survey, applications and perspectives. *arXiv preprint arXiv:2209.00517*.

Jing, B., Corso, G., Berlinghieri, R., and Jaakkola, T. (2022). Subspace diffusion generative models. *arXiv preprint arXiv:2205.01490*.

Kim, H., Mnih, A., Schwarz, J., Garnelo, M., Eslami, A., Rosenbaum, D., Vinyals, O., and Teh, Y. W. (2019). Attentive neural processes. In *International Conference on Learning Representations*.

Lázaro-Gredilla, M. and Figueiras-Vidal, A. (2009). Inter-domain gaussian processes for sparse inference using inducing features. In Bengio, Y., Schuurmans, D., Lafferty, J., Williams, C., and Culotta, A., editors, *Advances in Neural Information Processing Systems*, volume 22. Curran Associates, Inc.

Lim, J. N., Vollmer, S., Wolf, L., and Duncan, A. (2022). $\mathcal{F}$-ebm: Energy based learning of functional data. *arXiv preprint arXiv:2202.01929*.

Louizos, C., Shi, X., Schutte, K., and Welling, M. (2019). The functional neural process. *Advances in Neural Information Processing Systems*, 32.

Mishra, S., Flaxman, S., Berah, T., Pakkanen, M., Zhu, H., and Bhatt, S. (2020). πvae: Encoding stochastic process priors with variational autoencoders. *arXiv preprint arXiv:2002.06873*.

Rasmussen, C. E. (2003). Gaussian processes in machine learning. In *Summer school on machine learning*, pages 63–71. Springer.

Ravuri, S., Lenc, K., Willson, M., Kangin, D., Lam, R., Mirowski, P., Fitzsimons, M., Athanassiadou, M., Kashem, S., Madge, S., Prudden, R., Mandhane, A., Clark, A., Brock, A., Simonyan, K., Hadsell, R., Robinson, N., Clancy, E., Arribas, A., and Mohamed, S. (2021). Skilful precipitation nowcasting using deep generative models of radar. *Nature*, 597(7878):672–677.

Rissanen, S., Heinonen, M., and Solin, A. (2022). Generative Modelling With Inverse Heat Dissipation.

Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D. J., and Norouzi, M. (2021). Image super-resolution via iterative refinement. *arXiv preprint arXiv:2104.07636*.

Singh, G., Yoon, J., Son, Y., and Ahn, S. (2019). Sequential neural processes. *Advances in Neural Information Processing Systems*, 32.

Song, Y. and Ermon, S. (2019). Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*.

Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. (2021). Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*.

Sun, S., Zhang, G., Shi, J., and Grosse, R. (2019). Functional variational bayesian neural networks. *arXiv preprint arXiv:1903.05779*.

Vincent, P. (2011). A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674.

Wynne, G. and Duncan, A. B. (2022). A kernel two-sample test for functional data. *Journal of Machine Learning Research*, 23(73):1–51.

Yang, L., Zhang, D., and Karniadakis, G. E. (2020). Physics-informed generative adversarial networks for stochastic differential equations. *SIAM Journal on Scientific Computing*, 42(1):A292–A317.

A Introduction to the supplementary material

We provide details on the implementation of SP-SGM in Appendix B. The proof of Proposition 2 is presented in Appendix D. The proof of Proposition 3 is presented in Appendix E. The likelihood computation is given in Appendix C. Experimental details are provided in Appendix F.

B SP-SGM in practice

B.1 SP-SGM algorithm

Algorithm 1 Spectral Process Score-Based Generative Model (SP-SGM)

Require: $T, \mathcal{D}, \theta_0, N_{\text{iter}}, \varepsilon, K, \{(\lambda_m, e_m)\}_{m=0}^M$
/// TRAINING ///
Get $\mathcal{D}^M$ from $\mathcal{D}$ ▷ Dataset projection
for $n \in \{0, \dots, N_{\text{iter}} - 1\}$ **do**
  Get $\{\mathbf{Y}_{m,0}\}_{m=0}^M$ mini-batch from $\mathcal{D}^M$
  $t \sim U([\varepsilon, T])$ ▷ Uniform sampling between ε and T
  $\mathbf{Y}_{m,t} = \mathrm{e}^{-t}\mathbf{Y}_{m,0} + (1 - \mathrm{e}^{-2t})^{1/2} G, \quad G \sim \mathrm{N}(0, \mathrm{Id})$ ▷ Diffuse
  Get DSM loss $\ell(\theta_n)$ ▷ Compute score matching loss
  $\theta_{n+1} = \texttt{optimiser\_update}(\theta_n, \ell(\theta_n))$ ▷ ADAM optimiser step
end for
$\theta^\star = \theta_{N_{\text{iter}}}$
/// SAMPLING ///
$\{\bar{Y}_{m,0}\}_{m=0}^M \sim \mathrm{N}(0, \mathrm{Id})$ ▷ Sample from Gaussian distribution
$b_\theta^\star(t, y) = \mathbf{s}_{\theta^\star}(T - t, \bar{y})$ for any $t \in [0, T]$, $\bar{y} \in \mathbb{R}^{M+1}$ ▷ Reverse process drift
$\{\bar{Y}_{m,n}\}_{m=0,n=0}^{M,N}$ Euler-Maruyama with drift $b_\theta^\star$ ▷ Approximate reverse diffusion
return $\theta^\star, x \mapsto \sum_{m=0}^M \lambda_m^{1/2} \bar{Y}_{m,N} e_m(x)$

B.2 Karhunen-Loève coefficients

In this section we show that the random coefficients from the KL basis expansion, given for any $m \in \mathbb{N}$ by

$$Z_m = \langle \mathbf{Y}, \lambda_m^{-1/2} e_m \rangle_{\mathrm{L}^2(\mathcal{X})} = \int_{\mathcal{X}} \langle (\mathbf{Y}(x) - \mu(x)), \lambda_m^{-1/2} e_m(x) \rangle \mathrm{d}x,$$

satisfy $\mathbb{E}[Z_m] = 0$ and furthermore if $\mathrm{K} = \mathrm{K}_{\mathbf{Y}}$ then $\mathbb{E}[Z_m Z_{m'}] = \delta_{m,m'}$. First, we have that for any $m \in \mathbb{N}$

$$\mathbb{E}[Z_m] = \int_{\mathcal{X}} \langle \mathbb{E}[\mathbf{Y}(x) - \mu(x)], \lambda_m^{-1/2} e_m(x) \rangle \mathrm{d}x = 0.$$

Then, we have for any $m, m' \in \mathbb{N}$

$$\begin{aligned}
\mathbb{E}[Z_m Z_{m'}] &= \mathbb{E}[\textstyle\int_{\mathcal{X}} \langle (\mathbf{Y}(x) - \mu(x)), \lambda_m^{-1/2} e_m(x) \rangle \mathrm{d}x \int_{\mathcal{X}} \langle (\mathbf{Y}(x') - \mu(x')), \lambda_{m'}^{-1/2} e_{m'}(x') \rangle \mathrm{d}x'] \\
&= \lambda_m^{-1/2} \lambda_{m'}^{-1/2} \textstyle\int_{\mathcal{X}} \int_{\mathcal{X}} \mathrm{Tr}(\mathbb{E}[(\mathbf{Y}(x) - \mu(x))(\mathbf{Y}(x') - \mu(x'))^\top] e_m(x) e_{m'}^\top(x')) \mathrm{d}x \mathrm{d}x' \\
&= \lambda_m^{-1/2} \lambda_{m'}^{-1/2} \textstyle\int_{\mathcal{X}} e_m(x)^\top \mathrm{K}_{\mathbf{Y}}(x, x') e_{m'}(x') \mathrm{d}x \mathrm{d}x' \\
&= \lambda_m^{-1/2} \lambda_{m'}^{-1/2} \textstyle\int_{\mathcal{X}} \lambda_m \langle e_m(x'), e_{m'}(x') \rangle \mathrm{d}x' = \lambda_m^{-1/2} \lambda_{m'}^{-1/2} \lambda_m \delta_{m,m'} = \delta_{m,m'}.
\end{aligned}$$

B.3 Computation of eigenfunctions

We focus on the case where $\mathbf{Y}$ takes values in $\mathbb{R}$ but the computations can be extended to the case where $\mathbf{Y}$ takes values in $\mathbb{R}^d$ with $d \in \mathbb{N}$. Recall from Theorem 1 the $\mathrm{L}^2(\mathcal{X})$ operator T_{K} associated with kernel K, given by $\mathrm{T}_{\mathrm{K}} f(x_2) = \int_{\mathcal{X}} \mathrm{K}(x_1, x_2) f(x_1) \mathrm{d}x_1$. In the following we describe how we approximate the orthonormal system of eigenfunctions $\{e_m\}_{m \in \mathbb{N}}$ of T_{K} from a dataset $\mathcal{D} = \{\mathbf{Y}^i\}_{i=1}^L$.

Firstly, we note that in practice we have access to the functions in our dataset at a finite set of query points, that is $\mathcal{D} = \{\{\mathbf{Y}^i_{x_{i,n}}\}_{n=1}^{N_i}\}_{i=1}^L$. For ease of notation, in the following we denote $\{x_s\}_{s=1}^S = \{x_{i,n} : i \in \{1, \dots, L\},\ n \in \{1, \dots N_i\}\}$, with $s = \sum_{i=1}^L N_i$.

The eigensystem $\{e_m\}_{m\in\mathbb{N}}$ satisfies $\mathrm{T_K} e_m(x) = \lambda_m e_m(x)$ for all $x \in \mathcal{X}$, which we approximate by

$$\lambda_m e_m(x) = \mathrm{T_K} e_m(x) \approx \tfrac{1}{S} \sum_{s=1}^{S} \mathrm{K}(x, x_s) e_m(x_s). \tag{3}$$

Evaluating the above at the entire set of evaluation points $x = \{x_s\}_{s=1}^S$ results in the matrix eigenproblem

$$\tfrac{1}{S} K u_m = \lambda_m^{\text{mat}} u_m,$$

where K is an $S \times S$ matrix and has entries $K_{i,j} = \mathrm{K}(x_i, x_j)$, $\lambda_m^{\text{mat}} \in \mathbb{R}$ is the matrix eigenvalue and $u_m \in \mathbb{R}^S$ is the corresponding matrix eigenvector, normalised such that $u_m^\top u_m = 1$. By Baker's theorem (Baker, 1979, Theorem 3.4), $\lambda_m^{\text{mat}} \to \lambda_m$ as $S \to \infty$. Note that if we take $e_m(x_s) = (u_m)_s$ then we have

$$\textstyle\int_{\mathcal{X}} e_m(x) e_m(x) p(x) \mathrm{d}x \approx \tfrac{1}{S} \sum_{s=1}^{S} e_m(x_s) e_m(x_s) = \tfrac{1}{S} u_m^\top u_m = \tfrac{1}{S}$$

where $p(x)$ is the density of observation points on $\mathcal{X}$. So, in order to have an orthonormal eigensystem $\int_{\mathcal{X}} e_m(x) e_{m'}(x) p(x) \mathrm{d}x = \delta_{mm'}$ we must rescale according to $e_m(x_s) = \sqrt{S}(u_m)_s$. Finally, rearranging (3) we get

$$\hat{e}_j(x) \approx (S\lambda_j^{\text{mat}})^{-1} \textstyle\sum_{s=1}^{S} k(x, x_s) e_j(x_s) \approx (\sqrt{S}\lambda_j^{\text{mat}})^{-1} \sum_{s=1}^{S} k(x, x_s) u_j(x_s).$$

Having obtained an approximate eigensystem as above, it is simple to obtain the spectral dataset $\mathcal{D}^M = \{\{Z_m^i\}_{m=0}^M : i \in \{1, \ldots, L\}\}$ using

$$Z_m^i \approx \tfrac{1}{N_i} \textstyle\sum_{n=1}^{N_i} (\mathbf{Y}_{x_n} - \mu(x_n))(\lambda_m^{\text{mat}})^{-1/2} \hat{e}_m(x_n).$$

where $\mu(x) = \mathbb{E}[\mathbf{Y}_x]$ is the process mean.

B.4 Truncation order

For a given approximate eigensystem $\{(\hat{\lambda}_m, \hat{e}_m)\}_{m=1}^L$, we define the truncation order M as

$$M = \arg\min_{m\in\mathbb{N}} \textstyle\sum_{j=1}^{m} \hat{\lambda}_j / \sum_{j=1}^{L} \hat{\lambda}_j \geq \eta,$$

where $\eta \in [0, 1]$ is a threshold parameter, fixed to $\eta = 0.99$ for our experimental results on the 1D datasets. Alternatively, as we do on the MNIST dataset, one can treat the truncation order M as an hyperparameter.

B.5 Eigenvalue weighted objective

We propose the following weighted DSM loss to train the score network $\boldsymbol{s}_\theta : [0, T] \times \mathbb{R}^{M+1} \to \mathbb{R}^{M+1}$:

$$\mathbb{E}_t \left\{ \mathbb{E}_{\mathbf{Z}_0, \mathbf{Z}_t} \left[\left\| \tilde{\lambda}^\alpha \odot (\boldsymbol{s}_\theta(\mathbf{Z}_t, t) - \nabla_{\mathbf{Z}_t} \log p(\mathbf{Z}_t | \mathbf{Z}_0)) \right\|_2^2 \right] \right\}, \tag{4}$$

with $\tilde{\lambda} = [\lambda_0/\Lambda, \ldots, \lambda_M/\Lambda]^\top$[9], $\Lambda = \sum_{m=0}^M \lambda_m$, the vector of normalised eigenvalues, $\alpha > 0$ a tuneable hyperparameter, $t \sim U([0, T])$, $\mathbf{Z}_0 = \{\lambda_m^{-1/2} \int_{\mathcal{X}} \langle \mathbf{Y}_x, e_m(x) \rangle \mathrm{d}x\}_{m=0}^M$ and

$$\mathrm{d}\mathbf{Z}_t = -(1/2)\mathbf{Z}_t \mathrm{d}t + \beta_t \mathrm{d}\mathbf{B}_t,$$

where $t \mapsto \beta_t$ is an hyperparameter. The motivation behind (4) is to put more importance on fitting well the lower frequency (i.e. higher eigenvalue) components of the decomposition $\sum_{m=0}^M {\lambda_m}^{1/2} Z_m e_m(x)$ since these matter the most in terms of quality of reconstruction as per Theorem 1. With $\alpha = 0$, we have that $\lambda^\alpha = \mathbf{1}$, thus that all dimensions of the spectral space are weighted equally and we recover the standard DSM loss.

C Likelihood evaluation

Let's denote the Karhunen-Loeve expansion with $\phi : \mathbb{R}^{M+1} \to \mathbb{R}^d$ such that $(Z_1, \ldots, Z_M) \mapsto \mu(\cdot) + \sum_{m=0}^M Z_m \sqrt{\lambda_m} e_m(\cdot)$. We have that $\frac{\partial \phi}{\partial Z_m} = \sqrt{\lambda_m} e_m$, hence with $J_z \phi = \left(\frac{\partial \phi}{\partial z_1}, \ldots, \frac{\partial \phi}{\partial z_N} \right)$ we get $J_z \phi^\top J_z \phi = \left(\sqrt{\lambda_m} \sqrt{\lambda_n} \langle e_m, e_n \rangle \right)_{mn} = \left(\sqrt{\lambda_m} \sqrt{\lambda_n} \delta_{mn} \right)_{mn} = \mathrm{diag}(\lambda_m)$. Hence, denoting a realisation $F = \phi(\mathbf{Z})$, applying the change of variable formula we have

$$\log p(F) = \textstyle\sum_m \log p(Z_m) - \log |J_z \phi^\top J_z \phi|^{1/2} = \sum_m \left(\log p(Z_m) - 1/2 \log \lambda_m \right).$$

[9] The exponentiation in (4) is to be understood in a pointwise fashion.

D Proof of Proposition 2

We recall Proposition 2.

Proposition 2. *We denote by* π^M *the distribution of* $(\sum_{m=0}^{M} {\lambda_m}^{1/2} Z_m e_m(x))_{x\in\mathcal{X}}$ *and* $\mathrm{S} = \{\pi^M_{x_1,\dots,x_n} : (x_1,\dots,x_n) \in \mathcal{X}^n,\ n \in \mathbb{N}\}$. *Then we have that* S *is* exchangeable *and* consistent.

Proof. Note that this result is the trivial direction of the Kolmogorov extension theorem (Charalambos and Aliprantis, 2013, Theorem 15.26) since F_M is by definition a stochastic process. For completeness, we provide a proof of this result in our setting. We denote $Z = \{Z_m\}_{m=0}^M$ and π the model distribution on Z. For any $n \in \mathbb{N}$ and $f \in \mathrm{C}_b((\mathbb{R}^d)^n, \mathbb{R})$ we have

$$\begin{aligned}\int f(y_1,\dots,y_n)\mathrm{d}\pi^M_{x_1,\dots,x_n}(y_1,\dots,y_n) &= \int\int_{\mathbb{R}^{M+1}} f(y_1,\dots,y_n)\mathrm{d}\pi^M_{x_1,\dots,x_n}(y_1,\dots,y_n,Z=z)\mathrm{d}z\\ &= \int\int_{\mathbb{R}^{M+1}} f(y_1,\dots,y_n)\mathrm{d}\pi^M_{x_1,\dots,x_n}(y_1,\dots,y_n,|Z=z)\mathrm{d}\pi(z)\\ &= \int\int_{\mathbb{R}^{M+1}} f(y_1,\dots,y_n)\textstyle\prod_{i=1}^n \mathrm{d}\pi^M_{x_i}(y_i|Z=z)\mathrm{d}\pi(z).\end{aligned}$$

Hence, for any $n \in \mathbb{N}$, n-permutation $\sigma \in \mathfrak{S}_n$, $x_1,\dots,x_n \in \mathcal{X}$, and continuous and bounded function $f \in \mathrm{C}_b((\mathbb{R}^d)^n)$ we have

$$\int f(y_{\sigma^{-1}(1)},\dots,y_{\sigma^{-1}(n)})\,\mathrm{d}\pi^M_{x_{\sigma(1)},\dots,x_{\sigma(n)}}(y_1,\dots,y_n) = \int f(y_1,\dots,y_n)\,\mathrm{d}\pi^M_{x_1,\dots,x_n}(y_1,\dots,y_n).$$

Similarly, for any $n_1 \le n_2$, $x_1,\dots,x_{n_2} \in \mathcal{X}$ and $f \in \mathrm{C}_b((\mathbb{R}^d)^{n_1})$ we have

$$\begin{aligned}&\int f(y_1,\dots,y_{n_1})\,\mathrm{d}\pi^M_{x_1,\dots,x_{n_2}}(y_1,\dots,y_{n_2})\\ &\quad= \int f(y_1,\dots,y_{n_1})\int_{\mathbb{R}^{M+1}}\mathrm{d}\pi^M_{x_1,\dots,x_{n_2}}(y_1,\dots,y_{n_2}|Z=z)\mathrm{d}\pi(z)\\ &\quad= \int f(y_1,\dots,y_{n_1})\int_{\mathbb{R}^{M+1}}\textstyle\prod_{i=1}^{n_2}\mathrm{d}\pi^M_{x_i}(y_i|Z=z)\mathrm{d}\pi(z)\\ &\quad= \int f(y_1,\dots,y_{n_1})\int_{\mathbb{R}^{M+1}}\textstyle\prod_{i=1}^{n_1}\mathrm{d}\pi^M_{x_i}(y_i|Z=z)\left(\prod_{i=n_1+1}^{n_2}\mathrm{d}\pi^M_{x_i}(y_i|Z=z)\right)\mathrm{d}\pi(z)\\ &\quad= \int f(y_1,\dots,y_{n_1})\int_{\mathbb{R}^{M+1}}\textstyle\prod_{i=1}^{n_1}\mathrm{d}\pi^M_{x_i}(y_i|Z=z)\mathrm{d}\pi(z)\\ &\quad= \int f(y_1,\dots,y_{n_1})\,\mathrm{d}\pi^M_{x_1,\dots,x_{n_1}}(y_1,\dots,y_{n_1}).\end{aligned}$$

In this setting, the random variable Z is finite dimensional as it is supported in $\mathbb{R}^M$, but more generally this result is still true with an (infinite dimensional) process F. □

E Proof of Proposition 3

In this section, we provide two proofs of Proposition 3.

E.1 Partition definition

We recall Proposition 3.

Proposition 4. *Let* π^0 *be the distribution of* $\sum_{m=0}^{+\infty}\lambda_m^{1/2}\bar{Z}_{m,0}e_m$ *and* π *the target distribution. Denote* $\mathrm{GP}(\mathcal{X})$ *the space of Gaussian processes on* $\mathcal{X}$ *and assume that* K *is the covariance kernel. Then,* $\pi^0 \in \arg\min_{\pi_{\mathrm{GP}}\in\mathrm{GP}(\mathcal{X})}\mathrm{KL}\left(\pi|\pi_{\mathrm{GP}}\right)$. *In addition,* $\sum_{m=0}^{M}\lambda_m^{1/2}\bar{Z}_{m,0}e_m$ *is the projection of* $\sum_{m=0}^{+\infty}\lambda_m^{1/2}\bar{Z}_{m,0}e_m$ *on the subspace of* $\mathrm{L}^2(\mathcal{X})$ *spanned by* $\{e_m\}_{m=0}^M$.

Note that in this first approach we do not assume any regularity on the samples of the process. This proof is based on (Sun et al., 2019, Theorem 1) which considered a modified definition of the Kullback-Leibler divergence. Namely, for any μ,ν probability measures over a probability space $(\Omega,\mathcal{F})$ we define

$$\mathrm{KL}\left(\mu|\nu\right) = \sup\{\mathrm{KL}\left(\mu_P|\nu_P\right) : P \text{ measurable finite partition of } \Omega\},$$

where for any measurable finite partition $P = \{\Omega_1,\dots,\Omega_N\}$ we define $\mu_P = \{\mu(\Omega_1),\dots,\mu(\Omega_N)\}$, similarly for ν_P. With that definition (Sun et al., 2019, Theorem 1) holds. Namely, for any $\mu,\nu \in \mathscr{P}((\mathbb{R}^d)^{\mathcal{X}})$[10], we have

$$\mathrm{KL}\left(\mu|\nu\right) = \sup\{\mathrm{KL}\left(\mu_{x_1,\dots,x_n}|\nu_{x_1,\dots,x_n}\right) : (x_1,\dots,x_n) \in \mathcal{X}^n,\ n \in \mathbb{N}\}. \tag{5}$$

[10]The product space $(\mathbb{R}^d)^{\mathcal{X}}$ is a measurable space with the cylindrical sigma algebra, see Sun et al. (2019) for more details in that context.

Let π be the target distribution of the stochastic process. We define $\mathrm{GP}(\mathcal{X})$ the space of Gaussian processes over the input space $\mathcal{X}$. We have that

$$\inf\{\sup\{\mathrm{KL}\left(\pi_{x_1,\dots,x_n}|\pi^0_{x_1,\dots,x_n}\right) : (x_1,\dots,x_n)\in\mathcal{X}^n,\ n\in\mathbb{N}\} : \pi^0\in\mathrm{GP}(\mathcal{X})\}$$
$$\geq \sup\{\inf\{\mathrm{KL}\left(\pi_{x_1,\dots,x_n}|\pi^0_{x_1,\dots,x_n}\right) : \pi^0\in\mathrm{GP}(\mathcal{X})\} : (x_1,\dots,x_n)\in\mathcal{X}^n,\ n\in\mathbb{N}\}.$$

In addition, we have

$$\sup\{\inf\{\mathrm{KL}\left(\pi_{x_1,\dots,x_n}|\pi^0_{x_1,\dots,x_n}\right) : \pi^0\in\mathrm{GP}(\mathcal{X})\} : (x_1,\dots,x_n)\in\mathcal{X}^n,\ n\in\mathbb{N}\}$$
$$= \sup\{\inf\{\mathrm{KL}\left(\pi_{x_1,\dots,x_n}|\pi^0_{x_1,\dots,x_n}\right) : \pi^0_{x_1,\dots,x_n}\in\mathrm{GP}(\{x_1,\dots,x_n\})\} : (x_1,\dots,x_n)\in\mathcal{X}^n,\ n\in\mathbb{N}\},$$

where we emphasise that the set $\mathrm{GP}(\{x_1,\dots,x_n\})$ is simply the set of nd-dimensional Gaussian probability measures. Hence, for any $\{x_1,\dots,x_n\}\in\mathcal{X}^n$ and $n\in\mathbb{N}$, we have that

$$\mathrm{KL}\left(\pi_{x_1,\dots,x_n}|\pi^{0,\star}_{x_1,\dots,x_n}\right) = \inf\{\mathrm{KL}\left(\pi_{x_1,\dots,x_n}|\pi^0_{x_1,\dots,x_n}\right) : \pi^0_{x_1,\dots,x_n}\in\mathrm{GP}(\{x_1,\dots,x_n\})\},$$

where $\pi^{0,\star}_{x_1,\dots,x_n}$ is the Gaussian measure with same mean and covariance matrix as $\pi_{x_1,\dots,x_n}$. Using the Kolmogorov extension theorem (Charalambos and Aliprantis, 2013, Theorem 15.26), there exists $\pi^\star\in\mathscr{P}((\mathbb{R}^d)^{\mathcal{X}})$ such that for any $\{x_1,\dots,x_n\}\in\mathcal{X}^n$ and $n\in\mathbb{N}$, $\pi^\star_{x_1,\dots,x_n} = \pi^{0,\star}_{x_1,\dots,x_n}$. Therefore, we have that

$$\mathrm{KL}\left(\pi_{x_1,\dots,x_n}|\pi^\star_{x_1,\dots,x_n}\right) = \inf\{\mathrm{KL}\left(\pi_{x_1,\dots,x_n}|\pi^0_{x_1,\dots,x_n}\right) : \pi_{x_1,\dots,x_n}\in\mathrm{GP}(\mathcal{X})\}.$$

Hence, using (6) we get that

$$\sup\{\inf\{\mathrm{KL}\left(\pi_{x_1,\dots,x_n}|\pi^0_{x_1,\dots,x_n}\right) : \pi^0\in\mathrm{GP}(\mathcal{X})\} : (x_1,\dots,x_n)\in\mathcal{X}^n,\ n\in\mathbb{N}\}$$
$$= \sup\{\mathrm{KL}\left(\pi_{x_1,\dots,x_n}|\pi^\star_{x_1,\dots,x_n}\right) : \{x_1,\dots,x_n\}\in\mathcal{X}^n\} = \mathrm{KL}\left(\pi|\pi^\star\right).$$

Therefore, we have that

$$\inf\{\sup\{\mathrm{KL}\left(\pi_{x_1,\dots,x_n}|\pi^0_{x_1,\dots,x_n}\right) : (x_1,\dots,x_n)\in\mathcal{X}^n,\ n\in\mathbb{N}\} : \pi^0\in\mathrm{GP}(\mathcal{X})\} \geq \mathrm{KL}\left(\pi|\pi^\star\right),$$

which implies using (6) that

$$\mathrm{KL}\left(\pi|\pi^\star\right) \leq \inf\{\mathrm{KL}\left(\pi|\pi^0\right) : \pi^0\in\mathrm{GP}(\mathcal{X})\}.$$

The equality holds since $\pi^\star\in\mathrm{GP}(\mathcal{X})$. Finally, since $\pi^\star$ and π have the same covariance kernels, they have the same Karhunen-Loève eigensystems. Therefore, there exists $\{Z_m\}_{m\in\mathbb{N}}$ such that $\mathbf{Y}^\star = \sum_{m\in\mathbb{N}}\lambda_m^{1/2}Z_m e_m$ has distribution $\pi^\star$, which concludes the proof.

E.2 Sample continuous

In our second approach, we restrict ourselves to the case of sample continuous processes. Namely, we now longer consider $\mu\in\mathscr{P}((\mathbb{R}^d)^{\mathcal{X}})$ but $\mu\in\mathscr{P}(\mathrm{C}(\mathcal{X},\mathbb{R}^d))$.

Proposition 5. *Assume that $\pi\in\mathrm{C}(\mathcal{X},\mathbb{R}^d)$ and that there exists $\phi:\ [0,+\infty)$ such that for any $x_1,x_2\in\mathcal{X}$,*

$$\mathbb{E}[\|\mathbf{Y}_{x_1}-\mathbf{Y}_{x_2}\|^2]\leq\phi(\|x_1-x_2\|),$$

such that $\int_0^{+\infty}\phi(\exp[-t^2])\mathrm{d}t<+\infty$. Let π^0 be the distribution of $\sum_{m=0}^{+\infty}\lambda_m^{1/2}\bar{Z}_{m,0}e_m$. Denote $\mathrm{GP}(\mathcal{X})$ the space of Gaussian processes on $\mathcal{X}$ and assume that K is the covariance kernel. Then, $\pi^0\in\arg\min_{\pi_{\mathrm{GP}}\in\mathrm{GP}(\mathcal{X})}\mathrm{KL}\left(\pi|\pi_{\mathrm{GP}}\right)$. In addition, $\sum_{m=0}^{M}\lambda_m^{1/2}\bar{Z}_{m,0}e_m$ is the projection of $\sum_{m=0}^{+\infty}\lambda_m^{1/2}\bar{Z}_{m,0}e_m$ on the subspace of $\mathrm{L}^2(\mathcal{X})$ spanned by $\{e_m\}_{m=0}^M$.

In that case, we show that

$$\mathrm{KL}\left(\mu|\nu\right) = \sup\{\mathrm{KL}\left(\mu_{x_1,\dots,x_n}|\nu_{x_1,\dots,x_n}\right) : (x_1,\dots,x_n)\in\mathcal{X}^n,\ n\in\mathbb{N}\}. \tag{6}$$

More precisely, we have the following proposition.

Proposition 6. *Let $\mu,\nu\in\mathscr{P}(\mathrm{C}(\mathcal{X},\mathbb{R}^d))$*

$$\mathrm{KL}\left(\mu|\nu\right) = \sup\{\mathrm{KL}\left(\mu_{x_1,\dots,x_n}|\nu_{x_1,\dots,x_n}\right) : (x_1,\dots,x_n)\in\mathcal{X}^n,\ n\in\mathbb{N}\}.$$

Proof. First, note that for any $\{x_1, \ldots, x_n\} \in \mathcal{X}^n$, $n \in \mathbb{N}$ we have using the data processing theorem (Ambrosio et al., 2008, Lemma 9.4.5)

$$\mathrm{KL}\left(\mu|\nu\right) \geq \mathrm{KL}\left(\mu_{x_1,\ldots,x_n}|\nu_{x_1,\ldots,x_n}\right). \tag{7}$$

Since $\mathcal{X}$ is compact, for any $n \in \mathbb{N}$, there exists $\{x_1, \ldots, x_n\}$ such that $\mathcal{X} \subset \cup_{k=1}^n \mathrm{B}(x_k, 1/(n+1))$. For any $n \in \mathbb{N}$, let $\{\varphi_k\}_{k=1}^n$ be a smooth partition of unity associated with $\{\mathrm{B}(x_k, 1/(n+1))\}_{k=1}^n$. For any $\mathrm{C}(\mathcal{X}, \mathbb{R}^d)$-valued random variable $\mathbf{Y}$ and $n \in \mathbb{N}$, denote $\mathbf{Y}^n$ such that for any $x \in \mathcal{X}$

$$\mathbf{Y}_x^n = \textstyle\sum_{k=1}^n \varphi_k(x) \mathbf{Y}_{x_k}.$$

Since $\mathbf{Y}$ is continuous and $\mathcal{X}$ is compact we have that $\mathbf{Y}$ is uniformly continuous. Hence, for any $\varepsilon > 0$, there exists $n \in \mathbb{N}$ such that for any x_1, x_2, $\|x_1 - x_2\| \leq 1/(n+1)$, $\|\mathbf{Y}_{x_1} - \mathbf{Y}_{x_2}\| \leq \varepsilon$. Therefore, we have that

$$\|\mathbf{Y}_x^n - \mathbf{Y}_x\| \leq \textstyle\sum_{k=1}^n \varphi_k(x) \|\mathbf{Y}_{x_k} - \mathbf{Y}_x\| \leq \varepsilon.$$

Therefore, we have that $\lim_{n \to +\infty} \sup_{x \in \mathcal{X}} \|\mathbf{Y}_x^n - \mathbf{Y}_x\| = 0$. Therefore, for any $n \in \mathbb{N}$, denoting μ_n the distribution of $\mathbf{Y}^n$, we get that $(\mu_n)_{n \in \mathbb{N}}$ converges to μ in $\mathscr{P}(\mathrm{C}(\mathcal{X}, \mathbb{R}^d))$. For any $n \in \mathbb{N}$, let $\mathbf{a} = \{a_k\}_{k=1}^n \in (\mathbb{R}^d)^n$ and $f_n^{\mathbf{a}} : \mathcal{X} \to \mathbb{R}^d$ such that for any $x \in \mathcal{X}$, $f_n^{\mathbf{a}}(x) = \sum_{k=1}^n a_k \varphi_k(x)$. Denote $\mathrm{C}_n(\mathcal{X}, \mathbb{R}^d) = \{f_n^{\mathbf{a}} : \mathbf{a} \in (\mathbb{R}^d)^n\}$. Define $\varphi_n : (\mathbb{R}^d)^n \to \mathrm{C}_n(\mathcal{X}, \mathbb{R}^d)$ such that for any $\mathbf{a} \in (\mathbb{R}^d)^n$, $\varphi_n(\mathbf{a}) = f_n^{\mathbf{a}}$. We have that φ_n is a bijection. In addition, we have that $(\varphi_n)_{\#}\mu_n = \mu_{x_1,\ldots,x_n}$. Therefore, using the data processing inequality we have that for any $n \in \mathbb{N}$

$$\mathrm{KL}\left(\mu_n|\nu_n\right) = \mathrm{KL}\left(\mu_{x_1,\ldots,x_n}|\nu_{x_1,\ldots,x_n}\right).$$

In addition, since $(\mu_n, \nu_n) \to (\mu, \nu)$ we using that the Kullback-Leibler divergence is lower semi-continuous (Dupuis and Ellis, 2011, Lemma 1.4.3)

$$\lim_{n \to +\infty} \mathrm{KL}\left(\mu_n|\nu_n\right) = \lim_{n \to +\infty} \mathrm{KL}\left(\mu_{x_1,\ldots,x_n}|\nu_{x_1,\ldots,x_n}\right) \geq \mathrm{KL}\left(\mu|\nu\right).$$

Therefore, we get that

$$\mathrm{KL}\left(\mu|\nu\right) \leq \sup\{\mathrm{KL}\left(\mu_{x_1,\ldots,x_n}|\nu_{x_1,\ldots,x_n}\right) : (x_1, \ldots, x_n) \in \mathcal{X}^n,\ n \in \mathbb{N}\}.$$

Combining this result and (7), we conclude the proof. □

The rest of the proof is similar to Appendix E.1, expect that one needs to check that the obtained Gaussian process $\pi^\star$ is sample continuous (up to a modification). This is done using that there exists $\phi : [0, +\infty)$ such that for any $x_1, x_2 \in \mathcal{X}$,

$$\mathbb{E}[\|\mathbf{Y}_{x_1} - \mathbf{Y}_{x_2}\|^2] \leq \phi(\|x_1 - x_2\|),$$

such that $\int_0^{+\infty} \phi(\exp[-t^2])\mathrm{d}t < +\infty$, see Fernique (1975).

F Experimental details

In the following sections we provide details on our experimental procedures.

F.1 1D Datasets

Quadratic A synthetic dataset constructed to exhibit clear bi-modality. Samples consist of function evaluations of $f(x) := ax^2 + \epsilon$ with $a \sim \mathrm{Unif}(\{-1, 1\})$ and $\epsilon \sim \mathrm{N}(0, 1)$ on a uniformly sampled grid of 100 evaluation points in $[-10, 10]$. We allow unlimited samples during training, validation and testing phases.

Melbourne A real dataset recording the number of pedestrians on streets in Melbourne, each sample being a period of 24 hours with readings taken hourly (giving a grid of 24 evaluation points per sample). The dataset is sourced from `http://www.timeseriesclassification.com/description.php?Dataset=MelbournePedestrian`. We pre-processed the dataset by removing rows with unobserved values and rescaling the dataset to have unit variance. The pre-processed dataset contains 2319 samples and we split the dataset in the ratio of $[0.8, 0.1, 0.1]$ for training, validation and testing respectively.

Gridwatch A real dataset recording the energy demand of the UK grid, each sample being a period of 24 hours with readings taken every five minutes (giving a grid of 288 evaluation points per sample). The dataset is sourced from https://www.gridwatch.templar.co.uk/download.php by selecting the 'demand' field. We performed a number of pre-processing steps:

- Select days where the readings were taken exactly on the fifth minute,
- Remove days where the the difference between two consecutive readings was beyond the 99.5^{th} percentile of differences,
- Remove days where consecutive readings did not vary for half an hour or more,
- Centered the dataset so each sample has zero mean,
- Scaled the dataset so each sample has unit variance.

This resulted in 1013 samples which we split in the ratio of $[0.8, 0.1, 0.1]$ for training, validation and testing respectively.

F.2 Baselines

Gaussian Processes We used a Matern-1/2 kernel with lengthscale and noise variance learnt by maximising the marginal likelihood.

Neural Processes We modified a PyTorch implementation of Neural Proccesses available online at https://github.com/EmilienDupont/neural-processes. Our modifications adapted the implementation to unconditional training/sampling by implementing Equation 7 of Garnelo et al. (2018b). We set the dimension of the representation of context points and the dimension of the latent variable to be 512 (1024 for Gridwatch dataset). The encoder and decoder were three-layered FCNs with 512 (1024 for Gridwatch dataset) units and ReLU activations, trained to 1000 epochs.

F.3 Functional MMD and two-sample test power

Our quantitative results are based off the functional MMD and kernel two-sample hypothesis test of Wynne and Duncan (2022).

The hypothesis test considered uses the null hypothesis that samples from the model and dataset are from the same distribution, while the alternative states that they are not. We know that the null is false as the distribution of samples from the model will never truly match the data distribution. Therefore, we expect a perfectly powerful test to always reject the null hypothesis. By choosing a specific test with finite power (for example by restricting the number of samples used when computing the test statistic), we can compare the quality of samples from different models by computing the power of the specific test on each of the sets of samples. If the test exhibits a lower power on a set of samples, it indicates those samples are harder to distinguish from the data distribution and thus can be considered of higher quality.

The specific test we choose is the kernel two-sample test with the ID kernel computed on 10 samples. We perform 1000 tests to estimate the power. Confidence intervals are obtained by repeating the procedure on models trained from different seeds.

F.4 Uncurated samples

We visualise samples from all models on the 1D datasets and from SP-SGM on MNIST.

Dataset	SP-SGM	NP	GP
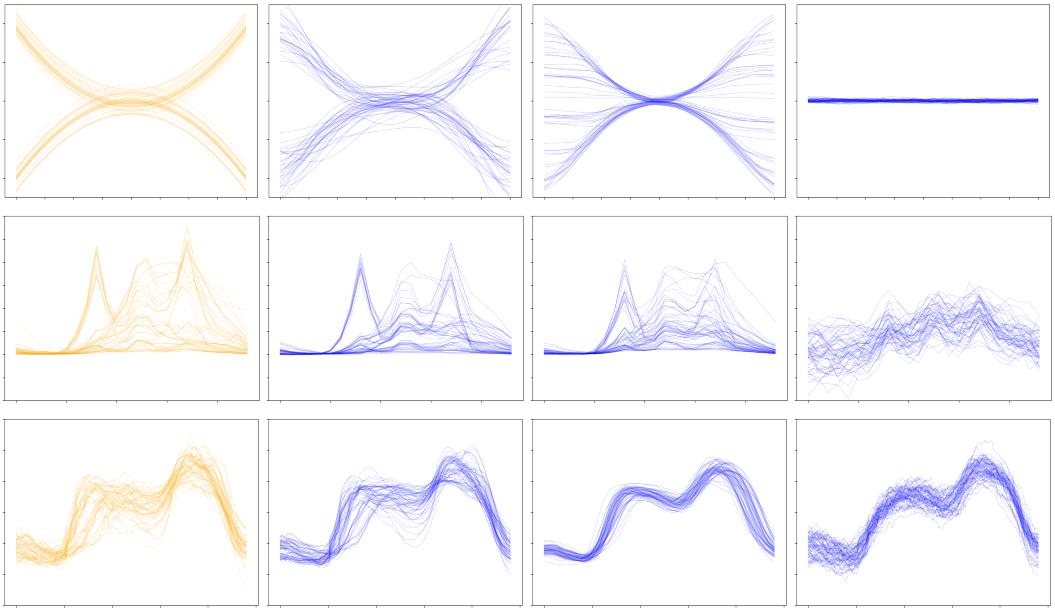			

Table 2: Samples from trained models on 1D datasets; Quadratic [top], Melbourne [middle] and Gridwatch [bottom].

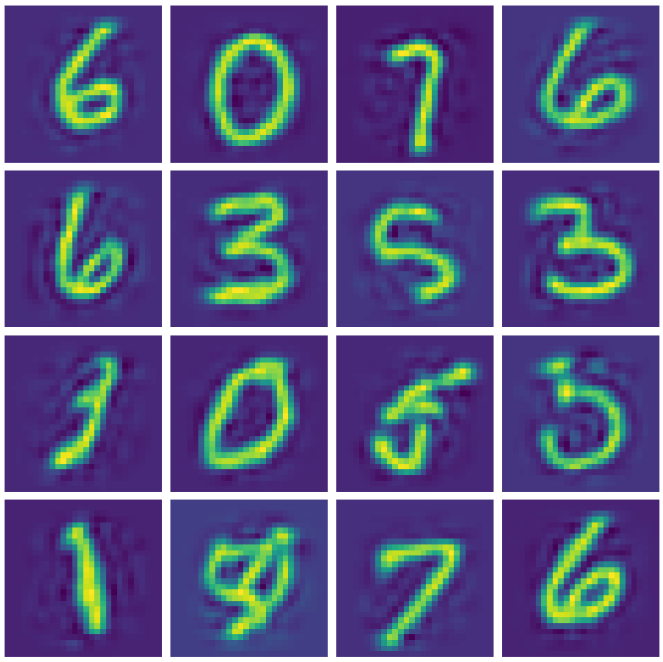

Figure 6: Uncurated MNIST samples, covariance kernel, $M = 100$

F.5 Further ablation studies

$M = 20$	$M = 100$	$M = 200$

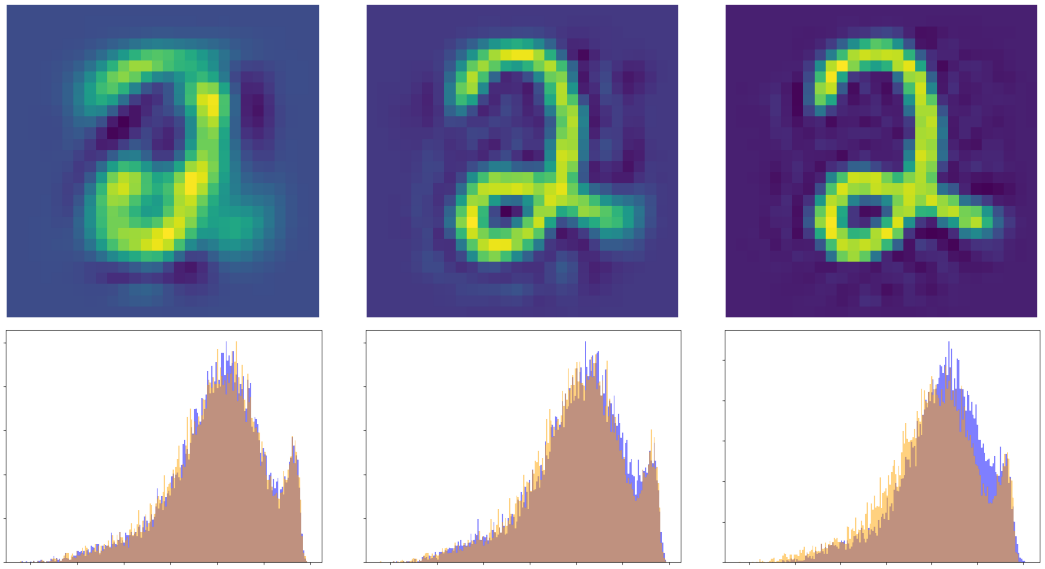

Figure 7: Effect of increasing truncation order M. Increasing reconstruction quality on MNIST digits [top], decreasing quality of fit in spectral space [bottom].

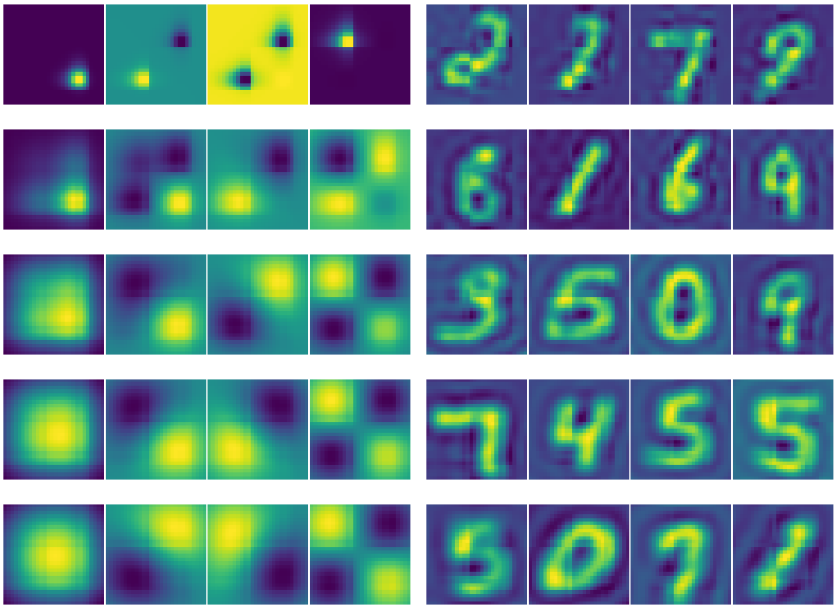

Figure 8: Eigenfunctions [left] and samples from SP-SGM [right] for RBF kernel with lengthscale in [1,2,3,4,5] [top to bottom respectively]

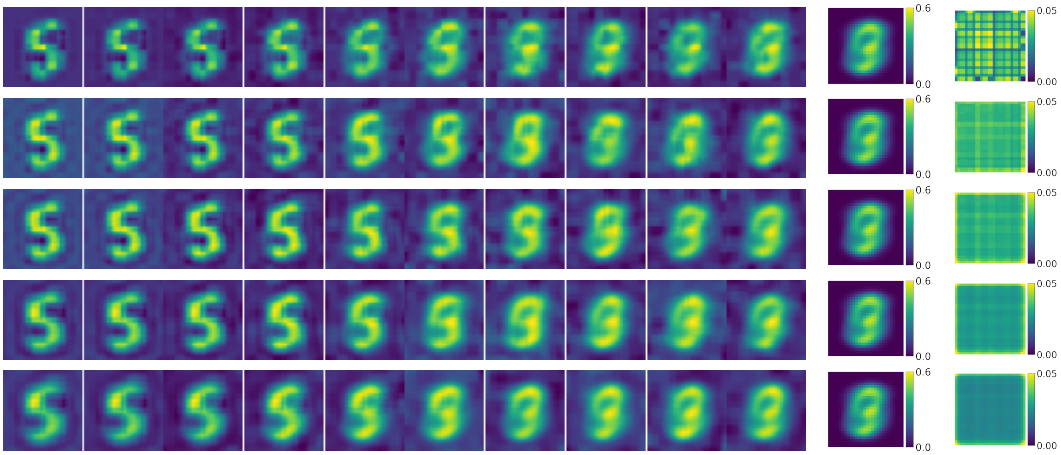

Figure 9: Forward noising process of SP-SGM on MNIST digits using RBF kernel with lengthscales in [1,2,3,4,5] [top to bottom respectively]. Rightmost columns show the pixel-wise mean and standard deviation of the reference measure respectively.

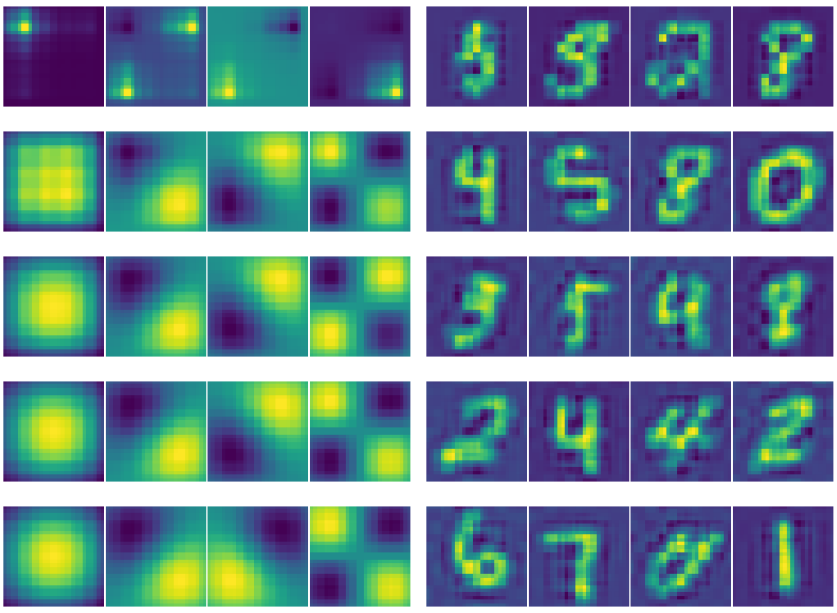

Figure 10: Eigenfunctions [left] and samples from SP-SGM [right] for Matern-3/2 kernel with lengthscale in [1,2,3,4,5] [top to bottom respectively]

Figure 11: Forward noising process of SP-SGM on MNIST digits using Matern-3/2 kernel with lengthscales in [1,2,3,4,5] [top to bottom respectively]. Rightmost columns show the pixel-wise mean and standard deviation of the reference measure respectively.