

---

# pySigLib - Fast Signature Kernels on CPU and GPU

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Signature-based methods have recently gained significant traction in machine learning  
2 for sequential data. In particular, signature kernels have emerged as powerful  
3 discriminators and training losses for generative models on time-series, notably  
4 in quantitative finance. However, existing implementations do not scale to the  
5 dataset sizes and sequence lengths encountered in practice. We present pySigLib,  
6 a high-performance library offering optimised implementations of signatures and  
7 signature kernels on CPU and GPU, fully compatible with PyTorch’s automatic  
8 differentiation. Beyond an efficient software stack for large-scale signature-kernel  
9 computation, we introduce a novel differentiation scheme that delivers accurate  
10 gradients at a fraction of the runtime of existing libraries. We believe our library  
11 unlocks the practical use of signature kernels for generative modelling and the  
12 distillation of next-generation sequence models.

## 13 1 Introduction

14 Most forms of sequential data such as financial time series, on sufficiently fine time scales, can be  
15 represented as a continuous path  $x : [0, 1] \rightarrow \mathbb{R}^d$ . It was first shown by [8], and then explored in  
16 greater detail and generality in the context of *rough path theory* in [32, 19, 5], that any path may be  
17 faithfully represented, up to reparameterisation, by the collection of its iterated integrals known as the  
18 *path-signature*. The path-signature  $S(x)$  is formally defined as the solution to the tensor differential  
19 equation  $dy_t = y_t \otimes dx_t$  in the *free tensor algebra*  $T((\mathbb{R}^d)) := \prod_{n=0}^{\infty} (\mathbb{R}^d)^{\otimes n}$ , and thereby can be  
20 interpreted as a non-commutative analogue of the exponential function. By applying Picard iteration,  
21 one recovers the familiar formulation of the path-signature as the sequence of iterated integrals  
22  $S(x) = (\int_{0 < t_1 < \dots < t_n < 1} dx_{t_1} \otimes \dots \otimes dx_{t_n})_{n \in \mathbb{N}}$ . The Stone–Weierstrass theorem ensures that linear  
23 functionals acting on the range of the path-signature are dense in the space of continuous real-valued  
24 functions defined on compact subsets of *unparameterised paths*. This makes the signature a powerful  
25 representation, enabling the approximation of path-dependent functionals using linear models.

26 Thanks to this *universal approximation property* and a variety of other algebraic features, signature  
27 methods have experienced a rapid rise in popularity in recent years, with applications across a wide  
28 range of data science domains, including quantitative finance [40, 13, 1, 13, 4], cybersecurity [11],  
29 information theory [43, 45, 48], and even quantum computing [12]. Signatures have also served as a  
30 theoretical foundation for proving universality properties of *neural differential equations* [34, 2] and  
31 more recently of *state-space models* (SSMs) [35, 36, 49]. In generative modelling, signatures have  
32 enabled new approaches to synthesizing financial time series in a model-independent manner [6], been  
33 used as universal nonlinearities in Seq2Seq architectures [24], and provided representation spaces for  
34 training *score-based diffusion models* on time series [3]. For a detailed and pedagogical introduction  
35 to the subject, the reader is referred to [7], while [16] offers a survey of recent applications.

36 In practice, paths are typically obtained via piecewise linear interpolation of discrete time series.  
37 Combining a non-trivial algebraic relation known as *Chen’s identity* with the elementary fact that the  
38 signature of a linear segment is the tensor exponential of its increment, yields a concise expression for  
39 the signature of a piecewise linear path  $x = x^1 * \dots * x^L$  as  $S(x) = \exp(\Delta x^1) \otimes \dots \otimes \exp(\Delta x^L)$ .

Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

40 This expression underlies the implementation of signature computations in standard Python libraries  
 41 such as `esig` [31], `iisignature` [42], and `signatory` [26]. It enables an efficient evaluation of the  
 42 signature with time complexity  $\mathcal{O}(Ld^n)$ , where  $n \in \mathbb{N}$  is the truncation level. However, due to the  
 43 exponential growth in the dimension  $d$ , the method becomes computationally prohibitive for large  $n$ .

44 A widely adopted solution to this curse of dimensionality is the use of *signature kernels* [27]. These  
 45 kernels take the form  $\langle S(x), S(y) \rangle$ , for suitable inner products  $\langle \cdot, \cdot \rangle$  on  $T((\mathbb{R}^d))$ , and can be computed  
 46 efficiently without explicitly evaluating the feature map  $S(x)$ . In particular, recent work has shown  
 47 that the signature kernel satisfies a Goursat PDE [44, 28]. Signature kernels have since found  
 48 applications in hypothesis testing [47, 30, 22], causality [33], kernel-based solvers for path-dependent  
 49 PDEs in derivative pricing under rough volatility [38], kernel formulations of deep hedging [37],  
 50 and have even been shown to arise as infinite-width limits of neural networks [10]. They have  
 51 also been used to train neural SDEs for time-series generation across diverse fields, including fluid  
 52 dynamics [46], computational neuroscience [21], and quantitative finance [23, 15, 20].

53 Despite these advances, **current software implementations of signature kernels do not scale**  
 54 **to large datasets of long time series, particularly common in quant finance**. Moreover, when  
 55 signature kernels are employed as loss functions, efficient and accurate backpropagation is crucial.  
 56 Existing implementations compute derivatives using a second PDE [29]; while natural, this approach  
 57 often yields inaccurate gradients, especially for short time series, leading to unreliable model training.

58 **Contribution** We introduce `pySigLib`, a **high-performance computational library offering**  
 59 **optimised CPU- and GPU-amenable implementations of signatures and signature kernels**  
 60 **in C++ and CUDA, fully compatible with PyTorch’s automatic differentiation**. `pySigLib`  
 61 is significantly faster than existing packages on both CPU and GPU thanks to algorithmic im-  
 62 provements, optimized memory access, and hardware-level parallelism via SIMD instructions.  
 63 It also supports efficient on-the-fly application of lead-lag and time-augmentation transforma-  
 64 tions, providing a substantial speed-up in financial applications. Furthermore, all functions  
 65 are fully backpropagatable through a PyTorch API. The library is open source and available at  
 66 <https://github.com/daniil-shmelev/pySigLib>, installable directly via `pip`, with documen-  
 67 tation hosted at <https://pysiglib.readthedocs.io>.

## 68 2 Computation of signature kernels

69 We begin with a brief account of signature kernels and numerical techniques to compute them.

### 70 2.1 Signature kernels as the solution of a Goursat PDE

71 A key result by [44] established that signature kernels can be computed as the solution of a Goursat-  
 72 type PDE. Notably, this result states that given two continuous paths  $x, y : [0, 1] \rightarrow \mathbb{R}^d$ , the  
 73 corresponding signature kernel  $k_{x,y}(s, t) := \langle S(x)_s, S(y)_t \rangle$  satisfies the hyperbolic PDE

$$\frac{\partial^2 k_{x,y}}{\partial s \partial t} = \langle \dot{x}_s, \dot{y}_t \rangle_V k_{x,y}, \quad k_{x,y}(u, \cdot) = k_{x,y}(\cdot, v) = 1.$$

74 An effective discretisation of the above PDE is given by [14, 44, 50]

$$\begin{aligned} \widehat{\mathbf{k}}_{i+1,j+1} &= \left( \widehat{\mathbf{k}}_{i+1,j} + \widehat{\mathbf{k}}_{i,j+1} \right) A(\Delta_{i,j}) - \widehat{\mathbf{k}}_{i,j} B(\Delta_{i,j}), \\ A(\Delta_{i,j}) &= \left( 1 + \frac{1}{2} \Delta_{i,j} + \frac{1}{12} \Delta_{i,j}^2 \right), \quad B(\Delta_{i,j}) = \left( 1 - \frac{1}{12} \Delta_{i,j}^2 \right) \\ \Delta_{i,j} &= \langle x_{t_{i+1}} - x_{t_i}, y_{s_{j+1}} - y_{s_j} \rangle \end{aligned} \tag{1}$$

75 over the dyadically refined grid  $P_{\lambda_1, \lambda_2} = \{(t_i, s_j) : 0 \leq i \leq 2^{\lambda_1} L_1, 0 \leq j \leq 2^{\lambda_2} L_2\}$  of order  
 76  $(\lambda_1, \lambda_2)$ , where  $L_1, L_2$  are the lengths of the discrete data streams  $x, y$  respectively.

### 77 2.2 An algorithm for signature kernels

78 Implementations of signature kernels exist for both CPU and GPU, most notably via the `sigkernel`  
 79 and `sigkerax` packages [44]. The resulting algorithm is outlined in 1. In our implementation,  
 80 several choices are notable: (1) as opposed to other packages, we allow  $\lambda_1 \neq \lambda_2$ , which can prove

81 useful when  $x$  and  $y$  are of significantly different lengths; (2) all values of  $\Delta_{i,j}$  required for (1) are  
 82 precomputed using a `matmul` operation. In cases where the dimension of the underlying path is  
 83 large, this operation accounts for almost all of the algorithm’s runtime, so it is critical that it is highly  
 84 optimized. In `pySigLib`, this is realised using PyTorch’s `bmm` function [39]; and (3) whereas other  
 85 packages implementing signature kernels precompute the dyadically refined path, `pySigLib` applies  
 86 the dyadic refinement on-the-fly, improving both computational and memory efficiency.

---

**Algorithm 1** CPU Algorithm for Signature Kernels

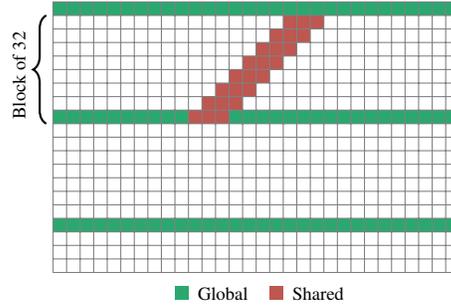
---

**Input:**  $\{x_i\}_{1 \leq i \leq L_1}, \{y_i\}_{1 \leq i \leq L_2}, \lambda_1, \lambda_2$  ▷ See (1) above  
 $dx = x_{1:L_1} - x_{0:L_1-1}; \quad dy = y_{1:L_2} - y_{0:L_2-1}$   
 $\Delta = \text{matmul}(dx^T, dy)$  ▷ See (2) above  
**for**  $s = 0, \dots, 2^{\lambda_1} L_1 - 1; \quad t = 0, \dots, 2^{\lambda_2} L_2 - 1$  **do** ▷ See (3) above  
 $p = \Delta[s \parallel 2^{\lambda_1}, t \parallel 2^{\lambda_1}]$   
 $\widehat{\mathbf{k}}[s, t] = (\widehat{\mathbf{k}}[s - 1, t] + \widehat{\mathbf{k}}[s, t - 1]) * (1 + \frac{1}{2}p + p^2) - \widehat{\mathbf{k}}[s - 1, t - 1] * (1 - p^2)$   
**end for**  
**return**  $\widehat{\mathbf{k}}[-1, -1]$

---

87 **2.3 Exploiting GPU parallelism**

88 As suggested in [44], the algorithm can be effectively  
 89 parallelised on a GPU, by noting that elements on  
 90 any given anti-diagonal of the PDE grid exhibit no  
 91 interdependencies, and so can be computed in parallel.  
 92 Following this approach, we note that it is not neces-  
 93 sary to store the entire PDE grid. Rather, one can store  
 94 only the current anti-diagonal and the two before it.  
 95 As the algorithm progresses, the three anti-diagonals  
 96 are “rotated”, so that the first anti-diagonal becomes  
 97 the second, the second becomes the third, and the  
 98 third is overwritten with new values to become the  
 99 first. This dramatically reduces the required memory allocation, allowing for the three anti-diagonals  
 100 to be kept in shared memory of the GPU, which has a significantly lower latency than the global  
 101 memory which a full PDE grid would be forced to reside in.



102 As of writing, when using other signature kernel libraries for GPU computations, the length of the  
 103 input streams is severely limited by the thread count of the GPU, as the algorithm attempts to assign  
 104 one thread per entry of the diagonal. We avoid this problem by computing kernels in blocks of 32. For  
 105 the first such block, a vector of initial conditions in global memory is populated with all ones, acting  
 106 as the first row of the PDE grid. As the block progresses, the initial condition is overwritten by the  
 107 final row in the block, which becomes the initial condition for the second block of 32. As such, one is  
 108 never limited by the thread count, as only 32 threads are allocated per kernel computation. Naturally,  
 109 this means the GPU is underutilized when computing a single kernel. However, when a batch of  
 110 kernels are computed, blocks associated with different kernels are allowed to run asynchronously  
 111 with respect to each other, meaning all available threads are used when the batch size is sufficiently  
 112 large. We anticipate this should be the case in almost all practical applications.

113 **2.4 Backpropagation through signature kernels**

114 Existing methods for backpropagation approximate derivatives as the solution to another PDE [29].  
 115 The downside to this is that the gradients are not exact. When either the length of the underlying path  
 116 or the dyadic order is too low, this difference is significant. To avoid this issue, we compute exact  
 117 gradients by differentiating through the operations of the solver (1). As we will see, the resulting  
 118 algorithm is significantly faster than the PDE approach, whilst producing exact gradients. Given  
 119 a scalar function  $F$  and its derivative with respect to the signature kernel  $\partial F / \partial \mathbf{k}_{L_1, L_2}$ , we aim to  
 120 compute  $\partial F / \partial \Delta_{a,b}$  for all indices  $a, b$ , which can then easily be used to compute the derivatives with  
 121 respect to the underlying paths  $x$  and  $y$  via a linear transformation. A naive approach may be to try  
 122 and differentiate the above relation directly with respect to  $\Delta_{a,b}$  to try and form an update rule for  
 123  $\partial \widehat{\mathbf{k}}_{i,j} / \partial \Delta_{a,b}$ . This approach would require traversing a PDE grid for each pair  $a, b$ , resulting in a

124 serial complexity of  $\mathcal{O}(2^{\lambda_1+\lambda_2} L_1^2 L_2^2)$ . One can, however, achieve the same result in only one traversal  
 125 of a PDE grid, by targeting instead  $\partial F/\partial \widehat{\mathbf{k}}_{a,b}$ . For simplicity, assume our grid is not dyadically  
 126 refined, i.e.  $\lambda_1 = \lambda_2 = 0$ . For convenience, define  $\partial F/\partial \widehat{\mathbf{k}}_{i,j} = 0$  if  $i > L_1$  or  $j > L_2$ . Then we  
 127 have the following relations, computable in one backward pass of a PDE grid:

$$\frac{\partial F}{\partial \widehat{\mathbf{k}}_{i,j}} = \frac{\partial F}{\partial \widehat{\mathbf{k}}_{i+1,j}} A(\Delta_{i,j-1}) + \frac{\partial F}{\partial \widehat{\mathbf{k}}_{i,j+1}} A(\Delta_{i-1,j}) - \frac{\partial F}{\partial \widehat{\mathbf{k}}_{i+1,j+1}} B(\Delta_{i,j}),$$

$$\frac{\partial F}{\partial \Delta_{i,j}} = \frac{\partial F}{\partial \widehat{\mathbf{k}}_{i+1,j+1}} \left[ (\widehat{\mathbf{k}}_{i+1,j} + \widehat{\mathbf{k}}_{i,j+1}) \left( \frac{1}{2} + \frac{1}{6} \Delta_{i,j} \right) + \frac{1}{6} \widehat{\mathbf{k}}_{i,j} \Delta_{i,j} \right].$$

128 The same approach applies to dyadically refined grids, noting that in this case there are multiple  
 129 indices  $a, b$  for which  $\widehat{\mathbf{k}}_{a,b}$  depends on  $\Delta_{i,j}$ .

---

**Algorithm 2** CPU Algorithm for Backpropagation through Signature Kernels

---

**Input:**  $x_i \in \mathbb{R}^d$ ,  $1 \leq i \leq L_1$  and  $y_i \in \mathbb{R}^d$ ,  $1 \leq i \leq L_2$

**Input:** Dyadic order for  $x$ ,  $\lambda_1$ , and dyadic order for  $y$ ,  $\lambda_2$

**Input:** Dyadically refined kernel grid  $\widehat{\mathbf{k}}$ , precomputed or computed on-the-fly.

**Input:** Derivative  $\partial F/\partial \widehat{\mathbf{k}}[s, t]$ .

$dx = x_{1:L_1} - x_{0:L_1-1}$ ,  $dy = y_{1:L_1} - y_{0:L_1-1}$ ;  $\Delta = \text{matmul}(dx^T, dy)$

$\Delta = \text{matmul}(dx^T, dy)$

**for**  $s = 2^{\lambda_1} L_1 - 1, \dots, 0$ ;  $t = 2^{\lambda_2} L_2 - 1, \dots, 0$  **do**

$p_{i,j} = \Delta[i \parallel 2^{\lambda_1}, j \parallel 2^{\lambda_2}]$  for  $(i, j) = (s, t), (s-1, t), (s, t-1)$

$d_1[s, t] = (d_1[s+1, t] * A(p_{s,t-1}) + d_1[s, t+1] * A(p_{s-1,t}) - d_1[s+1, t+1] * B(p_{s,t}))$

$d_2[s \parallel 2^{\lambda_1}, t \parallel 2^{\lambda_2}] += d_1[s, t] * \{ (\widehat{\mathbf{k}}[s+1, t] + \widehat{\mathbf{k}}[s, t+1]) * A'(p_{s,t}) - \widehat{\mathbf{k}}[s, t] * B'(p_{s,t}) \}$

**end for**

**return**  $d_2$

▷ Result is  $\partial F/\partial \Delta$ , which can be transformed into  $\partial F/\partial x$  or  $\partial F/\partial y$

---

130 Table 1 and Figure 1 show the runtime of signature kernels and their backpropagation. A 14-core  
 131 i7-13700H CPU with 32GB of RAM and an NVIDIA RTX 4060 GPU are used. The minimum  
 132 runtime is taken over 50 runs. Results are compared to the `sigkernel` package, which is limited by  
 133 the thread count (1024) on GPU and memory on CPU.

(B, L, d)	Forward (CPU)		Forward (GPU)		Backward (CPU)		Backward (GPU)	
	sigkernel	pySigLib	sigkernel	pySigLib	sigkernel	pySigLib	sigkernel	pySigLib
(128, 256, 8)	0.3610	<b>0.0118</b>	0.0088	<b>0.0034</b>	2.3248	<b>0.0322</b>	0.1475	<b>0.0057</b>
(128, 512, 16)	1.3608	<b>0.0364</b>	0.0376	<b>0.0117</b>	15.0445	<b>0.1276</b>	15.3843	<b>0.0231</b>
(128, 1024, 32)	6.0874	<b>0.2325</b>	-	<b>0.0653</b>	-	<b>0.6019</b>	-	<b>0.1112</b>

Table 1: Runtime (seconds) for a batch of  $B$  paths of length  $L$  and dimension  $d$ .

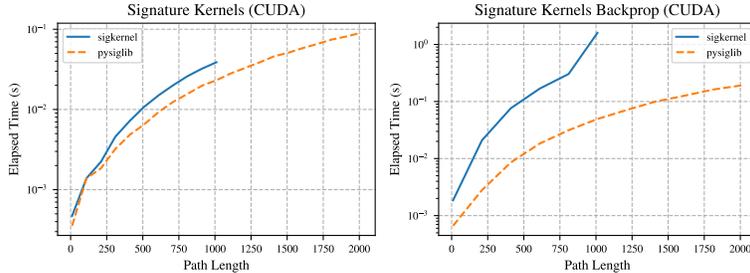


Figure 1: Runtime (seconds) for a batch of 32 paths of dimension 5.

### 134 3 Conclusion

135 We have introduced `pySigLib`, a computational library designed to overcome the challenges of  
 136 applying signature-based methods to large-scale data. Our results show significant improvements in  
 137 the computation of signature kernels, allowing for efficient processing of long time-series data. We  
 138 presented a novel approach to backpropagation through signature kernels, which achieves accurate  
 139 gradients at a fraction of the runtime of existing implementations. Concrete experiments integrating  
 140 `pySigLib` into deep learning pipelines are beyond the scope of this paper and left for future work.

## References

- [1] E. Abi Jaber and L.-A. Gérard. Signature volatility models: pricing and hedging with fourier. *SIAM Journal on Financial Mathematics*, 16(2):606–642, 2025.
- [2] I. P. Arribas, C. Salvi, and L. Szpruch. Sig-sdes model for quantitative finance. In *ACM International Conference on AI in Finance*, 2020.
- [3] B. Barancikova, Z. Huang, and C. Salvi. Sigdiffusions: Score-based diffusion models for time series via log-signature embeddings. In *The Thirteenth International Conference on Learning Representations*.
- [4] C. Bayer, L. Pelizzari, and J.-J. Zhu. Pricing american options under rough volatility using deep-signatures and signature-kernels. *arXiv preprint arXiv:2501.06758*, 2025.
- [5] H. Boedihardjo, X. Geng, T. Lyons, and D. Yang. The signature of a rough path: uniqueness. *Advances in Mathematics*, 293:720–737, 2016.
- [6] H. Buehler, B. Horvath, T. Lyons, I. P. Arribas, and B. Wood. A data-driven market simulator for small data environments. *arXiv preprint arXiv:2006.14498*, 2020.
- [7] T. Cass and C. Salvi. Lecture Notes on Rough Paths and Applications to Machine Learning, 2024.
- [8] K.-T. Chen. Integration of paths, geometric invariants and a generalized baker-hausdorff formula. *Annals of Mathematics*, pages 163–178, 1957.
- [9] I. Chevyrev and A. Kormilitzin. A primer on the signature method in machine learning. *arXiv preprint arXiv:1603.03788*, 2016.
- [10] N. M. Cirone, M. Lemerrier, and C. Salvi. Neural signature kernels as infinite-width-depth-limits of controlled resnets. In *International Conference on Machine Learning*, pages 25358–25425. PMLR, 2023.
- [11] T. Cochrane, P. Foster, V. Chhabra, M. Lemerrier, T. Lyons, and C. Salvi. Sk-tree: a systematic malware detection algorithm on streaming trees via the signature kernel. In *2021 IEEE International Conference on Cyber Security and Resilience (CSR)*, pages 35–40. IEEE, 2021.
- [12] S. Crew, C. Salvi, W. F. Turner, T. Cass, and A. Jacquier. Quantum path signatures. *arXiv preprint arXiv:2508.05103*, 2025.
- [13] C. Cuchiero and J. Möller. Signature methods in stochastic portfolio theory. *arXiv preprint arXiv:2310.02322*, 2023.
- [14] J. T. Day. A runge-kutta method for the numerical solution of the goursat problem in hyperbolic partial differential equations. *Comput. J.*, 9(1):81–83, 1966.
- [15] P. Díaz, T. Lozano, and J. Vives i Santa Eulàlia. Neural stochastic differential equations for conditional time series generation using the signature wasserstein-1 metric. *Journal Of Computational Finance*, 2023, vol. 27, num. 1, p. 1-23, 2023.
- [16] A. Fermanian, T. Lyons, J. Morrill, and C. Salvi. New directions in the applications of rough path theory. *IEEE BITS the Information Theory Magazine*, 3(2):41–53, 2023.
- [17] G. Flint, B. Hambly, and T. Lyons. Discretely sampled signals and the rough hoff process. *Stochastic Processes and their Applications*, 126(9):2593–2614, 2016.
- [18] L. G. Gyurkó, T. Lyons, M. Kontkowski, and J. Field. Extracting information from the signature of a financial data stream. *arXiv preprint arXiv:1307.7244*, 2013.
- [19] B. Hambly and T. Lyons. Uniqueness for the signature of a path of bounded variation and the reduced path group. *Annals of Mathematics*, pages 109–167, 2010.
- [20] M. Hoglund, E. Ferrucci, C. Hernández, A. M. Gonzalez, C. Salvi, L. Sánchez-Betancourt, and Y. Zhang. A neural rde approach for continuous-time non-markovian stochastic control problems. *arXiv preprint arXiv:2306.14258*, 2023.

- 187 [21] C. Holberg and C. Salvi. Exact gradients for stochastic spiking neural networks driven by rough  
188 signals. *Advances in Neural Information Processing Systems*, 37:31907–31939, 2024.
- 189 [22] B. Horvath, M. Lemerrier, C. Liu, T. Lyons, and C. Salvi. Optimal stopping via distribution  
190 regression: a higher rank signature approach, 2023.
- 191 [23] Z. Issa, B. Horvath, M. Lemerrier, and C. Salvi. Non-adversarial training of neural sdes with  
192 signature kernel scores. *Advances in Neural Information Processing Systems*, 36:11102–11126,  
193 2023.
- 194 [24] P. Kidger, P. Bonnier, I. Perez Arribas, C. Salvi, and T. Lyons. Deep signature transforms.  
195 *Advances in neural information processing systems*, 32, 2019.
- 196 [25] P. Kidger and T. Lyons. Signatory: differentiable computations of the signature and logsignature  
197 transforms, on both CPU and GPU. *arXiv preprint arXiv:2001.00706*, 2020.
- 198 [26] P. Kidger and T. Lyons. Signatory: differentiable computations of the signature and logsignature  
199 transforms, on both CPU and GPU, 2021.
- 200 [27] F. J. Király and H. Oberhauser. Kernels for sequentially ordered data. *Journal of Machine  
201 Learning Research*, 20(31):1–45, 2019.
- 202 [28] M. Lemerrier, T. Lyons, and C. Salvi. Log-pde methods for rough signature kernels. *arXiv  
203 preprint arXiv:2404.02926*, 2024.
- 204 [29] M. Lemerrier, C. Salvi, T. Cass, E. V. Bonilla, T. Damoulas, and T. J. Lyons. Siggpde: Scaling  
205 sparse gaussian processes on sequential data. In *International Conference on Machine Learning*,  
206 pages 6233–6242. PMLR, 2021.
- 207 [30] M. Lemerrier, C. Salvi, T. Damoulas, E. Bonilla, and T. Lyons. Distribution regression for  
208 sequential data. In *International Conference on Artificial Intelligence and Statistics*, pages  
209 3754–3762. PMLR, 2021.
- 210 [31] T. Lyons and al. Coropa computational rough paths (software library), 2010.
- 211 [32] T. J. Lyons. Differential equations driven by rough signals. *Revista Matemática Iberoamericana*,  
212 14(2):215–310, 1998.
- 213 [33] G. Manten, C. Casolo, E. Ferrucci, S. W. Mogensen, C. Salvi, and N. Kilbertus. Signature kernel  
214 conditional independence tests in causal discovery for stochastic processes. In *The Thirteenth  
215 International Conference on Learning Representations*.
- 216 [34] J. Morrill, C. Salvi, P. Kidger, and J. Foster. Neural rough differential equations for long time  
217 series. In *International Conference on Machine Learning*, pages 7829–7838. PMLR, 2021.
- 218 [35] N. Muça Cirone, A. Orvieto, B. Walker, C. Salvi, and T. Lyons. Theoretical foundations of deep  
219 selective state-space models, 2024.
- 220 [36] N. Muça Cirone and C. Salvi. Parallelflo: Parallelizing linear transformers via flow discretiza-  
221 tion, 2025.
- 222 [37] N. Muça Cirone and C. Salvi. Rough kernel hedging, 2025.
- 223 [38] A. Pannier and C. Salvi. A path-dependent pde solver based on signature kernels, 2024.
- 224 [39] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin,  
225 N. Gimsheine, L. Antiga, et al. Pytorch: An imperative style, high-performance deep learning  
226 library. *Advances in neural information processing systems*, 32, 2019.
- 227 [40] I. Perez Arribas. *Signatures in machine learning and finance*. PhD thesis, University of Oxford,  
228 2020.
- 229 [41] J. Reizenstein. Iterated-integral signatures in machine learning. 2019.
- 230 [42] J. Reizenstein and B. Graham. The iisignature library: efficient calculation of iterated-integral  
231 signatures and log signatures. *arXiv preprint arXiv:1802.08252*, 2018.

- 232 [43] C. Salvi. *Rough paths, kernels, differential equations and an algebra of functions on streams*.  
233 PhD thesis, University of Oxford, 2021.
- 234 [44] C. Salvi, T. Cass, J. Foster, T. Lyons, and W. Yang. The signature kernel is the solution of a  
235 goursat pde. *SIAM Journal on Mathematics of Data Science*, 3(3):873–899, 2021.
- 236 [45] C. Salvi, J. Diehl, T. Lyons, R. Preiss, and J. Reizenstein. A structure theorem for streamed  
237 information. *Journal of Algebra*, 634:911–938, 2023.
- 238 [46] C. Salvi, M. Lemercier, and A. Gerasimovics. Neural stochastic PDEs: Resolution-invariant  
239 learning of continuous spatiotemporal dynamics. *Advances in Neural Information Processing*  
240 *Systems*, 35:1333–1344, 2022.
- 241 [47] C. Salvi, M. Lemercier, C. Liu, B. Horvath, T. Damoulas, and T. Lyons. Higher order kernel  
242 mean embeddings to capture filtrations of stochastic processes. *Advances in Neural Information*  
243 *Processing Systems*, 34:16635–16647, 2021.
- 244 [48] D. Shmelev and C. Salvi. Sparse signature coefficient recovery via kernels, 2024.
- 245 [49] B. Walker, L. Yang, N. M. Cirone, C. Salvi, and T. Lyons. Structured linear cdes: Maximally  
246 expressive and parallel-in-time sequence models. *arXiv preprint arXiv:2505.17761*, 2025.
- 247 [50] A.-M. Wazwaz. On the numerical solution of the goursat problem. *Applied mathematics and*  
248 *computation*, 59(1):89–95, 1993.

## 249 A Runtime for signature kernels on CPU

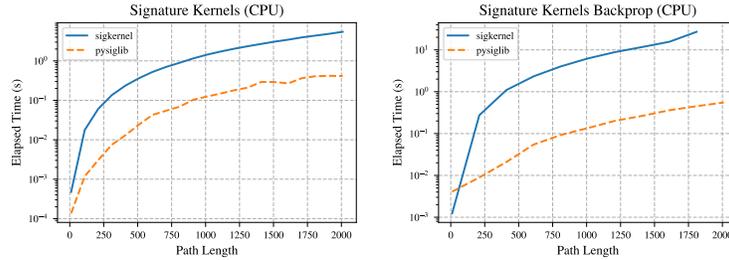


Figure 2: Runtime (seconds) for signature kernels and backpropagation, computed for a batch of 32 paths of dimension 5 on a i7-13700H CPU and an NVIDIA RTX 4060 GPU. The minimum time is taken over 50 runs. The `sigkernel` package is limited by the thread count (1024) on GPU.

## 250 B Path Transformations

251 Given a path, it is common to apply certain path-to-path transformations before computing a transform  
252 such as the signature. The most common path-to-path transformations are time-augmentation and the  
253 lead-lag transform. pySigLib provides backpropagatable implementations of these transforms. Where  
254 possible, pySigLib is capable of adapting algorithms internally to perform these transformations  
255 on-the-fly, which is often faster and more memory-efficient than pre-computing them. Given a path  
256 consisting of points  $x_{t_i} \in \mathbb{R}^d$ , the time augmented path is defined by  $\hat{x}_{t_i} := (x_{t_i}, t_i) \in \mathbb{R}^{d+1}$ .  
257 Time-augmentation is useful for introducing time-reparametrisation variance to the signature, in  
258 applications where the timing of events is critical. The lead-lag transformation is defined by  $X_{t_i}^{LL} :=$   
259  $(X_{t_i}^{\text{Lead}}, X_{t_i}^{\text{Lag}})$ , where [9, 18, 17]

$$X_{t_i}^{\text{Lead}} := \begin{cases} X_{t_k} & \text{if } i = 2k, \\ X_{t_k} & \text{if } i = 2k - 1, \end{cases}$$

$$X_{t_i}^{\text{Lag}} := \begin{cases} X_{t_k} & \text{if } i = 2k, \\ X_{t_k} & \text{if } i = 2k + 1, \end{cases}$$

260 The lead-lag transform is particularly useful for feature-extraction applied to financial data streams,  
 261 where the signature of the lead-lag path may be used as an approximation to the Itô-signature [17, 18].

## 262 C Truncated Signatures

263 We implement and test two algorithms for computing truncated signatures. The first of these is the  
 264 direct approach, as used by `iisignature` [42]. Suppose we are given as input a path consisting of  
 265 points  $x_i \in \mathbb{R}^d$ ,  $1 \leq i \leq L$ , and a truncation level  $N$  for the resulting signature. Suppose we have  
 266 already used the first  $\ell$  of these points to construct the signature  $S(x_{1:\ell}) = (A_0, A_1, \dots, A_N)$ . Then  
 267 the update rule for the next step of the algorithm is given by

$$A_k = \sum_{i=0}^k A_i \otimes \frac{z^{\otimes(k-i)}}{(k-i)!}$$

268 where  $z := x_{\ell+1} - x_\ell$ . The resulting algorithm is outlined in 3. Given the computational complexity  
 269 of signature computations, it is particularly important to minimise costly memory allocation and  
 270 access operations within the algorithm. Two major design choices allow us to do this:

- 271 1.  $(A_0, A_1, \dots, A_N)$  is stored as a single flattened contiguous array instead of using separate  
 272 containers for each level, improving cache locality.
- 273 2. Since the update for any given level  $A_k$  depends only on  $A_i$  for  $i < k$ , the levels are updated  
 274 in reverse order, starting at  $A_N$  and ending at  $A_1$ , allowing these to be written directly  
 275 into the existing memory for  $(A_0, A_1, \dots, A_N)$  without the need for storing intermediary  
 276 computations.

---

### Algorithm 3 Direct Algorithm for Truncated Signatures

---

**Input:**  $x_i \in \mathbb{R}^d$ ,  $1 \leq i \leq L$

**Input:** Truncation level  $N$

$z = x_1 - x_0$

$(A_0, A_1, \dots, A_N) = (1, z, \dots, z^{\otimes N}/N!)$  ▷ See (1) above

**for**  $\ell = 1, \dots, L - 1$  **do**

$z = x_{\ell+1} - x_\ell$

$(\tilde{A}_0, \tilde{A}_1, \dots, \tilde{A}_N) = (1, z, \dots, z^{\otimes N}/N!)$

**for**  $k = N, \dots, 1$  **do** ▷ See (2) above

**for**  $i = k - 1, \dots, 1$  **do**

$A_k = A_k + A_i \otimes \tilde{A}_{k-i}$

**end for**

$A_k = A_k + \tilde{A}_k$

**end for**

**end for**

---

277 Despite these optimisations, the direct algorithm remains complex, and is largely bound by costly  
 278 memory access operations. These issues are somewhat improved by Horner's algorithm for polyno-  
 279 mial multiplication, as used by `signatory` [25]. In Horner's algorithm, the update rule is rewritten  
 280 as

$$\begin{aligned} A_k &= \sum_{i=0}^k A_i \otimes \frac{z^{\otimes(k-i)}}{(k-i)!} \\ &= \left( \left( \left( \left( \left( \frac{z}{k} + A_1 \right) \otimes \frac{z}{k-1} + A_2 \right) \otimes \frac{z}{k-2} + \dots \right) \otimes \frac{z}{2} + A_{k-1} \right) \otimes z + A_k, \right. \end{aligned}$$

281 which minimises the number of multiplications necessary, and reduces the number of memory  
 282 accesses to the right-hand tensor. In our implementation of Horner's algorithm, it will be useful to  
 283 split the update rule as

$$\begin{aligned} A_k &= (B_k + A_{k-1}) \otimes z + A_k \\ B_k &:= \left( \dots \left( \left( \frac{z}{k} + A_1 \right) \otimes \frac{z}{k-1} + A_2 \right) \otimes \frac{z}{k-2} + \dots \right) \otimes \frac{z}{2} \end{aligned}$$

284 to allow for additional memory access and allocation optimisations, as we will see below. The  
 285 resulting algorithm is outlined in 4.

---

**Algorithm 4** Horner’s Algorithm for Truncated Signatures

---

**Input:**  $x_i \in \mathbb{R}^d, 1 \leq i \leq L$

**Input:** Truncation level  $N$

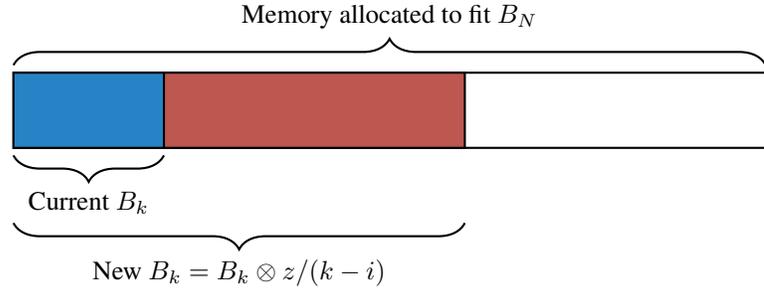
```

 $z = x_1 - x_0$ 
 $(A_0, A_1, \dots, A_N) = (1, z, \dots, z^{\otimes N}/N!)$                                 ▷ See (1) above
for  $\ell = 1, \dots, L - 1$  do
   $z = x_{\ell+1} - x_\ell$ 
  for  $k = N, \dots, 2$  do                                                    ▷ See (2) above
     $B_k = z/k$ 
    for  $i = 1, \dots, k - 2$  do
       $B_k = B_k + A_i$ 
       $B_k = B_k \otimes z/(k - i)$                                             ▷ See (3) below
    end for
     $B_k = B_k + A_{k-1}$ 
     $A_k = B_k \otimes z + A_k$                                               ▷ See (4) below
  end for
end for
  
```

---

286 The two design choices of Algorithm 3 also apply here. In addition, the following two choices are of  
 287 particular importance:

- 288 3. A single, continuous block of memory is pre-allocated to fit  $B_N$ , and re-used by all inter-  
 289 mediate computations involving  $B_k$ . The multiplication  $B_k = B_k \otimes z/(k - i)$  is written  
 290 directly into the same memory, but carried out in reverse order such that the new values of  
 291  $B_k$  (of which there are now  $d$  times as many) erase the old values only at the very end of the  
 292 multiplication, when they are no longer required.



- 293 4. The last multiplication and addition  $B_k \otimes z + A_k$  is written directly into the result,  $A_k$ .
- 294

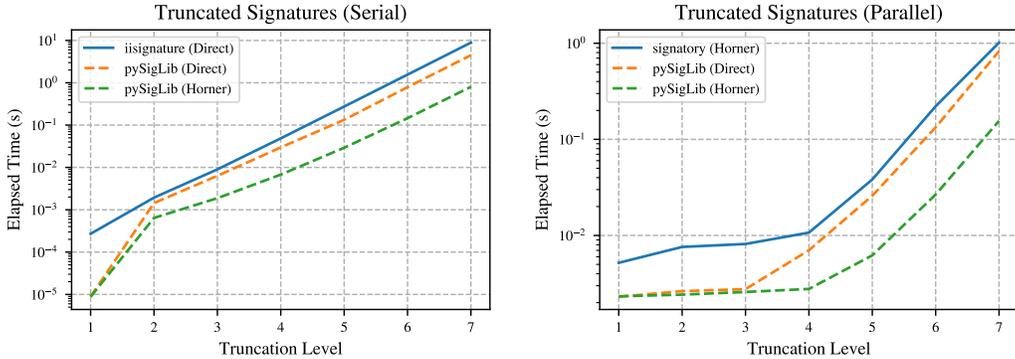


Figure 3: Timings for a batch of 32 paths of length 1024, dimension 5. The minimum time is taken over 50 runs. Ran on a i7-13700H CPU.

Elapsed Time (s)	CPU (Serial)			CPU (Parallel)	
	esig	iisignature	pySigLib	signatory	pySigLib
Signature	0.9093	0.2441	<b>0.0227</b>	0.0272	<b>0.0050</b>

Table 2: Timings for a batch of 32 paths of length 128, dimension 4, and a truncation level of 7. The minimum time is taken over 50 runs. Ran on a i7-13700H CPU.

Elapsed Time (s)	CPU (Serial)			CPU (Parallel)	
	esig	iisignature	pySigLib	signatory	pySigLib
Signature	9.0908	3.6238	<b>0.2805</b>	0.3423	<b>0.0590</b>

Table 3: Timings for a batch of 32 paths of length 1024, dimension 16, and a truncation level of 4. The minimum time is taken over 50 runs. Ran on a i7-13700H CPU.

### 295 C.1 Backpropagation through Truncated Signatures

296 Backpropagation is implemented efficiently by using the time-reversed path to deconstruct the  
 297 signature. We refer the reader to [41, Section 4.9] for details of the algorithm. Our optimisations  
 298 of the backpropagation are largely the same as for the forward pass of the signature. A slight  
 299 modification is made to the original algorithm presented in [41], whereby Horner’s algorithm is used  
 300 to deconstruct the signature.

## 301 **NeurIPS Paper Checklist**

302 The checklist is designed to encourage best practices for responsible machine learning research,  
303 addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove  
304 the checklist: **The papers not including the checklist will be desk rejected.** The checklist should  
305 follow the references and follow the (optional) supplemental material. The checklist does NOT count  
306 towards the page limit.

307 Please read the checklist guidelines carefully for information on how to answer these questions. For  
308 each question in the checklist:

- 309 • You should answer **[Yes]** , **[No]** , or **[NA]** .
- 310 • **[NA]** means either that the question is Not Applicable for that particular paper or the  
311 relevant information is Not Available.
- 312 • Please provide a short (1–2 sentence) justification right after your answer (even for NA).

313 **The checklist answers are an integral part of your paper submission.** They are visible to the  
314 reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it  
315 (after eventual revisions) with the final version of your paper, and its final version will be published  
316 with the paper.

317 The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation.  
318 While "**[Yes]**" is generally preferable to "**[No]**", it is perfectly acceptable to answer "**[No]**" provided a  
319 proper justification is given (e.g., "error bars are not reported because it would be too computationally  
320 expensive" or "we were unable to find the license for the dataset we used"). In general, answering  
321 "**[No]**" or "**[NA]**" is not grounds for rejection. While the questions are phrased in a binary way, we  
322 acknowledge that the true answer is often more nuanced, so please just use your best judgment and  
323 write a justification to elaborate. All supporting evidence can appear either in the main paper or the  
324 supplemental material, provided in appendix. If you answer **[Yes]** to a question, in the justification  
325 please point to the section(s) where related material for the question can be found.

326 **IMPORTANT**, please:

- 327 • **Delete this instruction block, but keep the section heading “NeurIPS Paper Checklist”,**
- 328 • **Keep the checklist subsection headings, questions/answers and guidelines below.**
- 329 • **Do not modify the questions and only use the provided macros for your answers.**

### 330 **1. Claims**

331 Question: Do the main claims made in the abstract and introduction accurately reflect the  
332 paper’s contributions and scope?

333 Answer: **[Yes]**

334 Justification: The claims made about the efficiency of the library accurately reflect the results  
335 presented in the paper (Section 2.4).

336 Guidelines:

- 337 • The answer NA means that the abstract and introduction do not include the claims  
338 made in the paper.
- 339 • The abstract and/or introduction should clearly state the claims made, including the  
340 contributions made in the paper and important assumptions and limitations. A No or  
341 NA answer to this question will not be perceived well by the reviewers.
- 342 • The claims made should match theoretical and experimental results, and reflect how  
343 much the results can be expected to generalize to other settings.
- 344 • It is fine to include aspirational goals as motivation as long as it is clear that these goals  
345 are not attained by the paper.

### 346 **2. Limitations**

347 Question: Does the paper discuss the limitations of the work performed by the authors?

348 Answer: **[Yes]**

349 Justification: Concrete experiments integrating the library into deep learning pipelines are  
350 beyond the scope of this paper and left for future work, as discussed in the conclusion of  
351 this work.

352 Guidelines:

- 353 • The answer NA means that the paper has no limitation while the answer No means that  
354 the paper has limitations, but those are not discussed in the paper.
- 355 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 356 • The paper should point out any strong assumptions and how robust the results are to  
357 violations of these assumptions (e.g., independence assumptions, noiseless settings,  
358 model well-specification, asymptotic approximations only holding locally). The authors  
359 should reflect on how these assumptions might be violated in practice and what the  
360 implications would be.
- 361 • The authors should reflect on the scope of the claims made, e.g., if the approach was  
362 only tested on a few datasets or with a few runs. In general, empirical results often  
363 depend on implicit assumptions, which should be articulated.
- 364 • The authors should reflect on the factors that influence the performance of the approach.  
365 For example, a facial recognition algorithm may perform poorly when image resolution  
366 is low or images are taken in low lighting. Or a speech-to-text system might not be  
367 used reliably to provide closed captions for online lectures because it fails to handle  
368 technical jargon.
- 369 • The authors should discuss the computational efficiency of the proposed algorithms  
370 and how they scale with dataset size.
- 371 • If applicable, the authors should discuss possible limitations of their approach to  
372 address problems of privacy and fairness.
- 373 • While the authors might fear that complete honesty about limitations might be used by  
374 reviewers as grounds for rejection, a worse outcome might be that reviewers discover  
375 limitations that aren't acknowledged in the paper. The authors should use their best  
376 judgment and recognize that individual actions in favor of transparency play an impor-  
377 tant role in developing norms that preserve the integrity of the community. Reviewers  
378 will be specifically instructed to not penalize honesty concerning limitations.

379 **3. Theory assumptions and proofs**

380 Question: For each theoretical result, does the paper provide the full set of assumptions and  
381 a complete (and correct) proof?

382 Answer: [NA]

383 Justification: The paper does not include theoretical results.

384 Guidelines:

- 385 • The answer NA means that the paper does not include theoretical results.
- 386 • All the theorems, formulas, and proofs in the paper should be numbered and cross-  
387 referenced.
- 388 • All assumptions should be clearly stated or referenced in the statement of any theorems.
- 389 • The proofs can either appear in the main paper or the supplemental material, but if  
390 they appear in the supplemental material, the authors are encouraged to provide a short  
391 proof sketch to provide intuition.
- 392 • Inversely, any informal proof provided in the core of the paper should be complemented  
393 by formal proofs provided in appendix or supplemental material.
- 394 • Theorems and Lemmas that the proof relies upon should be properly referenced.

395 **4. Experimental result reproducibility**

396 Question: Does the paper fully disclose all the information needed to reproduce the main ex-  
397 perimental results of the paper to the extent that it affects the main claims and/or conclusions  
398 of the paper (regardless of whether the code and data are provided or not)?

399 Answer: [Yes]

400 Justification: All parameters for results concerning runtime are specified, and code repro-  
401 ducing the results is available in the public repository of the library.

402 Guidelines:

- 403 • The answer NA means that the paper does not include experiments.
- 404 • If the paper includes experiments, a No answer to this question will not be perceived  
405 well by the reviewers: Making the paper reproducible is important, regardless of  
406 whether the code and data are provided or not.
- 407 • If the contribution is a dataset and/or model, the authors should describe the steps taken  
408 to make their results reproducible or verifiable.

- 409
- 410
- 411
- 412
- 413
- 414
- 415
- 416
- 417
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
  - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
    - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
    - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
    - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
    - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 434 5. **Open access to data and code**

435 Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

436 Answer: [\[Yes\]](#)

437 Justification: Code reproducing the results is available in the public repository of the library.  
438 Guidelines:

- 439
- 440
- The answer NA means that paper does not include experiments requiring code.
  - Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
  - While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
  - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
  - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
  - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
  - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
  - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

#### 441 6. **Experimental setting/details**

442 Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

443 Answer: [\[Yes\]](#)

444

465 Justification: The paper does not involve training or testing of models. All parameters for  
466 results concerning runtime are specified, and code reproducing the results is available in the  
467 public repository of the library.

468 Guidelines:

- 469 • The answer NA means that the paper does not include experiments.
- 470 • The experimental setting should be presented in the core of the paper to a level of detail  
471 that is necessary to appreciate the results and make sense of them.
- 472 • The full details can be provided either with the code, in appendix, or as supplemental  
473 material.

## 474 7. Experiment statistical significance

475 Question: Does the paper report error bars suitably and correctly defined or other appropriate  
476 information about the statistical significance of the experiments?

477 Answer: [No]

478 Justification: Error bars and statistical significance tests are not suitable for the results of the  
479 paper.

480 Guidelines:

- 481 • The answer NA means that the paper does not include experiments.
- 482 • The authors should answer "Yes" if the results are accompanied by error bars, confi-  
483 dence intervals, or statistical significance tests, at least for the experiments that support  
484 the main claims of the paper.
- 485 • The factors of variability that the error bars are capturing should be clearly stated (for  
486 example, train/test split, initialization, random drawing of some parameter, or overall  
487 run with given experimental conditions).
- 488 • The method for calculating the error bars should be explained (closed form formula,  
489 call to a library function, bootstrap, etc.)
- 490 • The assumptions made should be given (e.g., Normally distributed errors).
- 491 • It should be clear whether the error bar is the standard deviation or the standard error  
492 of the mean.
- 493 • It is OK to report 1-sigma error bars, but one should state it. The authors should  
494 preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis  
495 of Normality of errors is not verified.
- 496 • For asymmetric distributions, the authors should be careful not to show in tables or  
497 figures symmetric error bars that would yield results that are out of range (e.g. negative  
498 error rates).
- 499 • If error bars are reported in tables or plots, The authors should explain in the text how  
500 they were calculated and reference the corresponding figures or tables in the text.

## 501 8. Experiments compute resources

502 Question: For each experiment, does the paper provide sufficient information on the com-  
503 puter resources (type of compute workers, memory, time of execution) needed to reproduce  
504 the experiments?

505 Answer: [Yes]

506 Justification: The exact processor type and available RAM is specified for the results.

507 Guidelines:

- 508 • The answer NA means that the paper does not include experiments.
- 509 • The paper should indicate the type of compute workers CPU or GPU, internal cluster,  
510 or cloud provider, including relevant memory and storage.
- 511 • The paper should provide the amount of compute required for each of the individual  
512 experimental runs as well as estimate the total compute.
- 513 • The paper should disclose whether the full research project required more compute  
514 than the experiments reported in the paper (e.g., preliminary or failed experiments that  
515 didn't make it into the paper).

## 516 9. Code of ethics

517 Question: Does the research conducted in the paper conform, in every respect, with the  
518 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

519 Answer: [Yes]

520 Justification: The authors have reviewed and conform with the NeurIPS Code of Ethics.

521 Guidelines:

522 • The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.

523 • If the authors answer No, they should explain the special circumstances that require a

524 deviation from the Code of Ethics.

525 • The authors should make sure to preserve anonymity (e.g., if there is a special consid-

526 eration due to laws or regulations in their jurisdiction).

527 **10. Broader impacts**

528 Question: Does the paper discuss both potential positive societal impacts and negative

529 societal impacts of the work performed?

530 Answer: [NA]

531 Justification: There is no negative societal impact to the work. Positive societal impacts are

532 discussed in the introduction, regarding the efficient training of machine learning models.

533 Guidelines:

534 • The answer NA means that there is no societal impact of the work performed.

535 • If the authors answer NA or No, they should explain why their work has no societal

536 impact or why the paper does not address societal impact.

537 • Examples of negative societal impacts include potential malicious or unintended uses

538 (e.g., disinformation, generating fake profiles, surveillance), fairness considerations

539 (e.g., deployment of technologies that could make decisions that unfairly impact specific

540 groups), privacy considerations, and security considerations.

541 • The conference expects that many papers will be foundational research and not tied

542 to particular applications, let alone deployments. However, if there is a direct path to

543 any negative applications, the authors should point it out. For example, it is legitimate

544 to point out that an improvement in the quality of generative models could be used to

545 generate deepfakes for disinformation. On the other hand, it is not needed to point out

546 that a generic algorithm for optimizing neural networks could enable people to train

547 models that generate Deepfakes faster.

548 • The authors should consider possible harms that could arise when the technology is

549 being used as intended and functioning correctly, harms that could arise when the

550 technology is being used as intended but gives incorrect results, and harms following

551 from (intentional or unintentional) misuse of the technology.

552 • If there are negative societal impacts, the authors could also discuss possible mitigation

553 strategies (e.g., gated release of models, providing defenses in addition to attacks,

554 mechanisms for monitoring misuse, mechanisms to monitor how a system learns from

555 feedback over time, improving the efficiency and accessibility of ML).

556 **11. Safeguards**

557 Question: Does the paper describe safeguards that have been put in place for responsible

558 release of data or models that have a high risk for misuse (e.g., pretrained language models,

559 image generators, or scraped datasets)?

560 Answer: [NA]

561 Justification: The paper poses no such risks, as it does not involve any data or models.

562 Guidelines:

563 • The answer NA means that the paper poses no such risks.

564 • Released models that have a high risk for misuse or dual-use should be released with

565 necessary safeguards to allow for controlled use of the model, for example by requiring

566 that users adhere to usage guidelines or restrictions to access the model or implementing

567 safety filters.

568 • Datasets that have been scraped from the Internet could pose safety risks. The authors

569 should describe how they avoided releasing unsafe images.

570 • We recognize that providing effective safeguards is challenging, and many papers do

571 not require this, but we encourage authors to take this into account and make a best

572 faith effort.

573 **12. Licenses for existing assets**

574 Question: Are the creators or original owners of assets (e.g., code, data, models), used in

575 the paper, properly credited and are the license and terms of use explicitly mentioned and

576 properly respected?

577 Answer: [NA]

578 Justification: The paper does not use existing assets.  
579 Guidelines:

- 580 • The answer NA means that the paper does not use existing assets.
- 581 • The authors should cite the original paper that produced the code package or dataset.
- 582 • The authors should state which version of the asset is used and, if possible, include a  
583 URL.
- 584 • The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- 585 • For scraped data from a particular source (e.g., website), the copyright and terms of  
586 service of that source should be provided.
- 587 • If assets are released, the license, copyright information, and terms of use in the  
588 package should be provided. For popular datasets, `paperswithcode.com/datasets`  
589 has curated licenses for some datasets. Their licensing guide can help determine the  
590 license of a dataset.
- 591 • For existing datasets that are re-packaged, both the original license and the license of  
592 the derived asset (if it has changed) should be provided.
- 593 • If this information is not available online, the authors are encouraged to reach out to  
594 the asset’s creators.

595 **13. New assets**  
596 Question: Are new assets introduced in the paper well documented and is the documentation  
597 provided alongside the assets?  
598 Answer: [Yes]  
599 Justification: The library is documented and a URL to the documentation is provided in the  
600 introduction.  
601 Guidelines:

- 602 • The answer NA means that the paper does not release new assets.
- 603 • Researchers should communicate the details of the dataset/code/model as part of their  
604 submissions via structured templates. This includes details about training, license,  
605 limitations, etc.
- 606 • The paper should discuss whether and how consent was obtained from people whose  
607 asset is used.
- 608 • At submission time, remember to anonymize your assets (if applicable). You can either  
609 create an anonymized URL or include an anonymized zip file.

610 **14. Crowdsourcing and research with human subjects**  
611 Question: For crowdsourcing experiments and research with human subjects, does the paper  
612 include the full text of instructions given to participants and screenshots, if applicable, as  
613 well as details about compensation (if any)?  
614 Answer: [NA]  
615 Justification: The paper does not involve crowdsourcing nor research with human subjects  
616 Guidelines:

- 617 • The answer NA means that the paper does not involve crowdsourcing nor research with  
618 human subjects.
- 619 • Including this information in the supplemental material is fine, but if the main contribu-  
620 tion of the paper involves human subjects, then as much detail as possible should be  
621 included in the main paper.
- 622 • According to the NeurIPS Code of Ethics, workers involved in data collection, curation,  
623 or other labor should be paid at least the minimum wage in the country of the data  
624 collector.

625 **15. Institutional review board (IRB) approvals or equivalent for research with human  
626 subjects**  
627 Question: Does the paper describe potential risks incurred by study participants, whether  
628 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)  
629 approvals (or an equivalent approval/review based on the requirements of your country or  
630 institution) were obtained?  
631 Answer: [NA]  
632 Justification: The paper does not involve crowdsourcing nor research with human subjects  
633 Guidelines:

- 634 • The answer NA means that the paper does not involve crowdsourcing nor research with  
635 human subjects.
- 636 • Depending on the country in which research is conducted, IRB approval (or equivalent)  
637 may be required for any human subjects research. If you obtained IRB approval, you  
638 should clearly state this in the paper.
- 639 • We recognize that the procedures for this may vary significantly between institutions  
640 and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the  
641 guidelines for their institution.
- 642 • For initial submissions, do not include any information that would break anonymity (if  
643 applicable), such as the institution conducting the review.

644 **16. Declaration of LLM usage**

645 Question: Does the paper describe the usage of LLMs if it is an important, original, or  
646 non-standard component of the core methods in this research? Note that if the LLM is used  
647 only for writing, editing, or formatting purposes and does not impact the core methodology,  
648 scientific rigorousness, or originality of the research, declaration is not required.

649 Answer: [NA]

650 Justification: The core method development in this research does not involve LLMs as any  
651 important, original, or non-standard components

652 Guidelines:

- 653 • The answer NA means that the core method development in this research does not  
654 involve LLMs as any important, original, or non-standard components.
- 655 • Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>)  
656 for what should or should not be described.