

GENERAL GRAPH RANDOM FEATURES


Isaac Reid^{1*}, Krzysztof Choromanski^{2,3*}, Eli Berger^{4*}, Adrian Weller^{1,5}

¹University of Cambridge, ²Google DeepMind, ³Columbia University,

⁴University of Haifa, ⁵Alan Turing Institute

ir337@cam.ac.uk, kchoro@google.com

ABSTRACT

We propose a novel random walk-based algorithm for unbiased estimation of arbitrary functions of a weighted adjacency matrix, coined *general graph random features* (g-GRFs). This includes many of the most popular examples of kernels defined on the nodes of a graph. Our algorithm enjoys subquadratic time complexity with respect to the number of nodes, overcoming the notoriously prohibitive cubic scaling of exact graph kernel evaluation. It can also be trivially distributed across machines, permitting learning on much larger networks. At the heart of the algorithm is a *modulation function* which upweights or downweights the contribution from different random walks depending on their lengths. We show that by parameterising it with a neural network we can obtain g-GRFs that give higher-quality kernel estimates or perform efficient, scalable kernel learning. We provide robust theoretical analysis and support our findings with experiments including pointwise estimation of fixed graph kernels, solving non-homogeneous graph ordinary differential equations, node clustering and kernel regression on triangular meshes. 

1 INTRODUCTION AND RELATED WORK

The *kernel trick* is a powerful technique to perform nonlinear inference using linear learning algorithms (Campbell, 2002; Kontorovich et al., 2008; Canu and Smola, 2006; Smola and Schölkopf, 2002). Supposing we have a set of N datapoints $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^N$, it replaces Euclidean dot products $\mathbf{x}_i^\top \mathbf{x}_j$ with evaluations of a *kernel function* $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, capturing the ‘similarity’ of the datapoints by instead taking an inner product between implicit (possibly infinite-dimensional) feature vectors in some Hilbert space \mathbb{H}_K .

An object of key importance is the *Gram matrix* $\mathbf{K} \in \mathbb{R}^{N \times N}$ whose entries enumerate the pairwise kernel evaluations, $\mathbf{K} := [K(\mathbf{x}_i, \mathbf{x}_j)]_{i,j=1}^N$. Despite the theoretical rigour and empirical success enjoyed by kernel-based learning algorithms, the requirement to manifest and invert this matrix leads to notoriously poor $\mathcal{O}(N^3)$ time-complexity scaling. This has spurred research dedicated to efficiently approximating \mathbf{K} , the chief example of which is *random features* (Rahimi and Recht, 2007): a Monte-Carlo approach which gives explicitly manifested, finite dimensional vectors $\phi(\mathbf{x}_i) \in \mathbb{R}^m$ whose Euclidean dot product is equal to the kernel evaluation in expectation,

$$\mathbf{K}_{ij} = \mathbb{E} [\phi(\mathbf{x}_i)^\top \phi(\mathbf{x}_j)]. \quad (1)$$

This allows one to construct a *low-rank decomposition* of \mathbf{K} which provides much better scalability. Testament to its utility, a rich taxonomy of random features exists to approximate many different Euclidean kernels including the Gaussian, softmax, and angular and linear kernels (Johnson, 1984; Dasgupta et al., 2010; Goemans and Williamson, 2004; Choromanski et al., 2020).

Kernels defined on discrete input spaces, e.g. $K : \mathcal{N} \times \mathcal{N} \rightarrow \mathbb{R}$ with \mathcal{N} the set of nodes of a graph \mathcal{G} (Smola and Kondor, 2003; Kondor and Lafferty, 2002; Chung and Yau, 1999),

*Equal contribution.

¹Code is available at https://github.com/isaac-reid/general_graph_random_features.

enjoy widespread applications including in bioinformatics (Borgwardt et al., 2005), community detection (Kloster and Gleich, 2014) and recommender systems (Yajima, 2006). More recently, they have been used in applications as diverse as manifold learning for deep generative modelling (Zhou et al., 2020) and for solving single- and multiple-source shortest path problems (Crane et al., 2017). However, for these graph-based kernel methods the problem of poor scalability is particularly acute. This is because even computing the corresponding Gram matrix \mathbf{K} is typically of at least cubic time complexity in the number of nodes N , requiring e.g. the inversion of an $N \times N$ matrix or computation of multiple matrix-matrix products. Despite the presence of this computational bottleneck, random feature methods for graph kernels have proved elusive. Indeed, only recently was a viable *graph random feature* (GRF) mechanism proposed by Choromanski (2023). Their algorithm uses an ensemble of random walkers which deposit a ‘load’ at every vertex they pass through that depends on i) the product of weights of edges traversed by the walker and ii) the marginal probability of the subwalk. Using this scheme, it is possible to construct random features $\{\phi(i)\}_{i=1}^N \subset \mathbb{R}^N$ such that $\phi(i)^\top \phi(j)$ gives an unbiased approximation to the ij -th matrix element of the 2-regularised Laplacian kernel. Multiple independent approximations can be combined to estimate the d -regularised Laplacian kernel with $d \neq 2$ or the diffusion kernel (although the latter is only asymptotically unbiased). The GRFs algorithm enjoys both subquadratic time complexity and strong empirical performance on tasks like k -means node clustering, and it is trivial to distribute across machines when working with massive graphs.

However, a key limitation of GRFs is that they only address a limited family of graph kernels which may not be suitable for the task at hand. Our central contribution is a simple modification which generalises the algorithm to *arbitrary functions of a weighted adjacency matrix*, allowing efficient and unbiased approximation a much broader class of graph node kernels. We achieve this by introducing an extra *modulation function* f that controls each walker’s load as it traverses the graph. As well as empowering practitioners to approximate many more fixed kernels, we demonstrate that f can also be parameterised by a neural network and learned. We use this powerful approach to optimise g-GRFs for higher-quality approximation of fixed kernels and for scalable implicit kernel learning.

The remainder of the manuscript is organised as follows. In **Sec. 2** we introduce *general graph random features* (g-GRFs) and prove that they enable scalable approximation of *arbitrary* functions of a weighted adjacency matrix, including many of the most popular examples of kernels defined on the nodes of a graph. We also extend the core algorithm with *neural modulation functions*, replacing one component of the g-GRF mechanism with a neural network, and derive generalisation bounds for the corresponding class of learnable graph kernels (**Sec. 2.1**). In **Sec. 3** we run extensive experiments, including: pointwise estimation of a variety of popular graph kernels (**Sec. 3.1**); simulation of time evolution under non-homogeneous graph ordinary differential equations (**Sec. 3.2**); kernelised k -means node clustering including on large graphs (**Sec. 3.3**); training a neural modulation function to suppress the mean square error of fixed kernel estimates (**Sec 3.4**); and training a neural modulation function to learn a kernel for node attribute prediction on triangular mesh graphs (**Sec. 3.5**).

2 GENERAL GRAPH RANDOM FEATURES

Consider a directed weighted graph $\mathcal{G}(\mathcal{N}, \mathcal{E}, \mathbf{W} := [w_{ij}]_{i,j \in \mathcal{N}})$ where: i) $\mathcal{N} := \{1, \dots, N\}$ is the set of nodes; ii) \mathcal{E} is the set of edges, with $(i, j) \in \mathcal{E}$ if there is a directed edge from i to j in \mathcal{G} ; and iii) \mathbf{W} is the *weighted adjacency matrix*, with w_{ij} the weight of the directed edge from i to j (equal to 0 if no such edge exists). Note that an undirected graph can be described as directed with the symmetric weighted adjacency matrix \mathbf{W} . Now consider the matrices $\mathbf{K}_\alpha(\mathbf{W}) \in \mathbb{R}^{N \times N}$, where $\alpha = (\alpha_k)_{k=0}^\infty$ and $\alpha_k \in \mathbb{R}$:

$$\mathbf{K}_\alpha(\mathbf{W}) = \sum_{k=0}^{\infty} \alpha_k \mathbf{W}^k. \quad (2)$$

We assume that the sum above converges for all \mathbf{W} under consideration, which can be ensured with a regulariser $\mathbf{W} \rightarrow \sigma \mathbf{W}$, $\sigma \in \mathbb{R}_+$. Without loss of generality, we also assume

that α is *normalised* such that $\alpha_0 = 1$. The matrix $\mathbf{K}_\alpha(\mathbf{W})$ can be associated with a graph function $K_\alpha^\mathcal{G} : \mathcal{N} \times \mathcal{N} \rightarrow \mathbb{R}$ mapping from a pair of graph nodes to a real number.

Note that if \mathcal{G} is an undirected graph then $\mathbf{K}_\alpha(\mathbf{W})$ automatically inherits the symmetry of \mathbf{W} . In this case, it follows from Weyl’s perturbation inequality (Bai et al., 2000) that $\mathbf{K}_\alpha(\mathbf{W})$ is positive semidefinite for any given α provided the spectral radius $\rho(\mathbf{W}) := \max_{\lambda \in \Lambda(\mathbf{W})} (|\lambda|)$ is sufficiently small (with $\Lambda(\mathbf{W})$ the set of eigenvalues of \mathbf{W}). This can again be ensured by multiplying the weight matrix \mathbf{W} by a regulariser $\sigma \in \mathbb{R}_+$. It then follows that $\mathbf{K}_\alpha(\mathbf{W})$ can be considered the Gram matrix of a *graph kernel function* $K_\alpha^\mathcal{G}$.

With suitably chosen $\alpha = (\alpha_k)_{k=0}^\infty$, the class described by Eq. 2 includes many popular examples of graph node kernels in the literature (Smola and Kondor, 2003; Chapelle et al., 2002). They measure connectivity between nodes and are typically functions of the *graph Laplacian matrix*, defined by $\mathbf{L} := \mathbf{I} - \widetilde{\mathbf{W}}$ with $\widetilde{\mathbf{W}} := [w_{ij}/\sqrt{d_i d_j}]_{i,j=1}^N$. Here, $d_i := \sum_j w_{ij}$ is the weighted degree of node i such that $\widetilde{\mathbf{W}}$ is the *normalised* weighted adjacency matrix. For reference, Table 1 gives the kernel definitions and normalised coefficients α_k (corresponding to powers of $\widetilde{\mathbf{W}}$) to be considered later in the manuscript. In practice, factors in α_k equal to a quantity raised to the power of k are absorbed into the normalisation of $\widetilde{\mathbf{W}}$.

Name	Form	α_k
d -regularised Laplacian	$(\mathbf{I}_N + \sigma^2 \mathbf{L})^{-d}$	$\binom{d+k-1}{k} (1 + \sigma^{-2})^{-k}$
p -step random walk	$(\alpha \mathbf{I}_N - \mathbf{L})^p, \alpha \geq 2$	$\binom{p}{k} (\alpha - 1)^{-k}$
Diffusion	$\exp(-\sigma^2 \mathbf{L}/2)$	$\frac{1}{k!} (\frac{\sigma^2}{2})^k$
Inverse Cosine	$\cos(\mathbf{L}\pi/4)$	$\frac{1}{k!} (\frac{\pi}{4})^k \cdot \left\{ (-1)^{\frac{k}{2}} \text{ if } k \text{ even}, (-1)^{\frac{k-1}{2}} \text{ if } k \text{ odd} \right\}$

Table 1: Different graph functions/kernels $K_\alpha^\mathcal{G} : \mathcal{N} \times \mathcal{N} \rightarrow \mathbb{R}$. The exp and cos mappings are defined via Taylor series expansions rather than element-wise, e.g. $\exp(\mathbf{M}) := \lim_{n \rightarrow \infty} (\mathbf{I}_N + \mathbf{M}/n)^n$ and $\cos(\mathbf{M}) := \text{Re}(\exp(i\mathbf{M}))$. σ and α are regularisers. Note that the diffusion kernel is sometimes instead defined by $\exp(\sigma^2(\mathbf{I}_N - \mathbf{L}))$ but these forms are equivalent up to normalisation.

The chief goal of this work is to construct a *random feature map* $\phi(i) : \mathcal{N} \rightarrow \mathbb{R}^l$ with $l \in \mathbb{N}$ that provides unbiased approximation of $\mathbf{K}_\alpha(\mathbf{W})$ as in Eq. 1. To do so, we consider the following algorithm.

Algorithm 1 Constructing a random feature vector $\phi_f(i) \in \mathbb{R}^N$ to approximate $\mathbf{K}_\alpha(\mathbf{W})$

Input: weighted adjacency matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$, vector of unweighted node degrees (no. neighbours) $\mathbf{d} \in \mathbb{R}^N$, modulation function $f : (\mathbb{N} \cup \{0\}) \rightarrow \mathbb{R}$, termination probability $p_{\text{halt}} \in (0, 1)$, node $i \in \mathcal{N}$, number of random walks to sample $m \in \mathbb{N}$.

Output: random feature vector $\phi_f(i) \in \mathbb{R}^N$

```

1: initialise:  $\phi_f(i) \leftarrow \mathbf{0}$ 
2: for  $w = 1, \dots, m$  do
3:   initialise:  $\text{load} \leftarrow 1$ 
4:   initialise:  $\text{current\_node} \leftarrow i$ 
5:   initialise:  $\text{terminated} \leftarrow \text{False}$ 
6:   initialise:  $\text{walk\_length} \leftarrow 0$ 
7:   while  $\text{terminated} = \text{False}$  do
8:      $\phi_f(i)[\text{current\_node}] \leftarrow \phi_f(i)[\text{current\_node}] + \text{load} \times f(\text{walk\_length})$ 
9:      $\text{walk\_length} \leftarrow \text{walk\_length} + 1$ 
10:     $\text{new\_node} \leftarrow \text{Unif}[\mathcal{N}(\text{current\_node})]$  ▷ assign to one of neighbours
11:     $\text{load} \leftarrow \text{load} \times \frac{d[\text{current\_node}]}{1 - p_{\text{halt}}} \times \mathbf{W}[\text{current\_node}, \text{new\_node}]$  ▷ update load
12:     $\text{current\_node} \leftarrow \text{new\_node}$ 
13:     $\text{terminated} \leftarrow (t \sim \text{Unif}(0, 1) < p_{\text{halt}})$  ▷ draw RV  $t$  to decide on termination
14:   end while
15: end for
16: normalise:  $\phi_f(i) \leftarrow \phi_f(i)/m$ 

```




Figure 1: Schematic for a random walk on a graph (solid red) and an accompanying modulation function f (dashed blue) used to approximate an arbitrary graph node function K^G .

This is identical to the algorithm presented by Choromanski (2023) for constructing features to approximate the 2-regularised Laplacian kernel, apart from the presence of the extra *modulation function* $f : (\mathbb{N} \cup \{0\}) \rightarrow \mathbb{R}$ in line 8 that upweights or downweights contributions from walks depending on their length (see Fig. 1). We refer to ϕ_f as *general graph random features* (g-GRFs), where the subscript f identifies the modulation function. Crucially, the time complexity of Alg. 1 is subquadratic in the number of nodes N , in contrast to exact methods which are $\mathcal{O}(N^3)$.²

We now state the following central result, proved in App. A.2.

Theorem 2.1 (Unbiased approximation of K_α via convolutions). *For two modulation functions: $f_1, f_2 : (\mathbb{N} \cup \{0\}) \rightarrow \mathbb{R}$, g-GRFs $(\phi_{f_1}(i))_{i=1}^N, (\phi_{f_2}(i))_{i=1}^N$ constructed according to Alg. 1 give unbiased approximation of K_α ,*

$$[K_\alpha]_{ij} = \mathbb{E} [\phi_{f_1}(i)^\top \phi_{f_2}(j)], \quad (3)$$

for kernels with an arbitrary Taylor expansion $\alpha = (\alpha_k)_{k=0}^\infty$ provided that $\alpha = f_1 * f_2$. Here, $*$ is the discrete convolution of the modulation functions f_1, f_2 ; that is, for all $k \in (\mathbb{N} \cup \{0\})$,

$$\sum_{p=0}^k f_1(k-p)f_2(p) = \alpha_k. \quad (4)$$

Clearly the class of pairs of modulation functions f_1, f_2 that satisfy Eq. 4 is highly degenerate. Indeed, it is possible to solve for f_1 given any f_2 and α provided $f_2(0) \neq 0$. For instance, a trivial solution is given by: $f_1(i) = \alpha_i$, $f_2(i) = \mathbb{I}(i=0)$ with $\mathbb{I}(\cdot)$ the indicator function. In this case, the walkers corresponding to f_2 are ‘lazy’, depositing all their load at the node at which they begin. Contributions to the estimator $\phi_{f_1}(i)^\top \phi_{f_2}(j)$ only come from walkers initialised at i traversing all the way to j rather than two walkers both passing through an intermediate node. Also of great interest is the case of *symmetric modulation functions* $f_1 = f_2$, where now intersections do contribute. In this case, the following is true (proved in App. A.3).

Theorem 2.2 (Computing symmetric modulation functions). *Supposing $f_1 = f_2 = f$, Eq. 4 is solved by a function f which is unique (up to a sign) and is given by*

$$f(i) = \pm \sum_{n=0}^i \binom{\frac{1}{2}}{n} \sum_{\substack{k_1+2k_2+3k_3+\dots=i \\ k_1+k_2+k_3+\dots=n}} \binom{n}{k_1 k_2 k_3 \dots} (\alpha_1^{k_1} \alpha_2^{k_2} \alpha_3^{k_3} \dots). \quad (5)$$

Moreover, $f(i)$ can be efficiently computed with the iterative formula

$$\begin{cases} f(0) = \pm\sqrt{\alpha_0} = \pm 1 \\ f(i+1) = \frac{\alpha_{i+1} - \sum_{p=0}^{i-1} f(i-p)f(p+1)}{2f(0)} \end{cases} \quad \text{for } i \geq 0. \quad (6)$$

²Concretely, Alg. 1 yields a pair of matrices $\phi_{1,2} := (\phi(i))_{i=1}^N \in \mathbb{R}^{N \times N}$ such that $K = \mathbb{E}(\phi_1 \phi_2^\top)$ in subquadratic time. Of course, explicitly multiplying the matrices to evaluate every element of \widehat{K} would be $\mathcal{O}(N^3)$, but we avoid this since in applications we just evaluate $\phi_1(\phi_2^\top v)$ where $v \in \mathbb{R}^N$ is some vector and the brackets give the order of computation. This is $\mathcal{O}(N^2)$.

For symmetric modulation functions, the random features $\phi_{f_1}(i)$ and $\phi_{f_2}(i)$ are identical apart from the particular sample of random walks used to construct them. They cannot share the same sample or estimates of diagonal kernel elements $[\mathbf{K}_\alpha]_{ii}$ will be biased.

Computational cost: Note that when running Alg. 1 one only needs to evaluate the modulation functions $f_{1,2}(i)$ up to the length of the longest walk one samples. A batch of size b , $(f_{1,2}(i))_{i=1}^b$, can be pre-computed in time $\mathcal{O}(b^2)$ and reused for random features corresponding to different nodes and even different graphs. Further values of f can be computed at runtime if b is too small and also reused for later computations. Moreover, the minimum length b required to ensure that all m walks are shorter than b with high probability ($\Pr(\cup_{i=1}^m \text{len}(\omega_i) \leq b) > 1 - \delta$, $\delta \ll 1$) scales only logarithmically with m (see App. A.1). This means that, despite estimating a much more general family of graph functions, g-GRFs are essentially no more expensive than the original GRF algorithm. Moreover, any techniques used for dimensionality reduction of regular GRFs (e.g. applying the Johnson-Lindenstrauss transform (Dasgupta et al., 2010) or using ‘anchor points’ (Choromanski, 2023)) can also be used with g-GRFs, providing further efficiency gains.

Generating functions: Inserting the constraint for unbiasedness in Eq. 4 back into the definition of $\mathbf{K}_\alpha(\mathbf{W})$, we immediately have that

$$\mathbf{K}_\alpha(\mathbf{W}) = \mathbf{K}_{f_1}(\mathbf{W})\mathbf{K}_{f_2}(\mathbf{W}) \quad (7)$$

where $\mathbf{K}_{f_1}(\mathbf{W}) := \sum_{i=0}^{\infty} f_1(i)\mathbf{W}^i$ is the *generating function* corresponding to the sequence $(f_1(i))_{i=0}^{\infty}$. This is natural because the (discrete) Fourier transform of a (discrete) convolution returns the product of the (discrete) Fourier transforms of the respective functions. In the symmetric case $f_1 = f_2$, it follows that

$$\mathbf{K}_f(\mathbf{W}) = \pm (\mathbf{K}_\alpha(\mathbf{W}))^{\frac{1}{2}}. \quad (8)$$

If the RHS has a simple Taylor expansion (e.g. $\mathbf{K}_\alpha(\mathbf{W}) = \exp(\mathbf{W})$ so $\mathbf{K}_f(\mathbf{W}) = \exp(\frac{\mathbf{W}}{2})$), this enables us obtain f without recourse to the conditional sum in Eq. 5 or the iterative expression in Eq. 6. This is the case for many popular graph kernels; we provide some prominent examples in the table left. A notable exception is the inverse cosine kernel.

Name	$f(i)$
d -regularised Laplacian	$\frac{(d-2+2i)!!}{(2i)!!(d-2)!!}$
p -step random walk	$\left(\frac{p}{i}\right)$
Diffusion	$\frac{1}{2^i i!}$

As an interesting corollary, by considering the diffusion kernel we have also proved that $\sum_{p=0}^k \frac{1}{2^p p!} \frac{1}{2^{k-p}(k-p)!} = \frac{1}{k!}$.

2.1 NEURAL MODULATION FUNCTIONS, KERNEL LEARNING AND GENERALISATION

Instead of using a fixed modulation function f to estimate a fixed kernel, it is possible to parameterise it more flexibly. For example, we can define a *neural modulation function* $f^{(N)} : (\mathbb{N} \cup \{0\}) \rightarrow \mathbb{R}$ by a neural network (with a restricted domain) whose input and output dimensionalities are equal to 1. During training, we can choose the loss function to target particular desiderata of g-GRFs: for example, to suppress the mean square error of estimates of some particular fixed kernel (Sec. 3.4), or to learn a kernel which performs better in a downstream task (Sec. 3.5). Implicitly learning \mathbf{K}_α via $f^{(N)}$ is more scalable than learning \mathbf{K}_α directly because it obviates the need to repeatedly compute the exact kernel, which is typically of $\mathcal{O}(N^3)$ time complexity. Since any modulation function f maps to a unique α by Eq. 4, it is also always straightforward to recover the exact kernel which the g-GRFs estimate, e.g. once the training is complete.

Supposing we have (implicitly) learned α , how can the learned kernel K_α^G be expected to generalise? Let $\psi_{K_\alpha} : x \rightarrow \mathbb{H}_{K_\alpha}$ denote the feature mapping from the input space to the reproducing kernel Hilbert space \mathbb{H}_{K_α} induced by the kernel K_α^G . Define the hypothesis set

$$H = \{x \rightarrow \mathbf{w}^\top \psi_{K_\alpha}(x) : |\alpha_i| \leq \alpha_i^{(M)}, \|\mathbf{w}\|_2 \leq 1\}, \quad (9)$$

where we restricted our family of kernels so that the absolute value of each Taylor coefficient α_i is smaller than some maximum value $\alpha_i^{(M)}$. Following very similar arguments to Cortes et al. (2010), the following is true.

Theorem 2.3 (Empirical Rademacher complexity bound). *For a fixed sample $S = (x_i)_{i=1}^m$, the empirical Rademacher complexity $\hat{\mathcal{R}}(H)$ is bounded by*

$$\hat{\mathcal{R}}(H) \leq \sqrt{\frac{1}{m} \sum_{i=0}^{\infty} \alpha_i^{(M)} \rho(\mathbf{W})^i}, \quad (10)$$

where $\rho(\mathbf{W})$ is the spectral radius of the weighted adjacency matrix \mathbf{W} .

Naturally, the bound on $\hat{\mathcal{R}}(H)$ increases monotonically with $\rho(\mathbf{W})$. Following standard arguments in the literature, this immediately yields generalisation bounds for learning kernels $K_{\mathcal{G}}$. We discuss this in detail, including proving Theorem 2.3, in App. A.4.

3 EXPERIMENTS

Here we test the empirical performance of g-GRFs, both for approximating fixed kernels (Secs 3.1-3.3) and with learnable neural modulation functions (Secs 3.4-3.5).

3.1 UNBIASED POINTWISE ESTIMATION OF FIXED KERNELS

We begin by confirming that g-GRFs do indeed give unbiased estimates of the graph kernels listed in Table 1, taking regularisers $\sigma = 0.25$ and $\alpha = 20$ with $p_{\text{halt}} = 0.1$. We use symmetric modulation functions f , computed with the closed forms where available and using the iterative scheme in Eq. 6 if not. Fig. 2 plots the relative Frobenius norm error between the true kernels \mathbf{K} and their approximations with g-GRFs $\hat{\mathbf{K}}$ (that is, $\|\mathbf{K} - \hat{\mathbf{K}}\|_F / \|\mathbf{K}\|_F$) against the number of random walkers m . We consider 8 different graphs: a small random Erdős-Rényi graph, a larger Erdős-Rényi graph, a binary tree, a d -regular graph and 4 real world examples (karate, dolphins, football and eurosis) (Ivashkin, 2023). They vary substantially in size. For every graph and for all kernels, the quality of the estimate improves as m grows and becomes very small with even a modest number of walkers.

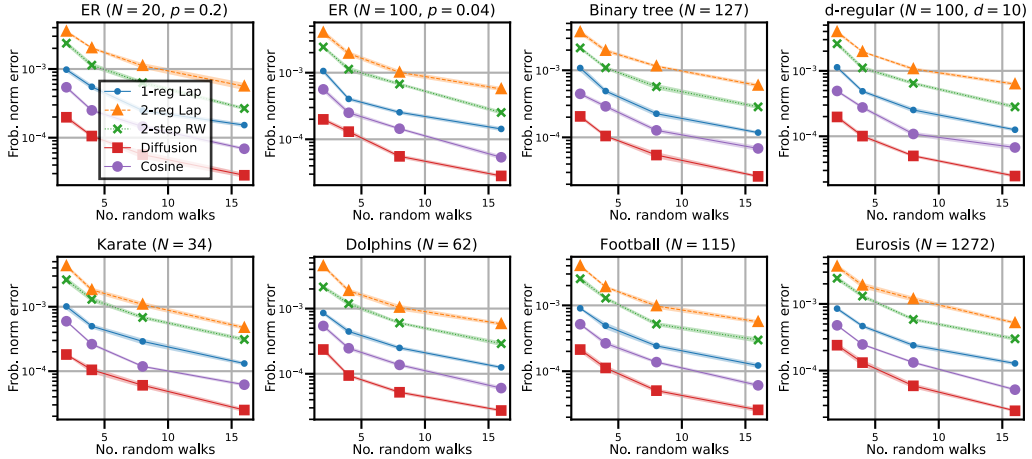


Figure 2: Unbiased estimation of popular kernels on different graphs using g-GRFs. The approximation error (y -axis) improves with the number of walkers (x -axis). We repeat 10 times; one standard deviation of the mean error is shaded.

3.2 SOLVING DIFFERENTIAL EQUATIONS ON GRAPHS

An intriguing application of g-GRFs for fixed kernels is efficiently computing approximate solutions of time-invariant non-homogeneous ordinary differential equations (ODEs) on graphs. Consider the following ODE defined on the nodes \mathcal{N} of the graph \mathcal{G} :

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{W}\mathbf{x}(t) + \mathbf{y}(t), \quad (11)$$

where $\mathbf{x}(t) \in \mathbb{R}^N$ is the state of the graph at time t , $\mathbf{W} \in \mathbb{R}^{N \times N}$ is a weighted adjacency matrix and $\mathbf{y}(t)$ is a (known) driving term. Assuming the null initial condition $\mathbf{x}(0) = \mathbf{0}$, Eq. 11 is solved by the convolution

$$\mathbf{x}(t) = \int_0^t \exp(\mathbf{W}(t - \tau)) \mathbf{y}(\tau) d\tau = \mathbb{E}_{\tau \in \mathcal{P}} \left[\frac{1}{p(\tau)} \exp(\mathbf{W}(t - \tau)) \mathbf{y}(\tau) \right] \quad (12)$$

where \mathcal{P} is a probability distribution on the interval $[0, t]$, equipped with a (preferably efficient) sampling mechanism and probability density function $p(\tau)$. Taking $n \in \mathbb{N}$ Monte Carlo samples $(\tau_i)_{i=1}^n \stackrel{\text{i.i.d.}}{\sim} \mathcal{P}$, we can construct the unbiased estimator:

$$\hat{\mathbf{x}}(t) := \frac{1}{n} \sum_{j=1}^n \frac{1}{p(\tau_j)} \exp(\mathbf{W}(t - \tau_j)) \mathbf{y}(\tau_j). \quad (13)$$

Note that $\exp(\mathbf{W}(t - \tau_j))$ is nothing other than the diffusion kernel, which is expensive to compute exactly for large N but can be efficiently approximated with g-GRFs. Take $\exp(\mathbf{W}(t - \tau_j)) \simeq \Phi_j \Phi_j^\top$ with $\Phi_j := (\phi(i))_{i=1}^N$ an $N \times N$ matrix whose rows are g-GRFs constructed to approximate the kernel at a particular τ_j . Then we have that

$$\hat{\mathbf{x}}(t) := \frac{1}{n} \sum_{j=1}^n \frac{1}{p(\tau_j)} \Phi_j \Phi_j^\top \mathbf{y}(\tau_j) \quad (14)$$

which can be computed in quadratic time (c.f. cubic if the heat kernel is computed exactly). Further speed gains are possible if dimensionality reduction techniques are applied to the g-GRFs (Choromanski 2023; Dasgupta et al. 2010).

As an example, we consider diffusion on three real-world graphs with a fixed source at one node, taking $\mathbf{W} = \mathbf{L}$ (the graph Laplacian) and $\mathbf{y}(t) = \mathbf{y} = (1, 0, 0, \dots)^\top$. The steady state is $\mathbf{x}(\infty) = \mathbf{W}^{-1}(-\mathbf{y})$. We simulate evolution under the ODE for $t = 1$ with $n = 10$ discretisation timesteps and \mathcal{P} uniform, approximating $\exp(\mathbf{W}(t - \tau_j))$ with different numbers of walkers m . As m grows, the quality of approximation improves and the (normalised) error on the final state $\|\hat{\mathbf{x}}(1) - \mathbf{x}(1)\|_2 / \|\mathbf{x}(1)\|_2$ drops for every graph. We take 100 repeats for statistics and plot the results in Fig. 3. One standard deviation of the mean error is shaded. $p_{\text{halt}} = 0.1$ and the regulariser is given by $\sigma = 1$.




Figure 3: ODE simulation error decreases as the number of walkers grows.

3.3 EFFICIENT KERNELISED GRAPH NODE CLUSTERING

As a further demonstration of the utility of our new mechanism, we show how estimates of the kernel $\mathbf{K} = \exp(\sigma^2 \mathbf{W})$ can be used to assign nodes to $k = 3$ clusters. Here, we choose \mathbf{W} to be the (unweighted) adjacency matrix and the regulariser is $\sigma^2 = 0.2$. We follow the algorithm proposed by Dhillon et al. (2004), comparing the clusters when we use exact and g-GRF-approximated kernels. For all graphs, $m \leq 80$. Table 2 reports the clustering error, defined by

$$E_c := \frac{\text{no. wrong pairs}}{N(N-1)/2}. \quad (15)$$

Table 2: Errors in kernelised k -means clustering when approximating the kernel $\exp(\sigma^2 \mathbf{W})$ with g-GRFs.

Graph	N	Clustering error, E_c
karate	34	0.08
dolphins	62	0.16
polbooks	105	0.12
football	115	0.02
databases	1046	0.10
eurosis	1272	0.09
cora	2485	0.01
citeseer	3300	0.04

This is the number of misclassified pairs of nodes (assigned to the same cluster when the converse is true or vice versa) divided by the total number of pairs. The error is small even with a modest number of walkers and on large graphs; kernel estimates efficiently constructed using g-GRFs can be readily deployed on downstream tasks where exact methods are slow.

3.4 LEARNING $f^{(N)}$ FOR BETTER KERNEL APPROXIMATION

Following the discussion in Sec. 2.1 we now replaced fixed f with a *neural modulation function* $f^{(N)}$ parameterised by a simple neural network with 1 hidden layer of size 1:

$$f^{(N)}(x) = \sigma_{\text{softplus}}(w_2 \sigma_{\text{ReLU}}(w_1 x + b_1) + b_2), \quad (16)$$

where $w_1, b_1, w_2, b_2 \in \mathbb{R}$ and σ_{ReLU} and σ_{softplus} are the ReLU and softplus activation functions, respectively. Bigger, more expressive architectures (including allowing $f^{(N)}$ to take negative values) can be used but this is found to be sufficient for our purposes.


We define our loss function to be the Frobenius norm error between a target Gram matrix and our g-GRF-approximated Gram matrix on the small Erdős-Rényi graph ($N = 20$) with $m = 16$ walks. For the target, we choose the 2-regularised Laplacian kernel. We train symmetric $f_1^{(N)} = f_2^{(N)}$ but provide a brief discussion of the asymmetric case (including its respective strengths and weaknesses) in App. A.5. On this graph, we minimise the loss with the Adam optimiser and a decaying learning rate (LR = 0.01, $\gamma = 0.975$, 1000 epochs). We make the following striking observation: $f^{(N)}$ does *not* generically converge to the unique unbiased (symmetric) modulation function implied by α , but instead to some different f that though biased gives a smaller mean squared error (MSE). This is possible because by downweighting long walks the learned $f^{(N)}$ gives estimators with a smaller variance, which is sufficient to suppress the MSE on the kernel matrix elements even though it no longer gives the target value in expectation. We then fix $f^{(N)}$ and use it for kernel approximation on the remaining graphs. The learned, biased $f^{(N)}$ still provides better kernel estimates, including for graphs with very different topologies and a much greater number of nodes: *eurosis* is bigger by a factor of over 60. See Table 3 for the results. $p_{\text{halt}} = 0.5$ and $\sigma = 0.8$.

Naturally, the learned $f^{(N)}$ is dependent on the the number of random walks m ; as m grows, the variance on the kernel approximation drops so it is intuitive that the learned $f^{(N)}$ will approach the unbiased f . Fig. 4 empirically confirms this is the case, showing the learned $f^{(N)}$ for different numbers of walkers. The line labelled ∞ is the unbiased modulation function, which for the 2-regularised Laplacian kernel is constant.

Table 3: Kernel approximation error with $m = 16$ walks and unbiased or learned modulation functions. Lower is better. Brackets give one standard deviation of the last digit.

Graph	N	Frob. norm error on \hat{K}	
		Unbiased	Learned
Small ER	20	0.0488(9)	0.0437(9)
Larger ER	100	0.0503(4)	0.0448(4)
Binary tree	127	0.0453(4)	0.0410(4)
d -regular	100	0.0490(2)	0.0434(2)
karate	34	0.0492(6)	0.0439(6)
dolphins	62	0.0505(5)	0.0449(5)
football	115	0.0520(2)	0.0459(2)
eurosis	1272	0.0551(2)	0.0484(2)

Figure 4: Learned modulation function with different numbers of random walkers m . It approaches the unbiased $f^{(N)}$ as $m \rightarrow \infty$



These learned modulation functions might guide the more principled construction of biased, low-MSE GRFs in the future. An analogue in Euclidean space is provided by structured orthogonal random features (SORFs), which replace a random Gaussian matrix with a **HD**-product to estimate the Gaussian kernel (Choromanski et al., 2017; Yu et al., 2016). This likewise improves kernel approximation quality despite breaking estimator unbiasedness.

3.5 IMPLICIT KERNEL LEARNING FOR NODE ATTRIBUTE PREDICTION

As suggested in Sec. 2.1, it is also possible to train the neural modulation function $f^{(N)}$ directly using performance on a downstream task, performing *implicit kernel learning*. We have argued that this is much more scalable than optimising \mathbf{K}_α directly.

In this spirit, we now address the problem of kernel regression on triangular mesh graphs, as previously considered by Reid et al. (2023). For graphs in this dataset (Dawson-Haggerty, 2023), every node is associated with a normal vector $\mathbf{v}^{(i)} \in \mathbb{R}^3$ equal to the mean of the normal vectors of its 3 surrounding faces. The task is to predict the directions of missing vectors (a random 5% split) from the remainder. Our (unnormalised) predictions are given by $\hat{\mathbf{v}}^{(i)} := \sum_j \hat{\mathbf{K}}^{(N)}(i, j) \mathbf{v}^{(j)}$, where $\hat{\mathbf{K}}^{(N)}$ is a kernel estimate constructed using g-GRFs with a neural modulation function $f^{(N)}$ (see Eq. 3). The angular prediction error is $1 - \cos \theta_i$ with θ_i the angle between the true $\mathbf{v}^{(i)}$ and approximate $\hat{\mathbf{v}}^{(i)}$ normals, averaged over the missing vectors. We train a symmetric pair $f^{(N)}$ using this angular prediction error on the small **cylinder** graph ($N = 210$) as the loss function. Then we freeze $f^{(N)}$ and compute the angular prediction error for other larger graphs. Fig. 5 shows the learned $f^{(N)}$ as well as some other modulation functions corresponding to popular fixed kernels. Note also that learning $f^{(N)}$ already includes (but is not limited to) optimising the lengthscale of a given kernel: taking $\mathbf{W} \rightarrow \beta \mathbf{W}$ is identical to $f(i) \rightarrow f(i)\beta^i \forall i$.

The prediction errors are highly correlated between the different modulation functions for a given random draw of walks; ensembles which ‘explore’ the graph poorly, terminating quickly or repeating edges, will give worse g-GRF estimators for every f . For this reason, we compute the prediction errors as the average normalised *difference* compared to the learned kernel result. Table 4 reports the results. Crucially, *this difference is found to be positive for every graph and every fixed kernel*, meaning the learned kernel always performs best.

Table 4: Normalised difference in angular prediction error compared to the learned kernel defined by $f^{(N)}$ (trained on the **cylinder** graph). All entries are positive since the learned kernel always performs best. We take 100 repeats (but only 10 for the very large **cycloidal** graph). The brackets give one standard deviation on the final digit.

Graph	N	Normalised $\Delta(\text{pred error})$ c.f. learned		
		1-reg Lap	2-reg Lap	Diffusion
cylinder	210	+0.40(2)	+0.85(4)	+0.029(3)
teapot	480	+0.81(5)	+1.78(8)	+0.059(3)
idler-riser	782	+0.52(2)	+1.12(1)	+0.042(2)
busted	1941	+0.81(2)	+1.60(4)	+0.063(2)
torus	4350	+2.13(5)	+5.3(1)	+0.067(2)
cycloidal	21384	+0.065(4)	+0.13(1)	+0.011(1)

It is remarkable that the learned $f^{(N)}$ gives the smallest error for all the graphs even though it was just trained on **cylinder**, the smallest one. We have implicitly learned a good kernel for this task which generalises well across topologies. It is also intriguing that the diffusion kernel performs only slightly worse. This is to be expected because their modulation functions are similar (see Fig. 5) so they encode very similar kernels, but this will not always be the case depending on the task at hand.

4 CONCLUSION

We have introduced ‘general graph random features’ (g-GRFs), a novel random walk-based algorithm for time-efficient estimation of arbitrary functions of a weighted adjacency matrix. The mechanism is conceptually simple and trivially distributed across machines, unlocking kernel-based machine learning on very large graphs. By parameterising one component of the random features with a simple neural network, we can further suppress the mean square error of estimators and perform scalable implicit kernel learning.




Figure 5: Fixed and learned modulation functions for kernel regression

5 ETHICS AND REPRODUCIBILITY

Ethics: Our work is foundational. There are no direct ethical concerns that we can see, though of course increases in scalability afforded by graph random features might amplify risks of graph-based machine learning, from bad actors or as unintended consequences.

Reproducibility: To foster reproducibility, we clearly state the central algorithm in Alg. 1. Source code is available at https://github.com/isaac-reid/general_graph_random_features. All theoretical results are accompanied by proofs in Appendices A.1, A.4 where any assumptions are made clear. The datasets we use correspond to standard graphs and are freely available online. We link suitable repositories in every instance. Except where prohibitively computationally expensive, results are reported with uncertainties to help comparison.

REFERENCES

- Zhaojun Bai, James Demmel, Jack Dongarra, Axel Ruhe, and Henk van der Vorst. *Templates for the solution of algebraic eigenvalue problems: a practical guide*. SIAM, 2000. URL <https://www.cs.ucdavis.edu/~bai/ET/contents.html>.
- Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002. URL <https://dl.acm.org/doi/pdf/10.5555/944919.944944>.
- Karsten M Borgwardt, Cheng Soon Ong, Stefan Schöner, SVN Vishwanathan, Alex J Smola, and Hans-Peter Kriegel. Protein function prediction via graph kernels. *Bioinformatics*, 21(suppl_1):i47–i56, 2005. URL <https://doi.org/10.1093/bioinformatics/bti1007>.
- Colin Campbell. Kernel methods: a survey of current techniques. *Neurocomputing*, 48(1-4): 63–84, 2002. doi: 10.1016/S0925-2312(01)00643-9. URL [https://doi.org/10.1016/S0925-2312\(01\)00643-9](https://doi.org/10.1016/S0925-2312(01)00643-9).
- Stéphane Canu and Alexander J. Smola. Kernel methods and the exponential family. *Neurocomputing*, 69(7-9):714–720, 2006. doi: 10.1016/j.neucom.2005.12.009. URL <https://doi.org/10.1016/j.neucom.2005.12.009>.
- Olivier Chapelle, Jason Weston, and Bernhard Schölkopf. Cluster kernels for semi-supervised learning. *Advances in neural information processing systems*, 15, 2002. URL <https://dl.acm.org/doi/10.5555/2968618.2968693>.
- Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*, 2020. URL <https://doi.org/10.48550/arXiv.2009.14794>.
- Krzysztof M Choromanski, Mark Rowland, and Adrian Weller. The unreasonable effectiveness of structured random orthogonal embeddings. *Advances in neural information processing systems*, 30, 2017. URL <https://doi.org/10.48550/arXiv.1703.00864>.
- Krzysztof Marcin Choromanski. Taming graph kernels with random features. In *International Conference on Machine Learning*, pages 5964–5977. PMLR, 2023. URL <https://doi.org/10.48550/arXiv.2305.00156>.
- Fan R. K. Chung and Shing-Tung Yau. Coverings, heat kernels and spanning trees. *Electron. J. Comb.*, 6, 1999. doi: 10.37236/1444. URL <https://doi.org/10.37236/1444>.
- Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Generalization bounds for learning kernels. In *Proceedings of the 27th Annual International Conference on Machine Learning (ICML 2010)*, 2010. URL <http://www.cs.nyu.edu/~mohri/pub/lk.pdf>.
- Keenan Crane, Clarisse Weischedel, and Max Wardetzky. The heat method for distance computation. *Communications of the ACM*, 60(11):90–99, 2017. URL <https://dl.acm.org/doi/10.1145/3131280>.

- Anirban Dasgupta, Ravi Kumar, and Tamás Sarlós. A sparse johnson: Lindenstrauss transform. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 341–350. ACM, 2010. doi: 10.1145/1806689.1806737. URL <https://doi.org/10.1145/1806689.1806737>.
- Michael Dawson-Haggerty. Trimesh repository, 2023. URL <https://github.com/mikedh/trimesh>.
- Inderjit S Dhillon, Yuqiang Guan, and Brian Kulis. Kernel k-means: spectral clustering and normalized cuts. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 551–556, 2004. URL <https://dl.acm.org/doi/10.1145/1014052.1014118>.
- Michel X. Goemans and David P. Williamson. Approximation algorithms for $\max\text{-}3\text{-}c_{\text{ut}}$ and other problems via complex semidefinite programming. *J. Comput. Syst. Sci.*, 68(2):442–470, 2004. doi: 10.1016/j.jcss.2003.07.012. URL <https://doi.org/10.1016/j.jcss.2003.07.012>.
- Vladimir Ivashkin. Community graphs repository, 2023. URL <https://github.com/vlivashkin/community-graphs>.
- William B Johnson. Extensions of lipschitz mappings into a hilbert space. *Contemp. Math.*, 26:189–206, 1984. URL <http://stanford.edu/class/cs114/readings/JL-Johnson.pdf>.
- Kyle Kloster and David F Gleich. Heat kernel based community detection. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1386–1395, 2014. URL <https://doi.org/10.48550/arXiv.1403.3148>.
- Vladimir Koltchinskii and Dmitry Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *The Annals of Statistics*, 30(1):1–50, 2002. URL <https://doi.org/10.48550/arXiv.math/0405343>.
- Risi Kondor and John D. Lafferty. Diffusion kernels on graphs and other discrete input spaces. In Claude Sammut and Achim G. Hoffmann, editors, *Machine Learning, Proceedings of the Nineteenth International Conference (ICML 2002), University of New South Wales, Sydney, Australia, July 8-12, 2002*, pages 315–322. Morgan Kaufmann, 2002. URL <https://www.ml.cmu.edu/research/dap-papers/kondor-diffusion-kernels.pdf>.
- Leonid Kontorovich, Corinna Cortes, and Mehryar Mohri. Kernel methods for learning languages. *Theor. Comput. Sci.*, 405(3):223–236, 2008. doi: 10.1016/j.tcs.2008.06.037. URL <https://doi.org/10.1016/j.tcs.2008.06.037>.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20, 2007. URL <https://dl.acm.org/doi/10.5555/2981562.2981710>.
- Isaac Reid, Krzysztof Choromanski, and Adrian Weller. Quasi-monte carlo graph random features. *arXiv preprint arXiv:2305.12470*, 2023. URL <https://doi.org/10.48550/arXiv.2305.12470>.
- Alexander J Smola and Risi Kondor. Kernels and regularization on graphs. In *Learning Theory and Kernel Machines: 16th Annual Conference on Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003, Washington, DC, USA, August 24-27, 2003. Proceedings*, pages 144–158. Springer, 2003. URL <https://people.cs.uchicago.edu/~risi/papers/SmolaKondor.pdf>.
- Alexander J. Smola and Bernhard Schölkopf. Bayesian kernel methods. In Shahrar Mendelson and Alexander J. Smola, editors, *Advanced Lectures on Machine Learning, Machine Learning Summer School 2002, Canberra, Australia, February 11-22, 2002, Revised Lectures*, volume 2600 of *Lecture Notes in Computer Science*, pages 65–117. Springer, 2002. URL https://doi.org/10.1007/3-540-36434-X_3.

- Yasutoshi Yajima. One-class support vector machines for recommendation tasks. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 230–239. Springer, 2006. URL https://dl.acm.org/doi/10.1007/11731139_28.
- Felix Xinnan X Yu, Ananda Theertha Suresh, Krzysztof M Choromanski, Daniel N Holtmann-Rice, and Sanjiv Kumar. Orthogonal random features. *Advances in neural information processing systems*, 29, 2016. URL <https://doi.org/10.48550/arXiv.1610.09072>.
- Yufan Zhou, Changyou Chen, and Jinhui Xu. Learning manifold implicitly via explicit heat-kernel learning. *Advances in Neural Information Processing Systems*, 33:477–487, 2020. URL <https://doi.org/10.48550/arXiv.2010.01761>.