

Motivation

Graph kernels are a key tool for machine learning on graphs and networks. They

- capture meaningful notions of node **similarity**.
- provide access to an enriched **feature space**.
- serve as an **upstream** input to many ML algorithms, e.g., for signal reconstruction, classification, or link prediction [1].

Challenge: Graph kernels face **scalability** issues, with prohibitive $\mathcal{O}(N^3)$ explicit computation time.

Contributions: Using tools from graph signal processing and randomized linear algebra, we provide

- **low-rank** graph kernel approximations, through
- explicit **random feature** node representations.
- **accuracy** and **scalability** guarantees.

Graph Kernels

Given an undirected weighted graph with N nodes and E edges, graph kernels take the form

$$\Gamma = h(L) \in \mathbb{R}^{N \times N}, \quad (1)$$

where L is the normalized graph Laplacian, and h a non-negative kernel function [2]. Using the graph Fourier basis (eigenvector basis of L),

$$\Gamma = \mathbf{V}h(\Lambda)\mathbf{V}^\top = \sum_{k=1}^N h(\lambda_k)\mathbf{v}_k\mathbf{v}_k^\top. \quad (2)$$

Interpretation: h preserves specific eigenvalues, or graph frequencies. Typical kernels emphasize (smooth) low-frequency components and suppress high-frequency ones, such that h is non-increasing.

Problem: Full eigendecomposition (ED) is prohibitive to compute ($\mathcal{O}(N^3)$ time complexity). Need of a method which bypasses explicit computation of the eigenvectors \mathbf{V} .

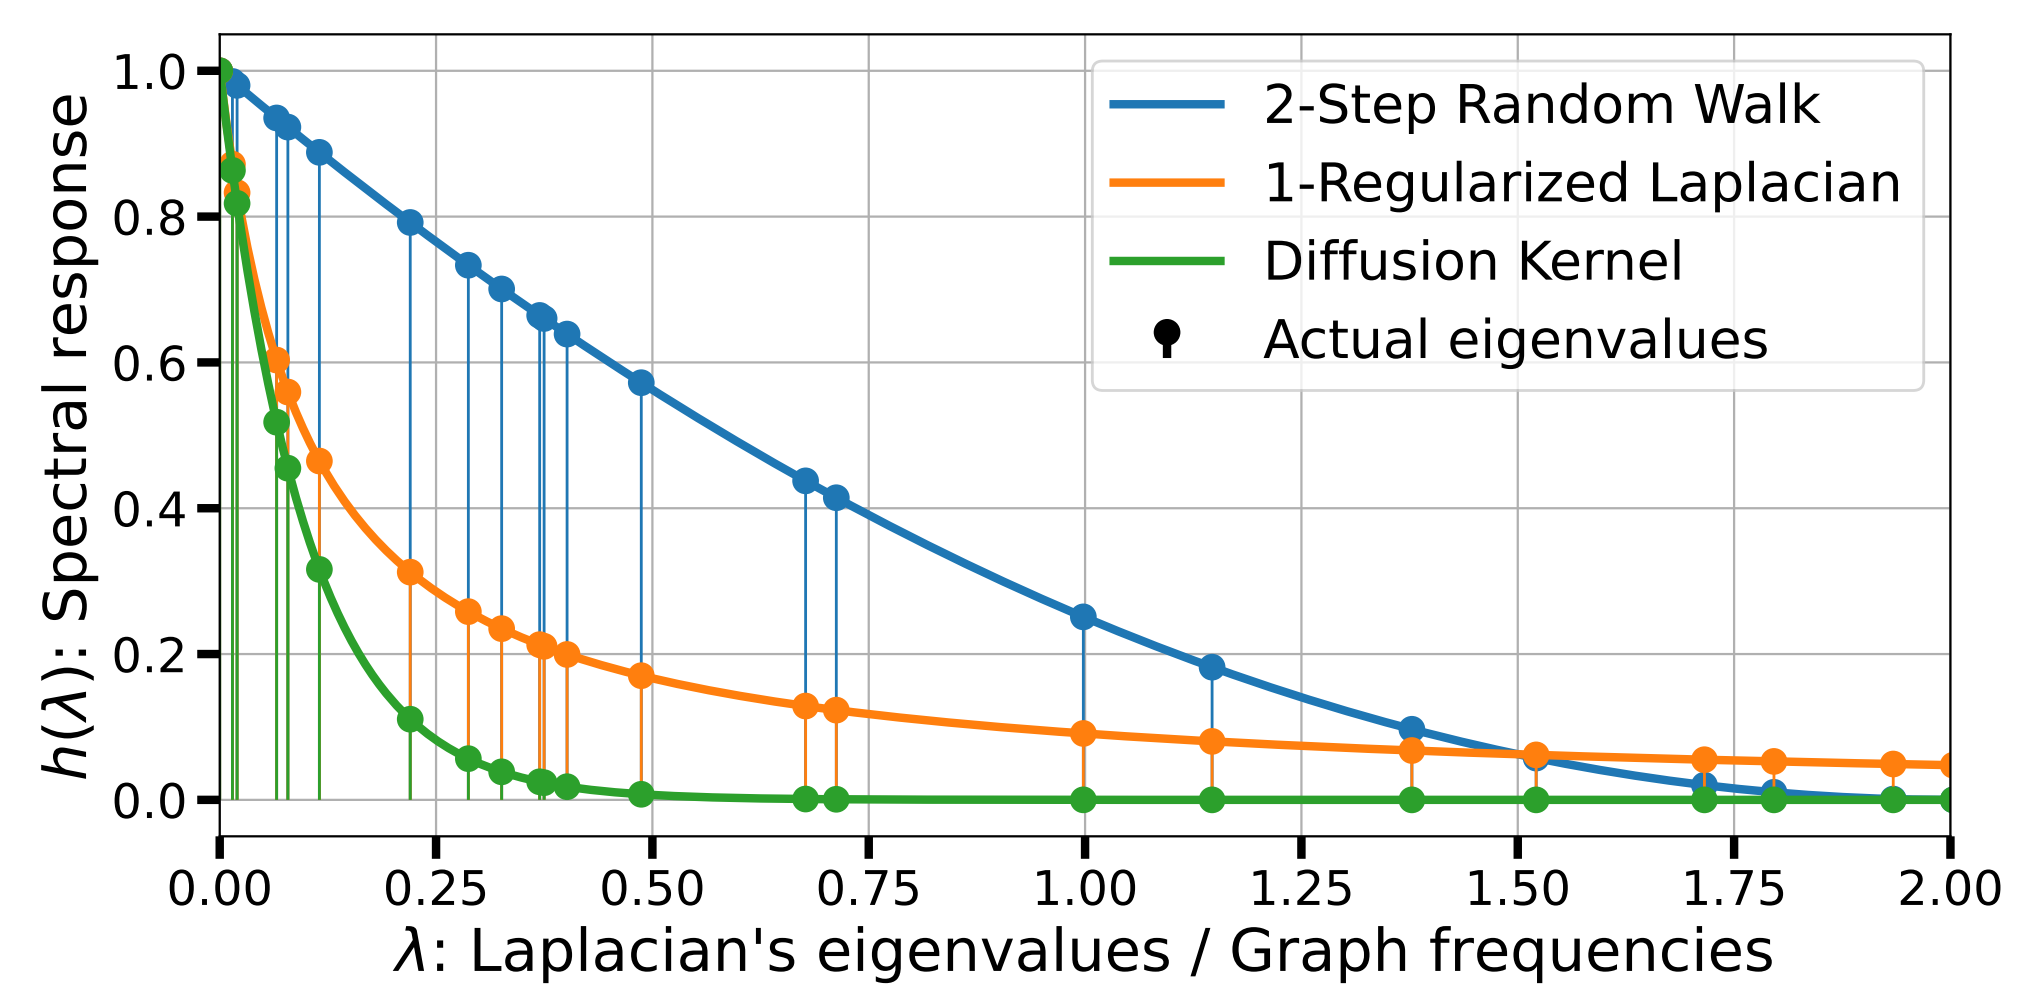


Figure 1: Some examples of graph kernels in the spectral domain (σ controls the bandwidth). Blue: $\Gamma = (I - L/\sigma)^2$, orange: $\Gamma = (I + \sigma^2 L)^{-1}$, green: $\Gamma = \exp(-\sigma^2 L)$.

Random Features

Our approach provides a random feature map

$$\phi : i \in \mathcal{V} \mapsto \phi_i = \Phi \delta_i \in \mathbb{R}^K, \quad (3)$$

with Φ a $K \times N$ feature matrix with columns $\{\phi_i\}_{i=1}^N$, such that kernel values can be approximated by simple inner products:

$$\Gamma_{ij} \approx \tilde{\Gamma}_{ij} = \langle \phi_i, \phi_j \rangle. \quad (4)$$

Benefits:

- **Explicit** node embeddings (feature vectors).
- **Fast** linear kernel evaluation.
- **Reduced** memory requirements.

Our approximation $\tilde{\Gamma} = \Phi^\top \Phi$ is rank- K . The best rank- K approximation to Γ for non-increasing h is

$$\Gamma^{(K)} = \mathbf{V}_{:K}h(\Lambda_{:K})\mathbf{V}_{:K}^\top = \sum_{k=1}^K h(\lambda_k)\mathbf{v}_k\mathbf{v}_k^\top. \quad (5)$$

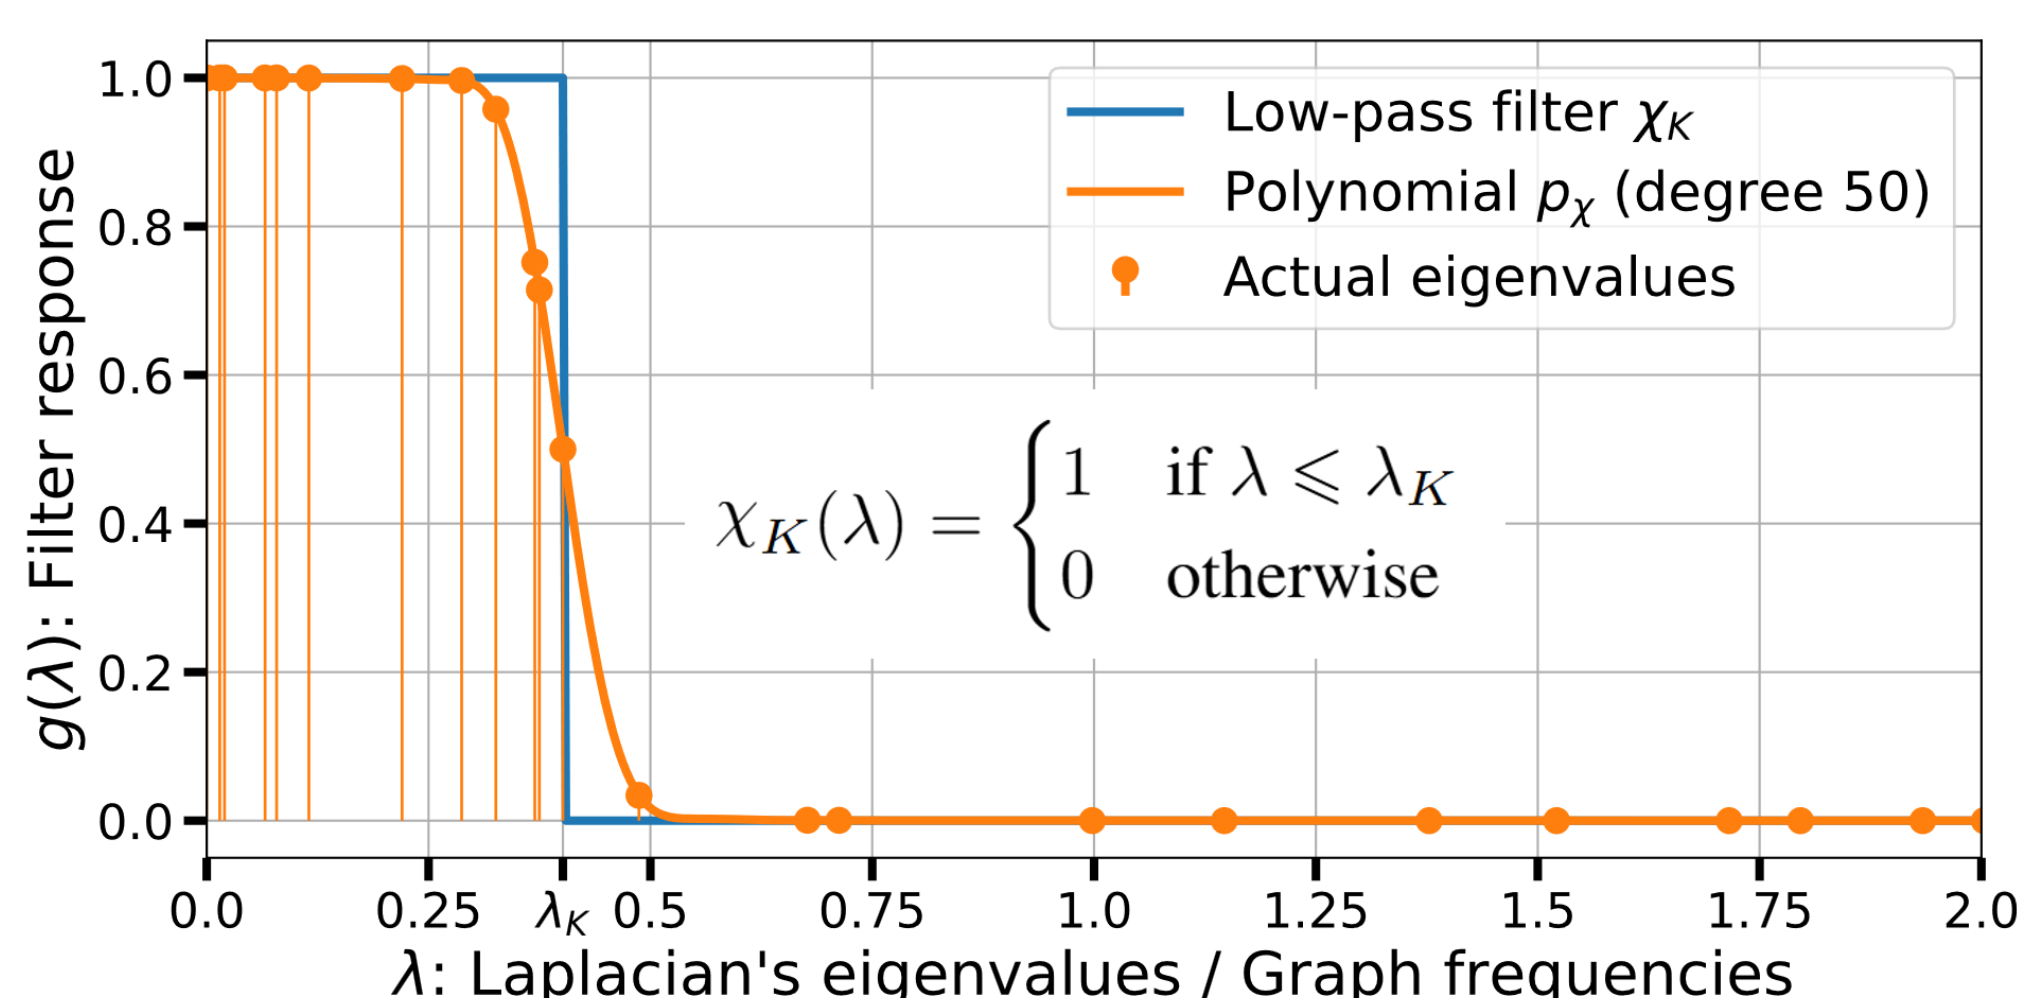


Figure 2: Perfect low-pass filter χ_K and its degree- M polynomial approximation p_χ .

Our Method, Part 1: Range Finding

Task: Approximately identify the subspace spanned by $\mathbf{V}_{:K}$, i.e., find an orthogonal matrix Q such that

$$\text{span}(Q) \approx \text{span}(\mathbf{V}_{:K}). \quad (6)$$

Method: Filter K Gaussian random signals $G \in \mathbb{R}^{N \times K}$ with a low-pass filter.

Idea: If $\chi_K := \mathbb{1}_{[0, \lambda_K]}$ is an ideal low-pass filter, then with probability 1,

$$B = \chi_K(L)G \implies \text{span}(B) = \text{span}(\mathbf{V}_{:K}), \quad (7)$$

and $Q = \text{ortho}(B)$ provides an exact solution to (6).

In practice: Constructing the $\chi_K(L)$ is infeasible without access to $\mathbf{V}_{:K}$. We replace χ_K with a degree- M polynomial approximation p_χ (see Figure 2), and obtain

$$Q = \text{ortho}(p_\chi(L)G). \quad (8)$$

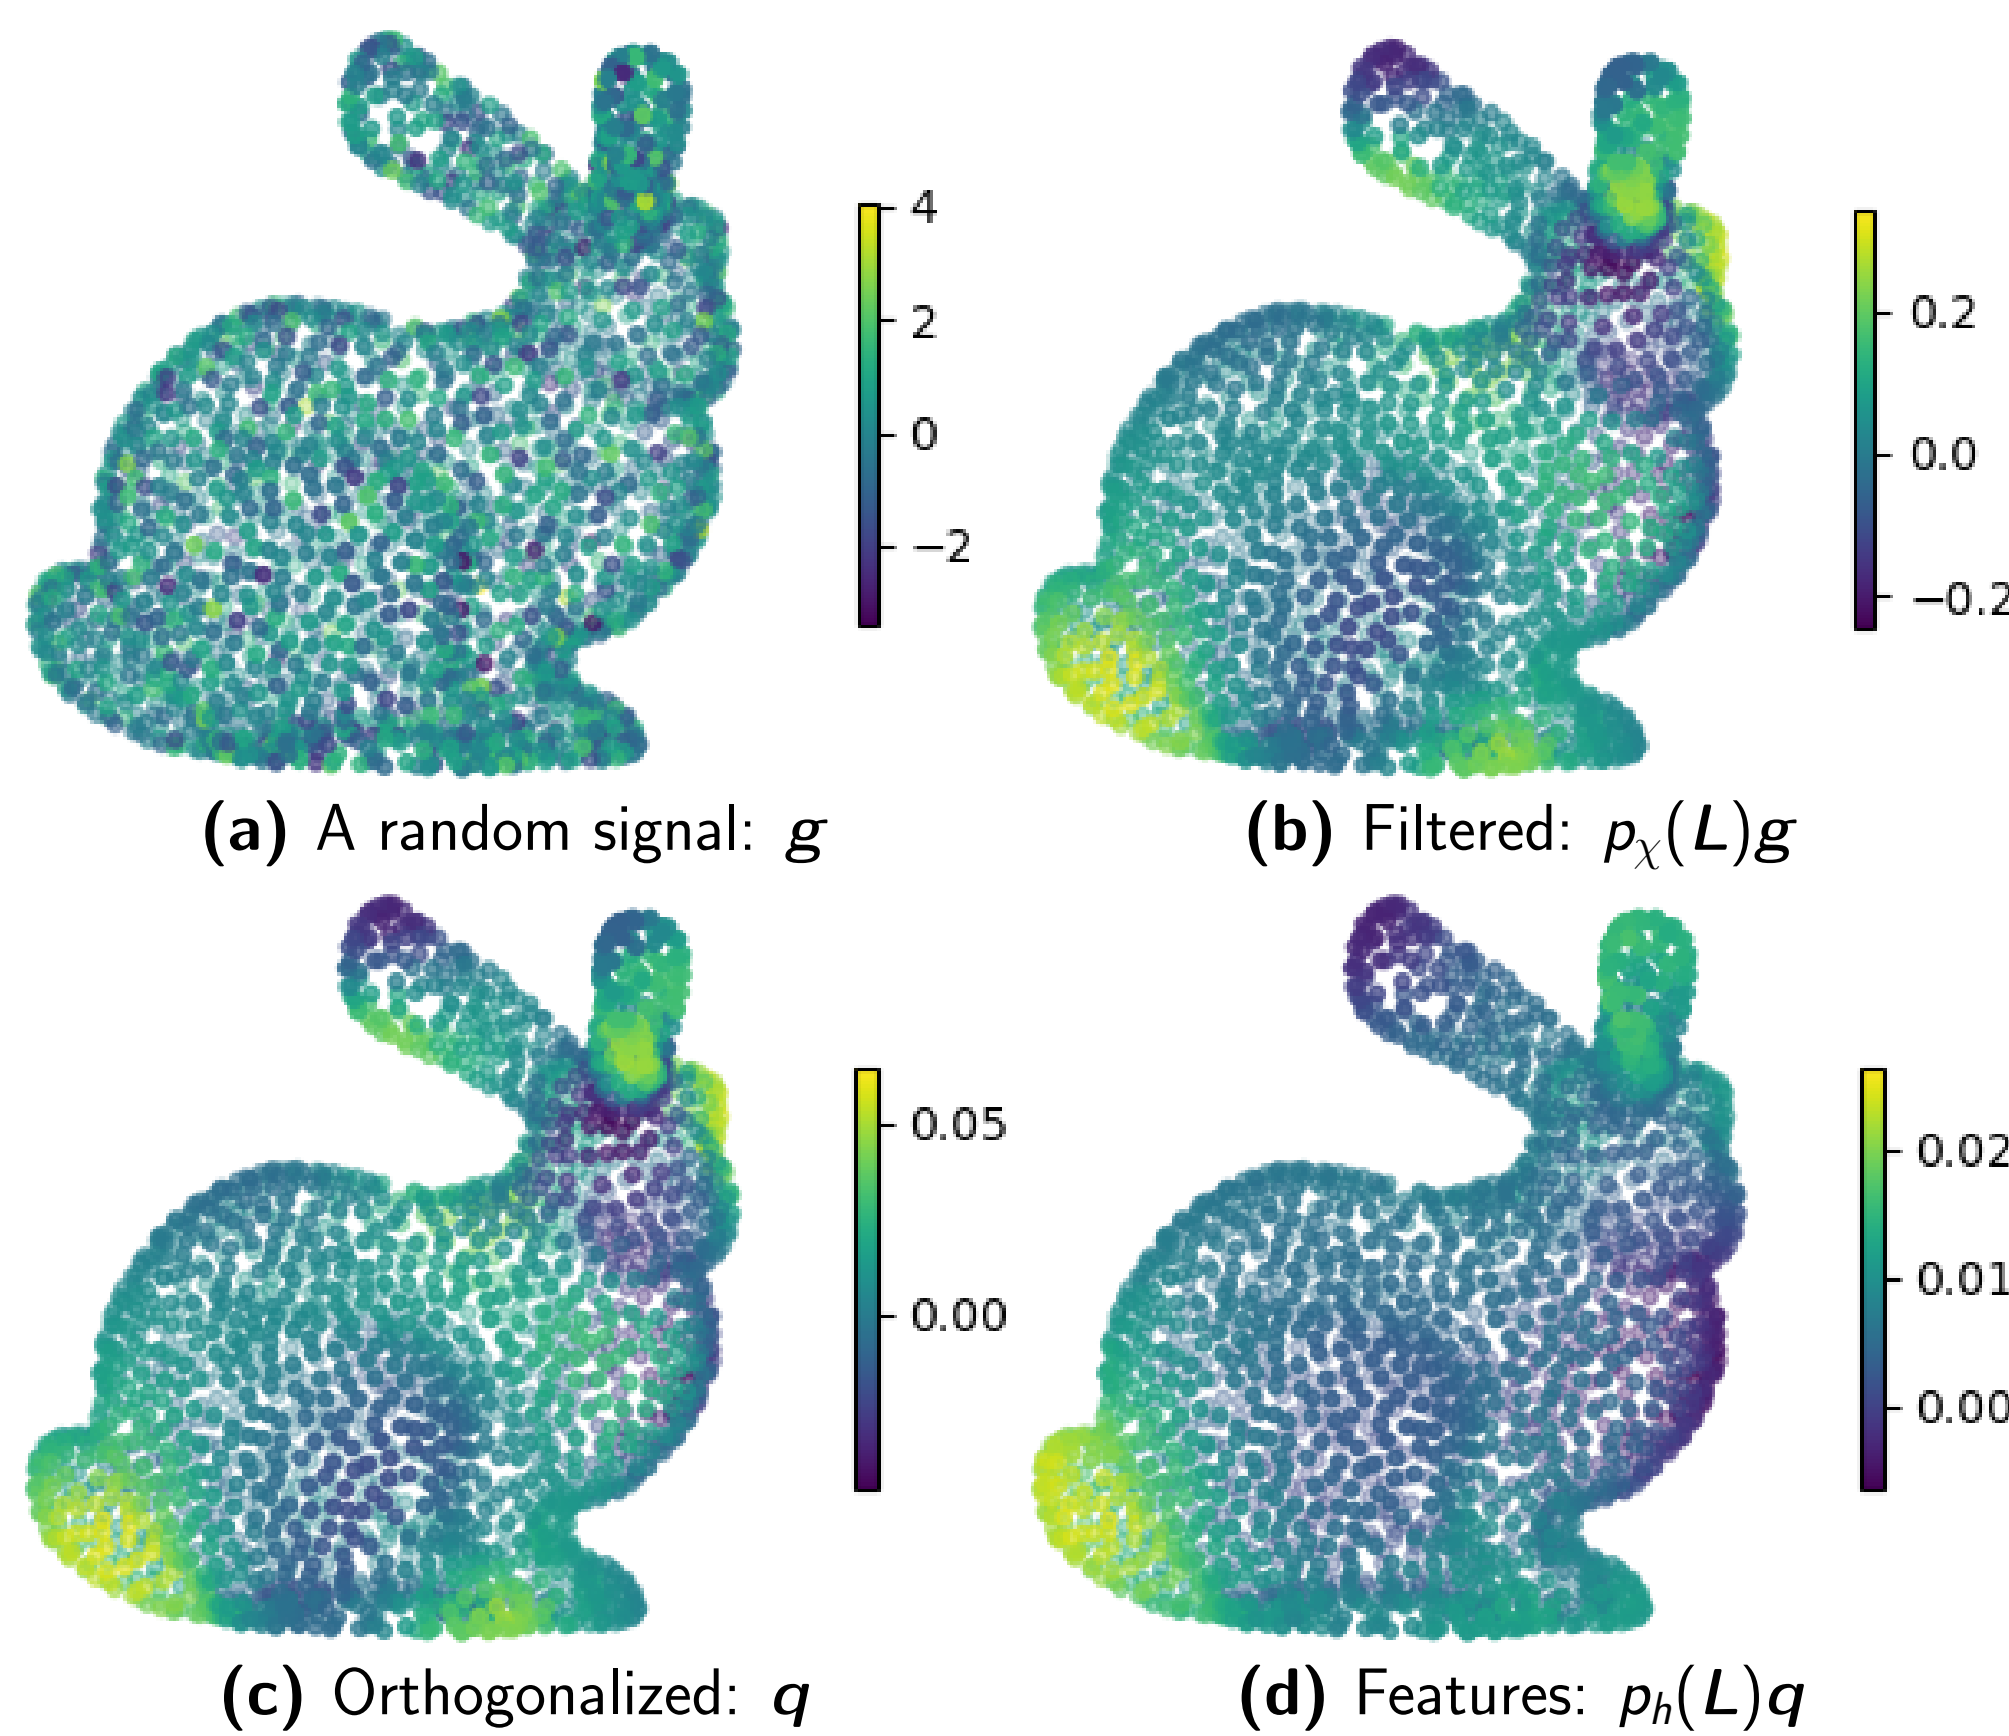


Figure 3: Example of random feature extraction.

Our Method, Part 2: Feature Construction

Task: Construct $\Phi \in \mathbb{R}^{K \times N}$ such that

$$\Phi^\top \Phi = \tilde{\Gamma} \approx \Gamma \quad (9)$$

Idea: The ideal feature matrix for kernel approximation is

$$\Phi^* = (h^{\frac{1}{2}}(L)\mathbf{V}_{:K})^\top, \\ \implies (\Phi^*)^\top \Phi^* = \sum_{k=1}^K h^{\frac{1}{2}}(\lambda_k)\mathbf{v}_k\mathbf{v}_k^\top h^{\frac{1}{2}}(\lambda_k) = \Gamma^{(K)}.$$

In practice: Replace $\mathbf{V}_{:K}$ with Q from (8), and approximate $h^{\frac{1}{2}}$ with a degree- M polynomial p_h , enabling the fast estimation of

$$\Phi = (p_h(L)Q)^\top \rightarrow \text{filter } Q \text{ with } p_h. \quad (10)$$

The full feature extraction process is illustrated on Figure 3.

Error Analysis

We can decompose the overall error $\mathcal{E} := \|\Gamma - \tilde{\Gamma}\|$ as

$$\mathcal{E} \leq \underbrace{\mathcal{E}_R^2}_{\text{Range Finding Error}} + \underbrace{\mathcal{E}_P}_{\text{Polynomial Approx. Error}}. \quad (11)$$

Proposition

Let $\epsilon_{h,M} := \max_{k \in [M]} |h^{\frac{1}{2}}(\lambda_k) - p_h(\lambda_k)|$. Then,


$$\mathcal{E}_P \leq 2\epsilon_{h,M} + \epsilon_{h,M}^2,$$

and

$$\mathbb{E}\mathcal{E}_R \leq h^{\frac{1}{2}}(\lambda_{K+1}) + 2\sqrt{\frac{K}{r-1}}p_\chi(\lambda_{K+1}) \\ + 2e^{\frac{\sqrt{K+r}}{r}}\left(\sum_{j>K} p_\chi(\lambda_j)\right)^{1/2}.$$

Here $r \ll K$ is an oversampling parameter such that $K+r$ is the effective embedding dimension. From the Weierstrass theorem, $\epsilon_{h,M}$ vanishes for large M , as does $p_\chi(\lambda_j)$ for $j > K$.

Notice that $h^{\frac{1}{2}}(\lambda_{K+1}) = \|\Gamma - \Gamma^{(K)}\|$ is the error committed by the best rank- K approximation.



Link to Paper.



Link to Code.

Results: Effect of Bandwidth

We compare our method to g-GRFs [3] on a Swiss-Roll graph ($N = 5000$), for the diffusion kernel $\Gamma = \exp(-\sigma^2 L)$ with varying bandwidth σ . Our approach

- achieves low approximation error for **spectrally localized** kernels.
- captures **global** geometric structure.
- is **complementary** to g-GRFs, which better captures localized behavior.

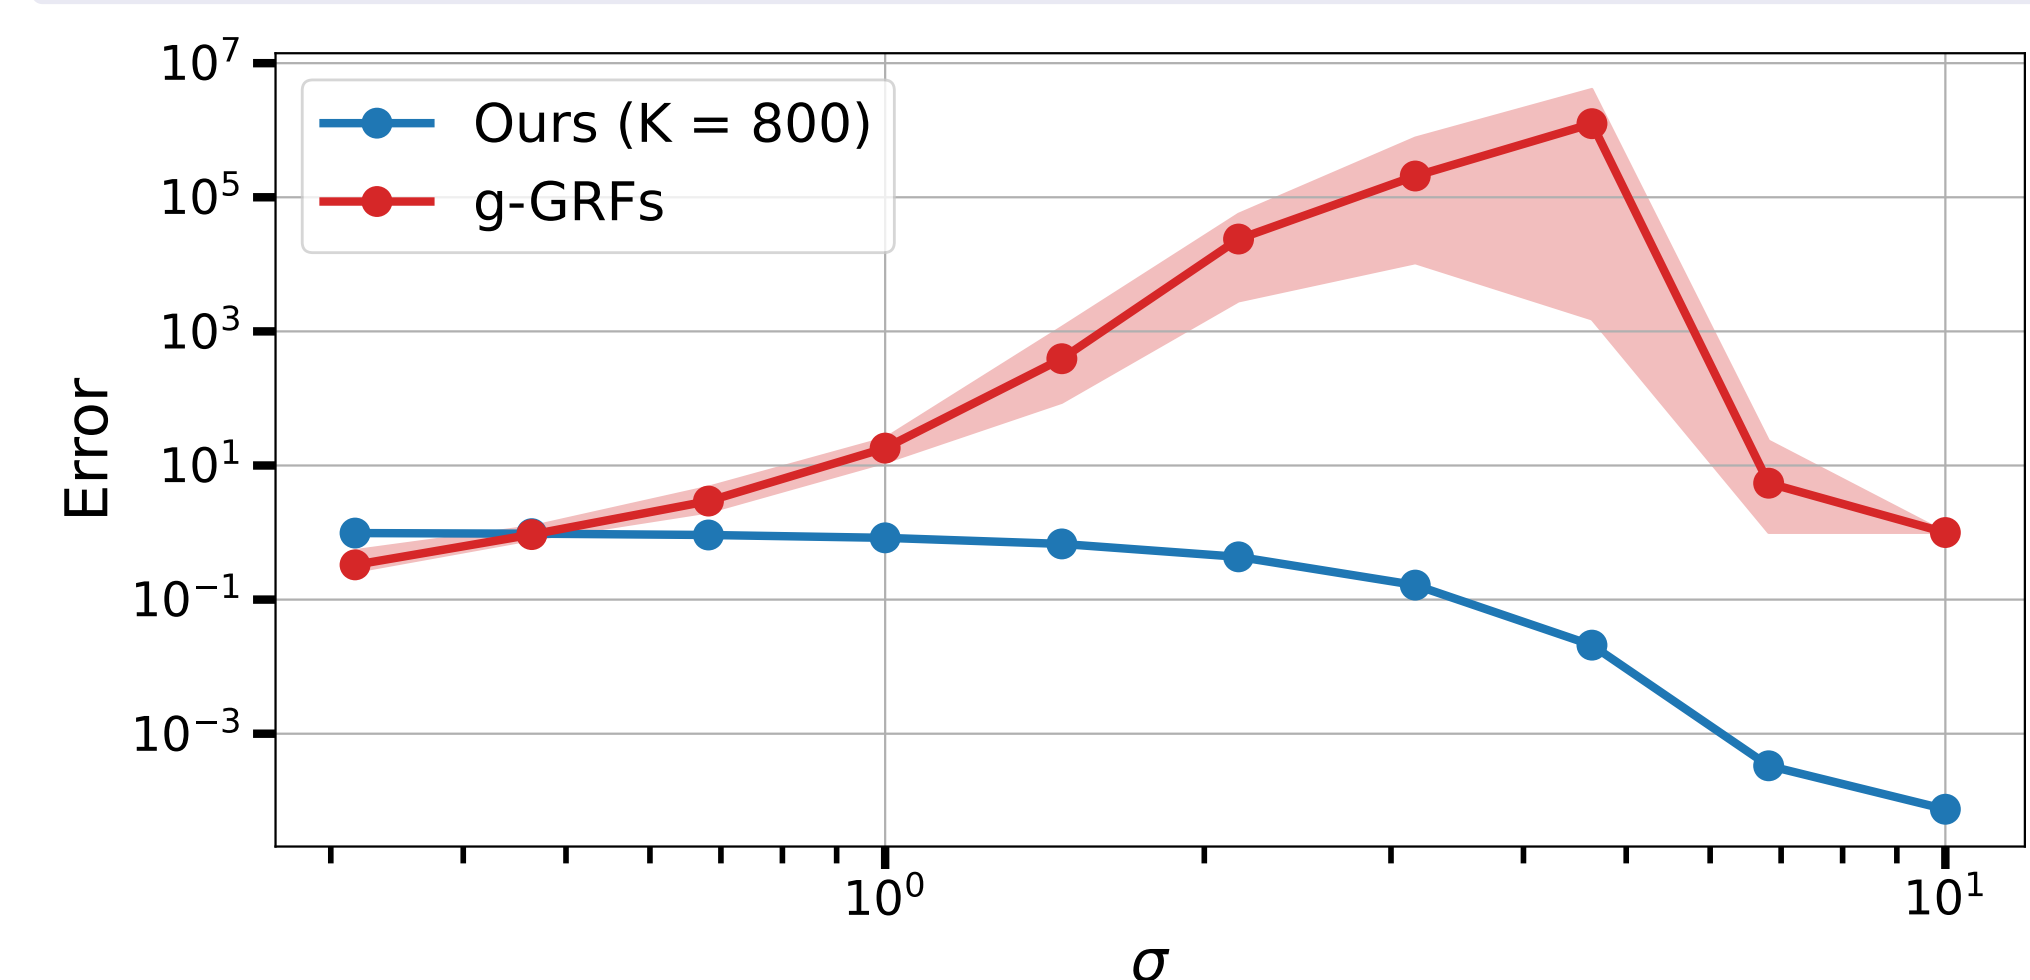


Figure 4: Average relative spectral norm of the error as a function of the bandwidth parameter σ .

Results: Effect of Target Rank

We examine the influence of the parameter K . We verify that

- the error **decreases** with K .
- the error **matches** that of $\Gamma^{(K)}$.
- our algorithm effectively recovers the **dominant kernel subspace**.

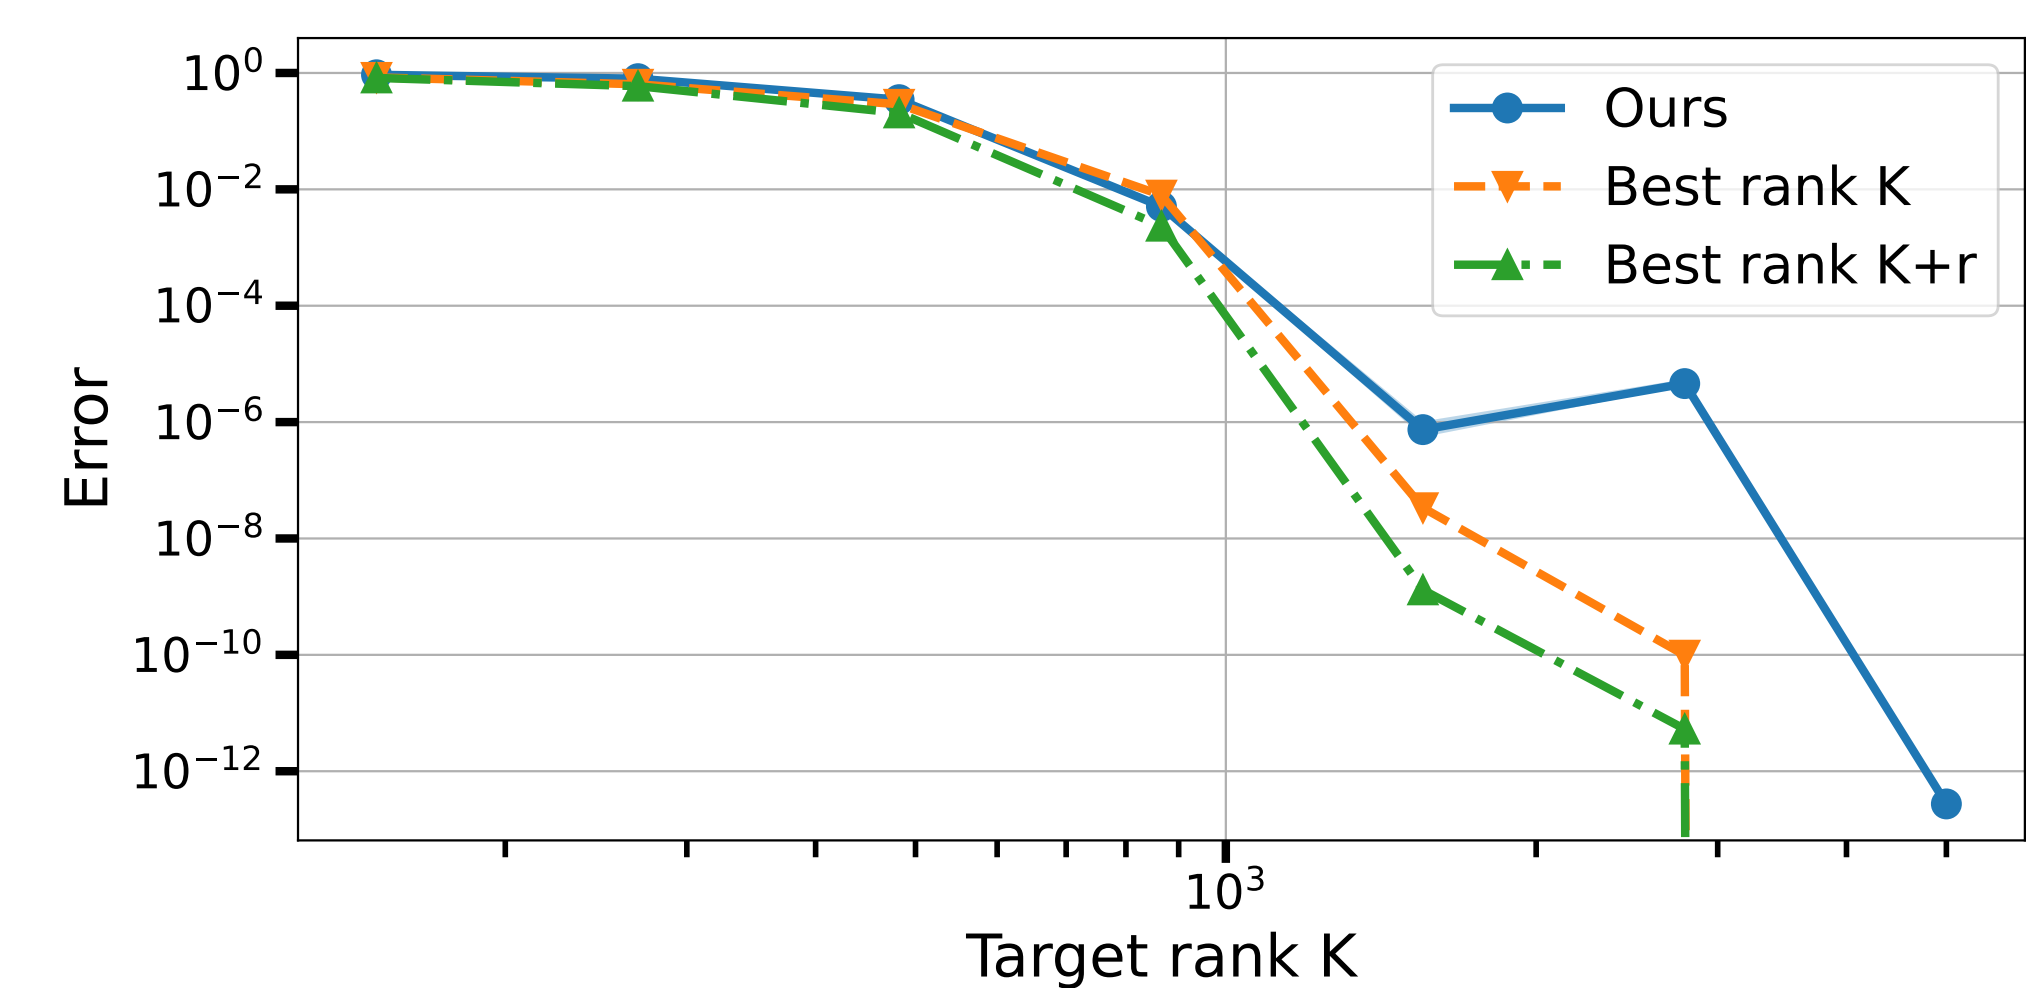


Figure 5: Average relative spectral norm of the error as a function of the target rank K , on a Swiss-Roll graph ($N = 5000$) and for the diffusion kernel ($\sigma = 5$).

Complexity Analysis

Our algorithm achieves **linear time** complexity for sparse graphs and **linear memory** complexity.

	Time	Memory
Full ED	$\mathcal{O}(N^3)$	$\mathcal{O}(N^2)$
Ours	$\mathcal{O}(EMK)$	$\mathcal{O}(NK)$

Table 1: Complexity comparison between explicit eigendecomposition and our approximation method. N : number of nodes, E : number of edges, M : order of polynomial approximations.

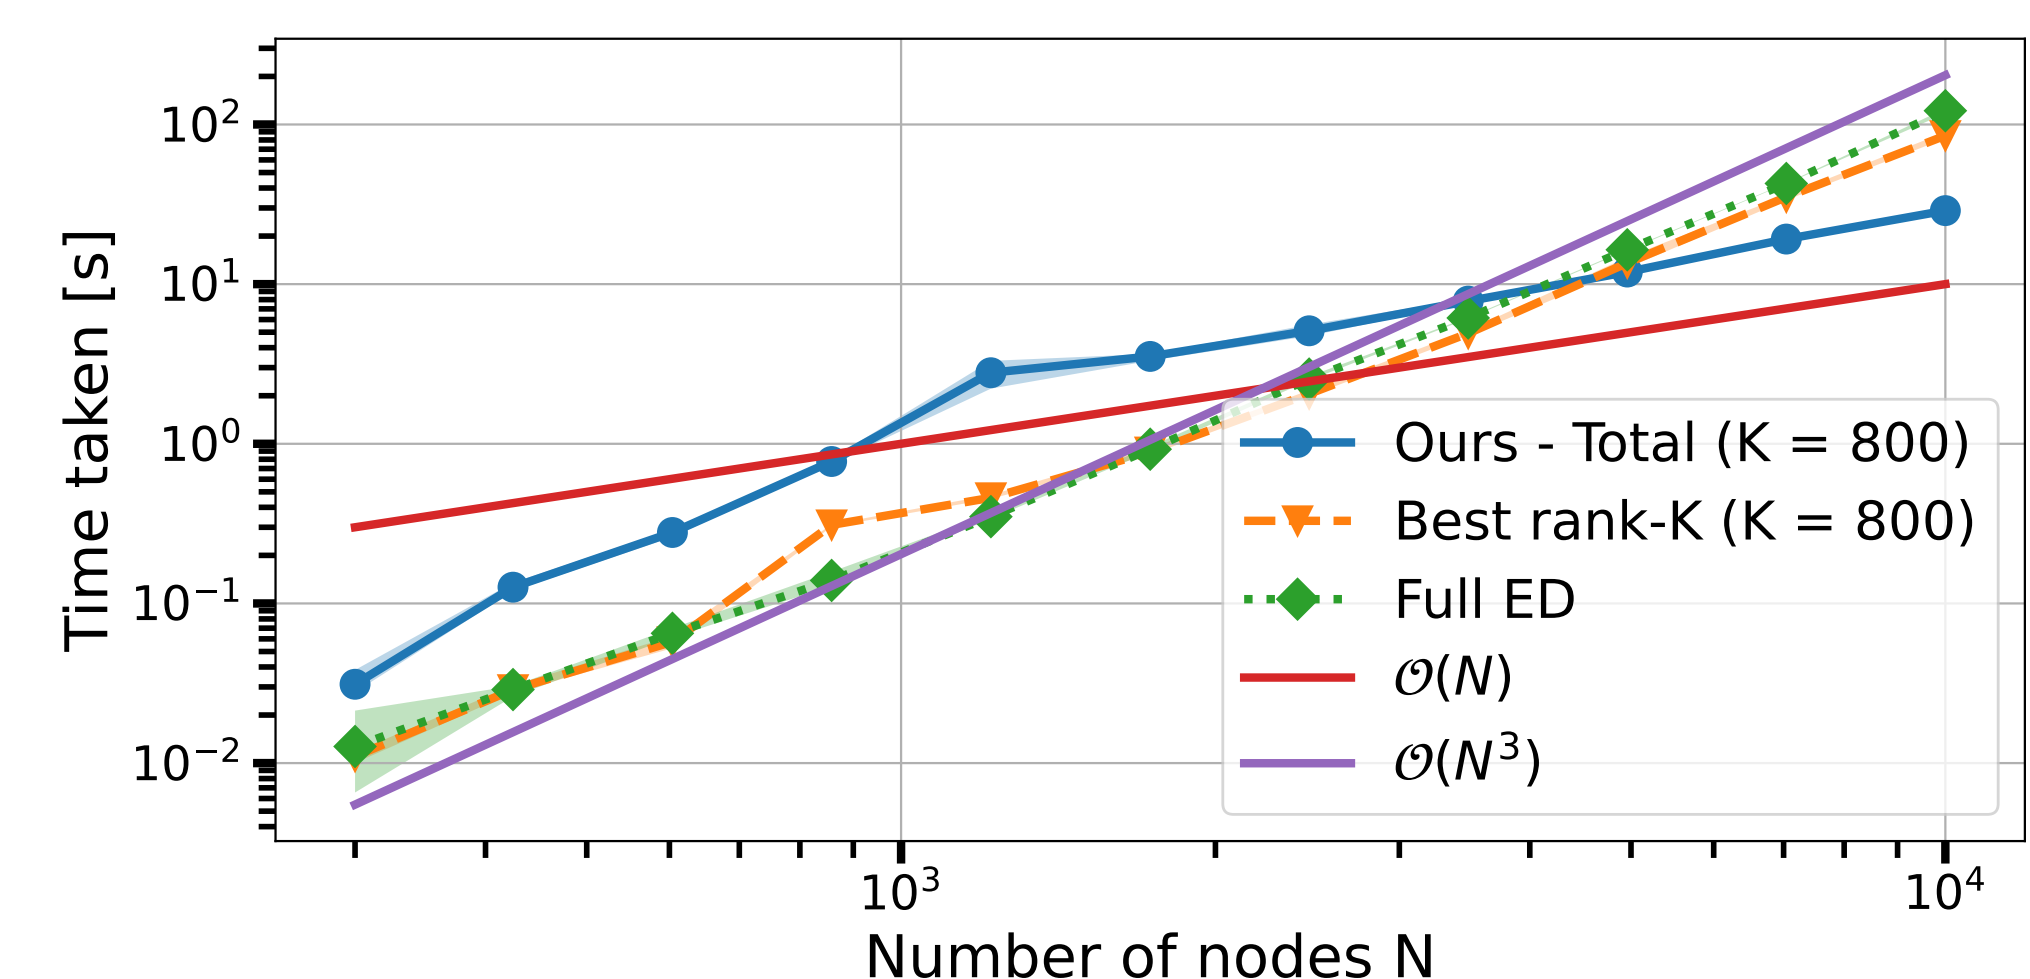


Figure 6: Average computation time as a function of the graph size N , on Swiss-Roll graphs.

References

- [1] A. Ortega, P. Frossard, J. Kovačević, J. M. Moura, and P. Vandergheynst, "Graph signal processing: Overview, challenges, and applications," in *Proc. IEEE*, vol. 106, no. 5, pp.808-828, 2018.
- [2] A. J. Smola and R. Kondor, "Kernels and regularization on graphs," in *16th annu. conf. learn. theory and 7th kernel wks.*, 2003, pp. 144-158.
- [3] I. Reid, K. Choromanski, E. Berger, and A. Weller, "General graph random features," *Int. conf. learn. represent.*, vol. 2024, pp. 28526-28543, 2024.